

DOCUMENT RESUME

ED 096 576

CG 009 204

AUTHOR Scandura, Joseph M., Ed.; And Others
TITLE 1974 Proceedings: Fifth Annual Interdisciplinary
Conference on Structural Learning.
INSTITUTION Pennsylvania Univ., Philadelphia. Mathematics
Education Research Group.
SPONS AGENCY Naval Personnel Research and Development Lab.,
Washington, D.C.
PUB DATE 74
NOTE 359p.
AVAILABLE FROM J. M. Scandura, MERGE Research Institute, 1249
Greentree Lane, Narberth, Pennsylvania 19072
(HC-\$5.00)

EDRS PRICE MF-\$0.75 HC-\$17.40 PLUS POSTAGE
DESCRIPTORS *Behavioral Science Research; *Cognitive Processes;
Concept Formation; Conference Reports;
*Interdisciplinary Approach; *Learning Processes;
*Learning Theories; Memory; Reading Processes

ABSTRACT

These proceedings emphasize multidisciplinary theory and research having implications for behavioral science and/or education. The contributions deal with one or more of the following questions: (1) How can one characterize the knowledge structures which underlie behavior associated with subject matter disciplines such as mathematics and language? How can one evaluate alternative characterization? (2) How can one determine the knowledge given individuals have at their command? (3) What are the mechanisms of performance, learning, memory, and perception as they involve structured knowledge? (4) What are the instructional conditions which control the acquisitions of structures? The proceedings contain 34 complete papers and summaries, including contributions by Banerji, Brainerd, Dienes, Fischer, Kopstein, Landauer, Merrill, Newell, Pask, Scandura, Shaw, Wilson, and Witz, and (edited) transcriptions of two discussion sessions, including Beilen, Fischer, and Minsky on "How Does Learning Take Place?" and "Theoretical Issues in Structural Learning." (Author)

ED 096 576

BEST COPY AVAILABLE

CG

1974 PROCEEDINGS: FIFTH ANNUAL
INTERDISCIPLINARY CONFERENCE ON STRUCTURAL LEARNING

Edited by
Joseph M. Scandura
John H. Durnin
Wallace H. Wulfeck II

U S DEPARTMENT OF HEALTH,
EDUCATION & WELFARE
NATIONAL INSTITUTE OF
EDUCATION

THIS DOCUMENT HAS BEEN REPRO-
DUCED EXACTLY AS RECEIVED FROM
THE PERSON OR ORGANIZATION ORIGIN-
ATING IT. POINTS OF VIEW OR OPINIONS
STATED DO NOT NECESSARILY REPRESENT
OFFICIAL NATIONAL INSTITUTE OF
EDUCATION POSITION OR POLICY

Multidiscipline &
Education
Research
Group
Enterprises

RESEARCH INSTITUTE

Reproduction in whole or part is permitted for any purpose of the United States Government.

Reproduction of the Proceedings was sponsored by the Personnel and Training Research Programs, Psychological Sciences Division, Office of Naval Research under Contract No. N00014-67-A-0399-0006, Contract Authority Identification Number NR No. NR 151-356X.

Approved for public release; distribution unlimited.

While the limited supply lasts, good quality bound copies of the 350 page proceedings may be obtained from J. M. Scandura, MERGE Research Institute, 1249 Greentree Lane, Narberth, Pennsylvania 19072. All orders must be prepaid at cost, \$5.00 per copy. Other copies may be obtained from the Defense Documentation Center (for Qualified Defense Requestors) or purchased from the National Technical Information Service, Springfield, Virginia.

CG 009 204
602 600

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER 1	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) 1974 PROCEEDINGS: FIFTH ANNUAL INTERDISCIPLINARY CONFERENCE ON STRUCTURAL LEARNING		5. TYPE OF REPORT & PERIOD COVERED Technical
7. AUTHOR(s) Editors Joseph M. Scandura John H. Durnin Wallace H. Wulfeck II		6. PERFORMING ORG. REPORT NUMBER B5
9. PERFORMING ORGANIZATION NAME AND ADDRESS MERGE Research Institute 1249 Green Tree Lane Narberth, Pa. 19072		8. CONTRACT OR GRANT NUMBER(s) N00014-67-A-0399-0006
11. CONTROLLING OFFICE NAME AND ADDRESS Personnel and Training Research Programs Office of Naval Research (Code 458) Arlington, Virginia 22217		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS 61153N RR042-04-02 NR151-356X
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office)		12. REPORT DATE June 30, 1974
		13. NUMBER OF PAGES
		15. SECURITY CLASS. (of this report) Unclassified
		15a. DECLASSIFICATION/DOWNGRADING SCHEDULE
16. DISTRIBUTION STATEMENT (of this Report) Approved for public release; distribution unlimited.		
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)		
18. SUPPLEMENTARY NOTES Prepared in cooperation with MERGE Research Institute.		
19. KEY WORDS (Continue on reverse side if necessary and identify by block number) Structural Learning, Cognition, Educational Design, Individual Differences, Content analysis, heterarchical, algorithmic, task analysis, human memory, information processing, cognitive development, language development, linguistics, mathematics education, perception, training research, artificial intelligence, computer simulation, structuralism, performance testing.		
20. ABSTRACT (Continue on reverse side if necessary and identify by block number) The proceedings cover three days of talks and discussions at the 1974 Fifth Annual Interdisciplinary Conference on Structural Learning. Although the contributions ranged widely, the emphasis throughout was on multi- disciplinary theory and research, having implications for behavioral science and/or education. The contributions dealt with one or more of the following questions: (1) How can one characterize the knowledge structures which underlie behavior associated with subject matter disciplines such as		

19. diagnosis, imagery, prose analysis, semantics, processing capacity, mathematics learning, human problem solving.
20. mathematics and language? How can one evaluate alternative characterizations? (2) How can one determine the knowledge given individuals have at their command? (3) What are the mechanisms of performance, learning, memory, and perception as they involve structured knowledge? (4) What are the instructional conditions which control the acquisition of structures?

The proceedings contain 34 complete papers and summaries, including contributions by Banerji, Brainerd, Dienes, Fischer, Kopstein, Landauer, Merrill, Newell, Pask, Scandura, Shaw, Wilson, and Witz, and (edited) transcriptions of two discussion sessions, including Beilen, Fischer, and Minsky on "How Does Learning Take Place?" and "Theoretical Issues in Structural Learning."

April 20 - 23, 1974
Stiteler Hall
208 S. 37th Street
University of Pennsylvania

Conference Director: Joseph M. Scandura
Associate Director--Arrangements: Wallace H. Wulfeck II
Proceedings: John H. Durnin

Saturday, April 20, 1974

Page

- 8:15-8:50 am REGISTRATION (Coffee & Donuts) - Stiteler Hall.
- 8:50-9:00 Introduction to the Conference Joseph M. Scandura, Conference Director Rm. B-26, Stiteler Hall.
- 9:00-10:55 Session 1. Rm. B-26.
- Discussant and Chairman: John W. Carr III, Moore School of Electrical Engineering, University of Pennsylvania
- 1 -- Perceiving the Pace of Change in Changing Faces: Toward a Theory of Shape Perception (35 min.)
Robert Shaw & John Pittenger, Center for Research in Human Learning, 205 Elliot Hall, Minneapolis, Minn. 55455.
- 10 -- A Formal Approach to Knowledge in Problem Solving (35 min.)
Ranan Banerji, Department of Information Sciences, Temple University, Philadelphia, Pa. 19122.
- 13 -- The Long-Term Memory Improvement Effect: Memory Improvement or Concept Development? (35 min.)
C.J. Brainerd, Department of Psychology, University of Alberta, Edmonton, Alberta, Canada.
- 10:55-11:10 Break (Coffee in Faculty Lounge, Stiteler Hall)
- 11:10-11:40 Session 2. Presentation Precie, Rm. B-26.
(Chairman: John H. Durnin, Villanova University)
- 23 -- Art Education - Means or End? (3 min.)
Boris Hlai Jr., Director of Research, Harcum College, Bryn Mawr, Pa. 19010
- 25 -- Cognitive Development as Problem Solving (3 min.)
Kurt Fischer, Department of Psychology, University of Denver, Denver, Colorado 80202.
- 32 -- Structuralism in Piaget's Intuitive Stage of Development (3 min.)
Joan Batt, Department of Psychology, New School for Social Research, 66 W. 12th St., New York, N.Y. 10011
- 37 -- A Study of Diagnostic-Test-Based Prescription in IPI Mathematics (3 min.)
Ann Graeber, Research for Better Schools, 1700 Market St., Suite 1700, Philadelphia, Pa. 19103.
- 40 -- Individualized Middle Mathematics: Extending and Improving the IPI Model (3 min.)
George Loverre, Research for Better Schools, 1700 Market St., Suite 1700, Philadelphia, Pa. 19104.
- 45 -- A Multidisciplinary Approach to Acceptability and Chomsky's Theory of Grammatical Hierarchy (3 min.)
Wayne Hreko & D. Kim Reid, Department of Special Education, Temple University, Philadelphia, Pa. 19122.
- 56 -- The Effect of Imaging and Text Structure on What is Learned from Reading a Passage (3 min.)
Sandra G. Koser, Cornell University, Stone Hall, Ithaca, New York 14850
- 11:45-12:30 Sessions 2A, B, C.
- | | Rm. B-26 | Rm. C-15 | Rm. C-18 |
|------------|---|---|---|
| Chairman: | Harry Gray, Moore School of Electrical Engineering, University of Pennsylvania. | Claude Mayberry, Associate Dean, University of Pennsylvania. | John H. Durnin, Villanova University |
| Papers by: | Kurt Fischer (20 min.)
Joan Batt (20 min.) | Boris Hlai Jr. (15 min.)
Ann Graeber (15 min.)
George Loverre (15 min.) | Hreko and Reid (20 min.)
Sandra G. Koser (20 min.) |
- 12:30-2:15 Lunch.
- 2:15-3:45 Session 3. Rm. B-26.
- Discussant and Chairman: Ranan Banerji, Department of Information Sciences, Temple University
- 63 -- Human Problem Solving: A Synthesis of Competence (A.I.), Psychological, and Measurement Approaches (40 min.)
Joseph M. Scandura, University of Pennsylvania, Philadelphia, Pa. 19174.
- 94 -- Some Aspects of Production Systems
Allen Newell, Carnegie-Mellon University, Pittsburgh, Pa. 15213.
- 95 -- Modeling the Meaning of "Au" (Out) in Laura (35 min.)
Klaus Witz, Department of Mathematics, University of Illinois, Urbana, Ill. 61801.
- 4:00-4:25 Session 4. Presentation Precie, Rm. B-26.
(Chairman: Wallace H. Wulfeck II, University of Pennsylvania.)
- 96 -- Representation, Abstraction, and Recognition of Structured Events (3 min.)
Frederick Hayes-Roth, Department of Psychology, University of Michigan, Ann Arbor, Mich. 48104.
- 112 -- Graphical Representation of Abstract Concepts in Textual Materials (3 min.)
John M. Yancey, School of Dental Medicine, University of Pennsylvania, Philadelphia, Pa. 19174.
- 114 -- The Effect of Identification of Semantic Relationships on Processing Load during Sentence Comprehension (3 min.)
Richard T. Isakson, Cornell University, Stone Hall, Ithaca, New York 14850.
- 121 -- Determination of Memory Load in Information Processing (3 min.)
Donald T. Voorhies, Structural Learning Program, University of Pennsylvania, Philadelphia, Pa. 19174.
- 134 -- An Approach to the Construction of the Contents, Objectives, and Sequences of Teaching within a Formalized Descriptive System (3 min.)
F. Schott, Seminar für Pädagogik, Technische Universität Braunschweig, 33 Braunschweig, T.G. Postfach 3339.
- 135 -- Demonstration of Teacher Planning as a Function of Management (3 min.)
Yehosh S. Dworkin, Abacus Associates, Inc., 443 Park Ave. S., New York, N.Y. 10016, and Nancy Dworkin, Hofstra University.
- 4:30-5:10 Sessions 4A, B, C.
- | | Rm. B-26 | Rm. C-15 | Rm. C-18 |
|------------|--|--|--|
| Chairman: | Ruzena Bajecy, Moore School of Electrical Engineering, University of Pennsylvania. | Wallace H. Wulfeck II, University of Pennsylvania | James Diamond, University of Pennsylvania |
| Papers by: | Hayes-Roth (20 min.)
John Yancey (20 min.) | Richard Isakson (20 min.)
Donald Voorhies (20 min.) | Dworkin & Dworkin (20 min.)
F. Schott (20 min.) |

Sunday, April 21, 1974

- 8:30-9:00 am Registration (Coffee & Donuts) - Stiteler Hall.
 9:00-10:30 Session 5. Rm. B-26.
 Discussant and Chairman: Zoltan Domotor, Department of Philosophy, University of Pennsylvania.
- 140 -- Heterarchies as an Alternative to Hierarchies for Sequencing Content in Instruction (40 min.)
 M. David Merrill, Division of Instructional Research, Development, and Evaluation, Brigham Young Univ. Provo, Utah.
- 148 -- A New Type of Test of Cognitive Functioning: An Intelligence Test(?) Based on Finite Affine Geometries (40 min.)
 Zoltan P. Dienes, Centre de Recherches en Psycho-Mathematique, University of Sherbrooke, Sherbrooke, Quebec, Canada.
- 10:30-10:45 Break (Coffee in Faculty Lounges, Stiteler Hall)
 10:45-11:10 Session 6. Presentation Precs, Rm. B-26.
 Chairman: George Lowarre, Research for Better Schools.
- 151 -- Criteria for Cognitive Structures: Reflections on Piaget's Theory of Genetic Epistemology and Introspection (3 min.)
 John Limber & Michael Chiavri, Department of Psychology, University of New Hampshire, Durham, N.H. 03824.
- 160 -- A Preliminary Proposal About a Piagetian Stage that Might Follow the Formal Operations Stage (3 min.)
 John M. Kennedy, Scarborough College, University of Toronto, West Hill, Ontario, Canada.
- 181 -- Higher Order Rule Characterization of Heuristics for Compass and Straightedge Constructions in Geometry (3 min.)
 John H. Durnin, Department of Mathematics Education, Villanova University, Villanova, Pa. 19085 (with Joseph M. Scandura and Wallace H. Wulfeck II).
- ?09 -- Diagnosis and Instruction of Higher Order Rules for Solving Geometry Construction Problems (3 min.)
 Wallace H. Wulfeck II, Structural Learning Program, University of Pennsylvania, Philadelphia, Pa. 19174
 (with J. M. Scandura, J. H. Durnin, & W. Ehrenpreis).
- Algorithms: Applications in Instructional Research and Development (3 min.)
 Vernon Gerlack, Director of Instructional Research, University of Minnesota, Minneapolis, Minn. 55455, and
 Fritz Brecke, Arizona State University, Tempe, Arizona 85281
- 222 -- Individualized Instruction and Multidiscipline Structural Learning (3 min.)
 Robert I. Starr, University of Missouri, St. Louis, 8001 Natural Bridge Road, St. Louis, Missouri 63221

11:15-11:55 Session 6A, B, C.

	Rm. B-26	Rm. C-15	Rm. C-18
Chairman:	Harry Beilen, Department of Psychology, Graduate Center, City University of New York.	George Lowarre, Research for Better Schools, Philadelphia, Pennsylvania.	Petar Jessel, Moore School of Electrical Engineering, University of Pennsylvania.
Papers by:	Limber and Chiavri (20 min.) John M. Kennedy (20 min.)	John Durnin (20 min.) Wallace Wulfeck II (20 min.)	Gerlack and Brecke (20 min.) Robert Starr (20 min.)

11:55-1:30 Lunch.

1:30-3:00 Session 7. Rm. B-26.

Discussant and Chairman: Marvin Minsky, Artificial Intelligence Laboratory, M.I.T., Cambridge, Mass.

- 227 -- How Much Structure is There in Memory (40 min.)
 T. K. Landauer, Bell Telephone Laboratories, Murray Hill, N.J. 07974.
- 231 -- Practical Methods for Building Extendable Knowledge Structures Based on the Idea that Concepts are Procedures for Reconstructing Relations and that Memories are Procedures for Reconstructing Concepts: Empirical Data and Theoretical Discussion (40 min.)
 Gordon Paik, System Research Ltd., Richmond, Surrey, England.
- 3:00-3:15 Break.
 3:15-3:40 Session 8. Presentation Precs, Rm. B-26.
 Chairman: Fred Davis, Center for Research in Evaluation and Measurement, University of Pennsylvania.
- 252 -- The Study of Syntax and Cognition: An Idea whose Time has Come (3 min.)
 Barbara A. Hutson, Department of Educational Psychology, State University of New York, Albany N.Y. 12222
- 264 -- A Systems and Set Theory Model for the Empirical Investigation of Perception and Structured Knowledge (3 min.)
 M. Sue Hogner, Department of Psychology, Catholic University of America, Washington, D.C. 20017.
- 274 -- An Interdisciplinary Theory of Adapting Behavior (3 min.)
 Doreen Steg and Rowalind Schulman, Drexel University, Philadelphia, Pa. 19104.
- 290 -- Application of Behavior Modification to Learning Center Choices in a Kindergarten Open Education Classroom (3 min.)
 Thomas D. Yaukey, Director, Wisconsin Early Childhood Study Center, University of Wisconsin, Madison, Wisc. 53706
- 293 -- The Organization and Acquisition of Formally Defined Concepts (3 min.)
 Jed Lewis, Center for Research in Human Learning, 205 Elliot Hall, Minneapolis, Minnesota 55455
- 310 -- A Test-Theoretical Approach to Structural Learning (3 min.)
 Wilhelm F. Kempf, Institut für die Pädagogik der Naturwissenschaften, Universität Kiel, Kiel, West Germany.

3:45-4:20 Sessions 9A, B, C.

	Rm. B-26	Rm. C-15	Rm. C-18
Chairman:	Donald Voorhies, University of Pennsylvania.	Charles Dwyer, Philosophy of Education, University of Pennsylvania.	Fred Davis, Center for Research in Evaluation and Measurement, University of Pennsylvania.
Papers by:	Barbara Hutson (20 min.) M. Sue Hogner (20 min.)	Steg and Schulman (20 min.) Thomas D. Yaukey (20 min.)	Jed Lewis (20 min.) Wilhelm Kempf (20 min.)

4:20-4:30 Short break.

* Unable to submit paper due to sponsor regulations.

Page Sunday, April 21, 1974

323

4:30-5:30 Roundtable Discussion. Rm. B-26.

-- How Does Learning Take Place?

Chairman: Joseph M. Scandura

Discussants: Ranan Benerji, Temple University
Zoltan P. Dienes, Centre de Recherches
en Psycho-Mathematique
Herry Beilen, C.U.N.Y.
Kurt Fischer, University of Denver
T.K. Landauer, Bell Telephone Laboratories

M. David Merrill, Brigham Young University
Marvin Minsky, M.I.T.
Gordon Pask, System Research Ltd.
Kalloo Wilson, University of Alberta
Felix Kopstein, Alexandria, Va.

5:30-6:30 WINE AND CHEESE - Faculty Lounge, Stiteler Hall.

Monday, April 22, 1974

9:00-9:30 Coffee and Donuts for Registrants. Faculty Lounge, Stiteler Hall.

9:30-12:00 Postsession I. Rm. C-24, Stiteler Hall.

335

-- Theoretical Issues in Structural Learning

Panel: R. Benerji, Z.P. Dienes, Z. Domotor, H.J. Gray, T.K. Landauer, M.D. Merrill, F. Kopstein,
M. Minsky, G. Pask, J.M. Scandura, K. Wilson, and other participants.

Panelists will have an opportunity to raise theoretical issues of concern followed by open-ended discussion and elaboration. Participants will have ample opportunity to comment, raise issues, and otherwise challenge the speakers.

12:00-2:00 Lunch.

2:00-4:30 Postsession II. Rm. C-24, Stiteler Hall.

Workshops - Theoretical Issues

Organization (10-15 min.)

Workshops: One or more intensive discussion workshops will be held depending on the participants' interests. The workshops may center about the work of one individual or group, or two or more groups in coordination. Possible topics and discussion leaders include: Relationships between production systems, hierarchies, entailment structures, and learning mechanisms (Minsky, Newell, Pask), Perception and cognition in the epistemic subject and in the individual (Beilen, Dienes, Limber, Kennedy, Witz), Criterion referenced and normative theories of testing creativity (Bracke, Davis, Dienes, Durnin, Gerlack, Kamf, Schott), Memory and thinking (Iskeon, Landauer, Pask, Voorhies), and Theory and research in Structural Learning (Scandura, Durnin, Voorhies, Wulfack).

Tuesday, April 23, 1974

9:00-9:30 Coffee and Donuts for Registrants. Faculty Lounge, Stiteler Hall.

9:30-12:00 Postsession III. Faculty Lounge, Stiteler Hall.

Educational Design in the U.S.A., England, Germany, and the World

Panel: Z. P. Dienes, W. Kamf, J.M. Scandura

Panelists will have an opportunity to give brief overviews of educational design in their respective countries (of origin or interest), followed by open-ended discussion.

12:00-2:00 Lunch

2:00-4:30 Postsession IV. Faculty Lounge.

Collaborative Planning Sessions and Workshops -- Educational Design

Organization (10-15 min.)

Workshops: Group workshops in educational design (as in Postsession II), and planning sessions to foster multidisciplinary collaboration.

PERCEIVING AGE IN FACIAL PROFILES

John B. Pittenger

Robert E. Shaw

University of Minnesota

Although people and animals are readily recognized by their facial characteristics, no theory of perception explains how this is accomplished. The studies to be reported provide support for a new theory explaining this phenomenon. Theories of object perception usually attempt to explain the perception of objects and patterns which do not change their shape over time. However, a truly adequate theory must also explain the perception of events where object configurations or the shape of objects undergo dynamic change. Shaw, McIntyre, and Mace, following Gibson,¹ argue that a successful theory of event perception must solve two problems. First, it must specify the change that constitutes the dynamic aspect of the event. Second, it must specify that which provides unity or identity to the structures involved in the event. The former requirement must be satisfied if we are to explain the classification of events (e.g., as running, rolling, growing, smiling, etc.), while the latter requirement must be satisfied if we are to explain the identification of the subject of the event (e.g. John runs, the ball rolls, the flower grows, Mary smiles, etc.). Thus, the perception of any event has two components: the detection of invariant information specifying the nature of the change involved, the transformational invariant, and the detection of invariant information specifying the structure which undergoes the change, the structural invariant. In the research to be reported we have attempted to treat the perception of faces, since they change over time due to growth, as a problem for a theory of event perception.

The human face is radically changed by growth; while the neo-natal head exhibits an exaggerated cranium and diminutive face, the face grows more rapidly than the cranium thereby altering not only the size of the head but its proportions as well. Consequently, Enlow² describes growth as a "remodeling" transformation which not only changes the size of the object (a similarity transformation) but its shape as well (a non-rigid transformation). Growth of the human face therefore can be considered to be a "viscal-elastic" event defined over the cranio-facial complex. We hypothesize that the perception of aging of faces depends upon the ability to detect the style of change in the shape of the face and head due to growth (the transformational invariant). Furthermore, we suppose that perception of sex, race, and individual identity depends upon recognition of those proportions of the face and head which are not changed by growth (the structural invariants). While feature and template theories treat only the perception of non-change, our approach treats perception of both change and non-change in a unified way.

¹R. Shaw, M. McIntyre and W. Mace in Studies in Perception: Essays in Honor of J. J. Gibson, R. MacLeod and H. Pick Jr., Ed. (Cornell University Press, Ithaca, New York, in press). J. Gibson, Vision Res., 1, 253-262 (1961).

²D. Enlow, The Human Face, An Account of the Postnatal Growth and Development of the Craniofacial Skeleton (Hoeber Medical Division, Harper & Row, New York, 1968), pp. 57-82.

The primary goal of the present studies was to test our hypothesis regarding the perceptual effects of the elastic component of the remodeling transformation. To do this we first had to discover a mathematical formulation of the elastic component of the growth process which, when applied to a basic facial profile, generates a sequence of faces reliably perceived to be ordered by age differences. Secondly, we had to show that while a given face was indeed aged by this transformation, it did not lose its individual identity.

The biological literature suggests three classes of transformations for the specification of the transformational invariants of skull growth: strain, shear, and radial growth. Arguments by D'Arcy Thompson and Cerasimov³ suggest that the strain imposed upon bony tissue by stresses caused by the growth of soft integuments provides the primary shape of the head. Consequently, we assume that the general shape of the head constitutes the primary perceptual information for age. The class of mathematical transformations that seems to characterize best the result of stress forces is also termed a strain, which stretches or compresses a structure. In addition to the strain component in growth, Albrecht Durer⁴ suggested that an identifying characteristic of faces is the degree to which the main angle of the profile is oblique with respect to the perpendicular. Indeed, as a person grows from infancy to adulthood the profile of face changes its angle obliquely toward being more prognathic. The mathematical transformation which best characterizes this change is called a shear. (Where a one-dimensional strain transforms a square into a rectangle, a shear transforms a square into a rhombus.)

The stimuli for our first age estimate experiment were produced by applying combinations of these transformations globally to a two-dimensional Cartesian space in which the profile of a ten year old boy had been placed so that the origin was at the ear-hole and the y-axis was perpendicular to the Frankfurt horizontal. The global transformations of the two dimensional space of the profile follows the technique of coordinate transformation suggested by D'Arcy Thompson. All calculations were performed by computer and the profiles drawn by a computer driven plotter.⁵

³D. A. W. Thompson, On Growth and Form (Cambridge University Press, Cambridge, 2nd ed., 1942) pp. 1027-1095. M. M. Gerasimov, Osnovy Vosstanovlenia Litsa Po Cherepy, (Basics of Restoration of the Face According to the Skull), (State Publishing House-Soviet Science, Moscow, 1949.)

⁴Dürer's work cited by Thompson, *ibid.*

⁵The formula used for the shear transformation in producing the stimuli expressed in rectangular coordinates was $y' = y$, $x' = x + \tan \theta y$ where $\tan \theta$ is the angle of shear and x', y' are new coordinates. The formula for the strain transformation used expressed for convenience in polar coordinates was $\theta = \theta$, $r' = r(1 - k \sin \theta)$ where r is the radial vector and θ is the angle specifying direction from the origin. Here k is a constant determining the parameter value of the strain. Thus in producing the stimuli $\tan \theta$ and k are the values to be manipulated for varying the amount of shear and strain, respectively. The detailed procedure for using these formulae can be found in Thompson (*ibid.*).

The initial outline profile was transformed by all 35 combinations of 7 levels of strain ($k = -.25, -.10, 0, +.10, +.25, +.35, +.55$) and 5 levels of shear ($\theta = -15^\circ, -5^\circ, 0^\circ, +5^\circ, +15^\circ$). These transformations are commutative. Shear was applied first. Slides of these profiles, projecting as black lines on a white background were presented to four groups of five subjects. The subjects were instructed to rate the ages of the profiles by choosing an arbitrary number to represent the age of the first profile and assigning multiples of this number to represent the age of succeeding profiles relative to the age of the first. The set of profiles was shown twice with a brief rest period between presentations. Different random orders of the profiles were used for each group of subjects and for each presentation. The full set of profiles is shown in Figure 1. Extreme values at both ends of the two variables were chosen deliberately to test the supernormal stimuli hypothesis.⁶

The mean magnitude estimates for the two presentations are shown in Table 1. Since the initial numbers chosen by subjects were arbitrary, the numbers do not represent age estimates in years. The hypotheses of monotonic increases of relative age estimates with increases in the strain parameter and with increases in shear angle were tested by Monte Carlo methods. For the strain parameter, we predict that for each individual subject the mean magnitude estimate for each strain level collapsed over levels of shear, will have a higher value than the mean for the next lower level of strain. Since the number of order predictions possible for N values of the experimental variables is $N-1$, we make 6 predictions for each subject and a total of 120 predictions for all 20 subjects. There were 109 (91%) correct predictions for the first presentation and 111 (93%) for the second. Testing the first result by a Monte Carlo simulation showed that with a distribution resulting from 5000 runs the probability of 109 or more successful predictions occurring by chance is far less than .001. A similar analysis was performed for the mean ratings for the levels of shear collapsed over the levels of strain. Four predictions are made for each subject, yielding a total of 80 predictions. There were 53 (66%) correct predictions for the first presentation and 52 (65%) for the second. A 5000 trial Monte Carlo simulation shows that the chance probability of 52 or more correct is less than .01.⁷

⁶Supernormal stimuli are produced by exaggerating some relevant aspect of a stimulus. Ethologists claim that such stimuli lead to exaggerated responses. (See N. Tinbergen, The Study of Instinct, Clarendon Press, Oxford, 1951). Gardner and Wallach (Perceptual Mot. Skills, 20, 135-112) used radial transformations to produce supernormally young and supernormally old facial profiles and found that observers perceived these profiles to be extremely young and extremely old.

⁷A distribution of the total number of correct predictions for 20 subjects responding randomly to strain was computed by the following Monte Carlo procedure. If a subject is responding randomly to strain, then the order of the 7 mean ratings for strain will be random. To simulate such a subject, the letters A to G were randomly ordered and the number of times A was to the right of B, B to the right of C, etc. through F to the right of G were counted. The number in the count represents the number of correct predictions for that pseudo-subject. The total count over 20 runs of the pseudo-subject

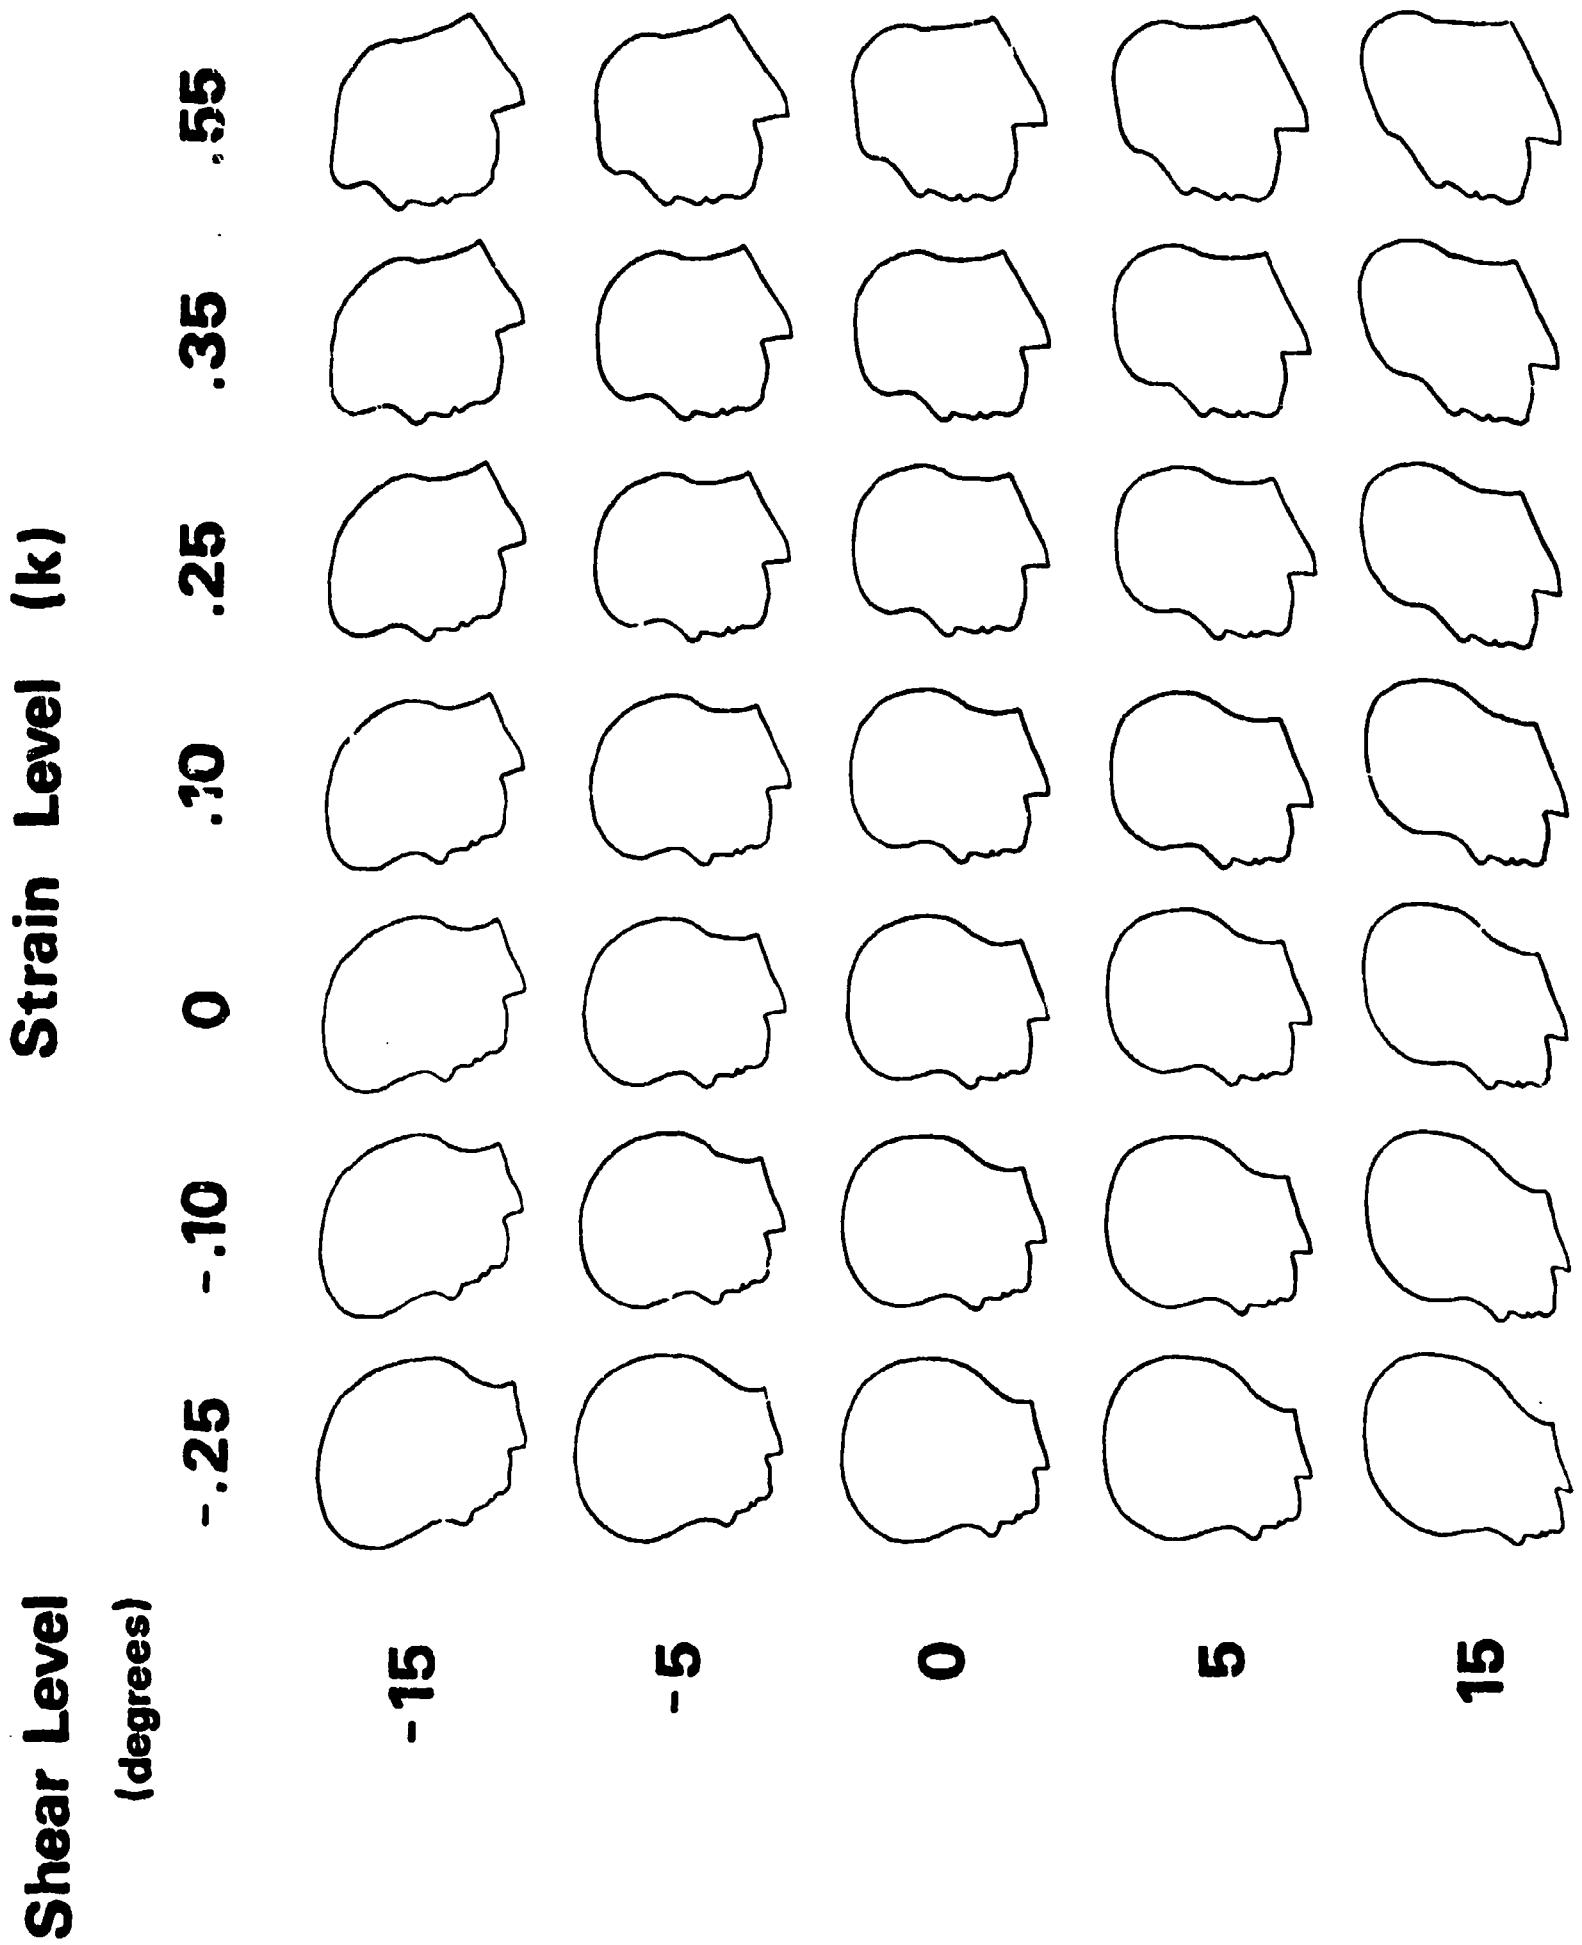


Figure 1

While both transformations have significant effects on perceived age, strain is considerably more powerful. In the first presentation, the range of mean estimates over levels of strain is almost an order of magnitude while for the shear transformation the increase from smallest to largest is less than 50%. Furthermore, strain leads to a considerably larger percentage of correct predictions than does shear. This difference was confirmed by a full replication of the experiment using different stimuli and 20 new subjects.⁸

Sensitivity to the shape changes produced by the strain transformation was assessed in a second experiment by presenting pairs of profiles produced by different levels of the transformation and requiring subjects to choose the older profile in each pair. A series of profiles was produced by applying strain transformations ranging from $k = -.25$ to $+.55$ to a single profile. Eighteen pairs of profiles were chosen; three for each of six levels of difference in k (.04, .08, .12, .16, .20, and .24). At each level, one pair came from the low end of the series, one from the middle, and one from the high end. The pairs were presented twice to four groups of 10 subjects. Different random orders were used for each presentation and each group. Subjects were informed that the study concerned the ability to make fine discriminations of age and that for each pair they were to choose the profile which appeared to be older. During the experiment they were not informed whether or not their responses were correct. By correct response we mean the choice of the profile with the larger k as the older.

An analysis of variance on percent errors of each presentation as a function of difference in k showed a typical psychophysical result; a decline in accuracy with smaller physical differences (for Difference in k , $F(5, 195) = 45.1$, $p < .001$) and an increase in sensitivity with experience in the task (for Presentations, $F(1, 39) = 21.9$, $p < .001$).

However, two other aspects of the results are more important for the question at hand. First, subjects do not merely discriminate the pairs consistently but choose the profile with the larger k as the older profile with greater than chance frequency; in the first presentation the larger k was selected on 83.2% of the trials and in the second, on 89.2% of the trials. In each presentation, each of the 40 subjects selected the profile with the larger k as older more than 50% of the trials. A sign test showed the chance probability of this last result to be far less than .001. Thus the conclusion of the first experiment is confirmed in a different experimental task. Second, sensitivity to the variable is surprisingly fine. Figure 2 shows three pairs of profiles; the pair with the fewest errors, the pair

represents the total number of correct predictions in a 20 subject experiment in which all subjects respond randomly to strain. To estimate the distribution of number of correct predictions in the experiment, the 20 subject count was performed 5000 times. The resulting distribution shows that the probability of 109 correct predictions occurring by chance is far less than .001. The Monte Carlo for shear is analogous.

⁸In the replication experiment there were 105 correct predictions (88%) in each presentation (chance probability less than .001). For shear there were only 45 successful predictions (56%) in the first presentation and 48 (60%) in the second. The chance probability of each of these is greater than .05.

6 Perceiving Age in Facial Profiles

Table 1

Mean Magnitude Estimates of Age - Experiment I

Shear Angle (in degrees)	Strain Parameter (k)							Mean
	-.25	-.10	0.00	+.10	+.25	+.35	+.55	
	First Presentation							
-15	21.3	24.5	31.5	60.5	66.9	75.4	188.6	66.9
- 5	19.7	27.7	36.2	47.7	64.9	152.8	170.4	74.2
0	17.1	25.8	32.6	49.4	71.6	178.5	166.3	77.3
+ 5	17.3	25.1	37.1	48.5	98.5	180.8	199.8	86.7
+15	19.0	33.7	50.0	63.3	147.5	158.0	206.7	96.9
Mean	18.9	27.3	37.5	53.8	89.9	149.1	186.3	
	Second Presentation							
-15	15.5	25.1	28.5	40.5	64.3	69.2	100.5	49.1
- 5	15.8	26.0	30.6	42.1	55.2	75.4	100.1	49.3
0	18.8	22.3	29.6	37.4	60.4	80.4	104.2	50.4
+ 5	18.0	23.7	34.3	39.3	61.7	79.2	108.9	52.1
+15	16.2	23.1	32.5	43.1	72.3	87.1	107.9	54.6
Mean	16.8	24.0	31.1	40.5	62.7	78.3	104.3	


Note - It was predicted that the magnitude estimates will increase from left to right in each row and from top to bottom in each column.

with most errors and an intermediate pair. Even for the pair with the most errors, the profile with the larger k was selected more often than predicted by chance, $\chi^2(2) = 13.2$, $p < .01$.⁹ As can be seen in Figure 2, this perceptual effect is produced by an extremely small change in the relative shapes of the profiles.


As a preliminary test of preservation of identity under the strain transformation, profile views of the external portions of the brain cases of 6 different skulls were traced from x-ray photographs and subjected to five levels of strain ($k = -.15, 0, +.15, +.30, +.45$). Five pairs of transformed profiles were selected from each individual sequences; the k for members of three pairs differed by .30 and those of the other two pairs by .45. A profile of a different skull was assigned to each of the above pairs which had the same k as one of the members of the pair. Slides were constructed of the profile triples such that the two profiles from distinct skulls which had the same level of k appeared in random positions at the bottom. Thirty subjects were presented the slides and asked to select which of the two profiles at the bottom of the slide that appeared most similar to the profile at the top. Subjects were given no feedback on the correctness of their answers during the experiment. The probability of a correct choice (choosing profiles that were transforms of the same skull) was a function of the similarity of the profiles paired at the bottom and the level of difference in k , a greater difference leading to more incorrect choices. The overall percentage of errors was low: For the 30 sets of stimuli presented to 30 subjects, the mean error was less than 17%, with no subject making more than 33% errors. Since no subject made 50% or more errors, a sign-test on the hypothesis of chance responding (binomial distribution) by each subject yields a probability of far less than .001. The results of these studies provide support for two important hypotheses: The strain transformation due presumably to growth, not only provides the major source of the relevant perceptual information for age-level, but also leaves invariant sufficient perceptual information for the specification of the individual identity of the person by the shape of the head alone.

⁹Recall that the pair was shown two times in each presentation of the set, once with the profile with the larger k on the right and once on the left. In the second presentation, 19 of the 40 subjects saw the profile with the larger k as older both times, 18 saw it once as older and once as younger, while only 3 saw it as younger both times. Since guessing would give an expected distribution of 10, 20, and 10 subjects, a χ^2 test shows that the obtained distribution is significantly different from chance.


¹⁰Preparation of this paper was supported in part by a Career Development Award to Robert E. Shaw from the National Institute of Child Health and Human Development (1 K04-HD24010) and by grants to the University of Minnesota, Center for Research in Human Learning, from the National Science Foundation (GB-35703X), the National Institute of Child Health and Human Development (HD-01136 and HD-00098), and the Graduate School of the University of Minnesota.



— K = .37
 --- K = .41
 Difference = .04
 Per cent Errors -
 P1 - 40
 P2 - 30



— K = .09
 --- K = .17
 Difference = .08
 Per cent Errors -
 P1 - 28
 P2 - 16



— K = .23
 --- K = .03
 Difference = .20
 Per cent Errors -
 P1 - 0
 P2 - 0

(P1 is the first presentation, P2 is the second)

Figure 2

References

- Enlow, D. The human face, an account of the postnatal growth and development of the craniofacial skeleton. New York: Harper and Row, 1968.
- Gardner, B. T. & Wallach, L. Shapes of figures identified as a baby's head. Perceptual and Motor Skills, 1965, 20, 135-142.
- Gerasimov, M. M. Osnovy Vosstanovleniya litsa po cherepy (Basics of restoration of the face according to the skull). Moscow: State Publishing House - Soviet Science, 1949.
- Gibson, J. J. Ecological optics, Vision Research, 1961, 1, 253-262.
- MacLeod, P. & Pick, Jr., H. (Ed.) Studies in perception: Essays in honor of J. J. Gibson. Ithaca, New York: Cornell University Press, in press.
- Thompson, D. A. W. On growth and form. Cambridge, England: Cambridge University Press, 1942.

A FORMAL APPROACH TO KNOWLEDGE IN PROBLEM SOLVING

Ranan Bihari Banerji
Temple University, Philadelphia, Pa.

and

Case Western Reserve University, Cleveland, Ohio

The purpose of this paper is to consider the similarities and differences between the methods of problem representation used in three different centers of research. In the process we hope to be able to bring the terminologies somewhat in line. In addition to problem representation I want to discuss the way knowledge is defined in the three activities in so far as it deals with problem solution.

The three specific centers I have in mind are the ones at i) Carnegie-Mellon University, spear-headed principally by Allen Newell and Herbert Simon ii) University of Pennsylvania, led principally by Joseph Scandura and iii) Case Western Reserve University, led by George Ernst and myself. The Carnegie-Mellon activity has been summarized by Simon and Newell in their book, "Human Problem Solving" published in 1972 by Prentice Hall. The work at University of Pennsylvania is described in Scandura's "Structural Learning", published by Gordon and Breach in 1973. The latest state of our work is probably best summarized by the report, "Mechanical Discovery of a Heuristic" by Ernst, Banerji, Hookway, Oyen and Shaffer, published by the Jennings Computation Center, CWRU. My own book, published in 1969, is somewhat back dated.

The motivations of these schools are in no way the same. Newell and Simon are contributing basically to Psychology, Scandura to educational evaluation while Ernst and I have been trying to study the abstract theory of non-numerical model of a problem. The Case and Carnegie models are even more similar since we share also our definition of a solution. As a matter of fact, the coincidence of these two models is no accident - we at Case borrowed it from Carnegie, or at least were strongly influenced by it.

Let us start with what Newell and Simon call a State Space (I called it "situation space" in my book to avoid the wrath of the haters of systems theory, but have since found it perfectly safe to hide behind what has now become accepted nomenclature). Each element of a state space is a "problem state". Each element of this space is described by a conjunction of statements in some language and so a set of predicates are defined on this set. Each predicate defines a statement form whose truth or falsity can be easily determined for each state. Scandura calls these the Atomic Decision Making capability. The other two schools had not given these any name. Also we have a set of unary functions which transform states to states. At Carnegie-Mellon they call these transformations, at Penn they are atomic operating rules. Carnegie-Mellon makes it clear that not all rules are applicable to all states - the applicability of a rule is testable in terms of the decision making capabilities. The choice of rules in the final pro-

gram is also guided by this.

The exact analogy ends here. At Carnegie-Mellon and at Case we define a subset of problem situations which at Carnegie-Mellon they call Goal States and we at Case call "Winning states". In addition we often specify a "Starting state" from which a goal state is to be reached. So our "goal" is a pair - an initial state and a set of goal states.

Scandura's "Goal environment" is much more general than the other two schools. A "goal situation" is a set of our goals, for each initial state a set of acceptable goal states. Thus a goal situation for Scandura is a relation on state space.

At this point another non-essential difference appears between Scandura and the rest of us which may be a source of confusion to the unwary. In his formulation he calls the set of initial states the stimuli and the goal states (or output-coded versions of them) the responses. Hence those states which are neither stimuli nor responses are only implicit in Scandura's model. Nevertheless, they are there, since many of the rules have these "limbo" states in their domain, as some of his examples show.

One of the major differences between Scandura's model and the State space model as presently conceived, however, is that in Scandura's Theory of Knowledge some rules are allowed to act on rules. The class of rules are, therefore expandable.

One of the rules that Scandura uses to obtain new rules from old rules yields programs - rules obtained by composition and recursion of other rules. The other two schools only use composition.

The last statement needs explication. Both Scandura and Simon-Newell consider programs with loops. However, as Scandura points out the execution of a program actually constructs a loopless path through state space - the same rule may be applied repeatedly, but every time it is applied, it is applied to a different state.

In Scandura's theory of knowledge rules grow by application of other rules till one is found which accounts for the goal situation i.e. it yields the proper output for each input. In Simon and Newell's model a solution is constructed by composition of rules so that the initial state is converted into a goal state.

Here the difference between the motivations of the three schools becomes important. In Scandura's work, the correct behavior by the subject is achieved by instruction and erroneous behavior is explained by assuming that it is due to a change in the program which describes the rules accounting for the correct behavior. The purpose is prediction and training. The knowledge is achieved when the proper rule is produced.

Simon and Newell are interested in the way the composition of rules is constructed among the plethora of possible ones: instruction of the subject stops at describing the problem environment and the goal. This construction process is carried out by the GPS which uses knowledge other than the problem specification. This knowledge, for the General Problem Solver, is the difference orderings and the table of connection.

So the growth of knowledge in Scandura's experiments can be considered to be guided by the knowledge in Newell and Simon's table of connections - as long as this growth is obtained by composition of rules. In Scandura's model, knowledge guides the behavior. In Simon and Newell's model, the difference transformation table guides the growth of knowledge.

At Case Western Reserve, we have been interested in constructing the table of connections i.e. studying the construction of knowledge at a level of abstraction removed one more level.

In what follows we shall discuss principles behind the program which constructs the table of connections.

The table of connections constructed by our method is always triangular and hence leads to efficient construction of correct behavior. However, the requirement of triangularity acts as a strong guiding force for the efficient working of the program which constructs the table itself.

It will be recalled that the connection table indicates how each transformation or atomic rule affects the result of the "differences" which can be considered (only informally - the GPS is recursive, Scandura behaviors are not) to be the analogs of the Scandura dmc's. For a table of connections to be triangular, it is necessary that there be a dmc which is affected by only one of the atomic rules.

As said before, each state is a conjunction of Predicates which naturally define certain dmc's. However, for the purpose of table construction one has to discover combinations of the properties defining the predicates. This way for each transformation one obtains a set of basic invariant properties.

Since the winning states are defined in terms of the original predicates, one can search for sentences which are implied by winning states and are expressible in terms of the invariant properties. This chain of implications can be continued. Since a move which keeps a property invariant also keeps invariant any implied property, the triangularity is assured.

Limitations of space and time precludes a detailed discussion here. The method will be exemplified at the presentation. The interested reader may write to Professor Ernst for a copy of the report.

THE LONG-TERM MEMORY IMPROVEMENT EFFECT:
MEMORY IMPROVEMENT OR CONCEPT DEVELOPMENT?

Charles J. Brainerd

University of Alberta

In this paper, I should like to report some selected preliminary findings from a recently completed longitudinal study of children's memory for simple visual stimuli. For the most part, these findings pertain to the so-called "long-term memory improvement effect" which Piaget first reported some seven years ago. On the basis of these findings, I shall propose a very elementary alternative interpretation which seems to explain both Piaget's data and my own. Essentially, the interpretation stipulates that, whatever else the long-term memory improvement effect is concerned with, it probably has nothing whatsoever to do with memory. Before discussing the data, however, a few remarks about the history and theoretical significance of the long-term memory improvement effect are in order.

In 1967, Piaget delivered the Heinz Werner Memorial Lectures at Clark University. During the first of the two lectures, he presented a rather startling empirical datum. He claimed that children's recall memories of certain special visual stimuli tend to be better six months after encoding than immediately after encoding. The research on which this claim was based involved three steps. First, groups of children were shown a simple two-dimensional stimulus and were asked to remember it because the experimenter would return sometime to question them about it. Second, the experimenter returned one week later to test the children's recall of the stimulus. The children were asked to draw a picture of the image. Occasionally, they also were asked to describe the stimulus verbally. Third and finally, the experimenter returned six months later and repeated the same memory tests administered one week after encoding.

Following the second recall test, the one-week and six-month responses were compared. Surprisingly, dramatic improvements in recall accuracy apparently had occurred during the six month interval between the two tests. In his lecture, Piaget reported detailed evidence (number of Ss, amount of change, etc.) on two stimuli: first, an array of 10 sticks of differing length seriated from shortest to longest and, second, the water level in a closed bottle tilted at various angles. The second type of stimulus is shown in Figure 1. Items 2 through 8 in Figure 1 show a closed bottle tilted at some possible angles. In all cases, the water level that Piaget's subject's were shown was parallel to the base line appearing just




FIGURE 1

below the bottle. The seriated array of 10 sticks was administered to 3- to 8-year-olds. The water bottle stimuli were administered to slightly older children.

As I already have noted, the six month improvements that Piaget reported were nothing short of dramatic. With the seriated array, the findings ran as follows: for roughly three-quarters of the Ss, the six-month drawings were more accurate reproductions of the array than the one-week drawings; the drawings of roughly one-quarter of the Ss neither improved nor deteriorated during the six-month interval; finally, not a single drawing was judged to be worse six months after encoding than one week after encoding. In other words, "forgetting" simply did not seem to occur with this particular stimulus. The results for children's recall of the water level of tilted bottles were similar if somewhat less dramatic: recall of the water level improved across the six month interval in roughly one-third of the Ss; slightly less than 60% of the Ss evidenced no change; only 10% of the subjects evidenced deterioration.

In view of our common sense expectation that the only possible change which should occur in recall across so long an interval as six months is decay, it was perhaps inevitable that the first response evoked by Piaget's data was skepticism about the reality of the data themselves. Less than a year after Piaget had reported his memory improvement effect, Robert Altemeyer, Daniel Fulton, and Kent Berney began a virtual point-for-point replication study which focused narrowly on the first of the two stimuli mentioned earlier-- the seriated array of 10 sticks. Although the Altemeyer-Fulton-Berney findings were less dramatic than Piaget's, there was overwhelming agreement with Piaget's data: recall improved in roughly half the Ss; recall did not change in roughly one-quarter of the Ss; recall deteriorated in roughly one-quarter of the Ss. The investigators concluded that "our data substantially confirm Piaget's results and further explorations yield results which in no way oppose his explanation (Altemeyer, Fulton, & Berney, 1969, p. 846)." Thus, the initial skepticism about the data themselves seemed unwarranted.

A year after the Clark address, Piaget published a comprehensive volume on memory (Piaget, Inhelder, & Sinclair, 1968). A year after that, Inhelder (1969) published an overview of the main theoretical points in the book plus illustrative empirical findings. We shall now briefly consider some of the key theoretical proposals. Before doing so, however, I should like to emphasize that the validity of Piaget's theory of memory appears to turn principally on the long-term memory improvement effect--or, more precisely, on his interpretation of the effect. None of the other research that Piaget has reported on memory to date poses even a slight problem from more traditional

theories of memory. It is only the long-term memory improvement effect that seems to require a new theory of the sort that Piaget proposes.

The point of departure for Piaget's analysis of memory is an assumption he has advanced many times before in conjunction with cognitive capacities other than memory (e.g., language, imagery, logic), namely, during each period of cognitive functioning, memory is strictly a dependent variable vis-a-vis the cognitive structures previously posited for that period. [In the case of the long-term memory improvement effect, the period is middle-childhood and the relevant cognitive structures are the so-called groupements.] Piaget actually seems to assume that there is no memory capacity as such--there are only the cognitive structures and they are responsible for our knowledge of the past. This brings us to the unique empirical prediction of the theory--the long-term memory improvement effect. To be perfectly accurate, the effect was not deduced from the theory and then tested as the term "prediction" implies. Instead, it appears to have been discovered quite by accident. In his book on memory, Piaget observes: "We...devised a number of preliminary tests to determine what our subjects remembered...of situations or states resulting from operational processes...This led to the (fairly obvious) conclusion that the memory of children is bound up with the way in which they interpret a model at various stages of their operational development...a chance encounter with one of our subjects persuaded us to look further into the matter. When this subject, who had been questioned about the memory test six months earlier, turned up for quite a different investigation, and was asked what he remembered of the first one, his reply showed that he did, in fact, recall it, but that he had schematized it further (Piaget, Inhelder, & Sinclair, 1968, my translation).

In any case, the general prediction of the theory is, in Inhelder's (1969) words, that long-term memory improvement will be observed for all visual stimuli which are "the (static) end result of operational transformations" *during the time interval when the cognitive structures that embody these transformations normally are acquired.* Returning to the two stimuli we discussed earlier, both the seriated array and the horizontal water level previously have been inferred to be the end result of the cognitive transformations embodied in the groupement structures of the concrete-operational period. [It should be noted parenthetically that it is a moot question as to whether or not such inferences are correct or, indeed, whether or not such a thing as a groupement structure even exists. For purposes of investigating the long-term memory improvement effect, we are forced to accept the structures and their behavioral correlates as givens.] Because the theory says that groupement structures normally emerge during the late preschool and early elementary school years, the expectation is that long-term memory improvement effects for seriated arrays and horizontal water levels should be observed during this age range. More generally

long-term memory improvement effects should be observed during this age range for *all* visual stimuli that the theory says are static end results of groupement transformations. [This leaves us with a wide assortment of alternative stimuli to choose from--i.e., double seriation matrices, two-dimensional projections of common three-dimensional objects, conservation transformations, and double classification matrices.] In view of the fact that some long-term memory improvement effects have already been observed for such stimuli during the appropriate age range, Piaget concludes that: (a) memory is indeed dependent on his *particular* cognitive structures and (b) the memory code must undergo structural reorganizations during the course of development which are dictated by these same cognitive structures.

Now, purely for the sake of argument, suppose we stipulate that Piaget has established two points beyond reasonable doubt. First, on the empirical side, his data are, in principle, replicable. That is, it is an established empirical fact that children's drawings and verbal descriptions of stimuli such as those we have mentioned are more accurate six months after encoding than one week after encoding. Second, on the theoretical side, let us accept the more dubious proposition that if long-term memory improvement effects satisfying the aforementioned conditions exist, then Piaget's cognitive structures are automatically implicated in any theory of memory. The question now becomes whether or not we can connect these two statements. That is, do the six-month changes necessarily reflect changes in *memory* or is something other than memory changing? Some rather elementary methodological considerations suggest that the second alternative is the more likely.

Methodologically, the most striking feature of the Genevan long-term memory design is its complete lack of control for extraneous influences other than memory. Of course, lack of controls for alternative interpretations is hardly a new phenomenon with Genevan research. In the case of the long-term memory studies, however, the possible extraneous factors that might have produced the changes attributed to memory are particularly heinous. For example, given the nature of the dependent variables (drawing and extemporaneous verbal description), the contaminating influences of language development and fine motor coordination immediately suggest themselves. During the age range in question, massive improvements in both of these capacities are known to occur.

A third contaminant also suggests itself: concept development. It is this extraneous influence that was the focus of the research I wish to report and it may be summarized as follows. The stimuli for which the long-term memory improvement effect already has been observed and other stimuli for which the effect should be observed all involve an underlying concept of some sort. Moreover, these underlying concepts are known to be acquired during precisely the same age range as the one examined in Piaget's memory studies. In

the case of the 10-stick array, the underlying concept is serial order. In the case of the water level stimulus, the underlying concept is horizontality. If one were to ask elementary schoolers or preschoolers *without prior encoding*, simply to draw a picture of 10 different sticks on two occasions six months apart, the degree of improvement between the two occasions no doubt would be considerable. Similarly, if we show children the bare outline of any of the tilted bottles in Figure 1 on two occasions six months apart and ask them to draw in the water level, there would also be improvement.

Thus, we simply cannot say just how much of Piaget's "memory" improvement data reflects actual improvement in memory and how much must be written off as improvement in other extraneous variables. There can be very little doubt that concept development, in particular, accounts for some of the observed improvement. The general aim of the study I shall now report was to separate these two factors and to determine the relative contribution of each to children's scores at one week and six months.

The study focused on one of the types of stimuli for which Piaget reported detailed memory improvement findings in 1967: the water level stimulus illustrated in Figure 1. It turns out that only two additions must be made to Piaget's design to sort out the effects of concept development and memory improvement. First, at least one other condition must be run. This condition is one in which the Ss encode stimuli that are identical with those that Piaget's Ss encoded, except that the stimuli are not "correct" instances of the underlying concept. [To illustrate, consider stimulus No. 2 in Figure 1. In place of the horizontal water level that Piaget's Ss were shown, suppose the Ss were shown a water level tilted in the same direction as the bottle.] The prediction of Piaget's memory theory for such a condition is straightforward and has been explicitly formulated elsewhere (Inhelder, 1969): Subjects encoding such stimuli should evidence no long-term memory improvement because the stimuli are not "the (static) end result of operational transformations."

The second change that must be imposed on Piaget's original design is purely psychometric. We must change the scoring procedure slightly. We now have two encoding conditions. One of them is a replication of Piaget's condition in which the Ss encode visual exemplars of some concrete-operational concept; the other is a control condition in which the Ss encode otherwise identical stimuli that do not happen to be exemplars of the concept. We now may determine the relative contributions of concept development and memory improvement simply by *applying exactly the same scoring method to the second condition as we apply to the first condition*. That is, instead of evaluating the responses of the Ss in the second condition in terms of their approximation to the stimuli originally encoded, we evaluate them in terms of their approximation to the conceptually correct stimuli encoded by the Ss in the first condition. If concept development produced some of the change Piaget

reported, then the six-month responses of the Ss in the second condition should be closer approximations to the conceptually correct stimuli than the one-week responses. Similarly, if long-term memory improvement produced some of the change Piaget reported, then the absolute amount of response improvement over the six month interval should be greater for Ss in the first condition (upon whom the effects of both memory improvement and concept development are operating) than for Ss in the second condition (upon whom the effects of only concept development are operating). Finally, if long-term memory improvement produced none of the change Piaget reported, then the absolute amount of improvement in the two conditions should be the same.

The study that eventually was conducted may be summarized as follows. The Ss were 450 first, second and third graders (150 per age level). Each S was assigned to one of three encoding conditions (150 condition). In the first condition, the Piaget replication, the Ss were shown stimuli No. 2, 4, 6 and 7 in Figure 1. Inside each bottle, a correct horizontal water level appeared. In the second condition, the Ss also were shown stimuli No. 2, 4, 6 and 7. However, a correct horizontal water level appeared inside only two bottles. An incorrect tilted water level appeared inside the other two bottles. In the third condition, the Ss were shown the same 4 stimuli and an incorrect tilted water level appeared inside all the bottles. Except for the type of water levels encoded, the procedures for the three conditions were identical.


During the encoding phase, E showed the Ss two large closed bottles--one empty and the other half-filled with red colored water ("Kool Aid"). The E explained that he would tilt the empty bottle at various angles and then show the Ss what the Kool-Aid looked like. Next, the E tilted the bottle to one of the four predetermined positions and showed the S a slide of the tilted bottle with the water level drawn in. Each time the E showed the S a slide, he cautioned the S to remember it because he would return later to ask questions about it. One week and six months after encoding, the Ss performed a recall task and a recognition task. The recall task was the same as Piaget's. The Ss were given sheets of paper on which stimuli No. 2, 4, 6, and 7 from Figure 1 were printed and they were asked to draw in the water level they had seen for each one. In the recognition task, the Ss were administered four recognition arrays. Each array contained six bottles tilted at the same angle and a different water level appeared within each bottle. One of the six water levels was the one S had seen and another was the conceptually correct level. Although Piaget's original design included no assessment of recognition memory, we decided to include a recognition test primarily as a control for the possible motor coordination effect mentioned earlier.

After the six-month retest was completed, the recall and recognition data from both tests were scored in terms of the degree to

which they approximated the correct horizontal water level. The amount of positive or negative change between the two tests then was determined. The recognition arrays had been constructed in such a way that the water levels appearing in each array formed a natural six-point scale. This scale was applied to the recognition data. It did not prove difficult to construct a similar scale for the recall data. The recall scoring method produced a satisfactorily high interscorer agreement of .85.


Before looking at the data, let us recall what we should find. If Piaget's original findings were at least partially the result of *memory* improvement, then the differences between the six-month and one-week scores should be greater in the first condition than in the second and greater in the second than in the third. On the other hand, if Piaget's original data were purely the result of concept development, then the differences between the six-month and one-week scores should be about the same in all three conditions. In Figure 2, we see that the recognition scores after six months were about 50% better than at one week for all three conditions. The \bar{S}_s in the first condition evidenced almost precisely 50% improvement; the \bar{S}_s in the second condition evidenced slightly less than 50% improvement; and the \bar{S}_s in the third condition evidenced slightly more than 50% improvement. None of these differences even approached significance. Now, let us look at the recall data in Figure 3. Again, the amount of improvement is around 50% and none of the between-group differences approached significance.

The data suggest two conclusions. First and most important the original findings on which Piaget based his claims for the existence of a long-term memory improvement effect probably have nothing to do with memory per se but, rather, result primarily from correlated improvements in the underlying concepts presupposed by his stimuli. At the very least, this seems to be a fair conclusion with regard to the Geneva data on recall of water levels. Even at the risk of over generalizing, I also would contend that it is by far the most probable explanation of his findings with other stimuli. The virtual identity of our three conditions on both the recall and recognition measures suggests a further interesting conclusion: Not only does long-term memory improvement for horizontal water levels apparently fail to occur, the \bar{S}_s seem to remember virtually nothing about the original stimuli six months after they were encoded. If forgetting was not almost total, then we would expect at least some small difference between the scores of \bar{S}_s in the first condition (who saw only horizontal water levels) and the scores of \bar{S}_s in the second condition (who never saw a horizontal water level). Such differences were not observed. To evaluate this second conclusion more directly, we ran a follow-up condition containing 60 \bar{S}_s . These \bar{S}_s were simply administered the one-week and six-month tests at the appropriate times. Instead of being asked to draw or select the water level they previously had seen, the \bar{S}_s were asked to draw or select the correct



Obtained Data by Condition
 Total N = 450 (150 Ss per Condition & per Age Level)

FIGURE 2



Obtained Data by Condition
 Total N = 450 (150 Ss per Condition & per Age Level)

FIGURE 3

water level. The same scoring procedures were used for the follow-up condition. The results were in line with our second conclusion: *The amount of change in both recognition and recall after six months did not differ significantly from the amount of change observed for Ss in our first condition.*

To sum up briefly, we have seen that Piaget's theory of memory "predicts" long-term memory improvements in any and all visual stimuli which are "the (static) end result of operational transformations." We also have seen that the data on which Piaget predicates his claims for the existence of such effects are subject to the more reasonable alternative interpretation of concept development. The data reported herein provide support for the alternative interpretation. Therefore, to the extent that Piaget's key assumption that memory always depends on his operational structures turns on the existence of the long-term memory improvement effect, we must conclude that the present evidence is inconsistent with the assumption. So as not to be guilty of over generalizing from a single set of stimuli, we are examining other sets of stimuli for which the long-term memory improvement effect has been claimed via the design outlined above. To date, our findings on these other stimuli have been consistent with the data just reported.

REFERENCES

- Altemeyer, R. A., Fulton, D., & Berney, K. M. Long-term memory improvement: Confirmation of a finding by Piaget. *Child Development*, 1969, 40, 845-857.
- Inhelder, B. Memory and intelligence in the child. In D. Elkind & J. H. Flavell (Eds.), *Studies in cognitive development*. New York: Oxford University Press, 1969.
- Piaget, J. *On the development of memory and identity*. Barre, Mass.: Clark University Press, 1968.
- Piaget, J., Inhelder, B., & Sinclair, H. *Memorie et intelligence*. Paris: Presses Universitaires de France, 1968.

ART EDUCATION - MEANS OR END?

Boris Blai, Jr.
Harcum Junior College

This pilot study, supported by a grant from the Spencer Foundation, investigated, empirically, Sir Herbert Read's views relating to education. These are:

(1) Real art education - termed 'visual learning' - "represents a group of stimuli-response behaviors which direct us, shape us, and maintain our very existence - all the time."

(2) "Our policy of education through art is based on the hypotheses that the images we evoke in the course of any kind of cognitive activity have a universal significance, and correspond to something permanent and unchanging in the nature of man."

The specific objective of this exploratory inquiry was to 'test' two propositions:

Hypothesis #1 - That art education, offered at a collegiate level, demanding the intellectual effort expected for any academic subject, helps students achieve: (1) a greater degree of emotional well-being through improved self-image - (the latter being brought about by helping the student produce satisfying objects he can see); (2) improved study habits and attitudes towards study; and (3) greater academic progress (achievement) than those who are not exposed to a collegiate art curriculum.

Hypothesis #2 - That art education provides a means for learning problem-solving, for Sir Herbert hypothesizes that the visual language of real art education is a direct route to the student's mind in developing problem-solving skill for application of other subject-matter areas.

Clearly then - if corroborating evidence is found for either or both of these hypotheses - the potential educational-learning implications could be far-reaching; for if the effects of real art education upon student learning are as hypothesized by Sir Herbert, the curriculum-structuring impact could be considerable at all levels of formal schooling and learning.

The subjects of this investigation were a group of 28 young women who graduated from Harcum Junior College in May 1973, after completing associate degree programs of study. There were 14 Visual Arts majors, and 14 non-Visual Arts majors; 'matched' in terms of College Board Scholastic Aptitude Test verbal and math scores; that is - there was less than a 50-point 'spread' of scores among the subjects.

Measures of College-entry, and College-graduation levels were obtained for these two groups through pre- and posttest administration of: (1) The Watson-Glaser Critical Thinking Appraisal instrument, consisting of a series of test exercises which require the application of some important abilities involved in problem-solving skills; (2) the Emotional Stability scale of the Gordon Personal Profile inventory; and (3) the Survey of Study Habits and Attitudes inventory. The fourth area investigated - academic progress, or achievement - was 'measured' by comparing the graduation cumulative grade-point averages of the Visual Arts and non-Visual Art majors.

To determine if the results of these comparisons - that is, the obtained differences between the averages earned by these two groups of students were likely to be chance variations or statistically significant differences - t-ratios were computed. As we are aware, a t-ratio of 3 is virtual certainty (about 999 chances out of 1000) that a true difference exists between the means of the two populations from which the samples were drawn, and is not one that might be due to sampling fluctuations alone - or some other chance variations. And of course, a t-ratio in excess of 3 is that much more assurance that a true difference exists between the two populations involved.

With the exception of study skill and attitudes, statistically-significant differences were found. For the difference in average performance in the problem-solving dimension, the t-ratio was 4.37. In the area of emotional well-being, the t-ratio was 3.33, and for the difference in cumulative averages earned between the Visual Arts and non-Visual Arts majors, the t-ratio was 4.40.

In the area of study skills and attitudes, the averages earned by both groups were the same; therefore in this instance, no statistically significant evidence was found to support Sir Herbert's assertion that art education eventuates in improved study habits.

The implications and conclusions of this pilot inquiry suggest the following: In theory, there are two principal elements associated with student learning. These are: (1) difference in initial aptitude (or Nature), and (2) college characteristics contributing to student learning (or Nature). Therefore any study of college impact upon students must take into account differences in initial aptitude among students. It is also recognized that higher posttest scores may, in fact, be due to the "maturation" factor, as well as some residual "practice" effect. Students grow and develop - as a result of their everyday life-learning experiences - with or without participating in experimental programs! However, if statistically significant differences are found among groups exposed to differing academic environments (in this case - collegiate programs), the net effect is to hold constant these factors of "maturation" and "practice effect" - to which all are subject - and therefore identify such differences with the environmental "press" (i.e. the various college majors).

To attempt to control for the student's ability prior to college entrance the ability which has contributed to the student's academic growth while in college, this study considered the average improvements among several small groups of Harcum students in a series of academic-related skills and achievements, as well as personal characteristics areas. Through a "before" and "after" testing situation for two samples of freshmen who were retested just prior to completion of their 2-year program, the net effect was to hold constant the factor of differences in student aptitude, by combining the differences in improvement among the low, average, and high academic aptitude students. The remaining differences between their collective performance initially upon entrance to Harcum, and their performance just prior to graduation two years hence, may then reasonably be associated with the College's effect - the so-called environmental 'press'.

The results of this pilot investigation have generally shown statistically significant differential-gains for this small-sample Harcum group enrolled in the Visual Arts curriculum, in contrast with a 'control' group enrolled in various other Harcum curricula. It is generally acknowledged that the visual and graphic arts are playing an increasingly greater role in contemporary life. Therefore, it is believed that further study, through replication of this pilot inquiry, might well be fruitful; to ascertain whether the gains here reported also are found to exist at various schooling levels; e.g., elementary, secondary, undergraduate and graduate collegiate. Additionally - are such gains found among males as well as females? And finally - do such gains persist over time.

If such future inquiries are well-designed, replicated a number of times, and show consistent directional findings and meaningful strengths of association, then the collective results of such investigations may, within reason, be generalized for purposes of modifying existing educational programs and practices.

COGNITIVE DEVELOPMENT AS PROBLEM-SOLVING:
THE MEANING OF DÉCALAGE IN SERIATION TASKS¹

Kurt W. Fischer

University of Denver

For many years child psychologists have been describing stages of development--stages of crawling, stages of vocalization, stages of classification of blocks, and so on. These stages have been useful as descriptions of the development of particular skills, but beyond this descriptive utility their significance has usually been unclear. During recent decades Jean Piaget has used stage descriptions in a more meaningful way (for example, Piaget & Inhelder, 1969; Piaget, 1952). He has placed his descriptions within the context of a comprehensive theory of intellectual development, with the result that his stages usually have a definite significance. In a general way, they indicate the progressive, step-by-step structuring of intelligence in the child, and more specifically most of the stages portray the development of skills that are essential components of intelligence. In other words, Piaget treats his stages as a means of assessing and portraying the development of intelligence in the child.

While Piaget's use of stages is a definite improvement over previous usage, many of his stages nevertheless encounter a major conceptual difficulty. When he finds stages of performance on a particular task, such as seriation of sticks or classification of blocks, he attributes the stages to the development of the child's intelligence, yet the stages of performance on a given task turn out to have an uncertain, ambiguous correspondence with intellectual development. While the individual task itself makes an important contribution to the child's stage of performance, seemingly minor changes in the nature of the task can produce dramatic changes in the child's stage.

This fluidity of stages is called décalage, or temporal displacement, by Piaget. At the same age tasks that appear to be almost identical produce very different stages of performance, and consequently any specific stage occurs at a different age for

¹The germ of the idea on which this paper is based grew during discussions with Susan Carey-Block and David Wood. The research reported in this paper is being supported by a research grant from the Spencer Foundation and by a Faculty Research Grant from the University of Denver.

each task.²

My proposal for solving this enigma is a relatively simple one, deriving directly from Piaget's own assertion that intelligence can only be understood as the interaction of organism and environment. It is not enough to consider only the organism (that is, the child's intellectual level) when trying to explain Piaget's stages. The environment (that is, the task) makes an equally important contribution to the child's performance. I would propose that the best way to deal with this interaction is to view the child as trying to solve the problem posed by the task. His stages of performance on a particular task are then produced jointly by his cognitive level and the difficulty of the task, because these two factors together determine how well he can understand the problem posed by the task. Consequently stages of cognitive development are best understood as stages of problem-solving.

Décalage in Seriation Tasks

A good example of the difficulty with Piaget's descriptions of stages is the standard Piagetian task of seriation. A child is given ten sticks of differing lengths that if ordered properly will form a series from shortest to longest. The adjacent sticks in the series differ by only a few millimeters in length. Most four-year-olds in our society cannot place these sticks into the proper series, while most nine-year-olds can make the series correctly. The most fine-grained analysis of stages of seriation provided by Piaget (1957; 1970) depicts four stages of seriation³ and I have added to the description a zero stage, in which the child cannot seriate at all. These stages are described in the column labelled "Stages of Seriation of Sticks" in Table 1. The descriptions of the stages have been altered slightly from Piaget's original descriptions on the basis of the findings of an experiment being carried out by Karen Cohen and me that will be discussed below.

The difficulty arises in its most salient form as soon as seriation is tested with different types of materials. Many diverse things can be seriated besides sticks differing a few millimeters in length, and research to date indicates that the ability to seriate many of these other materials follows essentially the same four stages as seriation of sticks. Yet the stages for some of these other seriation tasks occur at very different ages from those for the sticks--so different in some

²Because Piaget presents décalage as a basic characteristic of all intellectual development, it does not literally contradict his theory. On the other hand, he has not adequately explained what décalage means or why it occurs.

³Because Piaget is using these stages primarily as an assessment tool and is not interested in the stages themselves, he differentiates different numbers of stages depending upon the purpose that he is writing for. The stages of seriation listed in this paper are based on the largest number of stages that he differentiates.

cases that they belong to entirely different periods of intellectual development within Piaget's framework. For example, the seriation of nesting cups develops through the stages between one and three years of age (Greenfield, Saltzman, & Nelson, 1972; Greenfield, 1972, personal communication), which places them firmly within what Piaget calls the pre-operational period. Piaget himself (1967, p. 50) states that by four or five years of age most children can seriate sticks that differ greatly in length (at least 25 mm. difference between adjacent sticks). The seriation of items described in propositions, such as "John is taller than Sam," "Sam is shorter than Herman," etc., apparently develops through the four stages of seriation between approximately ten and fifteen years of age (Inhelder & Piaget, 1958, Chapter 16), which places them within what Piaget calls the formal-operational period. Consequently, the stages of seriation cannot be used in any general way to indicate that the child has reached a particular period of intellectual development, because the same stages can occur in entirely different periods if different types of seriation tasks are used.

Décalage occurs in the same way in several other types of tasks used commonly by Piaget. Piaget himself has carefully documented the fact that in the many different kinds of conservation tasks, all show essentially the same stages of development, but those stages occur at different ages for each task (Piaget & Inhelder, 1969). Tasks involving classification probably show the same kind of décalage.

A Proposal for Solution of the Difficulty

Within Piaget's theory the significance of these stages is still unclear. Why do the same stages occur in tasks that develop at such different ages, and more generally what do the stages mean?

The proposal of this paper is that these stages depict more than just changes in intellectual capacity. They describe phases in the solution of a problem of a particular type. That is, a child trying to solve some specific seriation task must solve the problem posed by that task. The stages that he shows in the task are stages in the solution of seriation and are produced jointly by his cognitive capacities and the difficulty of the specific task. Intellectual development and problem-solving involve essentially the same process. Intellectual development can be treated as problem-solving, and problem-solving can be treated as a developmental process in miniature (termed "microgenesis" by Werner & Kaplan, 1963).

I am using the word "stage" to refer to changes that occur with the development of intelligence and the word "phase" to refer to changes that occur in any learning or problem-solving situation, irrespective of whether the development of intelligence is involved. For one type of problem, such as seriation, all the various kinds of seriation tasks produce the same stages because they all require solution of the same kind of problem, seriation. The sequence of steps in problem-solving are the

phases of seriation. Phases become stages when the steps are stretched out over a long period of time as a result of cognitive development; an advance in cognitive sophistication or cognitive capacity is required for a child to achieve a more advanced stage.

Not only can a set of phases be described for solution of seriation tasks, but another set can be described for solution of conservation tasks, perhaps still another set for solution of classification tasks, and so forth. Two tasks of the same type produce the same stages at different ages because they both involve solution of the same type of problem, and the problem must be solved by a set of phases required by the structure of the problem.

The proposed phases for seriation are shown in Table 1 in the column labelled "General Organization," and the observed stages of seriation of sticks are described in the column labelled "Stages of Seriation of Sticks."

The far-right column, labelled "Hypothesized Transition Rule," suggests a way of conceptualizing the problem-solving process that occurs as the child moves from phase to phase. The proposed model is a combination of analysis-by-synthesis and differentiation-and-integration analyses of problem-solving (Fischer, in preparation). To move from Phase 0, where he does not seriate at all, to Phase 1, where he makes a rough discrimination between large and small, the child must recognize that the problem involved in the task is one of size. In Phase 1 his understanding of size and how to use it is still relatively primitive, however. To move from Phase 1 to Phase 2, where he tries to order all or most of the items and succeeds in making a partial series, the child must define the general outline of the problem: All the items should be ordered in terms of size. In Phase 2, however, he still does not have the ability actually to seriate all the items. To move from Phase 2 to Phase 3, where he succeeds in making a complete series by successive adjustments (trial and error), the child must differentiate the components of the problem: Each item must be treated as both larger than some items and smaller than others. In Phase 3, however, he still cannot completely intercoordinate these components mentally. He deals with them by making successive adjustments of pairs of items. To move from Phase 3 to Phase 4, where he anticipates the entire series from the start, and succeeds in making a complete series, the child must integrate the components of the problem into a new scheme: All the items together form a step-by-step series, and each item has its own specific place in the series. At this point the problem-solving process is complete.

A number of research questions follow directly from the proposal that the stages of seriation actually reflect phases of problem-solving. I have recently begun a research project, sponsored by the Spencer Foundation, to pursue several of them. There is time to mention only a few of the more interesting ones here.

Table 1

The Hypothesized Four Phases of Seriation

<u>Phase</u>	<u>General Organization</u>	<u>Stages of Seriation of Sticks</u>	<u>Hypothesized Transition Rule</u>
0	No understanding	The child does not understand that the sticks should be arranged in order of length.	Person recognizes that there is a problem to be solved involving size.
1	Distinguishes large from small.	He understands that the problem involves length, but does not know how to treat a single stick as both longer than one stick and shorter than another: he distinguishes long sticks from short sticks.	The general outline of the problem is defined: all the items are to be ordered in terms of size.
2	Makes partial series by trial and error.	He understands that the problem requires ordering all the sticks in terms of length, but he still has difficulty treating a stick as both longer than one stick and shorter than another: he makes partial series of sticks through successive comparisons of pairs of sticks.	The components of the problem are differentiated: each item must be treated as both larger than some items and smaller than others.
3	Makes complete series by trial and error.	He understands that each stick must be treated as both longer than one and shorter than another, but still cannot completely intercoordinate these relations mentally: he can make complete series by successive adjustments of pairs of sticks.	The components of the problem are integrated into a new higher-level scheme: all the items form a step-by-step series, and each item has its own specific place in the series.
4	Anticipates full series from the start.	He fully understands how to seriate because he can mentally intercoordinate the relations of "longer than" and "shorter than": he can anticipate the entire series from the start.	

It follows most obviously from the model that an adult solving a new seriation problem will show phases in the solution of that problem that are similar to the stages shown by a child in seriation, because the phases are hypothesized to be general phases of problem-solving for seriation. We are still working on this prediction, trying to find an appropriate task. Most seriation tasks are inappropriate, because adults are very good seriators with familiar materials. What is probably required is a task using unfamiliar, abstract materials.

It follows from the model that a child performing a seriation task will commonly show several successive phases as he repeats the task a few times. If, as the model assumes, the child is trying to solve the problem of seriation with this particular task, then as he is first acquainting himself with the task he will perform at an earlier phase. Several repetitions of the task will allow him to understand the problem more fully and so to progress to a later phase. (Piaget usually reports only the best that the child can do, the highest phase that he can reach, which is his stage for this task.) Karen Cohen and I are testing this prediction in a study in progress. Thus far we have tested about a dozen preschool children with the stick-seriation task, and preliminary analyses support the prediction. Most of the children advance about one phase over three to five trials with the same seriation task.

It follows from the model that, if a child is given one form of a seriation task and after he has performed it several times he is given a slightly easier form of the same task, he should advance about one phase with the easier form (unless of course he has already reached the final phase with the more difficult form). This prediction derives from the idea that, if the second form of the task is slightly easier, the problem-solving process should be facilitated. Karen Cohen and I are testing this prediction in the same study mentioned above. The initial sticks that we give the child to seriate differ by about 4 mm. between adjacent sticks. Ms. Cohen thought of a simple, neat way to make the task slightly easier: give the child sticks which differ by 8 mm. instead of 4 mm. Preliminary analyses support the prediction once again. Most children advance about one phase when they are given the easier sticks. For example, a child who performs at Phase 1 on the first trial with the difficult sticks will progress during later trials to Phase 2 with those same sticks; when he is given the easier sticks, his performance will advance one additional phase--to Phase 3.

In our research project we are now extending this study to several additional seriation tasks that develop during different periods of intellectual development. If the model is correct, the same predictions should apply to seriation achieved in any period.

In summary, I have suggested a model that treats Piaget's stages of "cognitive development" in tasks such as seriation as phases of

problem-solving. The model explains the occurrence of décalage in different versions of these tasks as different occurrences of the problem-solving process specific to seriation. The stage of performance shown for any specific seriation task is the result of an interaction of the child's general intellectual level and the characteristics of the specific version of seriation that he is faced with. That is, the stage is produced by the child's attempt to solve the problem posed by the task. The nature of this interaction is elaborated more fully in a book in preparation entitled Piaget, Learning, and Cognitive Development, and a model for depicting the child's general intellectual level is presented both in the book and in a paper to be presented in a symposium at the convention of the American Psychological Association this year in New Orleans (Fischer, 1974).

References

- Fischer, K. W. The hierarchy of cognitive development: Piaget systematized. Paper to be presented at the Annual Convention of the American Psychological Association, 1974, as part of the Symposium "Structural Psychology: Convergence toward a New Paradigm."
- Fischer, K. W. Piaget, learning, and cognitive development, in preparation.
- Greenfield, P., Saltzman, E., & Nelson, K. The development of rulebound strategies for manipulating seriated cups: A parallel between action and grammar. Cognitive Psychology, 1972, 3, 291-310.
- Inhelder, B., & Piaget, J. The growth of logical thinking from childhood to adolescence. New York: Basic Books, 1958.
- Piaget, J. The origins of intelligence in children. New York: International Universities Press, 1952.
- Piaget, J. Logique et équilibre dans les comportements du sujet. Études d'Épistémologie Génétique, 1957, 2, 27-118.
- Piaget, J. Six psychological studies. New York: Random House, 1967.
- Piaget, J. Piaget's theory. In P. H. Mussen (Ed.), Carmichael's manual of child psychology. Volume 1. Pp. 703-732. New York: Wiley, 1970.
- Piaget, J., & Inhelder, B. The psychology of the child. New York: Basic Books, 1969.
- Werner, H., & Kaplan, B. Symbol formation. New York: Wiley, 1963.

STRUCTURALISM IN PIAGET'S INTUITIVE STAGE OF DEVELOPMENT

Joan Batt
New School for Social Research

Introduction: Piaget's theory is very complex and there are many issues on which to focus. Therefore, it is possible, for example, to consider, the theory of the development of thinking in children, or, to view the work as a theoretical paradigm of modes of mental operations, or "thinking," in general. Kessen has said of Piaget, "Piaget is more a student of thinking than of children." (Kessen, 1962) Elkind has, also, noted that Piaget should be considered primarily as a genetic epistemologist rather than as a child psychologist. (Elkind, 1962) One of Piaget's major collaborators, Inhelder, has described the complexity of Piaget's work. She said, "Piaget is a zoologist by training, an epistemologist by vocation and a logician by method." (Inhelder, 1962) Piaget's theory of the growth of knowledge in the child is particularly concerned with the development of the individual as a genetic, self-evolving and self-regulating mental system which operates according to the principle of equilibrium. Cognitive development consists of having new experiences which are continuously disrupted and provide the impetus for developing more satisfactory, abstract ways of thinking. Structuralist concepts are important in Piaget's theory since the development of thinking can be seen as the continual breaking down of inadequate structures of knowing so that the mind develops more abstract, logical and mature structures.

There are 4 main stages in Piaget's theory of development. The intuitive stage is the second sub-stage of the second stage, the pre-operational stage between 2 and 7 years. The intuitive stage emerges from the former pre-operational sub-stage, that of egocentricism which occurs between 2 and 4 years. During the egocentric stage, the child is preoccupied with the beginning of language development and symbolic thinking. Also, the child is unable to understand a situation or idea from any viewpoint other than his own and "believes himself to be the center of the Universe." (Piaget, 1955)

By the end of the egocentric stage, around 4 years, the child becomes capable of simple intuitive thinking. i.e., he is beginning to have an intuitive grasp of adult, abstract thought. He is in the process of developing further representative symbolic activities, which manifest structures of knowing. However, even though the child has an intuitive grasp of logic, his thinking is still pre-logical, unidirectional, perceptual, and imaginal. Other characteristics which Piaget has applied to the child's intuitive thinking are: artificial, egocentric, nominally realistic, momentary, and phenomenalistic.

By age 6 or 7 years, the child develops articulated intuitive thinking and is now ready to enter the next major stage, that of concrete operational thinking. In the concrete operational stage the child can think logically, but only in specific, concrete situations.

The development of the child's thinking in the intuitive stage will be studied in two areas: those of logico-mathematical concepts and linguistic concepts. A structuralist interpretation will be emphasized.

Piaget's Structuralist Concepts and Theory. Piaget postulates that all psychological functioning has a structural aspect, i.e., that it relates to underlying structures and that in this way it deals with knowledge. Knowledge is defined as the behavioral manifestation of structures within the organism. (Piaget, 1968) At all age levels, structures exist which are the prototypes of logic and which by progressive equilibration lead to logical structures. Structure is, thus,

an important concept in Piaget's theories as are many others, such as equilibrium, stage, operation, assimilation and reversibility.

It is helpful to begin with a definition of structure as written by Piaget. He says: "Structure is a . . . system which presents the laws or properties of a totality seen as a system. These laws of totality are different from the laws or properties of the elements which comprise the system." (Piaget, 1968, p. 143) He, also, says that structure is a global term, like that of the Gestalt and not easily reduced to quantitative terms.'

In a recent scholarly work, Structuralism (1970), Piaget elaborates on three defining marks of structuralism as an epistemological method of inquiry into the nature of the growth of knowledge in children. These marks are: wholeness, transformation and self-regulation. By wholeness is meant that structures are not just aggregates, but organized wholes. Also, "the logical procedures or natural processes by which the whole is formed are primary." That structures are systems of transformations reveals their dynamic quality. As systems, they are simultaneously structuring and structured. The self-regulatory aspect of structures refers to their rhythm, regulation and operation.

Piaget discusses examples of structuralism found in fields other than that of child psychology. Some fields are: mathematics, the physical and biological sciences, the social sciences, language and philosophy.

In order to further clarify the meaning of the term, structure, three Piagetian scholars will be presented with definitions. Elkind defines Piaget's use of structure as a "mental system or totality whose principles of activity are different from those of the parts which make it up." He, also, compares structures to schemes. (Elkind, 1968) Furth says that structure refers to the "general form, the interrelatedness of parts within an organized totality." Structure can be used interchangeably with terms such as organization, system, form and coordination. (Furth, 1970) Phillips gives us a definition of structure by contrasting it with function. He says: "The basic underlying idea is that functions remain invariant, but that structures change systematically as the child develops. This change in structure is development." (Phillips, 1969) Structure and function, thus work together. The biologically inherited modes of assimilation and accommodation are functional, not structural.

Intuitive Thinking: Logico-Mathematical Development: A structuralist view of the child's development during the intuitive stage will be given. Empirical studies done by Piaget and his followers will be cited. First, logico-mathematical concepts will be studied, followed by language development. Among the many types of experiments which Piaget and his followers have done on the child's intuitive grasp of logico-mathematical concepts are: order, seriation, classification, conservation of matter, weight and area, space, time, causality, the invariant horizontal, problem-solving and explanations of natural forces. The following concepts will be discussed in this paper: order, seriation, classification, conservation of matter and the invariant horizontal.

Concerning the development of the conceptual structures of order and classification, Piaget says: "The order relations, for example, which on the sensori-motor plane were altogether immersed in the sensori-motor schema now become dissociated and give rise to a specific activity of 'ranking' or 'ordering.'" The same is true for other concepts.

The seriation of sticks, A, B, C, etc. of different lengths placed side by side, also, reveals the child's thinking. Children at 5 years are able to construct uncoordinated pairs only and "he gropes his way from

step to step." (Piaget, 1960) Not until the operational level is he able to understand the logical principles involved, such as transitivity of length and classification.

Another example of early intuitive thinking, involving the concept of classification is shown in the following experiment: A child age 3 or 4 years, when given a group of heterogeneous objects to classify groups them along no, or very general, classificatory principles. The groups are made on the basis of momentary perceptual impressions such as putting some of the circles together, but then arbitrarily adding few squares. He is unable to deal with the quantifiers, "one", "some" and "all" or the logical principles of class inclusion-exclusion. It is not until he is 5 or 6 years that he begins to think logically.

Directly applying Piaget's structuralist concepts to the above experiments, we can say that the whole, transforming and self-regulating principles that determine the child's thinking are at this stage pre-logical. The knowledge which the child has is intuitive as are the cognitive structures, themselves.

Piaget has given a detailed analysis of the child's progression from simple to articulated thinking in the development of the concept of conservation of matter. He analyses the progression into 4 sub-stages. (Piaget, 1963)

The following experiment is used as an example: The child is shown two identical balls of clay which he judges to be equal. One ball is then elongated by the Experimenter so that it resembles a sausage. The child then immediately thinks that the longer piece of clay contains more clay. The early intuitive child around 4 years does not use reason based on an understanding of the logical principle of compensatory relationships. He only attends to the dominant perceptual characteristic, in this case, length. According to Piaget, this is the first of four strategies which the child uses in grasping the concept of conservation of matter.

During the second phase of the experiment, the child realizes that his answer was wrong. He sees that the piece of clay is thinner, as well as longer.

During the third phase, the child will reason about the two dimensions of length and breadth, at once. First, "he will oscillate between the two." He will say, "I don't know. It's more because it's longer. . . , no, it's thinner, so it's less." Gradually, he will discover the principle of compensatory relationships. It is at this point that the child is ready to understand the concept in a concrete operational way.

The child in the fourth phase now reasons in terms of the the two transformations and has grasped the principle of compensation. The structural understanding of the concept is now reversible, extra-temporal and equilibrated in a concrete operational way.

This description of the development of the concept of the conservation of matter in the intuitive child supports Piaget's view that "structures are built up from the simple to the complex, little by little, ad infinitum" and that, at all age levels structures exist . . . which by progressive equilibration lead to logico-mathematical structures." (Piaget, 1968) Piaget emphasizes the self-constructivistic approach toward cognitive growth. Rather, than being either innate or essentially pre-formed, structures are continually in the process of forming themselves according to logical laws of transformation. (Piaget, 1968)

Intuitive Thinking: Language Development. In discussing language development from a structural viewpoint, Piaget discusses various topics, such as, the epistemological basis of language, the relationship

between thinking and language, including use of linguistic syntax, degree of communication and the ontological meaning of language. Each of these topics will be briefly discussed.

In Structuralism, Piaget gives an account of the epistemological basis of language. He discusses such topics as: synchronic structuralism, transformational structuralism and the relationship between ontogenesis and phylogenesis, and the relationship between linguistic and logical structures. In all cases, Piaget traces language development using an equilibration model. Language is neither the result only of external social formations nor of innate ideas, but, rather, of their interaction in a continually developing cognitive system.

Piaget points to the interrelatedness of language and intelligence. Each reflects the other. The language of a child can be analysed in terms of his use of grammar and syntax as well as degree of communication and understanding of ontological principles.

Piaget describes experiments by H. Sinclair de Zwart (1967) on the use of quantitative adjectives as novel and precise. She found that pre-operational children rarely used any adjectives except scalar ones, i.e., "That one is big" and "That one is little" while concrete operational children used vector, comparative vocabulary, such as, "That one is bigger than the other." The older group also used binary connections.

Piaget has also noted that the pre-operational child has no systematic basis for using grammar. Pronouns, personal and demonstrative adjectives, such as "he," "she," or "that" are used "right and left without any indication of what they refer to." Also, "words are much nearer to action and movement than for the adult. . . the child is impelled, even when he is alone, to speak as he acts. . ." (Piaget, 1955)

Piaget, also, discusses the social aspect of language or degree of communication. Language, like other abilities and forms of communication is a "group institution." Like other kinds of social behavior, its rules are imposed on individuals. He explains: "The syntax and semantics of a language yield a set of rules to which any individual speaking that language must submit, not only when he wants to express his thoughts to others, but even when he expresses it 'internally.'" (Piaget, 1970, p.75)

The child's emergence from egocentric speech, involving repetition, the monologue and the collective monologue to socialized speech is traced by Piaget in three stages. (Piaget, 1955) The stages will be described.

During the first stage between 3 and 5 years, the child is capable of speaking primarily only in a monologue. The child speaks for and by himself and his words have no social function.

The second stage between 5 and 7 years is known as that of the collective monologue. Now, the child talks in the presence of others and hears others, but is not yet conversing or sharing ideas. This stage is the most social of the egocentric varieties of child language, since to the pleasure of talking it adds that of soliloquizing before others and of interesting them in oneself.

The third stage is called that of collaboration in abstract thought. Now, around 8 or 9 years, the child is able to hold an actual conversation with others in which ideas are shared. His thinking is socialized at the concrete operational level. Not until adolescence will the individual become capable of abstract thinking and mature social interaction.

Other important ontological characteristics of the child's thinking during the intuitive stage which are revealed in his language are animism, artificialism and syncretism. An example of animistic thinking is shown in two statements concerning the moon. When the child remarks that "the moon follows him when he goes out for a walk," he is revealing that he is unable to differentiate between logical implication and physical causation. Again, this is true when the child says, "the moon doesn't fall down because its very high up."

critical remarks. Any questions have arisen as to the practical value and theoretical consistency of Piaget's theory. A large body of general criticism seems to stem from the fact that Piaget gives not enough explicit attention to either environmental or cultural influences on the child. Elkind (Elkind, 1963), for example, has found shifts in age-norms as a function of socio-cultural status. Hunt (Hunt, 1961) has questioned the value of a rich environmental background on the growth of intelligence. A major debate has been going on for several years between two Piagetian experimenters (Braine, 1964; Smedslund, 1965). Each has done experiments on the development of the concept of transitivity of length, raising questions as to the constancy of age-norms and describing cognitive processes.

An intensive summary of experiments done by Piaget and his followers on the acquisition of logical operations between 3 and 10 years has been compiled by Beilin (Beilin, 1971). Among the many topics considered are: identity, cognitive conflict, reversibility, learning, verbal training methods, multiple-training strategies and classification.

In general, it seems that there are many problematic issues involved in interpreting Piaget's work and finding experimental procedures for testing it. It appears that a more flexible attitude is required, one which utilizes Piaget's principles and concepts, but, is not restricted to the bounds of the work already done.

Bibliography

- Beilin, Harry. The training and acquisition of logical operations. In Piagetian Cognitive-Developmental Research and Mathematics Education. Ed. by E. Roszkopf, et.al., Nat'l. Council of Teachers of Mathematics, Wash. D. C., 1971.
- Braine, Martin. Development of a grasp of transitivity of length: A reply to Smedslund. Child Dev't., 1964, 35, 300-310.
- Garth, Hans. Piaget for Teachers. Englewood-Cliffs, N.J.:Prentice-Hall, 1960.
- Hunt, John McV. Intelligence and Experience. N.Y.:Ronald Press, 1961.
- Wheeler, Barbel. Some aspects of Piaget's genetic approach to cognition. In Thought in the Young Child, Ed. by W. Kessen and C. Muhlman, N.Y.:D. Loro r., 27,2, 1962, 19-40.
- Williams, John. The Origins of Intellect: Piaget's Theory. San Francisco, C.A.:Freeman and Co., 1969.
- Piaget, Jean. The Language and Thought of the Child. N.Y.:World Pub. Co., 1954. La g n se des structures logiques  l mentaires: Classifications et s riations. Neuchatel:Delachaux and Niestle, 1959. Psychology of Intelligence, N.Y.:Littlefield, Adams, 1963. The Child's Conception of the World; N.Y.:Littlefield, Adams, 1963.
- _____. Six Psychological Studies (Ed. by David Elkind), N.Y.:Random House, 1968.
- _____. The Science of Education and the Psychology of the Child. N.Y.: Viking Press, 1969.
- _____. Structuralism. N.Y.:Harper and Row, 1970.
- Smedslund, Jan. The development of transitivity of length: A comment on Braine's reply. Child Dev't., 1965, 36, 577-580.
- de Waart, M. Sinclair. Acquisition du langage et d veloppement de la pens e. Paris: Dunod, 1967.

DIAGNOSTIC-TEST-BASED PRESCRIPTIONS IN IPI MATHEMATICS

Anna Olga Graeber

Research for Better Schools

The Individually Prescribed Instruction in Mathematics program includes diagnostic tests designed to determine the units of the program and the objectives within these units appropriate for an individual learner. At present, diagnosis of a pupil's competence of a specific component of a given objective is left to the classroom teacher. In IPI, the process of assigning specific self-instructional materials based on analysis of a pupil's work on diagnostic tests is called prescription writing.

Past studies indicate that IPI teachers need and express a desire for help in prescription writing. Cox (1966), in a report on the first two years of the use of IPI, noted that "a major concern in the evaluation of the IPI system is the process of assigning materials which, within IPI, becomes the development of pupil prescriptions." Cox analyzed the prescriptions written for pupils in grades one to five in the C-Fractions unit. His analysis led to the conclusion that although patterns of prescriptions do appear, "these patterns seem to be models associated with certain teachers rather than with certain types of pupils."

Bolvin (1967) evaluated teacher prescriptions written during a school year at Oakleaf Elementary School in the Baldwin-Whitehall School District near Pittsburgh, Pennsylvania. His analysis of prescriptions led, as did the study by Cox, to the identification of two predominant types of prescriptions. The first type of prescription included materials that introduced the skill as well as almost all the materials available for the skill. Only after the pupil had completed many materials was the curriculum embedded test prescribed. The second type of prescription included only a few select materials. If the pupil failed to achieve mastery criteria on the curriculum embedded test, a few more materials were prescribed. Two of Bolvin's conclusions are of particular interest:

1. The teacher, having only about two minutes a day to write each child's prescription, needed the necessary information updated regularly and easily accessible to her.
2. Although prescriptions varied from child to child, the rationale behind this variation did not follow the child's learning needs.

If the teacher's analysis of a pupil's work, examination of available instructional materials, and actual assignment of materials are done within two minutes, the second conclusion, while disturbing, is understandable.

In 1967 Research for Better Schools, a regional laboratory of the United States Office of Education that was instrumental in the early dissemination of IPI, held a conference for teachers and administrators using IPI Mathematics to elicit suggestions for improving the instructional system and the teacher training program. The conference participants identified diagnosis as one process with which they wanted additional help and direction.

Holste (1972) investigated whether the type of prescription the teacher used affected the rate, the retention, or the level of on-task behavior of children during mathematics instruction. Holste found that the least effective prescription, both in terms of rate of progress and of retention, were those written by the teachers who used the procedure advocated by the authors of IPI.

The present study was planned to investigate whether a diagnostic instrument, designed to test the components of a single objective, could be used effectively as a basis for prescription writing. The investigator used the method Resnick (1967) calls "component analysis" to identify components of the six objectives comprising the C-level Fractions unit. The procedure begins by examining a specific instructional objective and listing answers to the question: To perform this behavior, what prerequisite or component behavior must the pupil be able to perform? For each behavior listed the same question is repeated; thus generating a hierarchy of objectives based on prerequisites. Resnick employed this method not for sequencing skills but rather for determining the subskills, or components, of each C-Fractions skill. Once the components of each objective had been determined, the pages in the existing self-instructional booklet that provided instruction for each component were identified. The page numbers were listed on the diagnostic instrument to match corresponding items. Thus, if a pupil missed any of the items on the diagnostic instrument, pages appropriate for the pupil's prescription could readily be identified.

In a school using the IPI Mathematics program, pupils at each of four grade levels, three through six, identified as likely to work in the C-Fractions unit during the school year were ranked by IQ and grouped to form matched pairs. One pupil in each pair was randomly assigned to the experimental group and received his prescriptions in the C-Fractions unit from the results of the diagnostic-prescriptive instrument. The other pupil in each matched pair received his prescriptions in this unit from his teacher as usual. Throughout the school year, data were collected to test the following hypotheses:

Pupils receiving diagnostic-test-based prescriptions spend less time working on a skill before mastery than do pupils receiving the usual teacher prescriptions.

Pupils receiving diagnostic-test-based prescriptions complete fewer instructional pages in a skill booklet before mastery of a skill than do pupils receiving the usual teacher prescriptions.

Pupils receiving diagnostic-test-based prescriptions require fewer curriculum embedded tests before mastery of a skill than do pupils receiving the usual teacher prescriptions.

For each hypothesis, a one-tailed sign test was used to determine whether the number of within-pair differences in the hypothesized direction was a significant number of the signed within-pair differences. The sign tests indicated that the amount of time taken and the number of instructional pages

completed by subjects receiving the two types of prescriptions differed significantly and favored the pupils receiving diagnostic-test-based prescriptions. Although pupils receiving diagnostic-test-based prescriptions took fewer curriculum embedded tests and fewer posttests than did pupils receiving prescriptions from their teachers, the difference in the number of tests taken was not significant.

The prescriptions of the eight teachers who had prescribed for control group pupils were analyzed. The teachers tended to prescribe all, or almost all, of the available instructional pages before prescribing a curriculum embedded test. Although the majority of prescriptions reviewed prescribed pages related to each component for which the subject had not demonstrated mastery at the time of the pretest, the inclusion of pages related to specific weaknesses is likely to have been a consequence of the length of the prescriptions rather than a result of accurate diagnosis.

The study suggests that if the proposed benefits of IPI are to be realized, teachers need assistance in prescription writing -- either in the form of diagnostic instruments, more extensive teacher training, or both.

REFERENCES

- Bolvin, J.O. Evaluating teacher functions. Working Paper 17, Learning Research and Development Center, Pittsburgh, Pennsylvania, 1967.
- Cox, R.C., et. al. A description and interim evaluation report concerning the first two years of the individually prescribed instruction project. Unpublished manuscript, Learning Research and Development Center, Pittsburgh, Pennsylvania, 1966.
- Holste, D.E. The effect of different prescriptions used by teachers in an IPI Mathematics program. Unpublished doctoral dissertation, University of Illinois at Urbana-Champaign, 1972.
- Resnick, L. Design of an early learning curriculum. Working Paper 16, Learning Research and Development Center, Pittsburgh, Pennsylvania, 1967.

INDIVIDUALIZED MIDDLE MATHEMATICS --
EXTENDING AND IMPROVING THE IPI MODEL

George F. Lowerre

Research for Better Schools

The idea of making the school curriculum fit the abilities of the students is not a new one. Even Horace Mann, the arch-proponent of the common school, recognized the importance of suiting the instruction to the student when he wrote in his Fourth Annual Report (1840)

Lessons, as far as it is possible, would be so adjusted to the capacity of the scholar, that there should be no failure in a recitation, not occasioned by culpable neglect.

Mann was really concerned with the capacity of an entire class of students. But Anderson (1962) reports that as early as 1888 a plan was in use in Pueblo, Colorado that emphasized individual work and individual progress. A continuing interest in the idea of individualization is documented by the themes of the 1925 and 1962 yearbooks of the National Society for the Study of Education -- both were concerned with individualization. Shane (1962) lists a total of 35 plans for individualization that had been tried out by 1962. Most of these plans, however, were organizational models by which individualization could be enhanced in a traditional, teacher-taught classroom.

IPI Mathematics applied to individualization in mathematics instruction the concepts and technology that had been developed by the early nineteen sixties -- behaviorally specified objectives, criterion referenced testing, and programmed instruction. Originally written at the Learning Research and Development Center (LRDC) at the University of Pittsburgh, IPI Mathematics was first tried out in the Baldwin-Whitehall School District in the school year 1965-66 and at five Research for Better Schools (RBS) demonstration schools in 1967. For 1968, the materials were revised by the staffs at both LRDC and RBS with assistance from the staff at Appleton-Century-Crofts. In 1970-71, a complete revision was made by the staff at RBS; this revision became available commercially in 1972.

With the completion of the commercial version of IPI Mathematics, two avenues were open for further, related work in the field of individualized mathematics instruction. One approach involved production of a new individualized elementary school mathematics program, with new content and objectives, a new organizational structure, and new learning activities that provide alternate paths to the attainment of each objective. Development of this program, called Individualized Mathematics, has been in the hands of LRDC. Parts of this program have been tried out at two schools in Pittsburgh and two schools in Philadelphia. Currently, further development of this program has been suspended for lack of funds.

The second approach involves extending the concepts and techniques of IPI Mathematics upward to the junior high school, improving them wherever possible. This program, called Individualized Middle Mathematics (IMM), is currently being developed by RBS.

The goal of the IMM development program is as follows: To develop and evaluate an individualized, behaviorally specified junior high school mathematics program that will

1. be usable by any child who has completed a typical K-6 mathematics program, whether individualized or not.
2. be compatible with IPI Mathematics, i.e., permit use in the same classroom.
3. provide freedom of choice for students.
4. provide help in making prescriptions.

It is not the objective of the Individualized Middle Mathematics program to create new junior high school mathematics content. The objective is to teach the content that is typical of seventh and eighth grade mathematics in 1973 - 74, and to organize that content as part of a system providing for individualization on the basis of entry point into the program, rate of learning, and choice of order in which the objectives are studied. The content has been selected after an examination of nine current junior high school mathematics text series and two standardized tests. The content was then stated in terms of behaviorally specified objectives. Each objective has been stated in two forms: formal and normative. The formal form of each objective states the objective with as much precision as the English language allows; the formal objective statements are intended for the developers, reviewers, and any others interested in an exact (if sometimes lengthy) statement of what is required to meet each objective. The normative form of each objective states the essence of the objective in brief, easily understood language. It is this form of each objective that is for use by teachers and students, and which will be printed on all student materials. One feature of the IPI Mathematics objectives has been continued -- each objective specifies the limits on the problems that the student is required to do to demonstrate mastery.

The number of entry skills has been kept to a bare minimum -- essentially they include the ability to compute with whole numbers and a knowledge of the basic concepts of fractions.

In order that the IMM materials can be used in the same classroom with IPI materials, the objectives have been organized into areas and levels, similar to the areas and levels in IPI. There are seven areas: Foundations, Integers, Rationals and Reals, Geometry and Measurement, Probability and Statistics, Equations and Inequalities, and Applications. There are four levels, labeled K, L, M, N. Levels K and L are intended to cover typical seventh grade mathematics content; levels M and N are intended to cover eighth grade mathematics.

The Individualized Middle Mathematics program includes all of the diagnostic tests that are part of IPI: placement tests (one per level), pretests (one per objective), curriculum-embedded tests (two per self-instructional booklet), and posttests (one per unit).^{*} All of these tests are criterion-referenced measures. There is no use of normative tests.

^{*} A unit consists of all of the objectives at one level in one area.

As in IPI, the principal instructional materials in IMM are self-instructional booklets. However, unlike IPI, some skill booklets have accompanying packets of manipulatives, and there are provisions for group activities.

In order to allow student choices in selecting the order in which the objectives are studied, the objectives are not arranged in a fixed sequence. The objectives in IPI are arranged in a linear order, or continuum, in which the sequencing cannot be altered. The objectives in IMM are not arranged in such a fixed sequence; instead they are arranged in a two dimensional set, where only one of the dimensions is linearly ordered. I have called such a set a bicontinuum, which I have defined as follows:

BICONTINUUM: A two-dimensional set in which at least one of the dimensions is arranged in a fixed sequence.

Within each topic area, the levels are ordered. That is, a student must demonstrate mastery of the objectives for Rationals and Reals at level K before he can go on to the objectives for Rationals and Reals at level L. Also, he must demonstrate mastery of objective number 1 in any unit before proceeding to objective number 2. However, a student working at any level will have considerable freedom to choose the area he wishes to study. The only limitations are as follows:

At level K (only) the student must demonstrate mastery of the skills in K-Foundations before working in any other unit.

At all levels, the student must demonstrate mastery of the skills in all other units at that level before working in the Equations and Inequalities unit and the Applications unit. He must complete the Equations and Inequalities unit before working in the Applications unit.

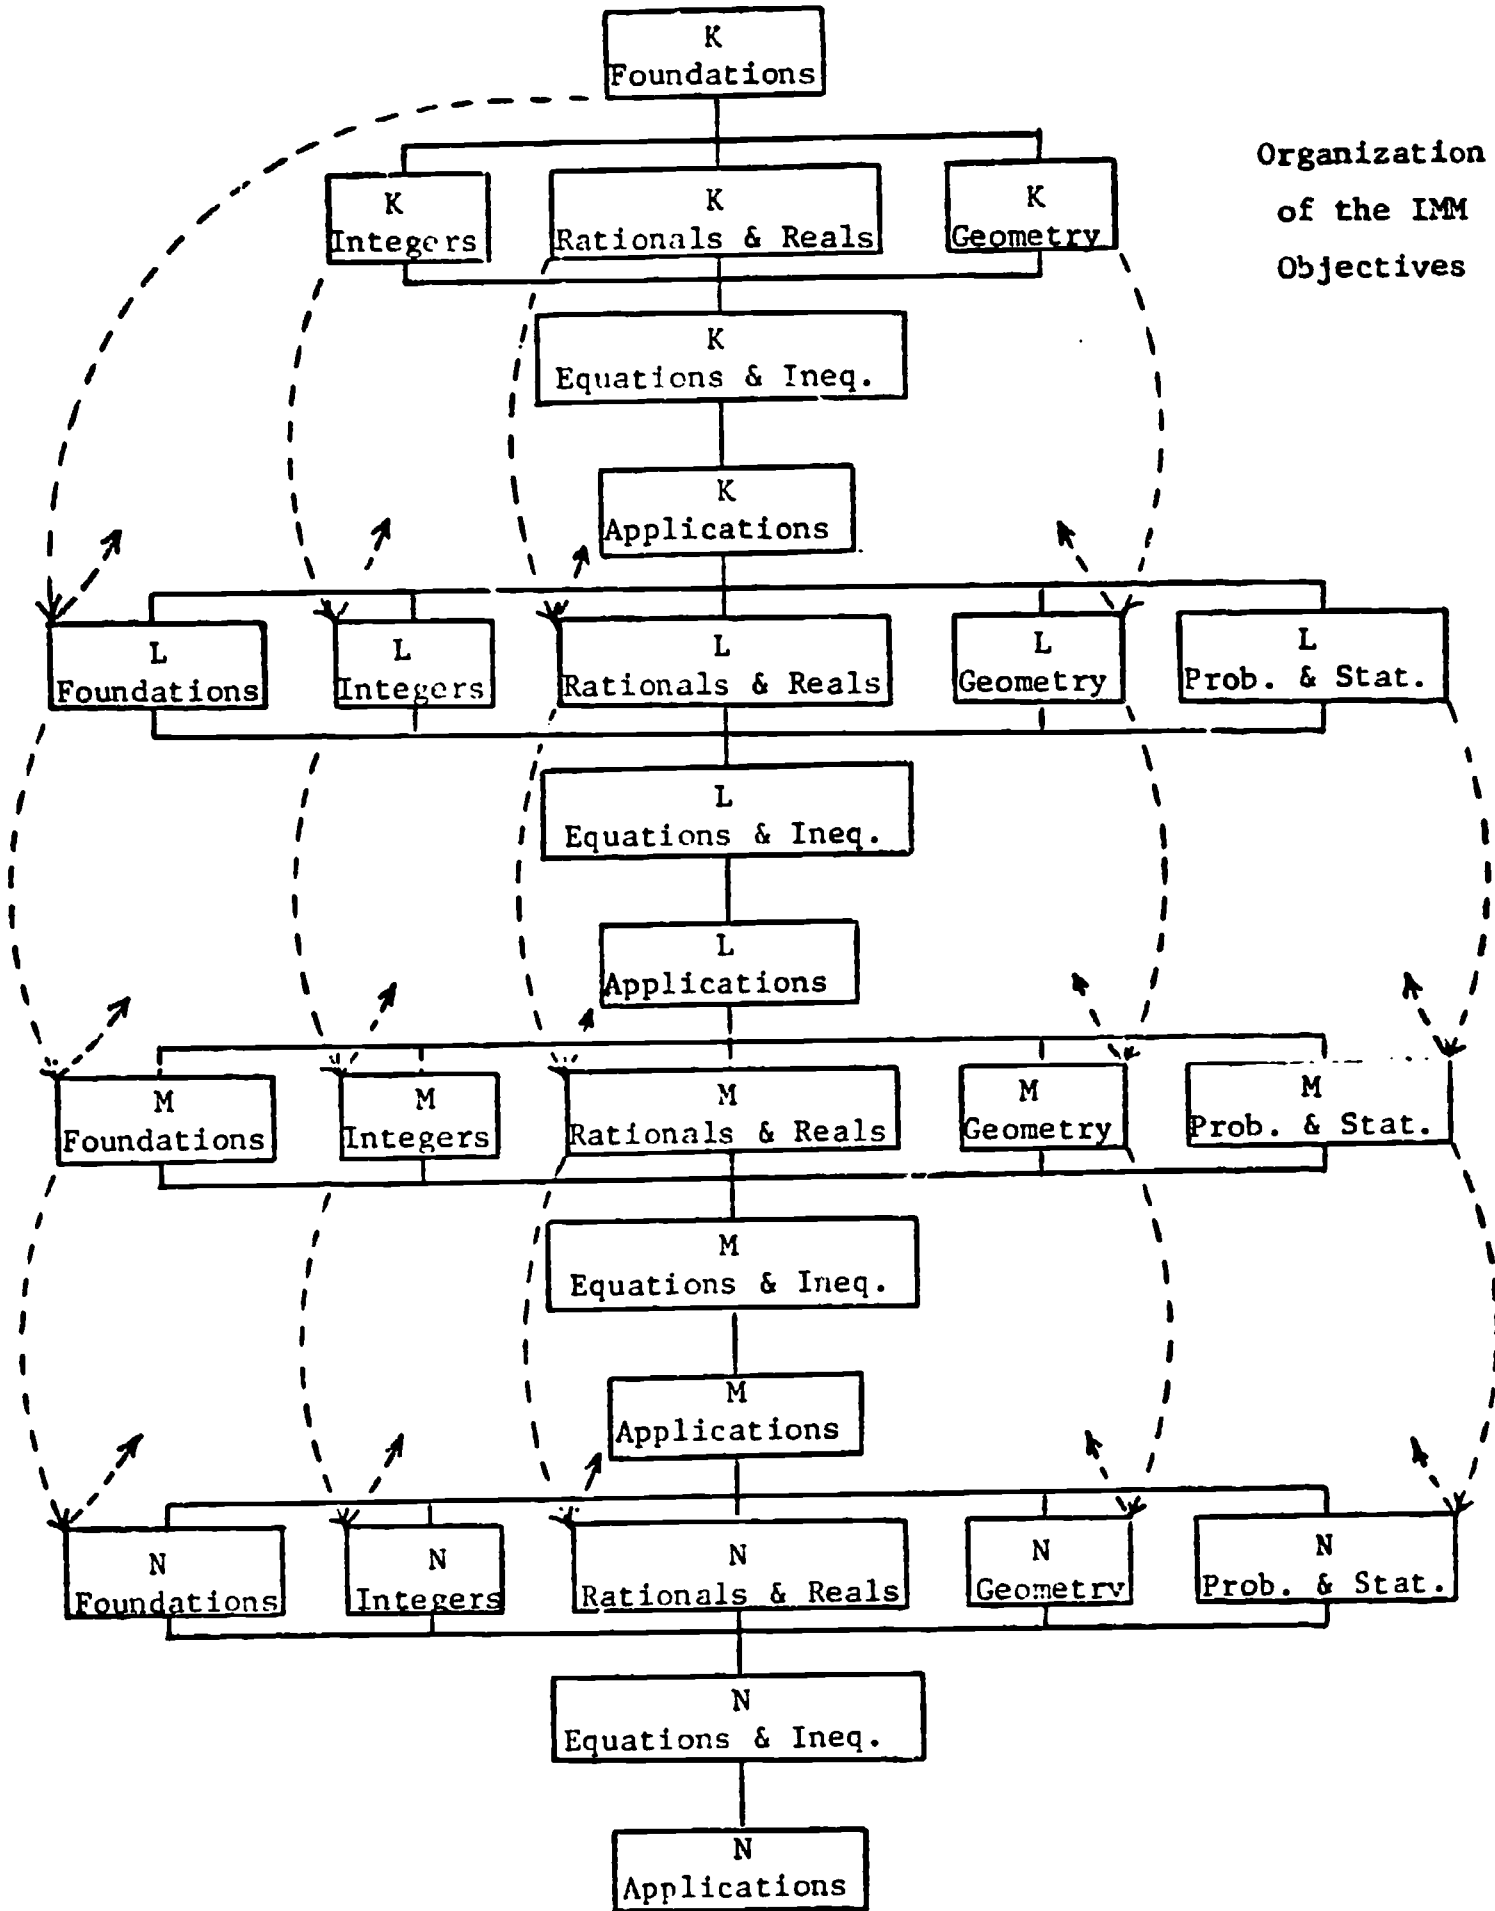
Except for these limitations, the student working at any level is free to choose the areas he wishes to study in any order.

An additional degree of freedom is built into the Bicontinuum. After mastering all of the objectives in any unit the student may, if he wishes, study the objectives at the next level in that area before completing the objectives in the other areas at the earlier level. For example, after demonstrating mastery of all of the objectives in K-Geometry, the student may study the objectives in L-Geometry before completing the rest of level K. He may not, however, progress to M-Geometry until he has mastered all of the units at level K.

Perhaps only a few students will choose the option of continuing to another level in one topic area before mastering all of the objectives in the other areas at the given level. However, this option (or second degree of freedom) is included for those students who may become especially interested in one particular topic area and wish to pursue it further before studying the other topics at the given level.

The organization of the proposed junior high bicontinuum is shown on the next page. Any units which lie on the same horizontal line can be studied in any

Organization of the IMM Objectives



order. The option of continuing to the next level in an area is indicated by the dotted lines. It must be emphasized that the student does not have a choice of whether or not to study any particular objective. His choices concern the order in which he studies the objectives.

To help the teacher prescribe the appropriate pages within each self-instructional booklet, each such booklet will contain two diagnostic-prescriptive instruments (D-P-I's). In the IMM program these are called Planning Aheads.

The Planning Ahead, which is taken before the student begins work in each booklet, is designed to identify the subskills within the booklet that the student has not mastered. Items on the Planning Ahead parallel the instructional sequence in the booklet and are scored on an all-or-none basis; if the student misses any one of the problems or any part of a problem related to a given subskill, it is assumed that he has not mastered that subskill and the corresponding pages should be prescribed. A second Planning Ahead will follow the Check Up in each self-instructional booklet; this will be administered to students who did not demonstrate mastery on the first Check Up in the booklet.

Graeber (1974) found that the use of diagnostic-prescriptive instruments with IPI booklets decreased both the number of booklet pages and the length of time required for mastery of each objective. Inclusion of diagnostic-prescriptive instruments (Planning Aheads) in IMM booklets should eliminate the problem of poor prescriptions and help shorten the length of time students work on each booklet. This is particularly important, since most IMM booklets will be longer than the typical IPI booklet. The teacher, however, remains the ultimate prescriber. He/she can give a prescription different from the one indicated by the Planning Ahead whenever it seems appropriate.

In summary, the Individualized Middle Mathematics program is designed to build on IPI Mathematics, using the same basic model. It should be easy to use the two programs simultaneously in the same classroom or to use IMM with junior high students who have never used IPI. At the same time, IMM has improved the IPI model by including student options and diagnostic-prescriptive instruments in the program.

REFERENCES

- Anderson, R.H. Organizing groups for instruction. Individualizing Instruction, Sixty-first Yearbook of the National Society for the Study of Education, Part I. Chicago: University of Chicago Press, 1962.
- Graeber, A.O. Diagnostic-test based prescriptions in individually prescribed instruction in mathematics. Unpublished doctoral dissertation, Teachers College, Columbia University, 1974.
- Mann, H. Fourth annual report (1840). In L.A. Cremin (Ed.), The Republic and the School Horace Mann on the Education of Free Men. New York: Teachers College, Columbia University, 1957.
- Shane, H.G. The school and individual differences. Individualizing Instruction, Sixty-first Yearbook of the National Society for the Study of Education, Part I. Chicago: University of Chicago Press, 1962.

ACCEPTABILITY AND CHOMSKY'S GRAMMATICAL HIERARCHY

Wayne P. Hresko and D. Kim Reid
Temple University

Chomsky (1965) postulates that in a given language a native speaker is capable of perceiving a hierarchy of grammatical violations. In brief, this hierarchy is composed of three structural rules comprising increasing levels of violation (Table 1). Chomsky states that "deviance would be greatest if the item substituted for frighten is a non-verb [lexical-rule violation], less

Table 1

An Example of Chomsky's Hierarchy

1. Sincerity may frighten the boy.
2. Sincerity may virtue the boy.
3. Sincerity may elapse the boy.
4. Sincerity may admire the boy.

great if it is a verb but a non-transitive verb [subcategorical-rule violation], and still less great if it is a transitive verb that does not take an abstract subject [selection-restriction rule violation] [Chomsky, 1965, p. 152]."

Chomsky's theory explicitly refers to what the speaker actually knows, not what he may report about his knowledge (Chomsky, 1965). When speaking, for example, people structure their language grammatically, but may be unable to report the basis for their structuring, that is, the rules involved. People have a tendency to speak in accordance with what is acceptable to them. According to Chomsky, "The notion 'acceptable' is not to be confused with 'grammatical'. Acceptability is a concept that belongs to the study of performance, whereas grammaticalness belongs to the study of competence.... Grammaticalness is only one of the many factors that interact to determine acceptability [1965, p. 11]."

Empirical investigation of the psychological reality of this grammatical hierarchy requires the use of performance-based models to infer levels of competence. Downey and Hakes (1968) and Moore (1972) judged subjects' perceptions of grammatical deviance by asking them to rank sentences on the basis of which appeared more acceptable. Moore also measured subjects' reaction time in labeling sentences as acceptable or unacceptable. Although the hierarchies as defined by Chomsky apply only to the verb position, Moore, in order to investigate the generalizability of the hierarchies to other sentence components, expanded the hierarchy to include violations in subject and object positions. In Moore's classification scheme, violations which appear in subject, verb, and object positions are labeled as lexical. Subcategorical violations are those which appear in the verb position. When a comparable violation appears in the subject or object position, it is termed as early selection restriction. Selection restriction in the verb position is called late selection restriction in subject and object positions.

None of the above studies indicated support for Chomsky's proposed hierarchies. The question arises whether the hierarchy does not exist or whether the rating of acceptability so confounds the grammaticality factor as to obscure it.

There is evidence that the hierarchy is psychologically real: of the studies undertaken to examine Chomsky's position, those specifically investigating grammaticality supported Chomsky's theory of hierarchies (Danks & Glucksberg, 1970, 1971; Stoltz, 1969). When the subject is specifically asked to rate according to grammaticality, perhaps the experimenter is focusing the subject's attention on some internal model of grammaticalness. When acceptability, however, is the basis for rating, subjects are free to use any combination of factors they may choose. Recent studies which have investigated the processing of sentences have evidenced the influence of semantic content on sentence processing (Bransford, Barclay, & Franks, 1972; Bransford & Franks, 1971; Sachs, 1967; James, Thompson, & Baldwin, 1973). The results of these studies reflect the notion that syntactic structure is neither the only nor the most important factor in processing sentences. Rather, the semantic structure of the sentence is of prime importance. Logically, a sentence which is ungrammatical but makes sense would be judged more acceptable than a sentence containing the same level of grammatical violation but does not make sense.

A sentence's acceptability would depend not only on its grammatical structure but also on its meaningfulness to the subject. No empirical verification, however, exists. A multidimensional scaling technique was, therefore, used to examine the dimensions underlying judgments of acceptability.

Method

Subjects

Subjects were 175 graduate students in the Department of Special Education, Temple University. Students were screened to exclude those who had taken coursework in the psychology of language. Of the 175 protocols, 14 were excluded from the analysis because of patterning of responses (all ones or nines circled).

Procedures

Nine sentences (Table 2) were selected from those used by Moore (1972) in which every level of violation of Chomsky's hierarchy appeared in subject, object, and verb position.

Each subject received a test booklet containing all possible pairs (36) of stimulus sentences. For each pair of sentences, in fifty-percent of the test booklets sentence one appeared above sentence two and in the remaining fifty-percent below it. To further control for order effects, the sequence of presentation of the pairs was randomized.

The instructions were:

The following is an experiment dealing with the psychology of language. We are interested in your perception of the acceptability of various sentences. By acceptable, we refer to sentences which may occur in normal, every day usage in standard English. At the top of each page of your test booklet are printed two sentences. You are asked to judge how similar each pair of sentences is with respect to their acceptability. Some sentences will appear more acceptable than others. This is a subjective task. There are no right or wrong answers. Participation in this task is not required by your instructor. If you do not wish to participate, either do not take a test booklet or place your test booklet face down on the desk. Your performance has no bearing on grading in this course. If you do participate, please answer all questions.

Table 2

Stimulus Sentences

Violation Location	Example	Level of Violation
Verb	1. Smart voters uncle honest politicians.	lexical violation
	2. Noisy dogs growl night animals.	subcategorical violation
	3. Catchy slogans believe unwary citizens.	selection restriction violation
Subject	4. Modern wanders improve factory efficiency.	lexical violation
	5. Sensible ideas distrust public officials.	early selection restriction violation
	6. Nosey ditches annoy suburb dwellers.	late selection restriction violation
Object	7. Large factories utilize efficient hesitates.	lexical violation
	8. Big corporations appoint many machines.	early selection restriction violation
	9. Factory foreman appreciate eager tools.	late selection restriction violation

The subjects indicated the degree of similarity for each pair of sentences with respect to acceptability by circling one value on a scale from one to nine, with nine indicating a high degree of similarity.

For clarity, consider that subjects are asked to compare cups of tea according to preference. The bases upon which they make their judgments may be the sweetness and strength of the tea. Although the cups of tea may also vary in temperature, subjects' ratings may not be affected by that dimension. Multidimensional scaling is a method by which a researcher may determine which dimensions underly subjects' judgments. The interested reader is referred to Green and Carmone (1972), and Shepard, Romney, and Nerlove (1972), for information concerning theory, methodology, and application.

The means of subjects' ratings of the similarities between pairs of sentences were scaled using the Kruskal-Shepard method (Kruskal, 1964a, 1964b). This scaling procedure gave multidimensional scales based on the rank orders of the perceived similarities among stimulus pairs. Sets of scale values were obtained for solutions with from one to five dimensions. For each solution, the operation was iterated until minimum stress was achieved for a least-squares fit using Kruskal's secondary scaling option. Kruskal's primary approach to ties was employed. Stress may be thought of as analogous to the amount of variance unaccounted for. The use of the term dimension is consistent with its use in multivariate behavioral research.

Prior to data collection, an external scale was generated to test the hypothesized grammaticality dimension. Tie procedures for rank order data were used to define Chomsky's hierarchy of grammaticality. Stimuli at each level of violation were considered tied. In this manner, subject, verb, and object violations in the (1) lexical position, (2) early selection restriction and subcategorical position, and (3) late selection restriction and verb selection restriction position were considered tied.

Results and Discussion

Stress values for one to five dimensions are shown in Figure 1. Both an acceptable level of stress (stress = .0827) and an elbow occurred at four dimensions (Kruskal, 1964a, 1964b). In accordance with the consideration of multidimensional scaling--visualizability, interpretability, and parsimony (Shepard, Romney, & Nerlove, 1972)--only pairs of dimensions occurring in three-space are plotted in Figures 2 through 4. The four dimensional solution was uninterpretable.

Stimuli cluster along dimension one according to meaningfulness. To the right of the configuration (See Figure 2) appear stimuli 5, 3, and 1: sentences having political overtones. In the center of the configuration occur stimuli 7, 4, 9, and 8: sentences containing references to industry. Stimuli 2 and 6, which are not meaningfully related to other stimuli, appear together in the upper left




Figure 1. Stress values in five dimensions.




Figure 2. Dimension 1 (meaning) plotted against dimension 2 (Chomsky).

of the configuration. The semantic content of the sentences adequately accounts for the clustering.

There appears, however, to be another level of meaning which explains the ordering of the stimuli along the dimension. The order is a function of the relevance of the stimuli to the subject. Clearly, politically slanted sentences and those having to do with work are more significant than "noisy dogs" and "nosey ditches". Thus, the results confirm Chomsky's belief that meaning is a factor in the judgments of sentence acceptability.

Dimension two (Figure 4) corresponds to Chomsky's hierarchy. To the left of the configuration occur the stimuli representing lexical-rule violations: stimuli 1, 7, and 4. Next appear stimuli 2 and 5, early selection and restriction and subcategorical-rule violations. Moving further to the right, the three stimuli representing verb selection-restriction and late selection-restriction violations cluster. Lexical violations dominate subcategorical and early selection-restriction violations which in turn dominate verb selection-restriction and late selection-restriction violations (with the exception of stimulus 8). The correlation of the external Chomsky scale and the points on dimension two is .69.

An advantage of using a multidimensional scaling technique is visualizability of the relationships among the variables. Because the researcher is able to dispense with the categorization of variables, the ordering among the subject, verb, and object within a level of violation becomes clear. Within each level of violation, an incorrect verb appears to be considered a more serious violation than an incorrect subject, which in turn appears to be more serious than violations within the object position. The importance of the verb in the processing of sentences is not a finding which is confined to this study. Fodor, Garrett, and Bever (1968), based on a series of experiments, concluded that the verb plays a crucial part in the recovery of the base structure in processing. Chomsky considered violations only within the verb position in his theory, but perceptions of violations within a given level of the hierarchy are colored by whether the violations occur within the subject, verb, or object position. These results support Moore's contention that violations within subject and object positions should be considered when examining grammatical hierarchies. Moore's theoretical interpretation of the subject-verb-object relationship, however, was not supported. He maintained that perceptions of grammaticality were based upon whether the disturbance occurred in the subject-verb-object sequence or in a subordinate part of the sentence.

Dimension three appears to indicate a temporal awareness of the occurrence of violations (See Figure 3). During post-testing subjects consistently reported that taken in an imaginative context, sentences 8 and 9 could be considered correct. In addition, many subjects viewed the word "wanders" in sentence 4 and "ditches" in sentence 6 as typographical errors (that is, "wonders" and "bitches") and therefore interpreted sentences 4 and 6 as being correct. This confusion potentially accounts for the clustering of sentences 4, 6, 9, and 8.

Other subjects, on the other hand, reported that they recognized the errors in sentences 4 and 6, which explains both their nearness to each other and to sentences 9 and 8. In both sentences 4 and 6 awareness of the error occurs before the verb. Centrally located along this dimension are stimuli 1, 3 and 5. Subjects reported that in these instances, they became aware of a grammatical violation at the mid-point of the sentence. In sentences 2 and 7 subjects were aware of a grammatical violation only when reaching the latter portion of the sentence.




Figure 3. Dimension 2 (Chomsky) plotted against dimension 3 (temporal).




Figure 4. Dimension 1 (meaning) plotted against dimension 3 (temporal).

What subjects are reporting as temporal awareness may have a relationship to the length of time required to process a sentence. Since subjects reported confusion on a semantic level as to the acceptability of sentences 4, 6, 8, and 9, it seems plausible that these sentences took the longest to process. Subjects also reported more difficulty in processing sentences 1, 3, and 5, than sentences 7 and 2. The reasons for this difficulty are unclear. One explanation may be that a sense of dissonance occurred in midstream rather than at the end of the sentence, interrupting the fluidity of processing. Admittedly this is speculation, but this dimension does appear to be related to a time factor. Research is now being planned to examine this hypothesis.

The results indicate that judgments of acceptability involve more than subjects' perceptions of grammatical deviance and are truly multidimensional in nature. So little structure is provided by an experimenter who asks a subject to rate sentences on the basis of acceptability that it is questionable whether the dimensions found in this study are the only ones which underly acceptability judgments. The presence of an uninterpretable four-dimensional solution lends further credence to the possibility of additional underlying dimensions.

It is not surprising that researchers using acceptability ratings failed to uncover subjects' perceptions of a grammatical hierarchy. The findings of these studies must be reexamined, since the presence of a perceived grammatical hierarchy may have been obscured by extraneous, confounding factors, such as the meaningfulness and possible temporal awareness dimensions indentified here.

References

- Bransford, J. D., Barclay, J. R., & Franks, J. J. Sentence memory: A constructive versus interpretative approach. Cognitive Psychology, 1972, 3, 193-209.
- Bransford, J. D., & Franks, J. J. The abstraction of linguistic ideas. Cognitive Psychology, 1971, 2, 331-350.
- Chomsky, N. Aspects of the theory of syntax. Cambridge: M. I. T. Press, 1965.
- Danks, J., & Glucksberg, S. Psychological scaling of linguistic properties, Language and Speech, 1970, 13, 118-140.
- Danks, J., & Glucksberg, S. Psychological scaling of adjective orders. Journal of Verbal Learning and Verbal Behavior, 1971, 10, 63-67.
- Downey, R. G., & Hakes, D. T. Some psychological effects of violating linguistic rules. Journal of Verbal Learning and Verbal Behavior, 1968, 7, 158-161.

- Fodor, J. A., Garrett, M., & Bever, T. G. Some syntactic determinants of sentential complexity, II: Verb structure. Perception and Psychophysics, 1968, 3, 453-460.
- Greene, P. E., & Carmone, F. J. Multidimensional scaling and related techniques in marketing analysis. Boston: Allyn and Bacon, Inc., 1970.
- James, C. T., Thompson, J. G., & Baldwin, J. M. The reconstructive process in sentence memory. Journal of Verbal Learning and Verbal Behaviors, 1973, 12, 51-63.
- Kruskal, J. B. Multidimensional scaling by optimizing goodness of fit to a nonmetric hypothesis. Psychometrika, 1964, 29, 1-27.(a)
- Kruskal, J. B. Nonmetric multidimensional scaling: A numerical method. Psychometrika, 1964, 29, 115-129. (b)
- Moore, T. E. Speeded recognition of ungrammaticality. Journal of Verbal Learning and Verbal Behavior, 1972, 11, 550-560.
- Sachs, J. S. Recognition memory for syntactic and semantic aspects of connected discourse. Perception and Psychophysics, 1967, 2, 437-442.
- Shepard, R. N., Romney, A. K., & Nerlove, S. B. (Eds.). Multidimensional scaling: Theory and application in the behavioral sciences. (Vols. 1 & 2) New York: Seminar Press, 1972.
- Stoltz, W. S. Some experiments with queer sentences. Language and Speech, 1969, 12, 203-219.

THE EFFECTS OF IMAGING AND TEXT STRUCTURE ON WHAT IS LEARNED FROM READING
A PASSAGE¹

Sandra G. Koser²
Cornell University

This paper will present some of the data which were collected in connection with two experiments on the role of imagery and text organization in what is learned from reading a passage.

This research was motivated primarily by the rather consistent finding that imaging can function as a powerful associative strategy for learning a wide variety of materials which range from verbal and pictorial paired-associates to sentences and paragraphs. It has been hypothesized that imagery's facilitative effects stem from its ability to function as a particularly good "relational organizer" (Bower, 1970). That is, it is hypothesized that associative imagery operates by linking or relating items so that they form unified wholes or higher order units. And, when one item is recalled, that item acts as a retrieval cue for the other items, thus regenerating the whole.

Imagery is also thought to have some unique properties which affect its associative functioning. Among these is imagery's ability to organize or represent spatial or figural information. This has been demonstrated by Brooks (1967, 1968) and Atwood (1971) who have shown that visual perception tasks interfere more with imagery than do verbal tasks. Thus, it would appear that imagery's particular power comes from its ability to relate, organize and represent spatial or figural information.

In most studies of visual imagery, the nature of the spatial or figural relationships between the items to be remembered by the subject has been left unspecified. In the Bower study, for example, the stimulus materials were paired-associates with no inherent figural relationship to one another. Rather, the subject was instructed to use imagery to create his or her own figural relationship. This raised the question of what the effect would be if these figural relationships were already clearly specified and presented to the Ss within a structured framework.

From research on text organization, we know that variations in the surface structure of concrete texts can affect an individual's ability to recall that text and to determine its underlying relationships (Fraser, 1969, 1972; Nathan and Moore, 1971). This seems to be due, in part, to differences between texts in the amount of surface restructuring or inter-sentence reorganization that an individual must do to clarify those texts' underlying relationships. In these terms, a well organized text would be one which required little or no sentence reorganization to clarify the text's relationships, while a poorly organized text would be one that required extensive reorganization of the sentences. If, as has been hypothesized, imagery does indeed function as an organizer of spatial or figural relations, then it should be possible to facilitate an individual's ability to reorganize poorly organized texts by instructing him or her to represent the text information visually. The two experiments to be described here were undertaken to test this hypothesis.

¹ This research was supported by Hatch funds, project # 406, under the direction of Dr. Marvin Glock, Department of Education, Cornell University.

² Address: Department of Education, Cornell University, Ithaca, N.Y. 14850.

Experiment I:**Method****Subjects**

Sixty Ss enrolled in an introductory psychology course participated. Of these 60 Ss, one S's data were eliminated because his materials were defective; one S's data were eliminated because she failed to complete the materials; and the data of ten Ss were eliminated because they indicated a failure to follow the strategy instructions. Of the remaining 48 Ss, 21 were in the repetition strategy group and 27 were in the imagery strategy group.

Materials

The three passages used in this experiment were similar in format to passages described by Frase (1972). Three groups of 16 sentences were written. Each group named four concrete objects and described four of the objects' attributes or characteristics. Each sentence group was presented to the Ss in one of three sequences or surface organizations, designated as referent, attribute or scrambled. In the referent organization, all the sentences pertaining to a particular object and its attributes were grouped together in a single paragraph. Each succeeding paragraph introduced a new referent with its attributes. For example, here are the first two paragraphs of a referent passage.

In the past few weeks, a well organized gang of thieves has committed four bank robberies. In each case, an unusual escape vehicle was used. In the first robbery, the bank manger identified the vehicle as a tank. The first vehicle was bright red. The red vehicle had a large spotlight mounted on its left side. The driver of the vehicle with the spotlight was a Black woman.

The second robbery vehicle turned out to be a jeep. The second vehicle was dark blue. The dark blue vehicle had a machine gun mounted at its rear. Observers said that the vehicle with the machine gun was driven by a White man.

In the attribute organization, all the sentences pertaining to a particular attribute were grouped together in a single paragraph. For example,

The red vehicle had a large spotlight mounted on its left side. The dark blue vehicle had a machine gun mounted at its rear. The yellow vehicle had an American flag mounted on its roof. The white vehicle had a battering ram mounted on its front.

Each succeeding paragraph described a different attribute and the last paragraph gave the sentences which included the vehicle names. Each paragraph maintained the same vehicle order. That is, the first sentence referred to vehicle one, the second sentence to vehicle two, and so on. In the scramble organization, the sentences were randomly ordered and arbitrarily divided into four sentence paragraphs.

Sixteen test questions were prepared for each passage. Four questions required information from one text sentence to be answered (level one); four required information from two text sentences (level two); four required information from three text sentences (level three); and four required information from four text sentences (level four). Thus, four questions

required only that the S remember the relations which were stated directly in the text, while 12 questions required the S to draw inferences from the text. Some examples of questions at each level follow.

Level one: What was the vehicle in the second robbery?

Level two: What color was the tank?

Level three: What was mounted on the tank?

Level four: Who was driving the tank?

Two sets of written strategy instructions were also prepared. One set directed the S to repeat the information in the passage paragraphs to him or herself, while the other set directed the S to form composite images of the objects described in the passages.

Design

The design was a 2 (strategy) x 3 (presentation order) x 3 (passage content) x 3 (organization) x 4 (question level) with repeated measures on the last two factors. A Latin square was used to balance passage content across organization levels.

Procedure

The Ss were tested in small groups. The materials, which had been assembled into individual packets, were randomly ordered and given to the Ss in the order in which they arrived for the experiment. The Ss were informed that they would receive different strategy instructions and that they were to apply those strategies as they read the passages. They were also informed that they would be tested on the passage information, and that their reading would be timed. The Ss were allowed 1' 30" to read a passage. After each passage, the Ss answered the test questions and then completed some questionnaire items which asked them to describe the strategy they had applied to the passage. The questionnaire item responses were later used by the E to determine whether or not the S had attempted to apply his or her assigned strategy.

Results

An unweighted means analysis of the number of questions answered correctly revealed a significant effect for strategy, $F = 8.11$, $df = 1/46$, $p < .01$; and for organization, $F = 82.27$, $df = 2/92$, $p < .001$. However, the strategy and organization main effects were qualified by a significant strategy by organization interaction, $F = 5.29$, $df = 2/92$, $p < .01$. A test of simple main effects of this interaction revealed a significant difference between the strategies for attribute text, $F = 74.10$, $df = 1/132$, $p < .001$. There were no significant differences between the strategies for referent text or for scrambled text, although the imagery strategy means were higher in both cases. The differences among text organizations were highly significant at both strategy levels. With imagery strategy, $F = 175.48$, $df = 2/92$, $p < .001$; with repetition strategy, $F = 174.44$, $df = 2/92$, $p < .001$. The organization means at each strategy level were compared using the Newman-Keuls procedure. Within each strategy level, the organization means are significantly different from one another ($p < .01$), with referent > attribute > scrambled. These differences are presented in figure 1.

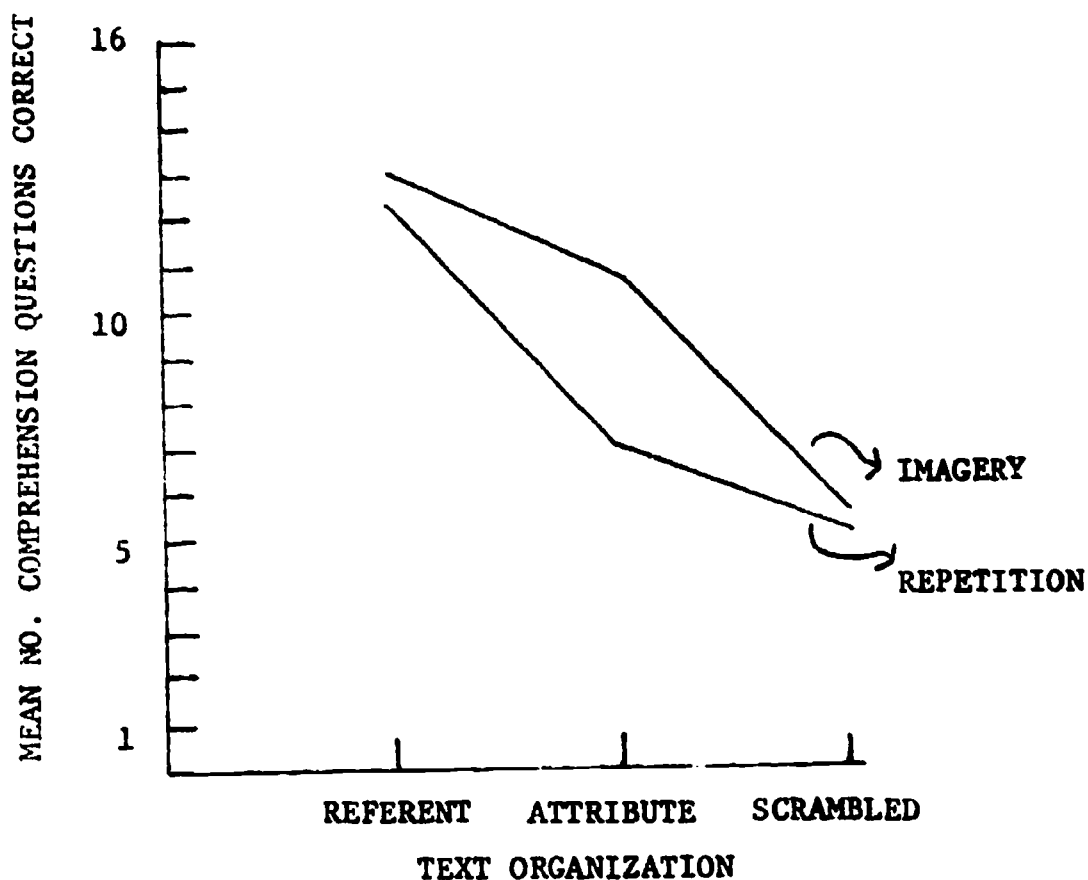


Fig. 1. MEAN NO. QUESTIONS CORRECT BY TEXT ORGANIZATION AND STRATEGY USED.

Discussion

It was originally hypothesized that variations in the surface structure of texts, that is, variations in text organization, would interact with the text processing strategy adopted by individuals to affect those individual's cognitive representations of the text information. It was further hypothesized that both the imagery strategy and the referent organized text would function similarly in that they would foster cognitive representations which allowed individuals to disambiguate the underlying attribution or relation structure of the texts. Since most of the text comprehension questions asked the Ss to draw inferences about what they had read, it was necessary for the Ss to have disambiguated the underlying structure of the texts in order to answer these questions correctly. It was, therefore, thought that these questions would provide a sufficiently sensitive measure of the role of both imaging and text organization in disambiguating and organizing underlying text relations.

From the data analysis, it is quite clear that referent organized text, regardless of which strategy is used to process it, fosters disambiguation of text attribute relations. Quite simply, subjects answer significantly more questions correctly when the text has a referent organization. This is not surprising, since the sequence of sentences in the referent organized text serves to group together all the attribution relations that exist between the objects discussed in the text and those objects' characteristics. Thus, the subject need not spend time regrouping the text sentences.

With attribute organized texts, Ss perform less well, but still better than they do with scrambled texts. In neither of these texts are the objects discussed grouped together with all their characteristics. In the attribute

text, however, the grouping is systematic or rule governed, while in the scrambled text, the grouping is random. With attribute text, each succeeding text paragraph always presents the attribute information in the same sequence. Apparently, the Ss are able to make some use of this sequence or surface information to disambiguate the text's relations. With scrambled texts, however, there is neither a facilitative sentence grouping (as there is with referent text), nor a facilitative organization rule (as there is with attribute text). Thus, the S must spend relatively more time disambiguating the text relations. Given the 1' 30" reading time constraint, then, it is no wonder that all Ss do poorly with scrambled texts. These findings and interpretations are consistent with what Frase (1969) reports.

It was hypothesized that the effect of using the imagery strategy would be to restructure non-referent text so that it was organizationally similar to referent text. It was, therefore, expected that Ss who claimed to have used imagery would be able to answer more questions correctly after reading attribute and scrambled texts than Ss who claimed to have used repetition. Thus, a strategy by organization interaction was predicted, with imagery Ss outperforming repetition Ss on attribute and scrambled texts. The predicted interaction was obtained; however, the imagery Ss outperformed the repetition Ss only on attribute text. Apparently, imagery served to reorganize attribute but not scrambled texts. This seemed an odd conclusion. Why should imagery function as a relational organizer in the one case but not in the other? Since this seemed a strange and strained explanation, reasonable alternatives were considered.

Perhaps the initial hypothesis, that imagery would disambiguate relations by reorganizing them was wrong. A plausible alternative seemed to be that while imagery would organize relations, it might play a more limited role in disambiguating them. That is, imagery might call attention to the existence of an underlying relational structure without actually providing a process for disambiguating it. This would account for the relatively poor showing of the imagery group with scrambled text. Since this text organization was in no way rule governed, it was very difficult for the Ss to disambiguate the underlying relations. Thus, the attempt to apply the imagery strategy would result in a grouping of objects with characteristics that was probably not correct. And this would lead to a poor inference making (question answering) performance. Unfortunately, the data of this experiment do not allow us to confirm or disconfirm this explanation since the dependent variable confounded the results of both disambiguation and organization processes.

A second experiment was, therefore, conducted to provide a measure of the S's organization of text relations which was independent of the correctness of those relations. It was again hypothesized that imagery would act as a relational organizer and would restructure or reorganize non-referent texts so that they exhibited a more referent like sentence organization. This hypothesis was tested by replicating experiment I, except that the dependent variable became the Ss' free-recall referent clustering scores. Based on these scores, an index of referent clustering was developed which described the degree to which the free-recall protocols expressed a referent organization. It was predicted that this index would be higher for Ss applying the imagery strategy than for Ss applying the repetition strategy.

Experiment II:

Method

Subjects

Forty-seven Ss enrolled in an introductory psychology course participated. Of these, the data of 11 Ss were eliminated because the Ss failed to follow

their strategy group instructions. As Ss arrived for the experiment, they were assigned to either the repetition or the imagery group. The assignment was random except that when a subject's data were eliminated from the study, the next S was asked to use that S's strategy. Complete data were collected on 36 Ss; 18 applied the imagery strategy and 18 applied the repetition strategy.

Materials

The materials were identical to those used in the first experiment except that no test questions were included.

Design

The design was identical to that of the first experiment except that there was no question level factor.

Procedure

All Ss were tested individually. The strategy instructions and the time limits imposed on the reading were the same as in experiment I. After time was called for each passage, the S was instructed to write down any information he or she could remember about the passage. They were also instructed to write in simple sentences and to avoid using pronouns. The free-recall protocols were scored for evidence of referent clustering. This was expressed by counting the number of consecutive relational sentences which referred to the same referent (R), subtracting one, and dividing by the total number of relational sentences (T) in that protocol minus the number of different referents referred to (K) in the protocol $((R - 1)/(T - K)) \times 100$. (This is essentially the same method of scoring that Frsae (1969) used.) For example, consider the following free recall excerpt.

The first vehicle was red. The red vehicle was a jeep.
The jeep was driven by a Black woman. The second vehicle
was dark blue. The fourth vehicle had a battering ram.
The second vehicle had a machine gun.

The sentences of this protocol would be coded as follows: AAABDB. Identical letters indicate sentences which refer to the same referent. The index of referent clustering for this protocol is $((3-1)/(6-3)) \times 100 = 67\%$.

Results

The mean referent clustering index was 84% for the imagery group and 44% for the repetition group; $F = 17.30$, $df = 1/18$, $p < .0008$. For referent text, it was 86%; for attribute text it was 53%; and for scrambled text it was 52%; $F = 15.36$, $df = 2/36$, $p < .0001$. The interaction of strategy and organization was not significant. The text organization clustering means were compared using the Newman-Keuls procedure. Referent text differed from attribute and scrambled text, which did not differ from one another ($p < .01$).

Discussion

Not surprisingly, there was significantly more referent clustering with referent organized text than with either attribute or scrambled text. And, as predicted, Ss applying the imagery strategy exhibited significantly more referent

clustering in their free recall protocols than did Ss using the repetition strategy (see fig. 2). These data demonstrate quite clearly that imagery does function to organize relations. These data are thus consistent with and tend to confirm the explanation put forth to account for the data collected in the first experiment.

That explanation suggested that the imagery strategy functions primarily as a relational organizer, and that its role in disambiguating relations is largely limited to calling attention to the existence of an underlying relational structure. Once attention has been focused on this underlying structure, its disambiguation seems to depend more on the characteristics of the text surface structure than it does on the imagery strategy.




Fig. 2. REFERENT CLUSTERING BY TEXT ORGANIZATION AND STRATEGY USED.

Atwood, G.E. An experimental study of visual imagination and memory. Cognitive Psychology, 1971, 2, 290-299.

Bower, G.H. Imagery as a relational organizer in associative learning. Journal of Verbal Learning and Verbal Behavior, 1970, 9, 529-533.

Brooks, L.R. The suppression of visualization in reading. Quarterly Journal of Experimental Psychology, 1967, 19, 289-299.

Brooks, L.R. Spatial and verbal components of the act of recall. Canadian Journal of Psychology, 1968, 22, 349-368.

Frase, L.T. Paragraph organization of written materials. Journal of Educational Psychology, 1969, 60, 394-401.

Frase, L.T. Maintenance and control in the acquisition of knowledge from written materials. In J.B. Carroll & R.O. Freedle (Eds.), Language Comprehension and the Acquisition of Knowledge, V.H. Winston & Sons, Washington, D.C., 1972.

Natkin, G.L. & Moore, J.W. The effects of instructional sequencing on learning from a simple knowledge hierarchy. Paper presented at the annual meeting of the American Educational Research Association, N.Y., February, 1971.

HUMAN PROBLEM SOLVING: A SYNTHESIS OF CONTENT, COGNITION, AND INDIVIDUAL DIFFERENCES

Joseph M. Scandura *
University of Pennsylvania

Why is it that some people can solve certain problems whereas others cannot? In one form or another this question has puzzled scholars almost from the beginning of civilization, and lies at the heart of the problem solving process. My goal this afternoon is to show how content (competence), cognition, and individual differences (measurement) can be synthesized in terms of the Structural Learning Theory and to show how this theory might be used to account for the findings of existing problem solving research.

(1) It is generally agreed that the successful problem solver must be able to understand the problems with which he is confronted (or which he defines for himself), he must be able to identify suitable subgoals, he must have at hand or be able to retrieve relevant information from memory, he must be able to distinguish relevant from irrelevant information, he must be able to derive solution procedures for each subgoal, he must be able to carry out these procedures correctly, and he must know how to verify both intermediate and terminal results. But, beyond this general outline, there is little general agreement on the actual nature of the underlying processes. What does it mean to say that a person understands a problem and how does he do it? How does he form subgoals? Retrieve relevant information? Derive solution procedures? Or, use them?

(2) In addition to questions concerning the general nature of problem solving, I will also consider the problem of identifying specific abilities (competencies) that are necessary for solving given classes of problems (content). As in research on computer simulation and artificial intelligence, this involves identifying specific procedures by which given classes of problems can be solved. Unlike most existing research in this direction, however, we have been concerned also with general methods of analysis by which such competencies may be identified. (John Durnin and Wallace Wulfeck describe some of our research on this problem.) More generally, we shall emphasize here the question of how analysis of specific problem domains can be combined with insights about the general nature of the problem solving process to yield a more comprehensive understanding of problem solving.

(3) Individual differences in problem solving ability are also critical. What is a difficult problem for one person, for example, might be quite simple for another. Are these differences due to genetics? Or, do they involve identifiable knowledge, knowledge that can be taught and learned? Undoubtedly, both factors are involved to some extent. But, where and how?

Any complete accounting of problem solving, as a minimum, will necessarily involve specification of the specific knowledge required in solving given classes of problems, an understanding of underlying psychological mechanisms, and some way to deal with individual differences. I believe that my structural learning theory (Scandura, 1973_a) is adequate for this purpose, especially with more recent refinements and extensions. Obviously, I cannot do justice to all that might reasonably be covered, in the space available. Rather than attempt a superficial overview of the structural learning theory, the empirical results, and implications for problem solving, I will concentrate on the more significant and relevant features of the basic theory. Further, to increase expository efficiency, instead of using a variety of real world examples, I shall try to "make the exposition concrete with a single, abstract example." In the concluding section, I shall suggest how the structural learning theory relates to other problem solving research. **

*This paper was prepared while the author was being supported by the MERGE Research Institute.

** Only material in sections 1 and 2 were actually presented at the Structural Learning Conference due to time limitations. Sections 3 and 4 constitute two thirds of the basic theory and are included here for completeness. Taken collectively, sections 1-4 constitute a draft of one chapter in my forthcoming book, Human Problem Solving: A Synthesis of Content, Cognition, and Individual Differences. This book will contain a more comprehensive treatment of the theory; a substantial amount of related empirical research on cognition in problem solving, content analysis, and the measurement of individual differences; and a more extensive discussion of relations to other problem solving research.

Any comments on the present article, especially critical ones, would be greatly appreciated.

1. OVERVIEW

In theoretical work on complex human behavior, there are three general kinds of questions that may be asked: (1) Competence (content): Given a class of tasks, however broad and general (e.g., value judgements) or narrow and specific (e.g., arithmetic computation), how can the competence required for performing these tasks best be characterized so as to reflect human behavior? (2) Cognitive theory: Given a specific task and the relevant knowledge available to a person, what are the mechanisms which determine how this knowledge is put to use and how new knowledge is acquired? (3) Measurement: What is the relationship between competence and knowledge? Or, equivalently, how can one find out what individual subjects know relative to given content domains.

Most existing theories deal with one or another of these themes with at best only passing concern for the others. Polya's (1962) analyses of heuristics, artificial intelligence (e.g., Minsky & Papert, 1972), task analysis (e.g., Gagne, 1970), and other forms of content analysis provide examples of question (1); traditional learning theories (e.g., Hilgard & Bower, 1966) provide examples of question (2); and mastery or criterion referenced testing (e.g., Bloom, 1973) comes closest to dealing with question (3). On the other hand, computer simulation (e.g., Newell & Simon, 1972), psycholinguistics (e.g., Carroll & Freedle, 1972), contemporary research on problem solving (e.g., Greeno, 1973), and semantic memory (e.g., Anderson & Bower, 1973; Kintsch, 1972; Rumelhart, Lindsay, & Norman, 1972; Quillian, 1968) tend to combine competence and psychological considerations. In addition, at a more informal level, educational psychologists (e.g., Gagne, 1970; Bloom, 1973) have been concerned with content, behavior, and (criterion referenced) measurement. With the possible exception of Pask's (e.g., Pask, Scott, & Kallikourdis, 1973) theory of conversations, however, very little has been done by way of (relatively) formal synthesis of all three concerns.

The structural learning theory, as it has evolved over the past several years (cf. Scandura, 1971, 1973) provides a unifying theoretical framework within which the concerns of the competence^a (e.g., artificial intelligence, linguistics, subject matter content) researcher, the psychologist, and the measurement specialist may be viewed. In the theory a sharp distinction is made between competence (content) and individual human knowledge, a distinction that is analogous but not identical to that which some linguists make between competence and performance. The theory also explicates the complementary relationship between competence and knowledge. This relationship takes the form of what essentially is a performance test theory. As a result, research in competence and behavioral domains can progress independently but, if it is to have relevance to human psychology, research in any one domain is necessarily subject to constraints imposed by the other.

The observer and the behaving subject both play a fundamental role in the theory, corresponding to the above distinction between competence and knowledge.* The observer determines the criteria to be used in judging both what count as stimuli (inputs) and what count as distinguishable behaviors. He is also responsible for identifying the competence (rules) against which individual behavior is to be judged.

In the theory, as in all generative theories, both competence and knowledge are represented as sets of rules (procedures or labeled directed graphs). Beyond this common feature, however, the characterization of competence differs from that in other generative theories (e.g., in linguistics, and artificial intelligence). In the structural learning theory, rules are allowed to operate on other rules to generate new ones in a way which allows for unanticipated, creative outcomes.

Regarding human behavior, whatever the individual learner does and can do in any particular situation is assumed in the theory to depend inextricably on what he already knows. In addition to whatever specific knowledge may be available to an individual, cognition and behavior are both assumed to be based on a specific control mechanism. This control mechanism is assumed to be common to all people and to determine how available knowledge is put to use and how new knowledge is created. Among other things it provides a basis for answering such questions as why it is that some people succeed in problem situations for which they do not explicitly know solution rules whereas others do not. In Scandura (1973_a) this mechanism is shown empirically to provide a viable

* Although a sharp conceptual difference was made, I used the term "knowledge" in my book (1973) to refer to both competence and behavior potential (knowledge). This usage was related to the terms "competence" and "knowledge" in an epilogue to Chapter 9.

account of learning (rule derivation) and of selection from among alternative rules. Subsequent research also has demonstrated its viability as an account of problem definition (forming subgoals), retrieval, and related phenomena.

In contrast to the control mechanism, specific knowledge may vary over individuals. Specifically, the theory shows how competence, identified by an observer, may be used to operationally define the knowledge had by individual subjects in a population. As it turns out, the rules in a competence theory serve as rulers for measuring human knowledge. Competence, however, may not be identified independently of behavior; indeed, behavior provides the ultimate test of adequacy for competence theories.

To summarize, although the structural learning theory provides a synthesis of the old, it also represents a basically new approach to the study of human behavior. Each of these concerns: competence, cognition, individual differences, and their relationships, is considered more fully in the following sections.

2. COMPETENCE

As indicated above, competence is represented in terms of rules introduced by an observer to account for behavior. To get some feeling for what this implies, let us take the role of the observer who wants to understand and be able to predict the behavior of some population of subjects. The first question which arises is what aspects of behavior are to be considered. Surely, if the observer is concerned with the subjects' problem solving behavior, then whether or not in the process a subject happens to scratch his head or get up to sharpen his pencil is not likely to be of great import. What are central are the types of problems and solutions in question. The first task of the observer, then, is to specify criteria for distinguishing between those aspects of behavior (i.e., those potentially observable inputs and outputs) that are of interest and those that are not. This is the observer's descriptive role.

These criteria, or equivalently, a class of potentially observable input-output pairs, against which actual behavior is to be judged, are predetermined. When the psychologist enters his laboratory, for example, he has a pretty good idea ahead of time what stimuli and what responses he is interested in. Whether or not the subject wiggles in his chair as he elicits the response "MUR" may not only be unanticipated, but typically also will be ignored. Similarly, in testing students to see whether they know some subject matter, the teacher usually determines in advance what are the stimuli and the corresponding acceptable responses.

So far so good. But, being able to identify relevant observables tells only part of the story. In order to have any hope of understanding the behavior, and especially of predicting it, the observer must have some idea as to how those observables can be generated. In effect, the observer, qua competence theorist, also determines how such behavior might reasonably be generated by members of the subject population(s) involved.* As in all competence theories, competence consists of a finite set of rules together with laws (control structures) governing the way in which these rules may interact in accounting for a given set of elements (in the present case, psychological observables).


In an important sense, the curriculum constructor (whether human or computer) finds itself in a situation which directly parallels that of the competence theorist. The curriculum constructor usually knows ahead of time what the student is to learn, if not in specific terms, then, at least generally. In addition to specifying expected behavior, for example, he must either specify or make "intelligent guesses" as to the rules students need to know to generate the desired behavior. As Glaser (1967) has pointed out, even where objectives appear to be difficult to pin down, as in some forms of open education, this is still a desirable goal.

2.1 Representation of Rules (Procedures, Algorithms).--Individual rules may be represented in terms of labeled directed graphs (see Figure 1), where the arrows represent operations (e.g., add, turn to the right, calm down) and the nodes represent decision making capabilities (e.g., If the numeral has one digit, do (operation) O_1 ; else, do O_2). The arrows are labeled to signify which operation goes with which condition (e.g., one-digit numerals might be operated on by O_1 ; others by O_2). Two or more of the nodes are

*Whereas perceptual as well as cognitive processes may be considered (cf. Scandura, 1973_a, Ch. 5), attention here is limited to the latter.

distinguished from the rest, one with the label "Start" and the other distinguished ones with "Stop." The reader, familiar with computer programming, may prefer to think of these graphs as schematic representations of flow diagrams. Further, although the mode of representation may be enriched to make explicit provision for perception and encoding, we prefer here to avoid this complication as it does not play an important role in the work described in later chapters.

Although formal distinctions exist among the terms "rule," "procedure," and "algorithm," I do not believe that these distinctions are critical in discussing human behavior and, throughout, I have used the terms more to connote concern rather than substance. "Procedure" and "algorithm" are used to connote concern for the internal structure of a labeled directed graph. (In formal treatments, algorithms necessarily terminate in a finite number of steps whereas procedures may not.) The term "rule" is used more generally, especially where internal structure is not at issue.



Furthermore, whereas all rules may be characterized in the same way, certain general distinctions are sometimes useful. For example, rules which operate on elements, which themselves may not act as rules, are termed "simple"; rules which operate on other rules are termed "nonsimple." Another distinction, which becomes useful in discussing cognition, depends on the use to which a rule is being put relative to another. A rule which operates on another rule is of a relatively "higher order," irrespective of whether the latter is simple or nonsimple. Higher order rules in this sense should not be confused with higher order rules as commonly used in the literature (e.g., Gagne, 1970; Restle, 1970). As used by most authors, higher order rules correspond to outputs of what are here called higher order rules. The distinction is conceptually important and directly parallels that in mathematics between a function value and the function itself.

Rules may serve a variety of purposes in accounting for potential behavior. In the role of higher order rules, they can be used to derive other rules, to select from among two or more alternative rules (cf. Taylor, 1960; Chapter 4*), and to retrieve information from memory (e.g., Bower and Anderson, 1972; Rumelhart, Lindsay, and Norman, 1972; Chapter 4*). Rules (including higher order ones) also may be used to assign meanings to mathematical (Scandura, 1973) and natural language (Suppes, 1974) descriptions, and to define problems or to break them into subgoals (e.g., Newell and Simon, 1972; Greeno, 1973; Chapter 4*).

In this regard, it is worth noting that rules need not operate on or generate observables. Thus, the meanings assigned to descriptions (by rules) and the rules derived by application of higher order rules are not themselves directly observable. Similarly, while some rules apply directly to observables, others do not. The distinction corresponds to that between classical and operant conditioning (Scandura, 1972, 1973).

As an example in which part of the stimulus (input) is external and part internal, consider the sequential task of putting an "X" in the next slot of the external stimulus . On the first presentation, the "X" goes in the first slot; on the second, it goes in the second slot, and so on, repeating the cycle beginning on the fourth presentation (and every first presentation, mod three). The point is that since the external stimulus is the same on each presentation, the only way the "X" can be properly placed is by "remembering" the last presentation. In effect, the effective input is the external , together with the (internalized) last placement of "X." From now on "input" means the effectively operating input, irrespective of whether all or part is observable or otherwise -- similarly, for "output."

One final point concerning representation. Rules need not be discrete. In particular, rules may be embedded in larger networks (labeled directed graphs). We return to this possibility below in discussing memory.

* i.e., Chapter 4 of my forthcoming book, Human Problem Solving: A Synthesis of Content, Cognition, and Individual Differences.

2.2 Interaction Among Rules.--In addition to the specification of individual rules, all competence theories assume either explicitly or implicitly a mode of interaction among these rules. A simple grammar, for example, consists of a finite set of rules, and is said to account for an input-output pair if some sequence of rules in the rule set can be found such that the successive application of these rules to the input generates the output. Although commonly assumed in all formal systems (e.g., Nelson, 1968) that rules (productions) are composable in this way, it is important to note that this mode of interaction is both very general and highly specific in comparison to human competence. Thus, while formal theories assume arbitrary composition, the fact that an individual is able to compose one pair of rules is no guarantee that he can necessarily compose another. On the other hand, forming composites of given rules is only one of many possible ways in which new rules may be derived from given ones.

In the structural learning theory, rules are allowed to interact in a more general manner. In particular, what is new in the theory is the idea of allowing rules to operate in a higher order fashion, that is, to operate on and to generate new rules. Suppose, for example, that a rule set contains rules for converting from certain English measures of length (e.g., yards, feet, inches) to their metric equivalents (e.g., meters, centimeters) and no others. According to the structural learning formulation, this rule set would be inadequate even for generating English equivalents of metric measures. The rule set has no provision for rules which operate in reverse direction. The whole situation changes, however, on addition of just one (higher order) rule to the set, a rule which maps English-to-metric rules (e.g., $p \text{ in.} \rightarrow 2.54 \times p \text{ cm.}$) into their inverses (e.g., $p \text{ cm.} \rightarrow p/2.54 \text{ in.}$). Not only will this inverse rule generate a metric-English rule for every English-metric rule in the set but with the addition of any similar conversion rule we get "free" the corresponding inverse.

Allowing rules to operate on rules in this fashion (to generate new rules which in turn can generate behavior) not only provides a great increase in explanatory power, but is consistent with how human beings use the knowledge they have (Scandura, 1973a). Thus, individuals do not need to know explicitly every rule that might be desired. Much of their knowledge is latent in a sense that it can be derived at will when needed from other information which is explicitly available.

2.3 Basic Example -- Finer Distinctions Between Structural Competence Theories and Traditional Grammars and Heterarchical Systems.--The formal distinction between traditional grammars and structural competence theories can be seen from the following basic example. (Variations of this example will be used throughout this paper for illustrative purposes.) Suppose that the given class of input-output pairs of interest consists of strings of the form $xB \rightarrow By$ where x is a string of a 's and y is the binary numeral representing the number of a 's (e.g., $aaaaaB \rightarrow B101$, $aaB \rightarrow B10$). A simple grammar which accounts for this class, includes the rules $r_1 = xxB_y \rightarrow xB0y$ and $r_2 = xxAB_y \rightarrow xBly$. To account for the pair $aaaaaB \rightarrow B101$, for instance, we see that

$$aaaaaB \xrightarrow{r_2} aaB1 \xrightarrow{r_1} aB01 \xrightarrow{r_2} B101$$

Notice that neither of the given rules is sufficient in itself to account for the given pair. It is necessary to assume that the rules may be applied successively as desired.

An equivalent way of accounting for this class is to explicitly include a generalized composition rule $*$ in the characterizing rule set (Scandura, 1973a), where $*$ operates on pairs of rules including itself [e.g., $*(r_1, r_2) = r_1*r_2$; $*(*, *) = **$ (adjoin one rule, then another); note: $** (r_1, r_2, r_1) = r_1*r_2*r_1$]. Denote this set $A = \{r_1, r_2, *\}$. Accounting for a given input-output pair, in this case, means either that there is a rule in A which generates the output on application to the input or that such a rule may be derived by application of rules in A to other rules in A . More generally, we say that A accounts for an input-output pair if there is a finite number n such that there is a rule in one of the following sets which generates the output from the input.

$$A = \{ *, r_1, r_2 \}$$

$$A^2 = A(A) = \{ *, r_1, r_2, **, r_1*r_1, r_1*r_2, r_2*r_1, r_2*r_2 \}$$

$$A^3 = A^2(A) = \{ *, r_{i_1}, **, r_{i_1}*r_{i_2}, ***, r_{i_1}*r_{i_2}*r_{i_3} \mid i_1, i_2, i_3 = 1, 2 \}$$

$$A^n = A^{n-1}(A) = \{ *, r_{i_1}, \dots, *1^2 \dots n, r_{i_1}*r_{i_2} \dots *r_{i_n} \mid i_1, i_2, \dots, i_n = 1, 2 \}$$

With respect to the above instance $aaaaaB \rightarrow B101$, for example, the rule $r_2*r_1*r_2$ in A^3 serves this purpose.

It is important to emphasize that these two formulations are mathematically equivalent insofar as computing power is concerned,* so in one sense we have nothing new. Mathematical equivalence, however, does not necessarily imply behavioral equivalence. As I have suggested above, for example, it is neither reasonable to assume that the composition rule $*$ (or any other rule) is necessarily universal in its applicability, or that composition is the only kind of higher order rule that might be included in a rule set. For example, the higher order rule

$$r_a \Rightarrow r_b$$

operates on rules involving a's and converts them into corresponding rules involving b's. The addition of just this one rule doubles the power of the given rule set to include an equivalent set of input-output pairs where the inputs involve b's instead of a's. More important, every time a new rule involving a's (in the domain of the higher order rule) is added to the rule set, we automatically get "free," because of this higher order rule, a corresponding rule involving b's.


Although the above structural learning formulation illustrates the major differences between it and simple grammars, the formulation is overly restrictive. Specifically, it is assumed that newly generated rules in A^n may only be applied to rules in A , hence the notation $A^n(A)$. A more general form of rule interaction allows for the derivation of new domain rules. For example, if the rules are allowed to interact at each level so as to generate new sets of domain rules, then the formulation would be generalized to

$$A^n(A^m) = \{ *, r_1, r_2, \dots, r_{i_1}*r_{i_2} \dots *r_{i_{nm}} \mid i_1, i_2, \dots, i_{nm} = 1, 2 \}$$

This rule set includes every rule that can be derived by applying rules in A^n to rules in A^m . Cognitive control mechanisms, corresponding to the above modes of rule interaction, are discussed in Section 3 (Goal Switching Mechanism and Extension to Multiple Processes, respectively).

For the computer scientist, the structural learning formulation may appear on first thought to be similar to modular, heterarchical systems in artificial intelligence (e.g., Minsky and Papert, 1972). In both cases, a finite number of modules (rules) is specified together with a control mechanism which governs the way in which the modules interact. (In the case of heterarchical systems, the modules are in the form of computer programs.) In actuality, there is a major difference. Structural competence theories, including both the rules and their assumed mode of interaction, are assumed to reflect idealized human behavior (under memory free conditions).** Although not necessarily inconsistent with human behavior, the major requirement of heterarchical systems is that they work with reasonable efficiency.

* Given the simplicity of this rule set, it is easy to avoid composition ($*$) altogether by allowing recursion on r_1 and r_2 . Thus, assuming suitable decision making capabilities, the rule denoted



has the same computing power as the collection of sets A^n ($n = 1, 2, \dots$).

** Rule sets, associated with individual knowledge, and control mechanisms, which correspond to the proposed mode of (rule) interaction, are discussed in Sections 3 and 2, respectively.

2.4 Semantics.-- The reader should not be misled by the syntactic character of our example. The proposed competence theory does a perfectly good job of handling meaning as well.

Any (composite) syntactic rule, for example, which accounts for a pair of the form $xB-By$, where x is a string of a 's and y is the binary numeral representing the number of a 's, may denote a semantic rule. In particular, it may denote a semantic rule which operates on unarranged piles of a given number of concrete objects (e.g., five 3" wooden dowels) and regroups them as a base 2 concrete embodiment of the corresponding binary numeral. For example, the syntactic rule $r_2*r_1*r_2$, as applied to the input $aaaaaB$ gives

$$aaaaaB \xrightarrow{r_2} a'a'B1_a \xrightarrow{r_1} a''B0_{a,1_a} \xrightarrow{r_2} B1_a''0_{a,1_a}$$

The corresponding semantic rule $rs_2*rs_1*rs_2$ regroups objects (e.g., five 3" dowels) at each step into sets of size 2^n , where n is the number of 0's and/or 1's in the input. Thus

$$///// \xrightarrow{rs_2} // // \phi \xrightarrow{rs_1} // // \circ \phi \xrightarrow{rs_2} \phi \phi \phi \phi$$

Notice that the a 's in the initial input ($aaaaaB$) denote single ($2^0=1$) objects (e.g., 3" dowels). Each a' in the string derived by application of rs_2 denotes two ($2^1=2$) objects and the " 1_a " to the right of B denotes one single element set, denoted ϕ (e.g., a 3" dowel with a rubber band around it). In the next string, the a'' denotes four ($2^2=4$) objects and the " 0_a " to the right of B denotes no sets containing two objects (denoted \circ). Similarly, the " 1_a " just to the right of B in the final output denotes one set of four ($2^2=4$) objects (denoted $\phi \phi \phi \phi$).

In short, each input-output element (whether terminal or intermediary), as well as each of the rules, may be assigned a semantic meaning. The a 's denote piles of objects, the sizes of the piles depending on the number of digits to the right of B . The digits 0 and 1, similarly, denote sets of size 2^n (e.g., groups of dowels with rubber bands around them). The " B " serves a strictly syntactic purpose, similar to that of the comma in ordinary English, and makes it easier to interpret each string.

2.5 Innate Bases and Direct Computing Power.--It is a well known fact that if a function (class of input-output pairs in which there is a unique output for each input) is computable (e.g., Rogers, 1967), then there is a countably infinite number of procedures which will account for the function. Equivalently, if one procedure will do the job, then it is possible to devise a countably infinite number of others which will also work. This is equally true of structural competence theories which consist of potentially large finite sets of rules. What keeps the number of feasible accounts within bounds, of course, is the added constraint that the rules in the rule set must reflect the "culture" of the population of subjects in question. This requirement of compatibility is made precise in the final section.

Even given these added constraints, viable, alternative rule sets may differ in important ways. Thus, a rule set K' is said to be an innate basis for another K , if there are finite numbers n and m such that every rule in K is contained in $K'^n(K'^m)$. For example, suppose $K' = \{ r_1, r_2, * \}$ and $K = \{ r_1, r_2, r_1*r_2 \}$. In this case, $K'(K') = \{ r_1, r_2, *, r_1*r_2, r_2*r_1 \}$ which contains each rule in K .

On the other hand, K has more direct computing power (than K') in the sense that the union of the domains of its simple rules is larger than (properly contains) the union of the corresponding domains ($\text{Dom } r_1 \cup \text{Dom } r_2$) in K' . Some implications of these relations will become clearer in the next section.

2.6 Construction of Structural Competence Theories.--In general, in order to devise a competence theory which adequately characterizes a given content domain and subject population, an observer must be intimately familiar with both the domain and the relevant "culture" and developmental level of the population. Unlike linguistic competence, however, structural competence is not concerned with a specific content domain but, generally, with all such domains. Given this breadth of concern, attention cannot be restricted to the evaluation of alternative rule accounts. Attention must also be

given to the general processes by which competence theories are constructed. In linguistics proper, the possibility of systematically devising grammars has not been taken seriously; indeed, it has been held unrealistic (e.g., Chomsky, 1957).

Nonetheless, general constraints imposed by the form of an acceptable competence theory, as is the case with structural competence, do make it possible to proceed in a quasi systematic manner. Furthermore, it is possible to proceed in an at least partially self correcting manner. Since the basic approach is detailed and illustrated in Parts III and IV*, we shall be brief.

To make our description definite, let us suppose that the domain of interest consists of the following set of input-output pairs $I-O = \{ xB-By \mid \text{where } x \text{ is a finite string of a's and } y \text{ is the binary numeral representing the number of a's} \} \cup \{ x'B-By' \mid \text{where } x' \text{ is a string of b's and } y' = y \text{ represents the number of b's} \} \cup \{ wT - Tz \mid \text{where } w \text{ and } z \text{ are strings consisting of the same numbers of a's and b's respectively} \} \cup \text{a complementary and ill defined set of additional input-output pairs involving a's and b's.}$

The first step in analyzing a content domain is to select a diverse (finite) sample of tasks from the domain. For example, suppose we select aaaaB-B101, aaB-B10, bbbB-B11, bbbbB-B100, and aaT-Tbb.

The next step is to devise solution rules for the sampled tasks which are consistent with how the subjects in question might be expected to solve them. The basic theory is neutral as to how this is done: via simple intuitive introspection or informal observation of typical subjects, or more systematic methods, such as described in Merrill (1974), for constructing procedures which simulate behavior on given classes of tasks. Admittedly, this imposes a heavy responsibility on the observer (theorist) but, if one wants to deal rigorously with both specific content domains and individual behavior, I see no viable alternative. Indeed, significant progress has already been made in this direction (e.g., Carroll, 1973; Polya, 1962; Minsky and Papert, 1972; Newell and Simon, 1972; Scandura, Durnin, Ehrenpreis, and Luger, 1971; Scandura, Durnin and Wulfeck, 1974) and it seems to me to present a fascinating challenge.

Given the five instances above, we get $r_2 * r_1 * r_2$, $r_1 * r_2$, $r_2' * r_2'$, $r_1' * r_1' * r_2'$, and r_{ab} , respectively, where r_1 , r_2 , and $*$ are as above, r_1' and r_2' parallel r_1 and r_2 with a's replaced by b's, and r_{ab} replaces each a to the left of T in the input with one b to the right of T in the output. Let K be this set of rules. Clearly, K accounts for some of the selected input-output pairs, but not all of them. It does not, for example, account for bbbbbB-B101.

The third step is to search for parallels among the solution rules to see whether one can identify an innate basis from which the initial set (K) can be generated. Observing the central role played by composition and noting that all of the solution rules may be accounted for in terms of a set of simpler rules, we get $K' = \{ *, r_1, r_2, r_1', r_2', r_{ab} \}$ where K' is an innate basis for K since there is a number n such that

$K \subset K^n = \{ \text{m-fold composites (e.g., } r_2 * r_1 * r_2 \text{) of } *, r_{ab}, r_1, r_2, r_1', r_2', r_{ab}; m=1, 2, \dots, n \}.$

K^n is clearly a far more powerful rule set. But, it still has important inadequacies. For example, it does not account for an input-output pair as simple as bbT-Taa.

Recognizing this, we can further increase the potential power of our set by introducing $r_a \Rightarrow r_b$, and eliminating r_1' and r_2' as redundant. This gives the second generation innate basis

$$K'' = \{ *, r_1, r_2, r_{ab}, r_a \Rightarrow r_b \}$$

Not only does K'' include all rules in K' but it accounts for all input-output pairs of the form $zT-Tw$ (where z is a string of b's and w is a string of a's). Further, the addition of any new rule involving a's is tantamount to also adding a corresponding rule involving b's.

Along with the increase in potential power obtained in this manner, it is important to notice that greater simplicity of the individual rules has also been achieved. Whereas the individual rules (in K) are composed of a number of atomic (unitary) rules, those in K'' are all atomic.

* Again, reference made to my forthcoming book: Human Problem Solving: A Synthesis of Content, Cognition, and Individual Differences.

This suggests an answer to the question of where higher order rules come from in the first place. Higher order ^(non simple) rules are no more complicated than simple ones; indeed, they may be quite simple (in the colloquial sense). All but the most confirmed environmentalists believe that people come into the world wired in with some capabilities (instincts, general arousal levels, basic perceptual capabilities, and the like)*. Following the present thesis to its natural conclusion, it is my current position that such capabilities provide the innate basis upon which knowledge grows by interacting with the environment. (It should be emphasized in this regard that this position is neutral as regards maturation. It is certainly reasonable to suppose that new capabilities may come about as a result of physical maturation, rather than learning, but the present theory remains open on this issue.)

On the other hand, the (simple) rules in K have more direct computing power in the sense that they account directly for more complex input-output pairs than the simple rules in K". (Strictly speaking, K does not have more direct computing power than K" since r_1 and r_2 in K" account for the special cases B-B0 and aB-B1, respectively.) In Section 3, which deals with human learning and performance, we shall see that direct computing power corresponds to the capability of responding in goal situations directly on the basis of available knowledge (rules) as opposed to first having to generate new knowledge.

2.7 Extensions of the Basic Method.--The preceding discussion sloughs over several matters that have considerable practical as well as theoretical importance. For one thing, nothing was said about competence accounts in which there are two or more rules which provide an equally viable account of some (sub)domain of tasks. For example, the subdomain of instances { $x''B-By''$ | where x'' is a string of s_n ($= 2s_{n-1} + 1$) a's where $s_0 = 1$ and y'' is the corresponding binary numeral } may be accounted for either by rule



or by rule (2) $\left. \begin{array}{l} \text{STOP} \\ r_2 \end{array} \right\}$ in which r_2 alone is used recursively.

As a more concrete example, recall that ordinary subtraction can be and frequently is performed by a variety of methods, the best known being the "borrowing" and "equal additions" algorithms.

The important point is that analogous pairs of rules may account for parallel subdomains. For example, a subdomain, corresponding to that above but involving b's instead of a's, would be solvable by rules similar to (1) and (2) above, with r_1 and r_2 replaced by r'_1 and r'_2 , respectively. When such parallels are observed among pairs (triples, etc.) of such rules, we construct higher order selection rules which distinguish among the rules in the various pairs. (Such higher rules may also allow for derivation as well as selection.) In this example, the selection rule "Choose the simpler rule" might be appropriate.

A second major concern has to do with whether the identified I-O inputs are "natural" in the sense that they correspond to unitary tasks (relative to the population). In the present context, for example, the instance aaaaB-B101 would be natural, whereas the instance aaaaB-2 would not be, even (especially) if we are told that the "2" denotes the numbers of 1's in the binary numeral representation of the number of a's. It would be more natural to think of the latter instance as a pair, aaaaB-B101 and B101-2.

Because the instance aaaaB-2, and the corresponding pair, both may be included in a given content domain, and because the relationship between them is likely to be shared by other instances, it is natural to think in terms of rules which operate on instances and generate other instances, in particular pairs (triples, etc.) of instances. We shall see in the following sections that such rules correspond to problem definition,

* In the structural learning theory, such capabilities are assumed to be over and above the control mechanism of Section 3 which is hypothesized to determine the way in which available rules interact.

or breaking down a problem into subproblems. For example, if presented with aaaaaB and asked to find the number of 1's in the binary numeral representing the number of a's, subjects would most likely first solve the task of generating the binary numeral B101, and then count the number of 1's in it. (This assumes, of course, that counting is either among the available rules or can be derived from such rules.)

In order to devise a competence account, which takes "problem definition" into account, careful attention must be given to the procedures which subjects in the population use to solve the sampled tasks (instances). In step one of the previous section, for example, we might want to pay closer attention to how subjects actually attack sampled problems. In particular, subjects who tend to solve parts of given problems, before attacking other parts, may be thought of as breaking the problem into subproblems. In this case, again, parallels in the way different problems are broken down are manifestations of underlying problem definition rules.

Because higher order derivation (selection) rules and problem definition rules may both enter into an account of any given instance, it may be helpful to make the distinction between them sharper and more operational. In those cases where subjects attack task instances by attempting to derive a total solution procedure, before actually applying it, higher order derivation rules may be assumed to be operating. On the other hand, where subjects actually solve part of the problem before attempting to derive a solution procedure for the rest, problem definition rules are involved. Given the instance aaaaaB-2, for example, a subject of the former type would attack the problem by generating a binary numeral rule and combining it implicitly or explicitly with a digit counting rule before actually applying (any part of) the derived rule to the input aaaaaB. A subject of the latter type, on the other hand, would first generate a binary numeral rule, then apply it, yielding B101, before attempting to find the number of 1's in the binary numeral. Notice, however, that higher order rules may be used to generate (partial) solution rules within any given subproblem.

Further, it should be clear that either type of account may be simulated by the other so that there are no analytic (procedural) differences between them. Any advantages of using both stem from greater flexibility insofar as accounting for the processes of behavior, rather than just end results. Use of this extended formulation in accounting for a realistic problem solving protocol is given in Chapter 14.*

It remains to note that the construction of structural competence theories may reasonably be expected to be a cumulative enterprise. Separate analyses of different content domains may be routinely combined as long as they pertain to the same population of subjects. In particular, the structural competence theory underlying a combined domain is essentially the union of the individual competence theories. In general, with closely related domains, it is reasonable to expect a significant number of rules to be common to each theory so that the combined rule set will contain fewer rules than the sum of the rules in the individual rule sets. In addition, although no empirical work in this direction has been completed, the combination of rule sets could suggest additional and more exacting analysis.

2.8 Global Competence.--In addition to rules of competence that tend to be associated with particular content domains, competence almost necessarily includes more global capabilities as well. Rather than being characteristic of particular content domains, they may enter into a wide variety of different domains and more aptly be thought of as characteristic of given subpopulations. The domains of spatial visualization and moral behavior, for example, tend to have this character in that, although presumably learned, they enter into a wide variety of behavior, and are constrained only by such subpopulation factors as sex and culture, respectively. That is, it is reasonable to suppose that some competencies (rules) are at once relatively independent of and relevant to particular content domains.

Most higher order rules, especially simple ones, tend to have this character. It would be unnatural, for example, to restrict the domain of a higher order composition rule to any particular content domain.

* i.e., Chapter 14 of my forthcoming book, Human Problem Solving: A Synthesis of Content, Cognition, and Individual Differences.

Undoubtedly, the most universally understood, as well as most studied, global competencies (rules) pertain to such concerns as morality, value judgements, logical reasoning, physical space, and the like. Although much of the research in these areas has had a largely empirical cast and has rarely been formulated rigorously in terms of rules, there are some notable exceptions. For example, the work of Taylor (1960) on motivation, Simon (1967) on emotions, and Newell and Simon (1972) on deductive reasoning appear to be formulated in compatible terms. Even the extensive work of the Piagetians is based on the schema -- a rule like, albeit relatively imprecise, construct. The latter research, particularly, clearly demonstrates that epistemic competence (i.e., associated with groups rather than individuals) tends to go through identifiable stages during early development with all that implies for devising competence theories for different subject populations.

Epistemic competence is complementary to that emphasized in the present volume in that it deals with general features of broader domains and populations. Furthermore, once specified, such competence provides additional constraints on analysis that may prove helpful in identifying the competence underlying specific content domains. There is no reason, of course, why one could not choose something like moral judgement itself as the area of specific concern. In this case, rather than developing descriptions of various stages of moral development (Kohlberg, 1971), say, theory development would consist of constructing rule sets that account for the kinds of moral judgements made by individuals at particular developmental levels. Presumably, of course, rule sets at one level should allow for growth consistent with rule sets at later levels. In competence theories of this type, notice also that what are considered to be inputs and outputs would likely have a rather molar character (e.g., inputs might consist of rather complex situations).

THE LEARNING AND BEHAVING SUBJECT

So far, nothing has been said about the actual behavior of people. The competence (rules or procedures) associated with a given content domain and subject population may be thought of as representing the knowledge had, relative to that content domain, by an all-knowledgeable member of the population. It reflects an ideal that may only rarely reflect individual knowledge.

What a particular subject does in a given situation, however, depends not on some ideal but on what he or she knows as an individual. Indeed, in order to talk meaningfully about the behavior of individual subjects in specific situations (Scandura, 1973^a) we must: (1) know how individuals use the specific knowledge that they do have and how this knowledge is affected and modified by interacting with the environment, and (2) have some way of determining what relevant knowledge individual subjects have available.

In section two, we take individual knowledge as given and consider the mechanisms which determine how this knowledge interacts with the environment (i.e., we deal with (1)). Section three deals with measuring knowledge (2), which involves specifying the relationships between competence and knowledge.

3. COGNITIVE MECHANISMS

Information processing psychology is predicated on the belief that human knowledge can be profitably represented in terms of processes (rules, procedures). Given a class of tasks, an information processing cognitive theory of human behavior (relative to that task domain) consists of a procedure which accounts for various facets of human behavior on these tasks. In computer simulation, for example, cognitive theory is equated with a computer program which reproduces the behavior of individual subjects on a given class of tasks (e.g., chess). In contemporary cognitive psychology, a large percentage of information processing research deals with simpler discrimination and short term memory tasks; latency provides one of the favorite indices of process form and complexity.

One of the main goals of contemporary psychological theory, as opposed to competence (e.g., grammar) theories, is to determine which aspects of performance are universally characteristic of human subjects and which aspects are more or less content and/or sub-population specific. Thus, for example, various parts of a program (theory) in computer simulation may pertain to content specific knowledge (and particular subpopulations of subjects) whereas other features may pertain to physiological constraints of the human

information processor (e.g., his limited capacity for processing information [Miller, 1956]; the time it takes to encode information [about 5 seconds], etc.). There are other features that, while neither common to all human beings nor content specific, are characteristic of various subpopulations (e.g., developmental levels, cultures, etc.). The distinction between learnable information and innate physiological capacities corresponds, in computer science, to "the logical equivalence of program and hardware (Gorn, 1973)."

Although there have been numerous attempts to incorporate learning in the information processing view (e.g., Feigenbaum & Feldman, 1963), existing process theories, for the most part, deal primarily with performance. The major goal of this section is to present a simple, cohesive information processing theory and to show how it provides a precise, comprehensive framework for viewing a wide variety of cognitive behavior.

The section is organized as follows:

(1) We propose and illustrate a (innate?) control mechanism that is assumed to govern the way in which known rules interact in producing behavior. This mechanism is just one of probably many "hardware" characteristics of the human information processor, and is assumed to operate with respect to arbitrary domains of knowledge and all subpopulations of subjects. In turn, we show how this mechanism accounts for simple performance, learning (rule derivation), motivation (rule selection), subgoal formation (defining problem situations), and retrieval. Relevant empirical data is presented in Scandura (1973).

(2) Another presumed "hardware" characteristic is the limited capacity for processing information. A theory is proposed which deals rigorously with the question of what counts as a unit (chunk) in processing available rules. Relevant empirical evidence and a stochastic extension of the deterministic theory, which accounts for certain difficult to control variables, is described in Voorhies and Scandura (this volume).

3.1 Representation of Knowledge.-- In the structural learning theory, knowledge is represented in precisely the same way as competence, that is, in terms of labeled directed graphs (rules). The term "knowledge," however, refers to an actual potential for behavior rather than to an account of some idealized, a priori class of potentially observable behavior. More specifically, rules are used to represent behavior potentials of individual subjects rather than to account for group behavior as is frequently the case in existing information processing theories in psychology.

Although all rules of knowledge denote a potential for behavior and are represented in terms of labeled directed graphs, recall that any given rule may play either of two distinct roles. Rules may act on other elements, in which case they serve as operators; or they may be operated upon, in which case they serve as states.

Decision making capabilities (nodes) are also rules, albeit rules which simply classify elements. (It should be emphasized, nonetheless, that decision making capabilities may involve any finite number of predicates [classifications] and, further, that each may be characterized in terms of a class of rules [cf. Landa, 1974].) A partition of strings into those that contain an even number of elements and those that contain an odd number of elements provides an example of a binary decision making capability. Goals in this view are represented as distinguished predicates (decision making capabilities). That is, goals are commands which tell what to do (i.e., what properties must be satisfied); they do not, however, tell how to do it. In the context of our illustration, "Find the binary numeral represented by the number of a's in x," is a goal.

It may also be worth noting that where internal structure is of no special interest, rules may be represented by single arrows, or, in the case of single, indivisible elements, as degenerate decision making capabilities (one element partitions). (That is, instead of having a partition on a set of elements, we have an undivided set which, for purposes of the particular analysis, may be viewed as an indivisible unit.)

Given this mode of representation, the question of central concern in the structural learning theory (of cognition) is how rules interact and are modified as the subject interacts with his environment. The structural learning theory is of the interactionist variety. What the subject does in a given situation depends not only on the situation but what the subject already knows and what he is trying to do. It also depends, of course, on behavioral constraints imposed by the physiological character of the human information processor.

Sections 3.2 through 3.10 are concerned primarily with behavioral constructs which reflect such physiologically based characteristics. Specific attention is given to a presumably universal control mechanism and to processing capacity. Among the constraints not considered are inborn instincts, general arousal level (e.g., Farley, 1974), time required to encode new input (cf. Newell & Simon, 1972), and decoding latencies characteristic of various body organs (e.g., hands, mouth).

My present position is that all such characteristics stem from presumably common "hardware" (human physiology). This does not mean, however, that I am committed to the view that all individual differences are due to environment. To the contrary, while I believe, for example, that all subjects have a fixed capacity for processing information, I also believe, that this capacity may vary over individuals. Such differences presumably are innate and may have a large cumulative effect on learning and behavior over the years. Furthermore, given the data accumulated to date, I take no firm position on whether particular physiological characteristics develop in the womb, or after birth. In the latter event, however, the kinds of characteristics we are considering here are more likely ascribable to physiological maturation than to learning in the usual sense.

3.2 Goal Switching Mechanism.--In the next seven sections, we consider the first hypothesized characteristic of the human information processor, that which pertains to the way in which learned rules interact. More particularly, a control mechanism is proposed and shown to provide a reasonable account of complex human behavior, in all its common ramifications: learning, problem solving, problem definition, motivation, and memory storage and retrieval. The basic idea rests on the assumption that human beings are goal directed information processors, and that control shifts among various higher and lower level goals automatically in a fixed, predetermined manner, according to requirements of the situation.

As a first approximation, the mechanism may be thought of, informally, as operating as follows: Given a task (stimulus and goal) for which the subject has a solution rule immediately available, the subject will actually apply the rule. (While this statement may appear almost tautological, it is an assumption. Rule use does not follow logically from availability.) On the other hand, when no such rule is available, control is assumed to automatically switch to the higher level goal satisfied by rules which do apply. With the higher level goal in force, the subject presumably selects from among available and relevant higher order rules in the same way as he would with any other goal. In effect, if the subject has an applicable (higher order) rule available, then he will use it. Where no such higher order rules are available, the theory assumes that control moves to still higher level goals. Conversely, once a higher level goal has been satisfied, by application of some rule, control is assumed to revert to the next lower level.

This mechanism is deceptive. While extremely simple, it is potentially quite powerful.

3.3 Learning (the acquisition of new rules).--To see how the mechanism might bring about learning and to obtain deeper insight into the mechanism itself, let us first describe the mechanism more precisely.

Given a goal situation (S, G) , the subject tests to see if the problem solution is immediately available (i.e., to see if he knows the solution, an element which satisfies G). If not, control shifts to the solution goal $(G^2=SG)$, consisting of the set of potential solution rules. The subject then tests each of the rules (r) available to him to see if r includes S in its domain $(S \in \text{Dom } r)$ and the range of r contains an element in the goal $(\exists x \in G \text{ such that } x \in \text{Ran } r)$. If such a rule is available, control reverts to the original goal G , and the rule is applied to S . If not, control shifts to the higher level goal $(G^3=HG)$, consisting of higher order rules (h) whose ranges contain a potential solution rule. With G^3 in force, the subject is assumed to test his available rules as before. If one (r') is found (i.e., if $\exists r \in G^2$ such that $r \in \text{Ran } r'$), control reverts to G^2 , and r' is applied to other available rules (in its domain), thereby generating a new (potential) solution rule. In turn, if SG is satisfied by the new solution rule, control reverts to the original goal G and the solution rule is applied to S . The problem is solved if the (potential, solution so generated satisfies the goal.

It should be noted that testing, goal shifting, and rule application are essentially the same at all goal levels. Generalization and formalization of the mechanism is given in Scandura (1973). Furthermore, as noted above, higher order rules are normally identical to other^a rules and obey the same laws of behavior. The descriptor "higher order" refers to the role of a rule in a particular instance of problem solving, not to its basic nature. (In computer science, this is referred to as unstratified control.)

For example, suppose a naive subject is confronted with the stimulus bbbbbB and asked to find the binary numeral which represents the number of b's. There are various alternative ways in which such a statement might be interpreted (e.g., consider the syntax and semantics of the situation). For present purposes, let us assume that the subject understands the problem in a syntactic sense, specifically, that he represents the goal G by

$$G = \text{Find a string } x, x \in \{x \mid x \text{ consists of} \\ \text{a B followed by a string of 0's and 1's}\}$$

Assume also that the only relevant rules available to the subject are

$$(1) r = r_2 * r_1 * r_2$$

$$(2) r_a \Rightarrow r_b$$

(The fact that [1] generates a goal string of the required type only when the input is "aaaaaB" is a special, but irrelevant, feature of this example.)

In this case, we may assume that the subject does not have a binary numeral string (i.e., B101) available which satisfies the goal. Hence, according to the postulated mechanism, control automatically shifts to

$$G^2 = \{r \mid S = \text{bbbbB} \in \text{Dom } r, \exists x \in G \text{ such that } x \in \text{Ran } r\}$$

Since neither of the available rules satisfy G^2 (e.g., (1) applies to strings involving a's rather than b's), control goes to

$$G^3 = \{r' \mid \exists r \in G^2 \text{ such that } r \in \text{Ran } r'\}$$

Rule (2) satisfies G^3 since its range involves rules involving b, one of which within our limited system applies to the stimulus bbbbbB and generates a binary numeral in G. Thus, control reverts to G^2 and Rule (2) is applied to Rule (1) (the only available rule to which it applies), thereby generating

$$(3) r'_2 * r'_1 * r'_2$$

where r'_i ($i = 1, 2$) operates on inputs involving b's but in every other way is identical to r_i . Now bbbbbB is in the domain of rule (3) and its range contains a binary numeral (i.e., a B followed by 0's and 1's) (i.e., rule (3) satisfies G^2). Thus, control reverts to the original goal G and rule (3) is applied to bbbbbB. Since the output B101 satisfies G, the problem is solved, and in the process the subject has learned the new rule (3).

As a second example, consider the situation where a subject has previously learned rules


$$(4) r_2 * r_2 * r_2$$

$$r_1 * r_2$$

These rules are adequate for solving certain of our illustrative tasks. After having learned a number of such specific rules (i.e., having solved a number of specific tasks of this type), it is reasonable to suppose that many subjects may discover (induce) a general method (rule) for solving any such task.

To see how this might come about, suppose the subject also knows rule h, a rule which operates on pairs of rules (like r_1 and r_2) in which the domains are complementary (i.e., where exactly one applies to each element in the given universe of elements [strings of a's]) and generates unitary, general recursive rules which combine the advantages of both. We shall not attempt to precisely formulate such an h but shall rather leave that as an exercise to the interested reader. Once formulated and applied to rules like 4,

h, as a minimum, should generate a general binary numeral generator (rule). One such general rule may be represented by the following labeled directed graph



In addition to illustrating how induction may enter into the theory (for details, see Scandura, 1973a), this example suggests how formerly discrete rules may be integrated into more comprehensive relational nets.*

3.4 Rule Selection.--As formulated, the goal switching mechanism tells what happens when a subject does not have any rule available for solving a given problem. But, it leaves open the question of what happens where two or more rules apply. Which of the alternatives does he select?

In order to account for rule selection, the rule derivation mechanism must be modified so that a rule is applied at a given goal level only where exactly one available rule applies. In particular, if two or more rules apply, control is assumed to move to a higher level goal. Otherwise, the mechanism is assumed to work as before.

To see the implications of this modification, it is only necessary to note that rules may act on two or more rules and output one of them. Higher order rules which operate in this way have been called selection rules (Scandura, 1973_a). As a simple example, consider the stimulus aaaaB and G as above. Assume further that the subject knows rules

$$(6) r_2 * r_1 * r_2$$

$$rg_2 * rg_1 * rg_2$$

where rg_i is a generalization of r_i ($i = 1, 2$) which applies not only to stimuli involving strings of little a's but to strings of any letter elements (e.g., it would apply to cccccB). Since both rules in (6) apply to the stimulus aaaaB and have ranges which contain an element of G, control would presumably shift to G^3 as above.

At this level, some selection rule would have to be available in order for the subject to make a choice. One such selection rule might be described loosely as "use the least general rule" (i.e., check the domains and if one domain is contained in the other, select that rule). Such selection rules may apply (potentially) to a broad class of rule pairs and, further, some empirical evidence suggests that they may be reflected in human behavior (Scandura, 1973a).

3.5 Rule Retrieval.--Although retrieval requires no essential change in the mechanism proposed, there is a connotative difference. Learning and problem solving correspond to operating on given rules (semantic nets, relational nets, labeled directed graphs) and generating new ones. Retrieval, on the other hand, refers to operating on a relational net and extracting a portion of it.** In effect, although goal switching operates in the same way, there is a difference as to whether the to-be-generated rule represents a unit of new knowledge or has previously been embedded in a more comprehensive cognitive structure.

* To my mind, one major open question concerning the goal switching mechanism involves the possible relationships in testing between $\text{Ran } r$ and G^n . In my original formulation (Scandura, 1973a), I had postulated that $\text{Ran } r \subset G^n$. Unfortunately, this condition is not fully adequate. For example, requiring the range of a composition rule to be contained in a goal is overly restrictive because the range of any realistic composition rule will include a wide variety of composite rules, including composite rules that have nothing to do with the task in question (e.g., see Scandura, 1974).

For this reason, I later proposed (Scandura, 1974) that $G^n \subset \text{Ran } r$. This seems to work and is operational in the sense that the relation can be tested algorithmically, but it is not the only possibility. Above, we have proposed still a third feasible relation between $\text{Ran } r$ and G^n . Ultimately, of course, the proper relation can only be determined empirically, and it could turn out to be content specific. Nonetheless, in the absence of critical, empirical evidence to the contrary, we may assume the above as a reasonable approximation.

For example, suppose that the subject has rule $r_2 * r_1 * r_2$ available and is presented with the task of generating the binary numeral representing the number of a's in aaaB? Given the above goal switching mechanism, what additional information (retrieval rule) would the subject need to know in order to succeed? One possibility would be to know a higher order retrieval rule which operates on chains of atomic rules (i.e., composite rules) and extracts subchains. Such a rule might, for example, check the domains of each component in turn to see if the given stimulus (aaaB) is in its domain and, if so, to remember it while checking in turn the ranges and domains of following components until there exists either a mismatch or the range contains an element in the desired goal.

3.6 Problem Definition.-- It is assumed above that the stimulus and goal adopted by the subject is known. In typical problem solving experiments, on the other hand, this information is not available directly. The subject is presented with a statement or description of the problem, to be sure, but this is not necessarily the same as ensuring that it will be interpreted as intended.

Problem definition (e.g., breaking a problem into a series of subgoals) is an important aspect of problem solving (e.g., Newell & Simon, 1972). Indeed, perhaps one of the first things a young child must learn in communicating with an adult about problem solving is that such statements as "solve the following problems" must be interpreted to mean: interpret each of the following problem statements, in turn, and then solve the problem so defined. In short, problem statements are invariably broken down into pairs of subproblems, both of which are important aspects of problem solving.

Although the context is new, no changes are required in the goal switching mechanism within any given subproblem. It remains to be shown how defined problems or series of subproblems can be generated by rules and to clarify the nature of problem definition goals (i.e., goals associated with interpreting subproblems). With regard to the latter, consider the problem statement: "Find the binary numeral represented by the number of a's in aaaaaB." The word "Find" in the statement is a sure cue that the subject is to break the problem into a pair of subgoals; first define the problem described by the problem statement; and then solve the problem so defined.

During the formative school years, and even at the frontiers of research, there are many different goal criteria which, for any given subject, may correspond to the word "define." For one subject, a problem may be adequately defined whenever a stimulus and a goal have been formulated which capture the meaning of certain key words (even though not all of the essential relationships are represented). For another, a well defined problem may involve logical consistency among the parts as well. For example, the process of understanding a problem statement may yield a defined problem that has no solution (e.g., find an albino with brown eyes). Two possible interpretations of the above statement might be: (1) Find (some) numeral obtained by applying r_1 's and r_2 's to aaaaaB and (2) Find (some) binary numeral which represents the number of a's in aaaaaB (i.e., check to see whether the numeral denotes the number five). The former goal would be satisfied by any numeral consisting of 0's and 1's; the latter requires also that the binary numeral represent the number five.

In addition, the process of defining problems (i.e., interpreting problem descriptions) frequently results in a series of subproblems. The problem statement "Find the number of 1's in the binary numeral representation of aaaaaB" provides a ready example.

** In most information processing theories of memory (e.g., Anderson & Bower, 1973; Kintsch, 1972; Rumelhart, Lindsay, & Norman, 1973), memory is also represented in terms of relational (semantic) nets, with retrieval being viewed as searching through the nets. In the present view, this corresponds to a situation where a rule (directed graph) is in the processor and the desired element is a node in that graph. Retrieval is similar to that in the Rumelhart, Lindsay, & Norman (1973) theory in which items (nodes) are "reconstructed" via higher order nodes (which correspond to subrules). The structural learning theory, however, explicitly allows for the retrieval of non-degenerate subrules.

Confronted with a statement of the form "Find C related to B related to A (where A is the given, C the desired output, and B an intermediary)," the natural thing to do (sub-goal to adopt) is to define the problem statement in terms of a pair of subproblems. Thus, given the illustrative problem statement, most adults would break it down into two subproblems: Find the binary numeral representation of aaaaaB, and then find the number of 1's in the obtained numeral. Once defined in this way, control would move automatically to the first subproblem.

Lest there be any confusion on the point, it should be emphasized that problems may be redefined at each given stage. The first subproblem, for example, implicitly allows for interpretation as a series of subproblems. In particular, each subproblem might be characterized by "reduce the number of a's to the left of B by r_1 or r_2 ." Indeed, this is precisely what Ernst and Newell's GPS (1969) and Newell and Simon's (1972) production systems would have the subject do.

There is another key feature of problem definition, however, which is not so easily captured by these other formulations (cf. Banerji, 1974). For example, suppose a subject is presented with the problem statement "Find the binary numeral representing the number of b's in bbbbbbB" but that he does not know any rules that operate on strings involving b's. He only knows the simple atomic rules r_1 and r_2 (which operate on strings involving a's) and the general fact that arbitrary strings having the same numbers of elements can be represented by the same numeral. In the present view, such a statement may be interpreted as (a) find the binary numeral represented by the number of a's in aaaaaB and then (b) assign the same numeral to the b's in bbbbbbB. In the GPS formulation, on the other hand, it is not automatically clear why aab1 (obtained by applying r_2 to aaaaaB), for example, is any closer to the binary numeral for bbbbbbB than the original stimulus itself.

Although we shall not attempt to specify specific rules for interpreting problem statements, it is easy to envision a system such as Winograd's (1972), Shank's (1973), or another recent question-answer system in artificial intelligence, serving this purpose. Further, it may be observed that where a subject knows two or more rules for interpreting a given class of problem statements, selections are assumed to follow as described above.

One additional point is worth mentioning. In empirical studies, it has been commonly observed that when the same problem is presented repeatedly over a series of trials, subjects tend to break it into progressively larger subproblems. This seems to be particularly true in well defined tasks where the atomic operations (rules) are specified to the subject in advance (e.g., The Tower of Hanoi [Klix & Sydow, 1967], missionaries and cannibals [Thomas, 1974; Greeno, 1973]). Given experience with a given problem, it seems likely on succeeding trials, that subjects may tend to derive more and more solution rules for the various subproblems. Correspondingly, the size of the subproblems identified may be expected to increase (and the number of such subproblems decrease).

3.7 Information Storage.--Although nothing new is involved, it seems appropriate, having discussed memory retrieval, to comment briefly on storage in memory. When presented with an item of information, whether it be a simple element, simple rule, or a complex, interrelated network, the item is necessarily presented in some context, either external and/or internal. Storage, in any case, involves deriving a rule which relates some portion of the context (retrieval cue) with the item to be stored. That is, the goal during storage is to derive a rule (network) which takes the context as input and contains the to-be-stored item in its range. In this view, when the subject derives a rule for getting from a given input (e.g., aaaaaB) to the desired output (e.g., B101), he has effectively "stored" the response with respect to the input. When the retrieval cue (input) is presented on another occasion, the subject applies the learned rule to generate (retrieve) the appropriate output.

3.8 Extension of the Mechanism to Multiple Processes.--All of the previous analyses pertain to single processes. Of perhaps greater ultimate interest is the study of more complex problem solving situations which involve two or more processes.

In the idealized control mechanism described above, it is assumed at each goal level that the domains and ranges of rules are tested simultaneously. To allow for multiple processes within a given (sub)problem, the mechanism must be generalized so that these tests are performed in series. Specifically, in the enriched mechanism (Scandura, 1973a), higher level goals are satisfied by rules whose ranges contain elements in next lower level goals.

Once a higher level goal has been satisfied (i.e., a rule satisfying it has been identified), control is assumed to revert to the corresponding domain goal. (The domain goal consists of elements which belong in the domain of the identified rule.) Control here is assumed to shift to higher and lower level (domain) goals as with any other initial goal. Once a domain goal has been satisfied, control reverts to the next lower level, the identified higher order rule is applied (to the domain elements generated), and the process continues. The composition rule (*) introduced above, for example, satisfies a higher level goal in the sense that its range contains a composite r_1*r_2 rule whose range contains a binary numeral output (in G). Hence, control is assumed to shift to the domain goal satisfied by pairs of compatible rules that are consistent with the higher order goal (i.e., S is in the domain of one rule in the pair and $\exists x \in G$ such that x is in the range of the other rule). The only pair that satisfies these conditions is (r_1, r_2) . Once the domain goal is satisfied, control moves to the next lower level goal; the composite rule is applied; and the process continues.

Although the above generalization makes it possible to handle a broader variety of problem situations (where needed domain rules are not immediately available), notice that the mechanism reduces essentially to the earlier one where a needed domain element is immediately available. In this case, the domain goals are satisfied directly; elements which satisfy them do not have to be derived.

3.9 Extensions and Further Comments.--Throughout the above discussion it was implicitly assumed that goals were structureless; they represented unitary tests related to specific tasks, ignoring the broader context (both external and internal) in which the tasks were presented. It is not hard to see limitations of such goals in problem solving reality. Consider, for example, a bizarre situation in which Goldfinger has placed James Bond on a ledge just above a snake pit and has presented him with the challenge of solving a difficult mathematical problem. If he succeeds, he goes free, else... While he is attempting to solve the problem, James must be extremely wary. His goal is a composite one; solve the problem but remain alive in so doing. Any rule used under such circumstances would likely have more to it than mathematics; it would include among its decisions making capabilities some sort of safety valve. If at any point a certain danger threshold is reached, Bond presumably would forget about the problem for a moment and avoid striking fangs. Perhaps a more common (and peaceful) occurrence is a situation scientists often find themselves in -- thinking about a problem while driving one's car to or from work. Under ordinary driving conditions, it is often possible to make reasonable progress. But, in an emergency, things change rapidly.

In addition to these examples, where shifting attention depends on real world events, over which the subject has little control, there are other situations where the triggering events may be a direct result of stimuli produced by the subject himself. Changes in internal stimulation, for example, might well account for such common observations as a person stopping in the middle of a task due to fatigue or hunger. Similarly, results obtained in previous attempts at solving a problem may lead a subject to suddenly take a new tact.

Once a shift from immediate to more general goal components is triggered, a subject is apt to engage global competencies (which may be incorporated into a solution rule as suggested above). Thus, for example, a given person may react in much the same way, in broad classes of problem situations, when tired. Stable individual differences in such cases, presumably, are a function of differences in relevant global capabilities and/or physiological predispositions.

As simple as it appears, the proposed control mechanism clearly has considerable power. It provides a potential basis for explaining various aspects of problem solving: learning, motivation, problem definition, storage and retrieval from memory, and apparently just about every conceivable kind of human behavior. There are some fine points which are still unresolved, of course, and the mechanism provides only a schema for the more specific competencies (rules) likely to be involved in any particular application.

The question of how to isolate the proposed control mechanism(s) for empirical study poses problems, problems not unlike those confronted in contemporary research on memory. When a subject performs on a memory task, for example, it is not always easy to determine whether information comes from short term or long term store (cf. Atkinson & Shiffrin, 1968; Waugh & Norman, 1965). Similarly, in most real world environments, it is not easy to determine which aspects of behavior are due to the proposed built-in control mechanism, which to specific learnable knowledge and, further, as we shall see in the next section, which to a person's limited capacity for processing information.

The approach my collaborators and I have taken has been to study behavior under idealized conditions, where it was known in advance exactly what relevant rules the subject knew and had immediately available and where explicit provision was made to ensure that the subject's processing capacity was not exceeded (e.g., where the subject had all of the time he needed and/or where paper and pencil or other aids were available for recording intermediate results). Relevant experiments are reported in Scandura (1973a, 1974).

3.10 Man: A Limited Capacity Information Processor.--It would be convenient if the control mechanism described above were all that were needed to explain human behavior. As noted above, however, this clearly is not the case.

Throughout the above discussion, no limits were imposed on the number of rules and elements that might be immediately available to the subject during cognition -- in effect, it was implicitly assumed that human capacity for processing information is essentially unbounded. One hardly need refer to the massive literature on information processing, however, to note that, in general, a subject's ability to cope with a given problem situation depends in substantial part on the memory load imposed by the task (e.g., Bruner, Goodnow, & Austin, 1956). Man is unquestionably a limited capacity information processor.

This section deals with the following questions. How does processing capacity interact with the proposed control mechanism in generating behavior? If the memory load imposed by a task is too great, the task may exceed the subject's capability even where the subject "knows how" to solve the task. What, however, constitutes memory load and how is it determined?

As a first approximation, memory load clearly depends on the number of things that must be kept in mind during processing. For illustrative purposes, recall once again our standard example: Let S be the stimulus $aaaaaB$, G the goal, $*$ a higher order composition rule, and r_1 and r_2 rules for operating on strings of even and odd numbers of a 's, respectively. Upon initial presentation, the subject must keep in mind each element in the set

$$A = \{S, G, *, r_1, r_2\}$$

At this point, the control mechanism takes over and introduces the higher level goal G^2 , thereby increasing memory load. With G^2 in control, the individual elements in set A are tested. Since none satisfies G^2 , the mechanism takes over again introducing G^3 , adding further to memory load. Upon subsequent testing, we find that G^3 is satisfied by $*$. Here, the control mechanism moves control to the domain (goal) of $*$. (The domain goal replaces G^3 , so memory load remains constant.) With this domain goal in force, the available rules are again tested, this time singling out the pair r_1 and r_2 . At this stage, $*$ is in control and operates on the pair r_1, r_2 , yielding the composite rule r_1*r_2 . By the control mechanism, control goes to G^2 and the available elements (including r_1*r_2) are tested. Since r_1*r_2 satisfies G^2 , r_1*r_2 is applied to stimulus S (with $G^2, *, r_1$, and r_2 no longer adding to memory load) yielding the desired output. As a final step, control goes to G and the output is tested against it.

This example provides a general, qualitative account of how memory load may change during processing. Because it deals exclusively with rule extensions, rather than their internal structure, however, it rather lacks detail and precision. As Miller (1956) has noted, memory load is not a simple function of the number of nominal elements that must be retained but rather of the number of "chunks" involved.

Although there are many examples of "chunks" in the literature (e.g., digits, letters, words, equivalences between base 10 and binary numerals, etc.), it is not clear exactly what a chunk is, or, to put it differently, what counts as a chunk in computing memory load. In addition to such entities as digits, words, and so on, my present view is that a chunk may also be, for example, a rule, a part of a rule, or a goal. Indeed, given the equivalence of data and process (e.g., Gorn, 1973), it must be so. In the current theory, what is a chunk depends on what elements (i.e., rules and goals) are available and the current stage of processing. There are three basic types of control that may obtain at any given point in time: (1) The control mechanism controls shifts among goals and rules. (2) Rules control generation. (3) Goals control testing (of rules).

Memory load is computed relative to what is in control (i.e., the mechanism, rule, or goal) at each moment. Let us suppose, for example, that a given rule is engaged. Then, memory load is determined as follows. At each stage of a computation (i.e., at each stage of an application of a rule to a given stimulus), a specific number of elements must be retained in the processor in order to determine future outputs and operations. At the conclusion of any given stage, it may be possible to eliminate certain of the elements required at a preceding stage, and new elements may be added. In adding $45 + 71$, for example, it is necessary to remember the two units digits 5 and 1 before summing them, but, afterwards, only the partial sum 6 must be retained.

The number of elements that must be retained in the processor at a given stage of a computation constitutes the memory load at that stage. The memory load of an entire computation is the maximum of the memory loads associated with the various stages of the computation. Details as to how memory loads may be determined with respect to individual rules, together with related empirical support, are given in Scandura (1973a) and Voorhies and Scandura (1974).

With respect to higher order rules, memory load depends on the degree to which the internal structure of the relevant input rules (and goals) must be taken into account in carrying out the individual decisions and operations of the higher order rule in question. The more the individual input rules act as units, the easier they are to operate on, and the lower the memory load. Conversely, memory loads will be relatively high where explicit attention in higher order rules is given to each component. In our basic illustration involving the control structure $*$, r_1 and r_2 were treated as wholes. No mention was made of internal structure.

In general, each rule imposes one unit of memory load on the processor just to the extent that atomicity reflects psychological reality in a given situation. Each additional subrule (subgraph) or subelement similarly increases memory load by one. Correspondingly, a unitary goal is also assumed to impose a load of one on the processor. With a composite goal, the load would be relatively larger (since a greater number of things would have to be kept simultaneously in mind).

Although it is beyond the scope of this treatment to delve into the matter, the distinction between unitary and composite goals may have some relation to the notion of intensity of concentration. In particular, the more unitary a goal, the greater the degree of concentration that may be possible, presumably due to the lower load imposed by the goal on the processor. In such a situation, however, the subject's behavior is apt to be less flexible. Perhaps this has something to do with the phrase "single minded."

4. INDIVIDUAL DIFFERENCES

In Section 2, we saw how arbitrary content domains can be analyzed to determine the underlying processes (higher and lower order rules) characteristic of given subject populations. In Section 3, we discussed the psychological characteristics of the human information processor and showed how they provide a comprehensive framework for viewing

much of human cognition. The specific cognitive content (rules) available to individual subjects was assumed.

In this section, we consider the question of how to identify (and characterize) the rules available to given individuals.

4.1 Background.--Current information processing approaches to this question have:

- (1) centered on attempts to identify processes which provide the best overall account of performance by some population (group) of individuals on some class of tasks, or
- (2) been concerned with identifying the processes used by particular individuals.

In the former case, underlying rules or processes are postulated, predictions based on these processes are made, and the predictions are compared with averaged data (usually averaged over items and groups, although sometimes over items within individuals). Such an approach ignores individual differences (involving individuals, items, or both). For example, suppose that we identify two processes by which the binary numeral representing the number of a's in x_B (where x is a string of a's) might be generated, namely rule

$$(1) \quad \text{START} \left. \begin{array}{l} \text{---} r_1 \\ \text{---} r_2 \end{array} \right\} \text{STOP}$$

and, collectively, $r_1, r_2, *$. In each case, (random) parameters can be introduced to represent the times required to perform the various parts of the respective processes. For example, $\alpha, \beta,$ and γ might be parameters representing the times required to carry out $r_1, r_2,$ and $*$, respectively. Ignoring the time required to determine whether the number of a's is odd or even (which would involve at least one additional parameter), the time required to apply rule (1) would be $x_1\alpha + x_2\beta$, where x_1 and x_2 denote the number of applications of r_1 and r_2 , respectively. (Presumably, by adding suitable parameters, one could also represent latencies involving the interactions among rules $r_1, r_2,$ and $*$.) Estimates of these parameters could be obtained by standard regression methods using empirical (group) data. The relative adequacy of the respective "theories," then, would be determined by substituting these estimates into the latency expressions and comparing the resulting latency predictions with the obtained average data. Finally, the process which provides the best fit would be chosen as the preferred account. (Note: This example is for illustrative purposes only. For an actual study along these lines, see Suppes and Groen, 1967.)

The main problem with this approach is that individual processes may be obscured by group data. The fact that a process provides the best overall account says nothing directly about the processes used by individuals. (Indeed, as became so clear as a result of the controversy concerning all-or-none versus incremental learning during the early 1960's, completely contrary statistical assumptions may lead to essentially the same predictions regarding average behavior. For further details and discussion along the lines of our example, see Scandura [1973a, pp. 259-262; 1967].)

The second approach is perhaps best typified by research on computer simulation. In computer simulation the goal is to devise explicit computer programs (rules) which perform on given classes of tasks (e.g., crytoarithmetic, chess) in the same way as individual subjects. That is, a given program may be thought of as representing the knowledge relative to a given task domain by a particular subject.

Although there is no entirely systematic method for devising simulation programs, the basic approach is not unlike that described in Section 2 on competence. Individual subjects may be observed while solving a variety of representative problems and programs devised to parallel these observed processes. (In existing simulations, such programs typically incorporate such performance constraints as a limited capacity for processing information [cf. Newell & Simon, 1972].) Finally, the outputs of trial programs are compared with new behavior, with subsequent refinements and evaluation where necessary.

To date, the approach seems to be useful primarily as a means of modeling performance. Little progress, for example, has been made in dealing with learning. Nonetheless, computer simulation research perhaps more than any other single movement is responsible for the mass shift in experimental psychology toward information processing.


From the present point of view, the simulation approach suffers two major limitations. (1) It can be extremely inefficient. Whereas rules designed to deal with averages often fail to account for individual behavior (or knowledge), the simulation approach too often fails from the other end. A simulation program may adequately represent the knowledge of particular individuals but not that of many (most) others in the population. In general, simulation programs must be constructed anew for each subject. (2) Simulation programs generally confound control mechanisms with rules representing specific knowledge. Where (easily) separable at all, as in Newell and Simon's (1972) production systems or Minsky and Papert's (1972) heterarchical systems, for example, there is little direct evidence or justification to suggest that the control mechanisms are characteristic of humans. As described in Section 3, more veridical separation (of control and knowledge) could have important advantages in accounting for cognitive activities (e.g., learning) other than simple performance.

4.2 Assessing behavior potential.--The structural learning theory combines many of the advantages of both of the above methods. The competence rules are used as an instrument of sorts with which to measure individual human knowledge. More specifically, the theory tells how, via a finite testing procedure, one can identify which parts of given rules in a competence theory individual subjects know -- that is, which rules or parts thereof accurately represent their potential for behavior. The rules in a competence theory in a very real sense serve as rulers of measurement, and provide a basis for the operational definition of human knowledge.

Let us briefly review how this may be accomplished (for details, see Scandura, 1973a; Durnin & Scandura, 1973). Note first that all rules may be represented as labeled directed graphs, and that the various components of any directed graph (e.g., the arrows and nodes) can always be broken down far enough (i.e., into simple enough components) so that each subject in a given population is able to perform each component perfectly or not at all (cf. Suppes, 1969; Scandura, 1970). In short, each component step of a procedure (rule) may be assumed to act in atomic fashion. (The above assumes, of course, that the subjects in the population share a common culture and that the rule of competence accurately reflects that culture relative to the task domain.)

In general, a path through a rule acts in atomic fashion if and only if each component acts in atomic fashion. Success on any path requires perfect performance on each component. Failure on any one part implies failure for the whole; the chain is only as strong as its weakest link. Furthermore, there are only a finite number of paths through any given rule since we do not distinguish paths according to the number of repetitions of loops.

For present purposes, let us assume that the component operations and decisions of rule (1) above act in atomic fashion. Then, there are three distinct paths through rule (1), represented by



Paths (1a) and (1b) are obviously atomic since they each involve only a single operation. Path (1c) is also atomic since each of its components is.

Collectively, the paths of a rule impose a partition on the domain of stimuli to which the rule applies. That is, each path makes it possible to generate responses to a uniquely specified equivalence class of stimulus items and to no others. For example (if we ignore for argumentative purposes the final step of generating a "1" to the left of a string of 0's, then), path (1a) applies to all strings which contain a number of a's equal to some power of two (e.g., aaaaaaaB) and path (1b) applies to all strings which contain a number of a's equal to $2s + 1$, where s is the number of a's in another string to which path (1b) applies (e.g., path (1b) applies to aaB, so it therefore also applies to aaaaaaaB because $7 = 2 \times 3 + 1$). Path (1c) applies to all other strings.

The fact that the paths of a rule partition its domain makes it possible to pinpoint through a finite testing procedure exactly what it is that each subject in the corresponding population knows relative to that rule. It is sufficient for this

purpose to test each subject on one item selected randomly from each equivalence class. Success on any one item, then, implies success on any other item drawn from the same equivalence class, and similarly for failure. In effect, individual knowledge (behavior potential) is represented in terms of rules, specifically in terms of subrules of given rules of competence.

Suppose, for example, that a subject is tested on the strings aaaaaaaB, aaaB, and aaaaB, corresponding to paths (la), (lb), and (lc), respectively. Also assume that the subject is successful on the first item but not on the second or third. In this case, the subject's knowledge would be represented by rule (la).

In constructing (sub)rules to characterize individual knowledge, it is important to keep in mind certain hierarchical relationships among the various paths. For example, if a subject succeeds on a test item corresponding to path (lc), it can reasonably be assumed that the subject will also succeed on paths (la) and (lb) because this path (now viewed as a rule) also applies to the items associated with paths (la) and (lb). Success on paths (la) and (lb), on the other hand, does not necessarily imply success on path (lc) because it is possible to learn repeated application of rule r_1 (path la) and/or rule r_2 (path lb) without necessarily having mastered the items corresponding to path (lc). Hence, the individual knowledge of a person who knows paths (la) and (lb), but not path (lc), would be represented by a discrete pair of rules, whereas a subject who also knows path (lc) might be represented by rule (l) (corresponding to path lc). In effect, the question of whether to represent knowledge as discrete networks (directed graphs) or as integrated networks depends on how the testee's performance relates to the hierarchical relationships among the various paths. If success on prerequisite paths implies success on a superordinate one, then knowledge may be represented in terms of the integrated network. Otherwise, discrete representation is indicated.

The above examples make clear that the knowledge attributed to different individuals may vary even though only one rule of competence may be involved. The idea is directly comparable to measuring different distances with the same ruler. Fortunately, none of this is idle speculation. Support stems from a wide variety of data collected by Durnin and myself (Scandura, 1973a; Durnin & Scandura, 1973), involving a large number of different tasks, with subjects ranging from preschool children to Ph.D. candidates. When run under carefully prescribed laboratory conditions, where it was known what the atomic rules were, it was possible to predict performance on new items, given performance on initially selected test items, with over 96% accuracy (Scandura, 1973a). When the testing took place under ordinary classroom conditions, where the subjects were run as a group and the competence theorist had to make judgements concerning the atomic rules, the predictions were accurate in about 84% of the cases (Durnin & Scandura, 1973).

4.3 Discussion.--Although the structural learning approach retains many features of the methods of average and simulation, it is not identical to either. Unlike the method of averages, where information processing models (rules) are designed to reflect the average performance of a group of subjects, rules of competence are designed to reflect idealized performance, which in some sense reflects the ultimate in performance to be expected of the subjects in the population in question. Furthermore, although accounts of average behavior are useful for many purposes (e.g., in devising overall instructional treatments), the structural approach, even here, allows more detailed characterization. For example, because the number of distinct paths through a given competence rule is finite, there are only a finite number of different combinations of paths (subrules) that might represent individual knowledge. Hence, it would be possible, by testing a broad sample of subjects, to determine percentages of subjects whose knowledge can be best represented by various subrules. Even single subrules may be used to summarize various aspects of overall behavior (much as with averages). For instance, the rule which accounts for the behavior of the largest proportion of subjects may be said to provide a "modal" account. Further, it takes little imagination to conceive of probabilistic rules which reflect average behavior directly. For example, instead of deterministic decisions, one could introduce probabilities.

Although the structural learning approach is more efficient than that based on simulation in that one competence rule can be used to characterize the knowledge of all

individuals in a given population, it has disadvantages in its simplest form. In particular, where two or more distinct "cultures" are combined in the same population, two or more distinctly different rules may be required to adequately represent the underlying competence. Ordinary subtraction, for example, is often carried out by quite different methods (rules), depending upon the culture in question. In contemporary American schools, the "borrowing" method is almost universal, whereas in an earlier day the "equal additions" method seemed to be favored. In Europe, even today, still a third method (of complements) is very common. Because simulation keys on the individual, such variations do not cause much difficulty.

In the structural learning formulation, the competence theorist identifies as many different rules as he feels are necessary to account for the various possibilities. The individual rules, then, may be routinely combined (by anyone familiar with programming) to form a single, unitary rule which combines the features of them all. To the extent that a combined rule is compatible with a subject's knowledge, it may be used as a basis for assessing behavior potential as before. Combined rules have the effect of distinguishing between test items that involve distinct paths of one or more of the various component rules. (For details, see Scandura, 1973a; Durnin & Scandura, 1973.)

At the other extreme, of course, a content domain may be too big. It may be so complex that it is extremely difficult, if not impossible, to devise a single rule of competence which accounts for the domain. Indeed, this is the norm in research on artificial intelligence.

4.4 Distinguishing Between Learning and Performance During Testing.--As it stands, the proposed approach to assessing behavior potential has an important, basic limitation. One can never be sure after testing, whether in fact the subject performed successfully because he knew a solution rule prior to testing, or because he derived a new one during the testing. In effect, test behavior may be based on direct performance or on learning.

Superficially, it might appear that the process of testing may affect what the subject knows in a manner analogous to the Heisenberg principle in physics. On tasks associated with particular solution rules, however, these alternatives are immaterial in so far as predicting behavior is concerned. Although there may be measurable differences (e.g., latency*) attributable to behavior based on learning versus performance, precisely the same success-failure predictions will be made.

A basic difference between performance and learning alternatives, however, may arise during testing when the given task domain is larger than that associated directly with the performance (solution) rule in question. In this case, the potential to learn may reflect itself in considerably broader behavior potential. The essential difference between performance and learning alternatives can be seen in the following example. Consider the task domain consisting of input strings of the form xB and $x'B$, where x is a string of a 's and x' is a string of b 's, and outputs of the form By , where y is the binary numeral representing the number of a 's or b 's in the corresponding input. Let those instances involving a 's as inputs constitute the (sub)task domain of primary interest. In this case, rule (1) provides a direct performance account. An alternative learning account of this (sub)task domain is provided by rules r_1 , r_2 , and a higher order recursion rule which operates on complementary rules like r_1 and r_2 and generates recursive combinations of the components (like rule [1]).

Clearly, any behavior profile associated with this (sub)domain can be accounted for either by rule (1) (the performance account) or by an alternative account based on learning. These alternative accounts, however, are not equivalent with regard to behavior on the rest of the task domain. Learning (account) subjects, for example, who are trained on rules r_1' and r_2' (which correspond to r_1 and r_2 and apply to inputs involving b 's) may be expected to perform better than performance subjects who are similarly trained. Specifically, in the former case, but not the latter, success would be expected on any input $x'B$, involving the same number of b 's as some success input involving a 's, and similarly for failure. As the reader may verify for himself, different success-failure profiles may be obtained by assuming different patterns of availability of initial (sub)rules.

* Consideration of latency data is beyond the scope of this paper, although it is not necessarily beyond that of the basic theory (cf. Scandura, 1973a).

In spite of these differences, it should be emphasized, for example, that the above learning account can be simulated by introducing a single, more complex rule which takes as inputs not only strings of a's or b's but pairs of individual component rules as well. Such a simulation rule would amount to combining the postulated learning mechanism with the individual lower and higher order rules. My personal preference for distinguishing between the presumably innate learning mechanism and rules representing specific knowledge is clear. If we assume this basic mechanism, then we can assess behavior potential, relative to arbitrary task domains, by testing separately with respect to each individual rule. This would not be possible where the control mechanism and the individual rules are incorporated into the same simulation rule.

The same type of analysis can be applied with task domains where two or more rules may account for a given (sub)domain. Although their joint effects may be determined by combining the individual rules, as indicated above, this may result in a loss of predictive power on other tasks in the larger domain where selections might have to be made. For example, consider the illustrative (sub)domain above together with rule (1) and, say, the degenerate special rule r which operates on aB , aaB , $aaaB$, and generates the outputs $B1$, $B10$, $B11$, respectively. Assume also that performance on these items via rule r is faster than would be expected by application of rule (1). In effect, the subject uses the special rule r on these three items and rule (1) on the others. Although such a performance profile may be accounted for adequately in terms of (some subportion of) a combined rule involving special rule r and rule (1), doing so would have a limitation similar to that described above involving learning. That is, if a subject's behavior is generated via application of a selection rule to alternative rules r and (1), rather than by direct application of the combined performance rule, then this would have important implications for performance beyond the immediate subdomain of interest. In particular, where alternative rules may be used, individuals tend to be systematic as to which type (of rule) is selected under which conditions. To see this, suppose, say, that the selection rule is defined by "Select the more specific rule," and envision special rules r' corresponding to rule r and rule (1) corresponding to rule (1), which involve b's as inputs.

To summarize, it would appear that whenever a given task domain is too big, so that it is not possible (or feasible) to identify a single performance rule underlying it, or is too small, in the sense that two or more rules account for (some part of) it, it is often desirable to proceed with the next step of the analysis. That is, prior to assessing behavior potential, the competence analysis should be extended (see Section 2) to identify the underlying higher order derivation and selection rules. In this case, if a subject happens, say, through special training, to know a particular (solution) rule for a certain subdomain, and not related higher order rules, then this will show up directly in the testing. The subject will succeed on tasks in the subdomain but will fail on other tasks solvable via the higher order rules.


Although it is beyond the scope of this paper to consider the matter, it would be desirable to extend this type of argument to higher order rules. For example, it makes a difference in predicting performance whether a solution rule is derived by direct application of a single higher order rule, or by a combination of problem definition (subgoal formation) rules and two or more simpler higher order rules (which apply individually to the subproblems).

4.5 Context.--Throughout the above analysis, we have assumed that all relevant knowledge is immediately available to the testee. Testing effects, due to the human processor's limited capacity for processing information, have been ignored.

In this section, we consider briefly how performance on given test items may be sensitive to context. Thus, for example, if the testee is presented with a novel item after solving a series of test items of a misleadingly similar type, he may fail even where he might otherwise succeed. The water jar problems of Luchins (1942) provide a well known prototype (see chapter 14 Scandura [in process] for more details). After solving a series of such problems using one rule, Luchins' subjects failed to see another simpler mode of solving a new water jar problem. Although Luchins did not run a parallel condition of testing on the novel problem before the initial series of problems was presented, it seems likely from an analysis of the problems that the

novel one would have been easier under these conditions.

The obvious explanation for such context effects is that the control processor may be filled with rules that are either adequate or inadequate to the task. Thus, after solving a series of tasks of a given type, the processor is apt to be filled with rules that are inadequate for solving new problems which, although they appear similar, require a different solution procedure. More generally, of course, the contents of the processor are determined by a variety of factors. In addition to being affected by attempts on previous test items, for example, the immediate context of a given test item may have an important effect. Suppose, for example, that a subject is shown how to cover the nine dots in



with five strokes (lines) of a pencil without lifting the pencil off the paper, and then is asked to do the same thing with only four strokes. If you don't know the answer, try the problem before reading the footnote.*

* Hint: Try the problem allowing some of the strokes to extend beyond the boundary of the dot array.

5. RELATIONSHIPS TO OTHER PROBLEM SOLVING RESEARCH

The proposed theory provides an intuitively simple, apparently viable mechanism by which individuals may be assumed to bring the knowledge they have to bear in problem solving. It also provides a basis for analyzing the specific competencies underlying given problem domains and for using such competencies to determine the relevant knowledge available to particular individuals.

It is still fair to ask how the theory fits into the current zeitgeist of problem solving research. (1) Does the structural learning theory provide reasonable explanations for the results of available psychological research on problem solving? (2) Does the theory provide a basis for explaining subject protocols obtained during the course of problem solving? Let me just consider a couple of examples.

5.1 Experimental Research -- Regarding Question (1), consider a result common to a number of earlier problem solving studies (e.g., Maier, 1930). Given a problem for which all of the components of a solution procedure are available, some subjects succeed on the problem whereas others do not. To the extent that memory is not a factor, explanation follows directly from the idealized theory. Those subjects who succeed can be distinguished from those who fail by the availability of an appropriate higher order rule.

In another well known study, Luchins (1942) found that, after solving a series of water jar problems using one solution procedure, subjects strongly tended to use the same procedure on a subsequent problem even though the problem could also be solved by a simpler method. On initial problems in the series, the subjects presumably had to derive the solution procedure in order to succeed. Once derived, however, the procedure could be assumed to be directly available (i.e., active) in the processor on subsequent problems. Since the procedure applied in all of the subsequent problem situations, including the one potentially solvable by a simpler rule, the theory would predict that the former procedure would be used. According to the structural learning theory, subjects are assumed to derive new solution procedures only where some previously learned solution procedure is not available. Prior constraints on a problem, of course, change the problem situation. To some extent, this may be what happened when Luchins told his subjects, "Don't be blind."

Before moving to protocol analysis, I would like to comment on the commonly made distinction between propositional and algorithmic (rule) knowledge. Although such a distinction can be made, in my opinion it has nothing to do with the nature of knowledge per se.

Whether knowledge (a rule) is to be considered as propositional or algorithmic depends not on any of its properties as such, but rather on the use to which the rule is being put by the subject at any given point in time. Knowledge that is actively being used, as an operator, corresponds roughly to the term "algorithmic," although a false distinction is sometimes made between ongoing operations involving symbol manipulation and, for example, manipulation of a geometric figure. Knowledge that is being acted upon or that is (to be) generated (e.g., a goal), acts as a state, and corresponds to "propositional" knowledge. Again, however, because of certain confusions in the literature the above correspondence is not rigorous (e.g., see Scandura, 1974).

I believe that a similar point has been made to psychologists by Alan Newell; indeed, it is fundamental throughout computer science where it goes under the rubric of "unstratified" control (e.g., Gorn, 1973).


In psychology, much of the confusion stems from the fact that the rules needed in problem solving may not stand alone but may be embedded in more encompassing relational nets of the sort discussed by Bower (1972), Kintsch (1972), Rumelhart, Lindsay, and Norman (1972), and Quillian (1968). (At a formal level, there is no distinction between a relational net, as used by these authors, and the labeled directed graphs used to represent rules in the structural learning theory [cf. Scandura, 1973a, 1973b].)

A complete theoretical account of the phenomena involved will necessarily require integration of a number of loose ends. Although it is not hard to conceive of (higher order) rules which extract subrules from available nets, however, very little research which bears directly on this problem has been done, either by us or by others. Nonetheless, the general nature of retrieval in this view almost certainly will have more in common with the reconstruction of items from memory (e.g., see Rumelhart et al., 1972) than it will with searching pathways from node to node (e.g., see Anderson in Bower, 1972),

where the nodes in the relational net are at the same level of abstraction. The real key in theory development, I think, will not only be the precise specification of retrieval rules but also extension of the proposed method of determining memory loads to determine forgetting priorities (re the limited capacity processor) during problem solving. As it stands, Voorhies and I (this volume) have reasonable confidence in the method only where a single rule is active (in the processor).

5.2 Problem Solving Protocols -- Granting that the results of problem solving experiments can be readily interpreted within the structural learning theory (as is the case with some other theories of problem solving), questions remain as to the theory's adequacy as a means of accounting for the process of problem solving -- for example, actual protocols of subjects in solving problems. This requirement is more stringent than in most experimental studies.

To see what might be involved, consider the problem of "circumscribing a circle about a given triangle using straightedge and compass."



The protocol goes as follows:

"What does this mean? Ah yes, that we have a certain triangle -- and we want to construct a circle that goes through the three vertices of the triangle. What should I do first?... If I knew the center of the circle, it would be easy. Let me see if I can find the center. How can I do this? Let's see, the center is equi-distant from the three vertices of the triangle. I don't know a procedure off hand. But perhaps I can figure one out.

What would happen if I bisected two of the angles?... The bisectors would meet at a point. Let's try it. Too bad! Their intersection isn't necessarily equi-distant from the three vertices.

What else can I do? Let's see -- I know how to find the locus of points equi-distant from two points. So! I can do this twice, each time with a different pair of vertices. Yes, that might work. Let me try it. Good! I have the center.

Now I need to construct the circle. That's easy. I can measure the distance from the center to one of the vertices and then I can use that distance as a radius to construct the circle. (The subject carries out the constructions and solves the problem.)"

Placing emphasis on the postulated theoretical mechanism, this particular protocol might be explained as follows. When confronted with the problem statement and asked to solve "the" problem, the subject apparently broke the problem defined by this statement into a series of subgoals: (1) Find a representation of the problem (i.e., identify the given, the goal, and any information about the solution) and then (2) solve the problem so defined. This method of breaking a problem down into subgoals appears to be broadly applicable and undoubtedly is learned very early. The presentation of almost any problem statement together with the direction "solve the problem" is almost universally interpreted to mean, first interpret the statement (i.e., identify the problem involved) and then solve the problem.

Once the problem was defined in this way, control according to the theory went to the first subproblem (where the stimulus is the statement and the goal is to find its meaning). Here, the subject appears to have a learned rule available by which the meaning can be determined. He seems to know, for example, what triangles and circles are, that "circumscribed" refers to the vertices of the triangle, and the necessary semantic grammar. It is beyond the scope of this analysis to get into the linguistic intricacies involved but it seems reasonable to suppose that analyses such as that of Winograd (1972) could be adapted for these purposes.

Once having determined the given problem, the subject appears to break the problem into a pair of subproblems: (1) Find the center of the goal circle (which the subject notes is equi-distant from the three triangle vertices). (2) Construct the circumscribing circle, given the center.

The subject apparently does not know a procedure for solving the first subproblem off hand so control shifts to the higher level goal of deriving one. The protocol becomes sketchy at this point but it appears, next, that a higher order conjunction rule (which generates conjunctions of rule pairs that apply to the same stimulus) may be either available or directly accessible (retrievable) to the subject. The next statement in the protocol suggests that the rule for bisecting an angle is available in the processor (along with possibly other elements). It is hard to say why this is so, without more information, but one possibility is that the subject may in the recent past have been confronted with a similar problem (e.g., inscribe a circle in a triangle) in which this rule was required for solution. In this case, the conjunction rule would be applied and the outputted, combined rule would be tested to see if it applies to S and has a range that contains G. Indeed, the subject seems according to the protocol to be checking to see if the point of intersection of the two bisectors is equidistant from the vertices. Rather than using a logical argument for this purpose, the subject's comments suggest that a quick sketch (or image) convinced him that the intersection so obtained is not equi-distant from the three vertices.

At this point, according to the theory, the subject starts the problem over. Although the protocol again is not explicit, it seems reasonable to suppose that the subject retained the higher order rule and searched for a domain rule which applies to the triangle (twice) and such that the cross product (intersection) of its range with itself contains the goal (i.e., the set of points equidistant from three arbitrary vertices).

This time, the subject recalls (retrieves or generates) a rule for constructing the locus of points equidistant from a pair of points, and generates a conjunction of the rule with itself. This conjunction satisfies the higher level goal, so control goes to the original goal and the subject actually uses the rule to construct an intersection of loci, namely, the center of the circle.

Next, control goes to the second subproblem, which appears to have been solved more directly. In this case, it would almost seem that the subject broke the second subproblem into even simpler subproblems. First, the subject measured the distance from the center to one of the vertices and, then, used that distance and center point to construct the goal circle. The solution to each subproblem appears to have been direct.

This example illustrates an important limitation of protocols. All protocols, even this relatively explicit one, have frequent gaps and provide only the barest suggestions concerning the processes actually used in problem solving. Even where a fairly explicit theory of problem solving is available, it is often difficult to distinguish between alternative explanations.

In spite of this general limitation, my interpretation does demonstrate the flexibility of the structural learning theory. It would appear, given any protocol, that some reasonable account of it can be devised that is consistent with the theory.

5.3 Further Comments -- The fact that the structural learning theory provides a way of explaining existing problem solving data tells only part of the story. Other theories purport to do the same thing. The computer simulation theories of Newell and Simon (1972) and Ernst and Banerji (1972) are particularly noteworthy in this regard.

There are, of course, important conceptual differences as well as similarities between these theories and the structural learning theory. The structural learning formalism (Scandura, 1973a), for example, is computationally (mathematically) equivalent to the problem space approach; both are bounded above by the power of a universal Turing machine (e.g., see Banerji, this volume).

Perhaps the major advantage of the structural learning theory is that it provides a more comprehensive picture of complex human functioning. Its potential power resides in its conceptual unity with regard to content, cognition (including learning, problem definition, motivation, memory), and individual differences.

References

- Anderson, J.R., and Bower, G.H. Human associative memory. New York: Halstead Press, 1973.
- Atkinson, R.C., and Shiffrin, R.M. Human memory: A proposed system and its control processes. In K.W. Spence and J. Spence (Eds.) The psychology of learning and motivation: Advance in research and theory. Vol. 2. New York: Academic Press, 1968.
- Banerji, R.B. A formal approach to knowledge in problem solving. This volume.
- Bloom, B.S. Recent developments in mastery learning. Educational Psychology, 1973, 10, 53-57.
- Bower, G.H. A selective review of organizational factors in memory. In E. Tulving & W. Donaldson (Eds.), Organization of memory. New York: Academic Press, 1972.
- Bruner, J.S., Goodnow, J.J., and Austin, G.A. A study of thinking. New York: Wiley, 1956
- Carroll, J.B. A program for sentence generation. Journal of Structural Learning, 1974, in press.
- Carroll, J.B. and Freedle, R.O. (Eds.) Language comprehension and the acquisition of knowledge. Washington, D.C.: V.H. Winston, 1972.
- Chomsky, N. Syntactic structures. The Hague: Mouton Press, 1957.
- Durnin, J.H. and Scandura, J.M. An algorithmic approach to assessing behavior potential: Comparison with item forms and hierarchical technologies. Journal of Educational Psychology, 1973, 65, 262-272.
- Ernst, G.W. and Banerji, R.B. Limitations in pattern recognition and problem solving. Proceedings of the ACM National Conference, Boston, 1972.
- Ernst, G.W. and Newell, A. GPS: A case study in generality and problem solving. New York: Academic Press, 1969.
- Farley, F. Basic process individual differences. Talk given at AERA Annual Meeting, April, 1974.
- Feigenbaum, E.A. and Feldman, J. (Eds.) Computers and thought. N.Y.: McGraw Hill, 1963.
- Gagne, R.M. The conditions of learning. (2nd edition) N.Y.: Holt, Rinehart, & Winston, 1968.
- Glaser, R. Objectives and evaluation: An individualized system. Science News, 1967, 1-3.
- Gorn, S. On the conclusive validation of symbol manipulative processes: How do you know it has to work? Journal of the Franklin Institute, 1973, 296, 499-518.
- Greeno, J.G. The structure of memory and the process of solving problems. In R. S. Iso (Ed.) Contemporary Issues in cognitive psychology: Loyola symposium. New York: Halstead-Wiley, 1973.
- Hilgard, E.R. and Bower, G.H. Theories of learning. (3rd edition) New York: Appleton-Century-Crofts, 1966.
- Kintsch, W. Notes on the structure of semantic memory. In E. Tulving and W. Donaldson (Eds.) Organization of memory. New York: Academic Press, 1972.
- Klix, F. and Sydow, H. The organization of information processing in problem solving behavior. Editions Du Centre National De la Recherche Scientifique, 15, quai Anatole-France-Paris VII, 1968.
- Kohlberg, L. From is to ought: How to commit the naturalistic fallacy and get away with it in the study of moral development. In T. Mischell (Ed.) Cognitive development and epistemology. New York: Academic Press, 1971.
- Landa, L.N. Algorithmization of learning and instruction. Englewood Cliffs, N.J.: Educational Technology Publications, 1974.
- Luchins, A.S. Mechanization in problem solving: The effect of Einstellung. Psychological Monographs, 1942, 54, No. 6.
- Maier, N.R.F. Reasoning in Humans. I. On direction. Journal of Comparative Psychology, 1930, 12, 115-143.
- Merrill, M.D. and Gibbons, A.S. Heterarchies and their relationship to behavioral hierarchies for sequencing content in instruction. 1974, This volume.

- Miller, G.A. The magical number seven, plus or minus two: Some limits on our capacity for processing information. Psychological Review, 1956, 63, 81-97.
- Minsky, M. and Papert, S. Artificial intelligence: Progress report. M.I.T., 1972.
- Nelson, R.J. Introduction to automata. New York: Wiley, 1968.
- Newell, A. and Simon, H.A. Human problem solving. Englewood Cliffs, N.J.: Prentice-Hall, 1972.
- Pask, G., Scott, B.C.E., and Kallikourdis, D.A. Theory of conversations and individuals (Exemplified by the learning process on CASTF). International Journal of Man-Machine Studies, 1973, 5, -43-566.
- Polya, G. Mathematical discovery. Vol. I. N.Y.: Wiley, 1962.
- Quillian, M.R. Semantic memory. In M. Minsky (Ed.), Semantic information processing. Cambridge, Mass.: M.I.T. Press, 1968.
- Restle, F. and Brown, E. Organization of serial pattern learning. In G.M. Bower (Ed.), The psychology of learning and motivation: Advances in research and theory. Vol. 4. New York: Academic Press, 1970.
- Rogers, H.J. Theory of recursive functions and effective computability. New York: McGraw Hill, 1967.
- Rumelhart, D.E., Lindsay, P.H. and Norman, D.A. A process model for long-term memory. In E. Tulving and W. Donaldson (Eds.) Organization of memory. N.Y.: Academic Press, 1972
- Scandura, J.M. The role of higher order rules in problem solving. Journal of Experimental Psychology, 1974, 6, 984-991.
- Scandura, J.M. Structural learning: Theory and research. N.Y.: Gordon & Breach, 1973a.
- Scandura, J.M. The structure of memory: Fixed or flexible? University of Pennsylvania, Structural Learning Report 73, 1973b.(Catalog of Selected Doc. in Psych., Spring, 1974)
- Scandura, J.M. What is a rule? Journal of Educational Psychology, 1972, 63, 179-185.
- Scandura, J.M. Deterministic theorizing in structural learning: Three level of empiricism. Journal of Structural Learning, 1971, 3, 21-53.
- Scandura, J.M. The role of rules in behavior: Toward an operational definition of what (rule) is learned. Psychological Review, 1970, 77, 516-533.
- Scandura, J.M. An alternative approach to research in mathematics learning--a reaction to Suppes' paper. Journal of Research and Development in Education, 1967, 1, 33-43.
- Scandura, J.M., Durnin, J.H., Ehrenpreis, W., and Luger, G. An algorithmic approach to mathematics: Concrete behavioral foundations. New York: Harper & Row, 1971.
- Scandura, J.M., Durnin, J.H., and Wulfeck, W.W.II. Higher order rule characterization of heuristics for compass and straightedge constructions in geometry. Artificial Intelligence, 1974, in press.
- Shank, R.C. Identification and conceptualization underlying natural language. In R.C. Shank and K.M. Colby (Eds.), Computer models of thought and language. San Francisco: Freeman, 1973. pp. 187-247.
- Simon, H. A. Motivational and emotional controls of cognition. Psychological Review, 1967, 1, 29-39.
- Suppes, P. The semantics of children's language. American Psychologist, 1974, 2, 103-115.
- Suppes, P. Stimulus-response theory of finite automata. Journal of Mathematical Psychology, 1969, 6, 327-355.
- Suppes, P. & Groen, G. J. Some counting models for first-grade performance data on simple addition facts. In J. M. Scandura (Ed.) Research in Mathematics Education, Washington: NCTM, 1967.
- Taylor, D. W. Toward an information processing theory of motivation. In M. R. Jones (Ed.) Nebraska symposium on motivation 1960, Lincoln, Neb.: Univ. of Nebraska Press, 1960.
- Thomas, J. C. An analysis of behavior in the Hobbits-Orcs problem. Cognitive Psychology, 1974, 2, 257-270.
- Waugh, N. C. & Norman, D. A. Primary memory. Psychological Review, 1965, 72, 89-104.
- Winograd, T. Understanding natural language. Cognitive Psychology, 1972, 3, 1-191.
- Voorhies, D. J. & Scandura, J. M. Determination of memory load in information processing, This volume.

SOME ASPECTS OF PRODUCTION SYSTEMS*

Allen Newell
Carnegie-Mellon University

* Due to illness in the family, Dr. Newell was unable to attend this year's conference. For your convenience, we are including a bibliography of Dr. Newell's work on production systems.

Bibliography

- Newell, A. Studies in Problem Solving: Subject 3 on the Cryptarithmic task: DONALD + GERALD = ROBERT, Carnegie Institute of Technology, 1967, 168 p.
Early version of production systems for cryptarithmic.
- Newell, A. "A theoretical exploration of mechanisms for coding the stimulus," in A. W. Melton and E. Martin (eds.) Coding Processes in Human Memory, Winston and Sons, Washington, D.C., 1972, pp 373-434.
Introduction to computer production system, PSG. (Contains a PS for cryptarithmic as an example.) Explores a particular task of stimulus encoding to see how production systems would handle it.
- Newell, A. "Production systems: Models of control structures," in W. C. Chase (ed.), Visual Information Processing, 1973a, pp 463-526.
Basic introductory paper on productions systems as models of the human immediate processor. Uses PSG. Contains detailed treatment of Sternberg paradigm.
- Newell, A. PSG Manual, System Version PSG2, Carnegie-Mellon University, Revised 15 August 1973b.
- Newell, A. PSGIP2: Immediate Processor Model 2, Carnegie-Mellon University, 9 Sep. 1973c.
A short working paper consisting entirely of production systems and example runs for a model described in a talk to Math Psychology Meeting, Montreal Sep. 73. Models Sternberg (as in Newell 1973a) and some elementary arithmetic tasks (inequalities, additions, multiplication). Done with PSG with STM-order for conflict resolution.
- Newell, A., and H. A. Simon, "An example of human chess playing in the light of chess playing programs," in N. Wiener and J. P. Schade (eds.) Progress in Biocybernetics, 2, Elsevier Publishing Co., Amsterdam, 1965, pp 19-75.
First use of production system as psychological model, in an analysis of chess.
- Newell, A., and H. A. Simon, Human Problem Solving, Prentice-Hall, 1972.
Contains an extensive general treatment of productions in the context of problem solving. Includes detailed treatment of cryptarithmic task. Much general background on information processing theories. Incorporates both Newell and Simon (1965) and Newell (1967).

MODELING THE MEANING OF "AU" IN LAURA *

Klaus Witz
University of Illinois

This talk was about the language production of a specific child, Laura, age 2. Twenty-five one hour video tapes were taken of Laura in her home over a period of 6 months. On tape twenty-four her maximum speech production consisted of 12 two-word utterances. On tape twenty-five, made 6 weeks later, she spoke over 200 four, five, and six word utterances.

The specific questions addressed were: (1) By what manner and mechanisms was Laura able, for the first time, to produce long utterances? (2) How can one distinguish production (output) from meaning processes which direct what words will be used?

A prosodic field is the totality of the aspects of intonation, pitch, frequency, stress, and loudness. It is a continuous space which represents at each moment the total attentive, interactive, and object directed state of the child. An important property of the field is historicity. Historicity means that the broad outline of the prosodic contour, the field as it stretches over time, that occurs on a particular occasion is already determined at the beginning of the contour, i.e. the first moment the contour is uttered.

Articulation is an area where there are located suitable functional systems to articulate individual words. Each individual word has a particular articulatory structure.

Traces from a graphic level record of Laura's speech output show that it is highly rhythmic, highly peaked. Each peak is a transient wave which momentarily unifies both the prosodic field and the articulatory region into a single audible output. The peaks may be explained in two ways: (1) A pacer or initiator causes the transients, and the prosodic contour adjusts to them. (2) The series of transients is fitted into the contour as a whole, in a time-global fashion. There are a number of phenomena within Laura's speech pattern which support the latter view. For example, she starts to say "silly putty" and she says, "Set up" and then repeats, "Set up putty". On the first there was a wrong stress distribution, the peak came out much too low. The second time it comes out nice, even, rhythmic, so this suggests that there is something time-global in the way stress is being put into the various peaks. The dynamic organization possessed by Laura on tape twenty-five was not present on tape twenty-four.

The second part of the talk concerned meaning. A special class of situations called expansions, where the second phrase of an utterance has an additional part (e.g. "Tastes icky, tastes icky silly putty"), was considered. Expansions are characterized by time lag, addition of words, expansion of the prosodic contour and usually a higher initial or maximum pitch on the second occurrence. It was found that the meaning process is relatively unchanged at the beginning of an utterance. The meaning process stays relatively unchanged in the first phase of an utterance, and then changes. One can think of the change as an enlargement of whatever the first phase is and while this is going on it has no effect on the actual production of the utterance. This has two implications: (1) The meaning process complex is a stationary, dynamic form that stays on level one for about two seconds and then enlarges. (2) The different submeanings corresponding to the subphrases are not really in evidence. Somehow one keeps the global entity going and the exact submeanings don't come into play.

* Summary of talk taken from tape recording.

REPRESENTATION, ABSTRACTION, AND RECOGNITION
OF STRUCTURED EVENTS

Frederick Hayes-Roth

The University of Michigan

ABSTRACT

Structured events are configurations of objects in logical, spatial, temporal or activity relations. A parameterized structural representation system for this class of events is proposed. Parameters in such representations are arbitrarily chosen symbols used to insure consistent references to the same object in diverse relations. All-or-none matching of two representations is the basis for pattern recognition. In this framework, descriptions of pattern or concept prototypes act as structural templates for stimuli. As a result, recognition can be performed in a natural and structural way and is unaffected by manipulations of irrelevant variables. A network embodiment of structural representations is described in which matching of any pattern can be performed associatively without reliance on intermediate memory or computation. Such a capability makes the associated recognition system a natural successor to Pandemonium as a paradigm for pattern recognition and learning. Partial matching of a set of representations is a basis for abstracting common subrepresentations. These abstractions can be used for the solution of general relational concept learning and recognition tasks and for the induction of schematic rules of behavior. Several efficient algorithms appropriate to these problems are described.

INTRODUCTION

Many prevalent questions in information science concerning perceptual, learning, classification, and retrieval functions share a common conceptual base. These functions all necessitate some more or less formal approach to the representation of information and techniques for comparing two or more information structures. As an example, the development of adaptive rules of behavior is easily understood by consideration of the application of mechanisms of abstraction to appropriate sets of training information. In this paper, the notion of a structured event, a discrete, integrated, relational information structure, is introduced. Structured events provide a desirable basis for the representation of a wide variety of knowledge. Further, transformations between events are easily described and can be used to represent general behavior rules.

The bulk of this paper addresses three basic questions which, for the sake of concreteness, are cast in the area of visual perception and classification learning: (1) How should a stimulus and a memory prototype (e.g. a Gothic capital letter "A") be represented? (2) How can a stimulus representation be matched to a memory representation to accomplish recognition? (3) How can representations of several exemplars of a single concept or pattern class be compared to abstract those characteristics common to each of them which are the basis of membership in that particular class? The answers to these questions will be easily seen to be generalizable to other problems involving the representation and comparison of

information structures. Before proceeding to a detailed discussion of these questions, however, a brief review of previous approaches and an overview of the proposed structural representation theory are provided.

PREVIOUS THEORIES OF REPRESENTATION AND MATCHING

Established theories for the representation and recognition of patterns fall into four categories. These will be called the graphic template model, spatial model, feature model, and generative model. Each class is considered in turn.

The graphic template model (see, for example, Lindsay & Norman, 1972) proposes a simple and seemingly desirable solution to both the representation and pattern matching problems. In brief, a memory prototype is a graphic template or replica of the type of item to be recognized. The memory representation of a Gothic capital "A" is conceived of as a stencil or photograph of that letter. Stimuli are recognized as "A"'s, roughly speaking, whenever the light which they emit is present where the template is light and absent where the template is dark. The advantages of such a system are that recognition is determined by a direct contact and comparison of homologous stimulus and memory structures and that all classification alternatives may conceivably be evaluated simultaneously. Its disadvantages are well known and are summarized by the statement that a template is an oversensitive representation and one which lacks any basis for generalization. Even insignificant deviations between the stimulus and template (e.g. differences in proportion, size, orientation, font) may completely inhibit proper recognition.

Spatial models underlie scaling theory approaches, signal detection theory, dimensional representations, correlational methods, discriminant analysis, and most decision theory approaches. In brief, a spatial model proposes that a prototype of a pattern class be represented as a single point in an n -dimensional (usually metric) space where each dimension reflects some multivalued attribute which takes a particular value for each stimulus. Any stimulus item is evaluated on each dimension and is represented by the point at the corresponding coordinates in the pattern space. A stimulus is matched to each memory prototype by some arbitrary distance function, and classification is usually performed by assigning the stimulus to the class associated with the nearest prototype.

These models have the simplicity of representation and matching operations as their outstanding qualities. Their primary disadvantages are a corollary of this simplicity. These models are generally incapable of representing criterial value dependencies among a subset of stimulus attributes. They are also unable to recognize that, although a particular set of features or feature values are relevant to one pattern class, those features may be completely irrelevant to the definition of other classes. One way in which this can be seen is by considering the pattern volume corresponding to each pattern class. In spatial models, the procedure of classifying all stimuli close to one prototype to the corresponding class defines a non-empty volume of points around each prototype which corresponds to the related class. If particular features values were irrelevant to some class (as "level of fatty acids" is irrelevant to the pattern class "Volvo"), a decision procedure should ignore "level of fatty acids" in considering the potential classification of objects as "Volvo." The appropriate pattern volume in such a case would then be a zero-volume hyperarea of some dimensionality less than that of the entire pattern space (Hayes-Roth, 1973, 1974(b); Michalski, 1973; Watanabe, 1973).

Feature models like that of Pandemonium (Selfridge, 1959) posit that each class is defined by the simultaneous presence of a particular set of feature values.

A prototype is represented as a set of criterial feature values and their necessary frequency of occurrence in a stimulus which is to be considered an element of the corresponding class. While these models solve the dependence and irrelevance problems attributed to spatial models, they too have difficulties. Chief among these is that stimuli possessing correct features in an inappropriate configuration are improperly recognized. This is a direct result of the matching process used with feature list representations. A measure of the degree of successful matching of stimulus and prototype is usually computed by determining the proportion of features prescribed by the prototype that are found in the stimulus. Classification is then performed by assigning the stimulus to the category associated with the best matched prototype.

Generative models include grammatical approaches to representation and matching, including linguistic pattern recognition (Uhr, 1971; Joshi, 1973; Thomason, 1973), analysis-by-synthesis (Neisser, 1967), and top-down language understanding systems (Winograd, 1973; Woods, 1970). These approaches employ syntactic and semantic rules to describe the kinds of stimuli that may be produced and encountered in the environment. A prototype is represented or conceived of in terms of a list of productions, transitions, or transformations which must be employed to generate it from a meaningless starting symbol. Matching of stimuli and prototypes is usually done in two steps. First, the stimulus is synthesized or "forged"; the rules of grammar are employed in some arbitrary manner until a particular sequence of transitions is found which leads to the production of an acceptably close replica of the stimulus. Second, when the transitions can be given an interpretation which is plausible in the problem domain, that interpretation is employed in some way to effect classification or understanding. One advantage of this type of approach is particularly noteworthy. Given sufficiently powerful grammars, the generative technique is capable of computing all definable recognition functions (Hopcroft & Ullman, 1969). In specific, any structural dependencies or feature irrelevancies can be appropriately incorporated into the rules of the grammar. On the other hand, the power of these approaches is intimately connected to their principal disadvantages. Among these is their usually enumerative, recursive, or non-deterministic searches of the space of possibilities (all sentences or stimuli which can be produced by the grammar). As a result of such search procedures, these programs do not provide plausible models of the apparently simple and immediate recognition that is apparently involved in the perception of a well known pattern like the letter "A". Furthermore, the performance of these systems is seriously degraded by any expansions of the possible set of rules applicable at each step.

In some sense, generative models seem related to natural mechanisms of recognition in the way Turing machines are related to human cognition. In both cases the strength of the relationship rests primarily on the ability of one machine to effect input-output transformations of a complexity equivalent to that of the natural system. Beyond that, there is little to support the notion that the machine provides a believable or otherwise promising model of the natural computations of interest.

STRUCTURAL REPRESENTATION THEORY

Structured events (Hayes-Roth, 1973) are configurations of objects in logical, spatial, temporal, or activity relations. To represent these configurations, several kinds of elements are used. A relation is a set composed of a predicate and a possibly ordered set of named predicate objects. A predicate is any property that is asserted to obtain among one, two, or more objects. The functional role of each element in a relation is identified by a preceding symbol, "p" for predicate, and completely arbitrary type names for the predicate objects. Parameters are arbitrarily

chosen symbols which name objects and are necessarily introduced in these representations to permit consistent multiple references to an object in diverse relations. Finally, a parameterized structural representation (PSR) is a two-tuple containing a set of parameters and a set of parameterized relations which are simultaneously true and which constitute the description of a single event of interest.




Figure 1. Structural template as prototype.

Although the current paper is primarily directed to the structure and formal properties of such representations, it will help to consider the concrete example in Figure 1, in which a Gothic capital letter "A" is illustrated and the five parameters used in its description are indicated. The corresponding PSR is labelled A and is given below:

A: ({a, b, c, d, e},
 {{p=node, name=a}, {p=node, name=b}, {p=node, name=c},
 {p=node, name=d}, {p=node, name=e},
 {p=line, node=a, node=b}, {p=line, node=a, node=c},
 {p=line, node=b, node=c}, {p=line, node=b, node=d},
 {p=line, node=c, node=d}, {p=line, node=c, node=e},
 {p=line, node=d, node=e}}) (1)

Each parameter in the parameter set {a, b, c, d, e} identifies a node of the line drawing of the "A". The body of the PSR contains relations asserting that each of these parameters is a "node" in the graphic structure and that several "line" relations obtain among these nodes. That is, any two nodes occurring in a single relation containing the predicate "line" are connected by a line in the graphic structure.

It is not claimed that the PSR (1) is an accurate model of what humans acquire when they learn the corresponding pattern. However, such a representation has several salutary features. The representation is insensitive to attribute of an "A" which are not essential to the basic concept. A similar description would be true of any "A" which reflected these structural dependencies regardless of its orientation, font, proportion, etc. As a result, the PSR labelled A is considered a structural template. It describes those properties--visual, in the current case--of any stimulus which are necessary for it to be classified as a proper instance of the class of block capital "A"'s. Although it is possible that additional features (unary relations) or other n-ary relations (n=2, 3, ...) might be present in any exemplar of this class, all of these may be considered extraneous for the purpose of representing a prototype "A".

Because the number of elementary relations may make the complete listing of a PSR tedious, a compact form is also used in this paper. In that form relations with redundant predicates and permutable objects are merged and represented by their set union. If this is done for the PSR A for Fig. 1, the following compact PSR is produced:

A: ({a, b, c, d, e},
 {{p=node, name=a, name=b, name=c, name=d, name=e},

{p=line, node=a, node=b, node=c}, {p=line, node=b, node=d},
 {p=line, node=c, node=d, node=e}}) (2)

Before considering the use of structural templates in recognition and more complicated problems, the general properties of the proposed structural representation system should be elaborated. Predicates may be chosen with complete freedom where the only consideration need be the task to be performed. Object type names, specified by the particular prefix symbols preceding objects in a relation, are also arbitrary and should be chosen to suit the specific task. Some predicates may be able to accept more objects than are of interest at any moment, and these others need not be specified at all. For example, one might conceive of a distinct predicate for each verb-sense in a language and associated predicate objects representing each of the cases associated with that verb-sense, e.g. agent, instrument, location, etc. In some instances, the order of a relation (the number of possible predicate objects) will not be the number of distinct object prefixes. For example, each of the line relations in (2) is actually a symmetric binary (second order) relation. Because it is tedious to distinguish first and second objects in symmetric relations, both are indistinguishably prefixed by the same name "node". The interpretation of such relations is straightforward.

Any set of relations may be named by assigning it an arbitrary parameter symbol, and the corresponding labelled event may be cited as the object of some predicate by referencing that symbol as the appropriate predicate object. This is particularly useful where a total configuration comprises several perceptually or logically distinct subevents which are related as unitary wholes in some way. Examples of such representations can be found in the translation of a linguistic phrase structure tree into a corresponding PSR, with each node in the tree associated with a parameter representing the set of descendants emanating from that node.

Although the PSRs thus far described have employed only parameters as predicate objects, in some sense this is not essential. When all objects are parameters and all content words are restricted to use as predicates, the PSRs are considered to be in explicit form. Otherwise, they are in implicit form. The relationship between the two forms is evident in the following semantically equivalent PSRs for the sentence, "John who is a very tall adolescent likes Mary to run with him."

{{t, u, v, w, x, y, z},
 {{p=name, person=w, word=x}, {p=name, person=y, word=z}, x:{{p=John}},
 z:{{p=Mary}}, {p=tall, how much=v, who=w}, v:{{p=very}}, t:{{p=run}},
 {p=adolescent, who=w}, u:{{p=do together, what=t, agent=w, agent=y}},
 {p=like, agent=w, object=u}}) (3)

{{u, w, y},
 {{p=name, person=w, word=John}, {p=name, person=y, word=Mary},
 {p=tall, how much=very, who=w}, {p=adolescent, who=w},
 u:{{p=do together, what=run, agent=w, agent=y}},
 {p=like, agent=w, object=u}}) (4)

The advantage of an implicit form like (4) is that it is easy to read because the content words which actually represent predicates may be used as if they were constants within relations. Some of the examples presented later in this paper will take advantage of this property. The value of an explicit form like (3) is that it clearly identifies every proposition predicated about the event. As a result, the comparison of any two structured events is most easily performed on explicit PSRs because irreconcilable differences between two patterns will always result from some lack of correspondence between predicates and will therefore be

easily interpreted. This will become clear when partial matching, the mechanism of abstraction, is explained later.

The representation of programmatic rules of behavior in this framework is achieved through the use of transformations between contingency and response event PSRs. The contingency PSR describes an internal condition or state of the computing system which must be satisfied before the transformation is invoked and the response occurs. This view of computing is like that in production systems (Newell, 1972; Becker, 1973) and pattern-directed procedure invocation (Hewitt, 1972; Rulifson, et al., 1972). The response PSR describes conditions which are to occur as a result of the successful invocation of the transformation. A computing system which is constructed in this way is called interrupt-driven or stimulus-driven, because computing is directly controlled by the detection of conditions of importance. Transformations whose contingencies are descriptions of stimulus encodings and whose responses describe appropriate classification or recognition behaviors are called categorical. Substitution transformations, which employ one or more parameters in both the contingency and response PSRs which are common between them, represent behaviors in which particular values are extracted from a stimulus and are incorporated directly into the related response configuration.

These two types of transformations alone are sufficient to represent all conceivable procedures. Although categorical transformations are really special cases of substitution transformations, they are induced in somewhat different ways as will be explained later. Transformations have been introduced at this point to complete the structural representation story and to facilitate the reader's appreciation of the generality of the later results.

Enough detail has now been provided on the representation question to permit the consideration of two types of matching: all-or-none matching, the logical basis of pattern recognition, and partial matching, the logical basis for abstraction of commonalities of several exemplars of a single class.

ALL-OR-NONE MATCHING: THE LOGICAL BASIS OF RECOGNITION

Presuming that known pattern prototypes are represented by PSRs, the task of recognizing the occurrence of a pattern exemplar is that of detecting when a stimulus contains all of the structural relations specified by the prototype. Consider the




Figure 2. All-or-none matching a stimulus and template.

example in Fig. 2 of determining whether the stimulus S matches the structural template A previously described (2). Two problems arise as a direct result of the introduction of arbitrarily chosen parameter symbols into the representations of both the stimulus and prototype. The binding problem concerns the determination of a distinct correspondent among the parameters in S for each parameter in the template A. Corresponding parameters represent objects which play equivalent roles in their respective patterns, those roles being defined by the relations in the

template. In the current example it is intuitively apparent that each node in A corresponds to the node in S which occupies a similar locus in the line drawings. These correspondences are denoted $u \equiv a$, $v \equiv b$, $w \equiv c$, $x \equiv d$, and $y \equiv e$.

The solution of the binding problem is, however, closely allied to a solution of another problem, the matching problem. The matching problem concerns the identification of a corresponding relation in S for each relation in A such that the two relations are completely comparable when the alphabetic differences between corresponding predicate object parameters are ignored. In the current example, it is clear that every relation in A is contained in S under the stated parameter equivalences. Thus S is said to match A, denoted $S(*)A$, and the interpretation is given that the pattern S contains the subpattern A.

In short, having solved the binding problem, the matching problem is to determine whether every relation present in A is also present in S. On the other hand, a desirable solution to the binding problem must necessarily provide a complete solution to the matching problem, too. In the next sections, two alternative algorithms for resolving these problems efficiently are presented.

INTERFERENCE MATCHING: A WIDELY APPLICABLE SERIAL PROCEDURE

The first algorithm, the interference matching procedure, is a three step serial procedure for simultaneously solving the binding and matching problems. The procedure is easily described in terms of the concept of models of two or more events. A model of a set of events with PSRs E_1, E_2, \dots, E_n comprises the following components: a set of parameter correspondences defining equivalent parameters in each of the E_i ; an abstraction E which is a set of relations common to each of the E_i when alphabetic differences between corresponding parameters are ignored; and a set of residuals, R_1, R_2, \dots, R_n which are sets of relations present in each original E_i , respectively, which are not represented in the common abstraction E. Distinct models of a stimulus and a template will describe alternative ways of matching the stimulus to the template (e.g. identifying more than one occurrence of a particular subpattern in a stimulus pattern). Distinct models of any set of events identify alternative ways in which each element in the set is similar to the others. The former interpretation will now be elaborated as the use of models in the all-or-none interference matching procedure for pattern recognition and information retrieval is explained.

The three steps of the procedure as applied to the question of determining whether S matches A are as follows:

Step 1. The stimulus and template are represented as explicit PSRs S and A, respectively. In parallel, all relations in S are compared with all relations in A. Each pair of relations from S and A which contain identical or semantically equivalent predicates are matched to produce a corresponding model. The abstraction in this model contains a single relation of the same type as in both compared relations. Each parameter in this relation is assigned a new symbolic name identifying the pair of parameters from S and A which must be assumed to correspond in the original patterns to allow the interpretation that the relations are exactly equivalent. The residuals of S and A in this model are all relations originally present in S and A, respectively, except for the matched relations represented in the new abstraction. In each model, the set of parameter correspondences contains each pair of parameters assumed to correspond by the forced equivalence of matched relations from S and A. At all times during the interference matching procedure, only models which entail consistent parameter bindings are allowed; that is, each para-

meter in A can be assigned at most one correspondent in S.

Step 2. In parallel, all pairs of consistent and mergeable models are combined to form new, more informative models. Two models are consistent if their parameter correspondences are jointly consistent in the previously described sense. They are mergeable if all residuals in each model contain the abstraction of the other model. That is each model to be merged must reduce the unexplained residuals of the other by identifying some relations which are common to both and which are complimentary to those already identified. When these conditions are satisfied, the models are combined by following these simple rules: the set of parameter correspondences of both models are merged by forming their set union; the abstractions of the two models are merged by forming their set union; and each residual in the new model is the set intersection of the corresponding residuals in the original models.

Step 3. Each model which contains a null residual for the template pattern A identifies one occurrence of the pattern A in the stimulus S. Each distinct model represents a unique occurrence of A in S. The correspondent in S for each parameter in A is specified by the appropriate element in the merged set of parameter correspondences.

Although it is apparent that this procedure is effective, two additional points should be made. First, every possible occurrence of the template A in the stimulus S is detected and represented by a corresponding model. Second, this algorithm is probably optimal for the task of identifying every occurrence of A in S. This conclusion is a direct result of the fact that only those models which may lead to a solution are ever manipulated, and manipulations cease as soon as sufficient evidence is obtained that any further operations will be futile. Further, the procedure rapidly reduces the set of mergeable models by reducing residuals whenever a merge occurs and, therefore, is guaranteed to halt.

A DIRECT MEMORY CONTACT METHOD: SANER NETS

The major drawback to the type of recognition process just described is that it involves a large amount of computation and the transition through several steps before recognition is achieved. It seems, therefore, seriously lacking as a potential model of the apparently immediate pattern recognition observable in natural systems. The reason that structured representations cannot be recognized as directly as items in a feature or spatial model is that the introduction of parameters into the representations entails a solution to the parameter binding problem whenever two patterns are compared. In this section one framework which banishes the parameter but completely retains the power of structural representations is introduced. In this framework, recognition can be easily and immediately achieved by the direct comparison of two superimposed homologous representations, one of the stimulus and one of the template.

The proposed method employs structured associative network event representation or, simply, SANER nets (Hayes-Roth, 1974(a)). The elements of these networks are nodes, corresponding to each distinct predicate object type name; arcs, directed according to lexicographic order, connecting any two nodes corresponding to loci occupied by the same parameter in the corresponding explicit PSR; event labels, used to relate all nodes and arcs arising from the same event representation; component frequencies, designating the number of times the same node or labelled arc arises in the net representation of one event; and permutation indices, auxiliary labels

attached to each arc event label designating that the connected node corresponds to the occurrence of a parameter object in a relation in position i in the k -th possible permutation of the complete list of the n relations in the PSR which are of the same predicate type. ($i=1,2,\dots,n; k=1,2,\dots,n!$)




Figure 3. A structured event and its SANER net representation. The label $C(1)$ on nodes in the SANER net indicates that that predicate object occurred one time in the PSR C . The label $C.1.1.1(1)$ on arcs in the SANER net indicates that only one position predicate and one geometric figure type (square, triangle) occurred in the PSR labelled C occupied by the same parameter (in one case f and in one case g).

An example of a structured event, its PSR, and a corresponding SANER net is given in Fig. 3. In this example, the SANER net is quite simple but reflects the significant attributes of this approach. Note that the representation is distributed; that is, the elements are connectable only logically, by the presence of a common event label (C). The content of the representation is a function of the predicate types employed, the multiplicity of references to each parameter (indicated by arcs), and the frequency with which such nodes and arcs occur. In the current example, all frequencies are one, as denoted by (1) in Fig. 3. The permutation indices appearing as auxiliary arc labels in Fig. 3 are interpreted as follows: the first index is the number of the permutation used to encode this net, 1 in this case where only one permutation is possible; the second index is the rank of the position relation in the permuted list of all position relations, 1 in this case because only one position relation occurs in the PSR C ; similarly, the third index is the rank of the square and triangle relations in their respective lists of all such relations in the PSR C .

Now suppose that the SANER net labelled C is one of the templates to be checked in subsequent recognition tasks, and a stimulus S is presented which happens to contain the event of interest as a subpattern. The stimulus is first encoded: it is represented as a set of alternative SANER nets, each arising from a distinct permuted ordering of classes of relations with identical predicates. One of these nets will necessarily satisfy the following recognition condition: each node and permutation indexed arc in the SANER representation of the template will be matched by an exactly indexed component in the stimulus net representation and the frequency of every template component will be equalled or exceeded by the frequency of the corresponding component in the stimulus SANER net.

The recognition condition is satisfied if and only if $S(*)C$. As a result, a theoretically simple computational device can be constructed for structured event recognition using only numerical threshold gates to compare corresponding component frequencies and compute an overall comparison conjunction. Because recognition is thus performed directly on the content of the representation and without intermediate computation and memory, it is appropriately considered an associative or direct memory

contact process. To my knowledge, SANER nets are the only framework yet designed in which direct memory contact recognition of structured events is possible. Because of the complete generality of PSRs and their correspondence to SANER nets, this approach appears to be adequate for all conceivable structured event recognition tasks. The major drawback to this technique is that the occurrence of several relations with a common predicate necessitates a potentially large number of permutations to generate the set of alternative SANER net encodings of a stimulus. Any one of these permutations may be chosen as the representation of a prototype template but, if recognition of stimuli is to be immediate, all possible permuted encodings of the stimulus relations must be generated before matching. This is a serious technological problem. The future solution of this problem would seem to be a goal of vast import because of the obvious practicality and generality of these techniques.

PARTIAL MATCHING: THE LOGICAL BASIS OF ABSTRACTION

An example of an abstraction problem is presented in Fig. 4 in which the stated task is to ascertain the rule by which all exemplars on the left are assigned to the category "A" and those on the right are assigned to "B". Such a rule is an example of a categorical transformational behavior. If the task seems too difficult, attend only to the "A" category because "B" contains anything that is not classifiable as "A".

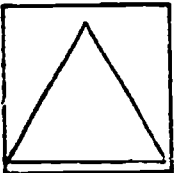
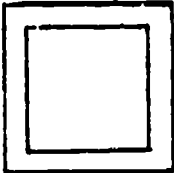
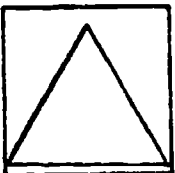
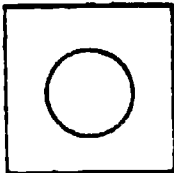
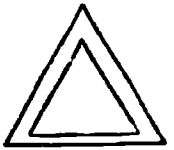
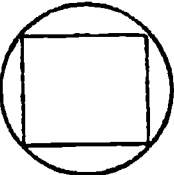
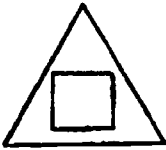
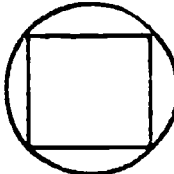
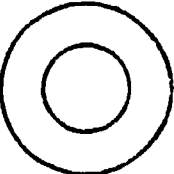
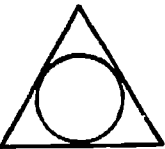
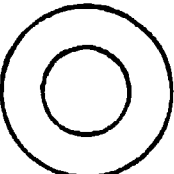
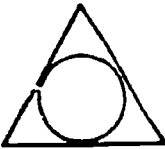
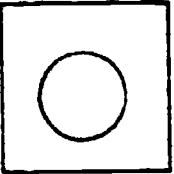
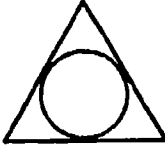
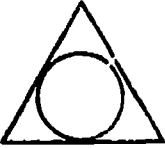

"A"		"B"	
A ₁ : 	A ₂ : 	B ₁ : 	B ₂ : 
			
A ₃ : 	A ₄ : 	B ₃ : 	B ₄ : 
			

Figure 4. An abstraction problem.

The solution to this problem is that "A" contains those events with two pairs of nested geometric forms in a vertical configuration such that the interior figures of both pairs are the same. A compact representation of the categorical transformational rules defined by this solution is given below in implicit form:

T: ({w, x, y, z},
 {{p=form, type=?, name=w, name=x, name=y, name=z},
 {p=contain, outer=w, inner=x}, {p=contain, outer=y, inner=z},
 {p=above, hi=w, hi=x, lo=y, lo=z},
 {p=same shape, name=x, name=z}})) (5)

U: ({w, x, y, z, T},
 {{p=not true, what=T}}) (6)

Say-A: ({w, x, y, z, T},
 {{p=transformation, contingency=T, response="A"}}) (7)

Say-B: ({w, x, y, z, T, U},
 {{p=transformation, contingency=U, response="B"}}) (8)

The rules Say-A (7) and Say-B (8) prescribe categorical (constant) responses to any stimuli which either match or do not match the criterial pattern specified by T, respectively. The "?" in the description T represents our indifference regarding exactly what type of geometric forms the four objects are.

This example usefully illustrates several points about the proposed approach to learning. First, the frequently encountered notion that exemplars of a concept or pattern are best interpreted as distortions of a pattern prototype (see for example, Franks & Bransford, 1971; Bransford & Franks, 1971; Reed, 1973; Thomason, 1973) is not adopted here. The preferred view holds that prototypes constitute propositions that are true of all exemplars of a concept (with appropriate interpretation for disjunctive concepts). If this view is accepted, the obvious solution to abstraction problems is to identify those properties common to each class and verify that they are not evidenced by the negative exemplars (see, for example, Hunt, 1959; Bruner *et al.*, 1956). The only two problems which arise in such an approach are (1) generating the possible unitary or compound propositions which are hypothetical bases for classification and (2) choosing appropriately among these. Solutions to these problems are now discussed.

Consider first the problem of producing hypotheses by identifying properties which are common to each element in a set of structured events. Suppose that the first two exemplars of category "A" in the previous example are given the following PSRs:

A₁: ({a, b, c, d},
 {{p=form, type=square, name=a},
 {p=triangle, name=b, name=c, name=d},
 {p=contain, outer=a, inner=b}, {p=contain, outer=c, inner=d},
 {p=above, hi=a, hi=b, lo=c, lo=d},
 {p=same shape, name=b, name=c, name=d}}) (9)

A₂: ({e, f, g, h},
 {{p=form, type=square, name=e, name=f, name=h},
 {p=form, type=circle, name=g},
 {p=contain, outer=e, inner=f}, {p=contain, outer=g, inner=h},
 {p=above, hi=e, hi=f, lo=g, lo=h},
 {p=same shape, name=e, name=f, name=h}}) (10)

The identification of properties common to both exemplars is easily operationalized in terms of models produced by the interference matching procedure previously described. If the corresponding explicit PSRs for A₁ and A₂ are matched using that

procedure, every distinct model produced represents a unique way in which the exemplars can be seen to be similar when objects identified by corresponding parameters are asserted to play functionally similar roles. In this use of the interference matching procedure, however, there is no distinction made between stimulus and template. Models are merged until it is no longer possible to reduce either of the two residuals. Such matching is called partial matching, because it determines the extent to which two PSRs are similar and identifies, in the residuals, associated irreconcilable differences.

One of the models so produced has the following compact implicit PSR as its abstraction:

```
{(ae, bf, cg, dh),
  {p=form, type=square, name=ae},
  {p=form, type=?, name=bf, name=cg, name=dh},
  {p=contain, outer=ae, inner=bf}, {p=contain, outer=cg, inner=dh},
  {p=above, hi=ae, hi=bf, lo=cg, lo=dh},
  {p=same shape, name=?, name=bf, name=dh}}
```

 (11)

This PSR can be given the following verbal interpretation, "Four geometric forms in two nested pairs in a vertical configuration such that the uppermost figure is a square and three forms, two of which are the interior figures, are the same shape." The parameters in this PSR are named by juxtaposing the two parameter symbols of corresponding parameters from A_1 and A_2 , respectively. The appearance of "?" in this PSR signifies a residual difference in the model of the compared explicit PSRs, an irreconcilable difference in the corresponding predicates occurring in the original representations.

This abstraction is related to the original concept learning problem in that it is a possible hypothesis for the criterion of membership in class "A". Moreover, any possible subrepresentation of (11) which could be produced simply by replacing parameter symbols by "?" or deleting whole relations is an equally plausible hypothesis about the true "A" prototype.

Because there are many such subrepresentations, one would be ill-advised to attempt to generate and validate each one at this point. Instead, attention should first be limited to the set of maximal abstractions, those abstractions which are not matched by any other abstraction yet produced. Maximal abstractions are easily identified in the interference matching procedure, because they are all contained only in models which were never merged into more informative models. Each of these maximal abstractions is subsequently partial-matched to the third exemplar of "A". The resultant maximal abstractions are identified and matched in turn to the last exemplar. At that point, the most informative maximal abstraction obtained represents the desired concept solution to the abstraction problem and is identical to T in (5). There may, however, be other abstractions which are also plausible bases for classification rather than the one asserted to be the "right" solution. Each of these will either be a maximal abstraction or a subrepresentation of one. In the next section, a statistical classification theory is presented which provides a method for selecting among these alternatives.

SCHEMATIC CLASSIFICATION THEORY

Within the proposed framework, the logical basis for classification decisions is the presence of one or more criterial characteristics in each item to be classified. Such a characteristic is called a criterial schema. Schemata are subrepre-

sentations, properties, relational attributes, or subpatterns of the structural descriptions of events. Schematic classification theory assigns to every possible schema matched by any exemplar of a pattern class a measure of its capacity to indicate membership in that class vis-a-vis the alternatives. By utilizing the operation of partial matching to identify subpatterns common to the set of exemplars of each pattern class and evaluating the predictive power or "cue validity" (Reed, 1973) of each of these schemata as indicators of membership in the associated class, each novel test item can be classified on the basis of the most powerful predictor which it matches.

A variety of measures may be used for the predictive power of a schema as an indicator that items which match it belong to a particular class. The Bayesian posterior expected net gain of weighted correct less weighted incorrect classifications resulting from the classification of all items matching a schema T as a member of class C --under the assumption that the conditional membership of an item in C which matches T is a Bernoulli process with some probability q to which we assign a uniform prior distribution--is one measure which is very general, widely applicable, intuitively reasonable, and easy to evaluate. If the gain attributable to each correct classification of an item in class C is denoted G_C , the loss attributable to an incorrect classification of an item in class C is denoted L_C , the proportion of elements which actually are elements of C is denoted p , the number of training exemplars of class C is denoted $|C|$, the number of training exemplars of class C matching T is denoted $|C(*)T|$, the number of non-exemplars of C is denoted $|\bar{C}|$ and the number of those matching T is denoted $|\bar{C}(*)T|$, the predictive power of the schema T as a basis for indicating that an item S is a member of class C given that $S(*)T$ is then defined to be:

$$U(T \Rightarrow C) = |G_C| p \frac{1 + |C(*)T|}{2 + |C|} - |L_C| (1-p) \frac{1 + |\bar{C}(*)T|}{2 + |\bar{C}|} \quad (12)$$

The only thing remaining to explain is how one might reasonably constrain the set of schemata considered as bases for classification. In even simple examples, the number of schemata is vast because it is the cardinality of the union of the complete sets of subrepresentations of each exemplar. This selection problem is greatly exacerbated when disjunctive concepts are encountered, because criterial schemata need not be matched by all exemplars of the same class, nor is it possible to know definitely how many disjunctive terms (independent schemata) are criterial or which exemplars manifest the same unknown criterial schema.

An algorithm has been developed which handles these difficulties in the same way for both conjunctive and disjunctive concepts. It is called the Space Limited Interference Matching (SLIM) procedure and is part of a fully implemented computerized learning and classification system. SLIM is chiefly constrained by the number of schemata it may maintain in working memory. Within the limitation imposed on memory space, the procedure performs the following actions in sequence: it successively partial-matches exemplars of each class; it extracts maximal abstractions; it evaluates the unconditional predictive power of each schema as in (12); it dynamically ranks schemata indicating the same class; it appropriately reduces (conditionalizes) the predictive power of schemata which are redundant with (are matched by) more informative and higher performing schemata or match less informative and higher performing schemata; and, finally, it eliminates from overcrowded storage the lowest conditionally performing schemata. As a result, the procedure dynamically optimizes the overall net performance of all schemata in storage. When the production of such maximal abstractions is complete, the schemata from all classes are merged and sorted so that test items may be classified according to the highest performing schema which they match.

RULE LEARNING: INDUCTION OF TRANSFORMATIONS

It should now be clear what is involved in learning categorical transformations, rules that assign a constant response like "A" or "B" to stimuli matching different specified contingencies. The schematic classification theory just discussed completely specifies how assignment of categorical responses should be made to each predictive schema which may be present in a stimulus. The implementation of these assignments as contingency-response transformations is straightforward.

To complete the discussion of rule learning it must be shown how substitution rules also can be induced from training exemplars. Fortunately, the solution of this problem requires no new operations. Suppose a set of exemplar (stimulus, response) event pairs is provided as the training information from which a substitution rule is to be abstracted. Substitution rules are induced in two steps. The first step is to partial-match each stimulus PSR with its related response PSR. The bindings of corresponding parameters in these matches identify where substitutions might be reasonably expected to occur. Each parameter in a pair of correspondents is renamed to be identical with its correspondent.

The second step in the rule learning process is to partial-match successively each (stimulus, response) pair with the others, ensuring that stimulus components are matched only with stimulus components, and likewise for response components. Each maximal abstraction obtained from this procedure will have one component arising from the matched stimulus events and one from the matched response events. The former becomes the contingency component and the latter the response component of a substitution transformation.

If this procedure is applied to (stimulus, response) pairs of (active, passive) equivalent sentences as an example, an active-to-passive transformation, illustrated in Fig. 5, is obtained. Parameters occurring in the contingency of a substitution transformation like that in Fig. 5 (NP_1 , NP_2 , V_1) represent universally quantified




Figure 5. An induced passive-to-active substitution transformation.

variables. Whenever a stimulus sentence matches such a contingency, the particular event which is labelled by the correspondent of each quantified variable is bound to that variable as its value and is substituted directly into the appropriate loci in the response configuration. As an example, the stimulus sentence "the boy hugged the dog" will match the contingency in Fig. 5 and will cause the binding of "the dog" to NP_1 , "hugged" to V_1 , and "the dog" to NP_2 . As a result, the sentence "the dog was hugged by the boy" will be produced as the response of this transformation.

CONCLUSIONS

Structural representations were introduced to permit a natural description of complex events. These relational representations necessitated the use of arbitrary parameters to name objects and events cited as predicate objects. Because the parameter symbols were chosen arbitrarily, even trivial questions like determining whether

two representations were identical would be very difficult to answer. An interference matching procedure was described which provides a basis for answering that question and other related ones of principal importance, including all-or-none matching for pattern recognition and partial matching for abstraction.

One approach toward direct memory contact methods of matching employed SANER nets to represent structured representations. These nets succeed in removing parameters from representations but require permuted stimulus encodings in many cases. Nevertheless, the discovery of such a network representation is a basis for optimism concerning the possibility that other direct memory contact methods may be found without comparable disadvantages.

The use of partial matching for generating abstractions is an essential component of pattern learning procedures based on schematic classification theory. In that framework, schemata which are highly predictive cues of class membership are abstracted and used to classify test items which match them. An algorithm for identifying a dynamically optimal set of schemata was described.

The prospects for fruitful application of these structured techniques are substantial. The manipulations of structured knowledge which these techniques accomplish appear to be basic to a wide variety of problems including pattern recognition and perception, comparison of current and goal states in problem solving, categorical response and substitution rule learning, and information matching and retrieval. In retrospect, the capabilities underlying representation, recognition, and abstraction of a pattern as simple as a Gothic letter "A" may also account for a wide variety of cognitive skills which appear, at first glance, wholly unrelated.

REFERENCES

- Becker, J.D. A model for the encoding of experiential information. In R.C. Shank & K.M. Colby, Computer models of thought and language. San Francisco: Freeman, 1973.
- Bransford, J.D., & Franks, J.J. Abstraction of linguistic ideas. Cognitive Psychology, 1971, 2, 331-350.
- Bruner, J.S., Goodnow, J.J., & Austin, G.A. A study of thinking. New York: Wiley, 1956.
- Franks, J.J., & Bransford, J.D. Abstraction of visual patterns. Journal of Experimental Psychology, 1971, 90, 65-74.
- Hayes-Roth, F. A structural approach to pattern learning and the acquisition of classificatory power. Proceedings of the First International Joint Conference on Pattern Recognition. New York: I.E.E.E., 1973.
- Hayes-Roth, F. An optimal network representation and other mechanisms for the recognition of structured events. Paper submitted for presentation at Second International Joint Conference on Pattern Recognition. 1974(a)
- Hayes-Roth, F. Schematic classification problems and their solution. MMPP 74-1. Ann Arbor: Michigan Mathematical Psychology Program, 1974(b).
- Hewitt, C. Description and theoretical analysis (using schemata) of PLANNER: a language for proving theorems and manipulating models in a robot. Cambridge: MIT Project MAC, 1972.
- Hopcroft, J.E., & Ullman, J.D. Formal languages and their relation to automata. Reading, Massachusetts: Addison-Wesley, 1969.
- Hunt, E.B. Concept learning: an information processing problem. New York: Wiley, 1962.
- Joshi, A.K. Remarks on some aspects of language structure and their relevance to pattern analysis. Pattern Recognition, 1973, 5, 365-382.

Lindsay, P.H., & Norman, D.A. Human information processing. An introduction to psychology. New York: Academic Press, 1972.

Michalski, R.S. AQVAL/1--computer implementation of a variable-valued logic system VL₁ and examples of its application to pattern recognition. Proceedings of the First Joint International Conference on Pattern Recognition. New York: I.E.E.E., 1973.

Neisser, U. Cognitive psychology. New York: Appleton, 1967.

Newell, A. A theoretical exploration of mechanisms for encoding the stimulus. In A.W. Melton & E. Martin (Eds.), Coding processes in human memory. Washington: Winston, 1972.

Reed, S.K. Psychological processes in pattern recognition. New York: Academic Press, 1973.

Rulifson, J.F., Derksen, J.A., & Waldinger, R.J. QA4: a procedural calculus for intuitive reasoning. Menlo Park: Stanford Research Institute, 1972.

Selfridge, O.G. Pandemonium: a paradigm for learning. In D.V. Blake & A.M. Uttley (Eds.), Proceedings of the Symposium on the Mechanisation of Thought Processes. London: H.M. Stationary Office, 1959.

Thomason, M.G. Finite fuzzy automata, regular fuzzy automata, and pattern recognition. Pattern Recognition, 1973, 5, 383-390.

Uhr, L. Flexible linguistic pattern recognition. Pattern Recognition, 1971, 3, 363-370.

Watanabe, S. Subspace method in pattern recognition. Proceedings of the First Joint International Conference on Pattern Recognition. New York: I.E.E.E., 1973.

Winograd, T. A program for understanding natural language. Cognitive Psychology, 1972, 3, 1-191.

Woods, W.A. Transition network grammars for natural language analysis. Communications of the ACM, 1970, 13, 591-606.

GRAPHICAL REPRESENTATION OF ABSTRACT CONCERNS
IN TEXTUAL MATERIALS

John M. Yancey, PhD
University of Pennsylvania

Graphical representations of abstract concepts (usually designated ideographs or ideograms) have been used throughout history. Ideographs are frequently employed in journals of science, and occasionally appear in student texts; yet the efficacy of this usage had not been previously investigated. The purpose of this study was to commence a line of inquiry into the graphical elaboration of textual materials intended to communicate abstract concepts.

Mental imagery research has shown relevant images to be positively associated with ease of learning, and accuracy of memory, for isolated verbal items and concrete textual materials. Several theoretical speculations have suggested that perceived structure also aids learning and memorial processes. It was, therefore, hypothesized that the graphical elaboration of a verbal text intended to communicate abstract concepts would elicit facilitating images, leading to superior learning retention. It was further predicted that ideographs depicting the inherent structure of the concepts described by the text would be associated with better performance than nonstructured graphics. Since imagery research has shown that abstract words seldom evoke images, it was hypothesized that the basic verbal text, even elaborated with redundant outlines, would not be associated with achievement levels as high as that demonstrated by subjects receiving the basic text elaborated with either structured or nonstructured ideographs. Subjects creating their own relevant graphics were expected to perform better than those provided with ideographs.

In this experiment, 345 undergraduates were randomly assigned to one of five levels of textual elaboration. Each of the five experimental groups received one of the following variations of reading materials: the basic verbal text only; the basic verbal text elaborated with redundant verbal outlines; the basic verbal text elaborated via graphical representations of redundant verbal outlines; the verbal text elaborated by means of graphical representations of the inherent structure of the abstract concepts contained in the basic text; or the basic verbal text, with graphical elaborations of the textual concepts to be constructed by the learner.

Following the initial learning session, all subjects took an

objective achievement test. Two weeks later an equivalent retention test was administered. Following the delayed test, a set of verbal transfer materials was read by all subjects and a transfer test administered. A postexperiment questionnaire measured the quantity and quality of mental images employed by each subject. The quantity and quality of subject-produced diagrams on the learning and testing materials was also recorded. In addition, other individual difference measures were compiled.

The data collected confirmed the hypothesis that relevant mental images are positively associated with superior learning and retention of abstract concepts. This significant relationship remained even after the affects of scholastic aptitude were statistically removed. The basic verbal test read by all subjects, however, elicited relevant images in approximately the same proportion of subjects in each group. Indeed, most of the subjects in every group developed and used appropriate images in most learning and retention tasks. These two factors effectively eliminated any possibility for the variations in textual elaboration to differentially evoke facilitating images.

The type of graphical elaboration provided clearly influenced the structure of subject-produced diagrams, but the equality of group achievement indicated the underlying images to be functionally equivalent. This failure to clearly assess the relationship between subject-produced graphics and image structure precluded an unequivocal examination of the hypothesis that images reflecting the inherent structure of abstract concepts are superior to nonstructured images.

The results of this experiment suggested several areas for further study. Specific recommendations were made regarding choice of materials (to best test the influence of various types of textual elaboration), and several individual difference variables were proposed for investigation. Pedagogical implications included possible criteria for using ideographs, and the specific benefits to be expected. The literature reviewed suggested several possible advantages of using the graphical mode to communicate the inherent structure of abstract concepts: (1) Ideographs communicate more directly (they remain essentially unchanged across languages); (2) Ideographs communicate an all-at-once gestalt (which is presumably more readily grasped, stored and subsequently recalled); (3) Ideographs establish similar images in both the sender and the receiver (establishing a common basis for associated verbal communiques); and (4) Ideographs communicate less irrelevant or incorrect detail (than image-evoking analogues).

The inherent risks of using image-evoking graphical representations of abstract concepts (or image-evoking concrete analogues) were enumerated. The possible risks, however, were found to be more than compensated by the probable advantages.

THE EFFECT OF IDENTIFICATION OF SEMANTIC RELATIONSHIPS
ON PROCESSING LOAD DURING SENTENCE COMPREHENSION ^{1,2}

Richard L. Isakson
Cornell University

Psycholinguistic research has not provided a clear picture on how or when the meanings of the individual words in a sentence are integrated by the listener or reader to form an internal, cognitive structure representing the meaning of a sentence. Attempts have been made to apply Chomsky's (1957, 1965) theory of linguistic competence to this problem. To some degree, these attempts demonstrated variables which influence comprehension. However, they did not succeed in accounting for how or when the meaning of a sentence is pieced together from the individual words comprising the sentence (Gough, 1971; Fodor and Garrett, 1966). Indeed, Chomsky never intended his theory to be a model of performance which would explain how the deep structure of a sentence is recovered by the listener (Lyons, 1970).

Evidence has been found that cognitive processing activity varies at different points as a sentence is comprehended. In reviewing a portion of this literature, Bever (1972, p. 105) states: "Such operations do not apply homogeneously throughout a sentence, and the effects of preliminary word- and phrase- processing are reflected in minor variations in attention to nonspeech stimuli during a clause." In several studies, sentences have been presented aurally and Ss were required to react to some signal contained in, or presented during the sentence. The signal was either a click (Abrams and Bever, 1969; Holmes and Forster, 1970) or a particular phenome in the sentence itself (Foss, 1969; Hakes and Foss, 1970; Foss and Jenkins, 1973). These studies assumed that if the signal occurred at a point where processing load on the listener was high, reaction time to the signal would be longer than it would be at a point where processing load was relatively low. These studies indicate that differences in reaction time are found as a result of certain variables assumed to influence processing demands such as major constituent boundaries, frequency of usage of particular lexical items, and ambiguity. The above findings do not, however, make clear how sentence meaning processing takes place.

¹Paper presented at the Fifth Annual Interdisciplinary Meeting on Structural Learning, April 1974, Philadelphia, Pennsylvania. For copies write to Richard L. Isakson, Dept. of Education, Stone Hall, Cornell University, Ithaca, New York, 14850.

²This research was partially supported by Hatch Act funds from the U.S. Department of Agriculture. Dr. George W. McConkie, Dept. of Education, Cornell University, was the advisor to this research.

The purpose of the present research is to examine sentence processing activity with regard to the identification, on the part of the reader, of the underlying semantic structure of a sentence. The semantic structure of a sentence has been described by Fillmore (1968, 1970, 1971). Fillmore (1968, p. 21) states: "The sentence in its base structure consists of a verb and one or more noun phrases, each associated with a verb in a particular case relationship." Some of the more common case relationships are: agentive, experiencer, objective, instrumental, and locative. The action or main verb together with its case actors form what Fillmore terms a "proposition".

In light of the contribution made by Fillmore's theory, the process of understanding a sentence can be viewed as an activity in which the reader identifies the individual words and stores their meanings in an active memory until the case role can be assigned to a segment of the surface input. At this point a type of processing occurs which integrates the meanings of the words into the cognitive representation that the reader is forming of the sentence meaning. Thus, at the points where a case relationship becomes clear, the reader processes a portion of the meaning of the sentence and thereby pieces together a cognitive representation of the underlying proposition of the sentence. This cognitive activity will be referred to as partial sentence meaning processing (PSMP).

Recent research (Isakson, 1974) has provided evidence in support of this view of sentence processing. Sentence pairs were constructed such that in one member of the pair the case role of a preceding noun phrase becomes clear at the 5th word. It was hypothesized that at the 5th word PSMP could be initiated, thereby placing greater processing load on the reader. In the other sentence type the 5th word did not make clear the case role of the preceding noun phrase and, therefore, it was hypothesized that PSMP would not be initiated at that point. Cognitive processing load was indexed by means of a click response task similar to that employed by Abrams and Bever (1969) and Holmes and Forster (1970). Subjects heard and responded to clicks presented shortly after the onset of either the 5th, 6th, or 7th word in the experimental sentences.

Response times to clicks at the 5th word position were characteristically longer in the sentences in which case role was identifiable at that point compared to sentences in which case role was not identifiable at the 5th word. The purpose of the present study was to replicate the findings of the first experiment and to monitor processing load at the 4th word position in addition to the three word positions which were probed in the earlier study. The various word positions were examined in order to locate, at least generally, the point at which case identification may have the greatest effect on processing load. Since the members of the sentence pairs are identical up through the 4th word, there should be no difference in response times to clicks occurring there. If no differences were found, the processing load at this point could be used as a baseline against which to compare processing demands at later points in the sentence. In an effort to avoid the possible

problem of some Ss not allowing their response times to vary according to the demands of the primary, comprehension task, as seemed to be the case in the first experiment, the pay-off system in the present experiment was modified to place less emphasis on speed of response and greater emphasis on comprehension.

Subjects -- Twenty-four Cornell University undergraduates participated as Ss.

Materials -- Sixteen sentence pairs were utilized from among the set of sentences developed for the first experiment. Each sentence was 13 words in length. An example sentence pair is:

Type A: The serene mother comforted the children who trusted her when they were afraid.

Type B: The serene mother comforted by children who trusted her grew braver every minute.

Both sentence types had identical words in the first nine word positions, except for position 5. At word position 5, Type A sentences always had the word "the" which served as a cue to identify the preceding word string as a main verb, or action of the sentence, and the animate or inanimate agent of the action. Type B sentences had either the word "with", "from", or "by" which left the case role of the preceding noun phrase unclear, and also failed to identify whether or not the verb was the main verb of the sentence.

Two 16mm films were prepared for presenting the sentences, with one word per frame. Each film contained eight Type A sentences and eight Type B sentences, plus 20 filler sentences of similar grammatical structure, and a number of blank frames during which base reaction times could be measured.

Procedures -- A subject was seated eight feet from the viewing screen. Each sentence was presented at a constant rate of approximately three words per second. Sometime during the presentation of each sentence, an audible click occurred in the S's earphones, at which time an electronic clock was started. When S pressed a microswitch the clock was stopped. The click always occurred 60 msec. after the onset of a word. It occurred at the 4th, 5th, 6th, or 7th word positions in the Type A and B sentences and at other positions in the filler sentences. After seeing a sentence the S attempted to recall it as accurately as possible. Subjects were given feedback after each sentence and were awarded points, later to be exchanged for money, for their sentence recall and click response performance.

Results -- The response time data are displayed in Table 1 and Figure 1. An analysis of variance of the data yielded a significant main effect for sentence type ($F = 7.46$, $df = 1/22$, $p < .012$). The mean response time to

clicks in Type A sentences across word positions 4, 5, 6 and 7 was 290.04 msec., while that of Type B sentences was 271.78 msec. The sentence type X click position interaction was not significant.

TABLE 1

Mean Click Response Time for Sentence Type A and B
at Word Position 4, 5, 6 and 7

Sentence Type	Word Position			
	4	5	6	7
A	308.29	299.88	283.50	268.50
B	264.08	283.38	275.13	264.54

An inspection of the graph in Figure 1 indicates that, rather than showing no difference, the response times to clicks occurring at the 4th word position were considerably longer in the Type A sentences. A t-test for correlated data between the two means at word position 4 indicated a significant difference ($t=2.47$, $df=23$, $p < .05$) in favor of the Type A sentences. The difference between the means at word position 5 approached significance ($t=1.45$, $df=23$, $p < .10$).

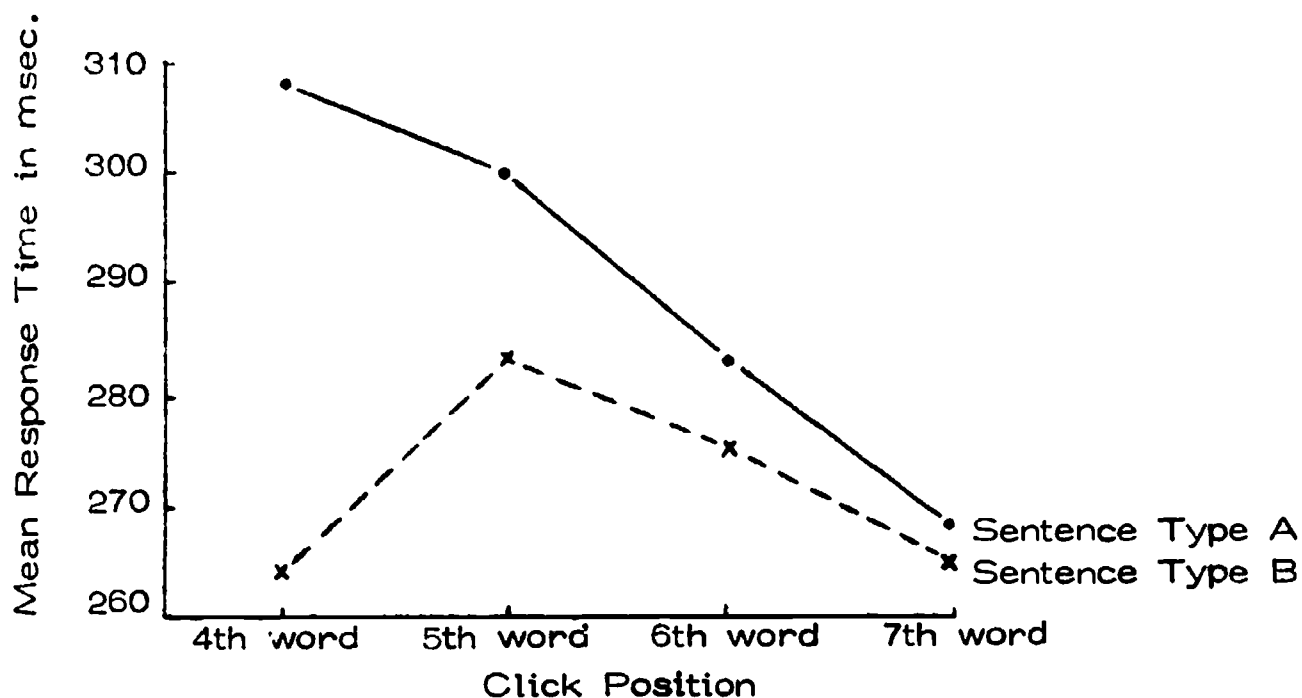


Figure 1

Mean Click Response Time in Type A and B Sentences
at Word Positions 4, 5, 6 and 7

Discussion -- The significant main effect for sentence type, with the Type A sentences resulting in a longer overall mean response time than the Type B sentences, supports the view that sentence processing activity is related to the identification of case relationships. The graph in Figure 1 indicates that the greatest difference between Type A and B sentences came at the 4th word position. This finding could be explained in terms of the possibility that Ss sometimes delayed their response to a click at word position 4 long enough for them to see the critical 5th word. Response times over 290 msec. would allow S to see word 5 while he was still in the process of responding to the click at word position 4. Upon seeing the 5th word in a Type A sentence, these Ss initiated PSMP which may have caused a delay in responding to the previously heard click. It should be noted, however, that in a subsequent experiment the difference between Type A and Type B sentences at the 4th word position was not found. However, mean response time at word position 5 in this more recent experiment was significantly longer for Type A than for Type B sentences. Thus, three experiments have found longer response times for Type A sentences at the 5th word position. These findings support the position that the meaning of a sentence is processed at points where the semantic structure of the sentence becomes clear to the reader such that he can assign a segment of the surface structure to a case role. An alternative view of sentence processing, in which the reader analyzes only the surface structure of the input string until after the entire sentence has been received, is called into question by the present results. Such a view is implied in the theory of the semantic component as put forth by Katz and Fodor (1963) and has more recently been suggested by Weisberg (1973) and Dell and Feldman (1974). If the reader were processing only the surface structure of the sentence at word position 5, one would expect greater processing load for the Type B sentence since its surface structure becomes considerably more complex at that point compared to the surface structure of the Type A sentence.

The picture of sentence processing which emerges from the present research is compatible with recent evidence (Suci and Hamacher, 1973; Shafto, 1973; and Wilson, 1973) which indicates that the memorial representation of a sentence is based on a case relationship structure. These findings and the findings of the present research indicate that as a sentence is understood, the reader, or listener, discovers the semantic or case relationships in the sentence's underlying structure and then forms a cognitive representation of the sentence's meaning based on these relationships.

REFERENCES

- Abrams, K. and Bever, T.G. Syntactic structure modifies attention during speech perception and recognition. Quarterly Journal of Experimental Psychology, 1969, 21, 280-290.

- Bever, T.G. Perception, thought, and language. In J.B. Carroll and R.O. Freedle (eds.), Language comprehension and the acquisition of knowledge. New York: Wiley, 1972.
- Chomsky, N. Aspects of a theory of syntax. Cambridge, Mass.: MIT Press, 1965.
- Chomsky, N. Syntactic structures. The Hague: Mouton, 1957.
- Dell, G.S. and Feldman, S. The role of semantics in sentence-processing. Paper presented at the Annual Meeting of the American Educational Research Association, 1974, Chicago, Illinois.
- Fillmore, C.J. The case for case. In E. Bach and R.T. Harms (eds.), Universals in linguistic theory. New York: Holt, Rinehart, & Winston, 1968.
- Fillmore, C.J. Types of lexical information. In D.D. Steinberg and L.A. Jakobovits (eds.), Semantics. Cambridge: Cambridge University Press, 1970.
- Fillmore, C.J. Some problems for case grammar. In monograph series on Languages and Linguistics. Washington: Georgetown University, 1971.
- Fodor, J.A. and Garrett, M. Some reflections on competence and performance. In J. Lyons and R.J. Wales (eds.), Psycholinguistic papers: The proceedings of the 1966 Edinburgh Conference. Edinburgh: Edinburgh University Press, 1966.
- Foss, D.J. Decision processes during sentence comprehension: Effects of lexical item difficulty and position upon decision times. Journal of Verbal Learning and Verbal Behavior, 1969, 8, 457-462.
- Foss, D.J. and Jenkins, C.M. Some effects of context on the comprehension of ambiguous sentences. Journal of Verbal Learning and Verbal Behavior, 1973, 12, 577-589.
- Gough, P.B. Experimental psycholinguistics. In W.O. Dingwal (ed), A survey of linguistic science. University of Maryland, 1971.
- Hakes, D.T. and Foss, D.J. Decision processes during sentence comprehension: Effects of surface structure reconsidered. Perception and Psychophysics, 1970, 8, 413-416.
- Holmes, V.M. and Forster, K.I. Detection of extraneous signals during sentence recognition. Perception and Psychophysics, 1970, 7, 297-301.

- Isakson, R.L. Cognitive processing load during reading. Paper presented at the Annual Meeting of the Eastern Psychological Association, 1970, 7, 297-301.
- Katz, J.J. and Fodor, J.A. The structure of a semantic theory. Language, 1968, 39, 170-210.
- Lyons, J. Noam Chomsky. New York: Viking Press, 1970.
- Shafto, M. The space for case. Journal of Verbal Learning and Verbal Behavior, 1973, 12, 551-562.
- Suci, G.J. and Hamacher, J.H. Psychological dimensions of case in sentence processing: Action role and animateness. International Journal of Psycholinguistics, 1972, 1, 34-48.
- Weisberg, R. Short-term representation of sentences. Journal of Experimental Psychology, 1973, 101, 381-383.
- Wilson, S. Personal Communication, Cornell University, 1973.

THE DETERMINATION OF MEMORY LOAD IN INFORMATION PROCESSING¹

Donald Voorhies and Joseph Scandura²
University of Pennsylvania

Miller's (1956) classic "The magic number 7 ± 2 " has inspired a considerable amount of experimental research on task memory load and its relation to human information processing capacity. To date, however, much of the most directly relevant research on the topic involves list learning of one sort or another. The reason is not hard to find. Given that to-be-remembered elements may contain more or less information, because of a process Miller (1956) called "chunking," the relation between list learning and memory load seems clear. The memory load imposed on a person retaining a list is simply the number of chunks.

Until recently, relatively little has been done by way of clarifying why, psychologically, a chunk is a chunk or, more generally, of developing general methods for determining memory loads on arbitrary tasks. As a first approximation, memory load clearly is a function of task complexity, and varies both across different tasks (e.g., memory span, addition, multiplication) and over different instances of given tasks (e.g., adding mentally two- versus three-digit numbers). Existing information theoretic (e.g., Posner, 1964) and regression (Suppes, 1967) models provide general quantitative measures of task complexity but they seem to work better with some tasks than others, and deal only indirectly with underlying cognitive processes.

Scandura (1973) has proposed a more general, deterministic-analytic method for determining the memory load imposed at each stage of processing by any given procedure (rule) as applied to any specific task. Data collected by Voorhies and Scandura (reported in Scandura, 1973; Voorhies, 1973) using this method showed that there are individual differences in processing capacity that are stable across different tasks (e.g., digit lists, addition). Not surprisingly, however, they also found that there were deviations from the deterministic ideal. For example, a subject with capacity eight sometimes recalled seven digits correctly and sometimes nine.

The main purpose of this paper is to introduce random variation into the Scandura (1973) theory as a means of accounting for obtained deviations. The extended stochastic model is applied both to previously reported and new data. The paper is organized as follows: (1) a review of relevant research pertaining to processing capacity; (2) a summary of the deterministic analytic model proposed (Scandura, 1973) for determining memory loads, along with relevant data; (3) the stochastic model proposed to account for deviations from ideal performance; and (4) application of the stochastic model both to existing data (Scandura, 1973) and to new data (Voorhies, 1973) collected under different conditions.

Related Research. The distinction between primary and secondary (permanent) memory has been with us since William James. It was not, however, until Miller's (1956) classic work that the notion of a fixed processing capacity came into prominence. Synthesizing research involving absolute judgements, span of attention, and span of immediate memory, Miller showed that humans have a definite limit in the amount of information they are able to process at one time, and that generally speaking the limit is 7 ± 2 elements.

Drachman and Zaks (1967) further demonstrated that there is a rather sudden decrement in performance -- a "memory cliff" -- at capacity. They administered the same task twice, using the data from the first administration to account and adjust for individual variation on the second. After first determining individual memory

-
1. This research was supported in part by Research Grant NSF GW-7901 from the National Science Foundation.
 2. Requests for reprints should be sent to Joseph M. Scandura, Structural Learning Program, University of Pennsylvania, Philadelphia, Pa. 19174.

spans, they administered additional memoranda (strings of digits) of various lengths. The dependent measure on the second test was a function of the digits recalled by subjects on strings of length span minus one, span, span plus one, etc. Pooling the data across subjects, and adjusting for individual spans in this way, there was a fairly sharp break in performance beyond the estimated memory span.

This limit appears to be independent of whether or not the elements being processed are familiar. Crowder (1967) and Conrad (1960), for example, found that requiring subjects to prefix to-be-remembered lists at recall with an extra, redundant digit (e.g., "1") causes a decrement in the percentage of correct responses. Savin (1968) further showed that prefixing a given digit to, say, an eight digit list causes the list to act at recall as a nine digit list. The extra digit, even though highly familiar to the subjects, apparently occupies as much processing space as any other.

Processing capacity, however, depends only indirectly on the number of elements. As Miller (1956) has shown, the essential factor is the number of effective units, which he called "chunks." In a similar vein, Bower (1969) has shown that the critical feature in recall is the number of functional units in memory, rather than the sheer number of units. He had subjects recall words from lists which included single words (one functional unit), three-word cliches (one functional unit), or three-word triples (three functional units). In all cases, although the number of words recalled varied, the number of functional units recalled was the same.

More important in view of the approach adopted below, there are indications in the literature that the number of chunks depends on the procedures subjects use as well as on the tasks themselves. Posner and Rossman (1965), for example, gave subjects number strings, and required them to transform the strings in various ways (e.g., add, give the highest number, give the lowest number, etc.). They found that both the number of transformations and their size had a significant effect on the amount of untransformed material that was lost from store. The more processing required of the subject, the less peripheral material he can continue to rehearse. This suggests that on the average the more complex a procedure, the more processing space is required to apply it. (In this regard, it should be noted that Posner and Rossman used an information theoretic measure of transformation difficulty (i.e., the size of reduction in "bits"), which on some tasks was not compatible with obtained difficulty. We comment on this again below.)

This dependency, clearly, is not a simple one; procedures can serve to decrease the number of chunks that must be retained, leave it unchanged, or possibly even increase the number. The results of Miller (1956), for example, suggest that a subject can increase the absolute number of elements (e.g., strings of binary digits) he can process by recoding them into a smaller number of chunks e.g., strings of decimal digits) in a way which makes it possible to regenerate the original elements on demand. Dalrymple-Alford (1967) overcame the highly variable and often idiosyncratic effects of rehearsal by building it (rehearsal) directly into his experimental procedure. Rather than attempt to prevent it, he tried to insure that rehearsal was of a known type and magnitude, and was uniform across all subjects. He had his subjects rehearse all previous digits aloud before the experimenter added a new digit. For example, suppose the first four digits were 2, 7, 3, and 9. Then, the subject had to say "two, seven, three, nine" before the experimenter presented the next digit. Following Brown (1958), Dalrymple-Alford reasoned that any increase in retention as a result of rehearsal is due to recoding and not rehearsal per se (i.e., the "strengthening of memory traces"). With this in mind, he also determined each subject's memory span in the conventional (Woodworth and Scholberg, 1954) manner. Surprisingly, the proportions of errorless repetitions were almost identical under both conditions indicating that no learning (recoding) occurred during rehearsal. Dalrymple-Alford further noted the absence of traditional serial position effects. First errors in repetition (breakdowns) tended to occur equally often at all positions of the number sequences. Meunier, Ritz, & Meunier (1972) have shown further that rehearsal is used (at least) primarily to maintain items in short-term memory.

The limited capacity hypothesis (cf. Broadbent, 1958) also has considerable explanatory power. It has been used as a basis for explaining and/or predicting the effects of a number of variables in short term memory experiments. Katz (1968), for example, found that primacy and recency effects on the serial position curve for short-term free recall are due to the subject's selective use of attention and storage and retrieval strategies operating under the constraint of a fixed short-term processing capacity. He had his subjects repeat random lists of two-digit numbers, and instructed one group to give the last pair before recalling the others. The other group was allowed to give the digits in any order they wished. Although he found no overall difference between the two groups in correct recall, the former group showed a relatively strong recency effect and the latter, a relatively strong primacy effect.

The results of several other studies suggest that processing capacity underlies the opposing effects of certain pairs of factors. Murdock (1965), for example, reasoned that if memory capacity is constant, then it should be possible to "trade off" the number of times a stimulus is presented with the exposure time on each trial. He constructed a list of six pairs of common English words, and presented them for study for a period of 24 seconds. Three conditions were used: (a) In the first condition, each pair was presented once, for a period of 4 seconds; (b) in the second, each pair was presented twice for a period of 2 seconds (on each trial); (c) in the third, each pair was presented four times for a period of 1 second (on each trial). Murdock found that there was essentially no difference between conditions a and c, with b only slightly better than both. In a study involving incidental learning, Somers (1967) obtained basically the same result. Sitterly (1968) used various combinations of digit and interdigit duration and also found that retention depended on total presentation time.

Other studies have been of a primarily empirical nature. Thus, for example, Williams and Fish (1965) showed that increasing the length of the individual items to be recalled, and increasing the number of symbols from which the items are constructed, both decrease the percentage of correct recall (of the items). Corballis (1966) varied the speed at which strings of nine digits were presented as well as the duration of time each digit was actually visible. He found that when stimulus durations were long, the number correct was higher the slower the presentation speed, but when stimulus durations were short, there was a tendency for this trend to be reversed. Although generally compatible with the limited capacity hypothesis, the results were relatively complex and no theoretical explanation was offered.

Use of the limited capacity hypothesis as a basis for explanation with more complex tasks has been limited largely to correlational studies. Whimbey, Fischhof, and Silikowitz (1969), for example, measured performance on a digit span task, a mental addition task, a vocabulary task, and a test of general intelligence. They found a generally high correlation between the digit span task and the mental addition task, suggesting that both are influenced by a limited processing capacity. Vocabulary and general intelligence both showed little relation. Whimbey and Leiblum (1967) found similar high correlations among simple repetition, repetition preceded by "0", and repetition with interspersed verbal activity. Later, Whimbey and Ryan (1969) found correlations between the digit span task and mental syllogistic reasoning problems.

By way of summary, three general conclusions may be drawn from the available literature. First, there is a definite limit on the number of effective units (chunks) that subjects may process at any one time; this limit is 7 ± 2 chunks. Second, the memory load imposed by a task depends in a nonsimple way on the procedure used by the subject as well as on the task itself. Third, the limited capacity hypothesis provides a useful basis for explanation, and, hence, it is not surprising that it plays a central role in most contemporary theories of memory (Anderson and Bower, 1973; Rumelhart, Lindsay, and Norman, 1972; Scandura, 1973).

Although more might be gleaned from the available literature, there are two important limitations inherent in it. For one thing, there is no way to specify precisely how a subject is to process given information, except with the simplest tasks. Savin (1968), for example, required his subjects to prefix a given digit

at recall but did not specify how this digit was to be processed in relation to the others. Other investigators have instructed their subjects not to organize the given digits but made no attempt to tell them what they should do. Under these conditions it is uncertain, for example, when information is in memory and when it is in the information processor (Waugh and Norman, 1965). Even on simple tests of memory span, recoding and rehearsal processes tend to be highly dependent on individual preference and susceptible, at best in only partly known ways, to even slight changes in experimental conditions. Indeed, in only one study reviewed (Dalrymple-Alford, 1967) was a serious attempt made to insure that the subjects used a particular procedure (rehearsal).

For another thing, even if procedures were to be specified, there is no a priori way available to determine the processing load they (the procedures) impose on the subject. Under these conditions it has been impossible to compare different tasks (procedures) with regard to memory load, short of direct empirical testing. As promising as the technique originally appeared, measures of information reduction (e.g., Posner, 1964) seem to work well only with certain tasks. Suppes (1967) proposed an analytic method for calculating memory loads that was used with some success in predicting latencies on simple addition problems. Suppes' (1967) regression model involved three structural variables, the magnitude of the sum (MAGSUM), the magnitude of the smallest addend (MAGSMALL), and the number of steps necessary to complete a problem (NSTEPS). In determining NSTEPS, specific account is taken of how many quantities must be kept in memory in the course of solving a problem. Although the predictions made by the model correlated .86 with the actual data of 24 fourth-grade students, the method used to determine memory load was too crude for present purposes.

Model for Memory Load.

In information processing, new elements may be encoded, thereby adding to memory load, and old elements may be decoded, thereby reducing memory load. In addition, internal operations themselves may serve either to generate new memory elements and/or to eliminate (erase) others.

Scandura (1973) has proposed an analytic model for calculating memory loads at each stage of processing, with arbitrary tasks and taking into account the particular procedures used. In the model, the memory load is characterized in terms of the number of distinct substimuli (defined below) which must be distinguished during the future course of a computation. In addition to specifying the number of substimuli at any given stage, which enter into the next operation, a count is taken of all substimuli which enter into any subsequent operation.

Directed graphs are used to represent underlying processes; with the nodes labelled so as to indicate memory loads at each stage of processing. In these directed graphs, nodes represent states and arrows between nodes represent operations between states.* The model rests on the principle that at any given state each substimulus (entity) that must be processed as a distinct unit, by the next or subsequent operation imposes a memory load of one (on the processor). The memory load at any given state, then, is the sum of such units, and the memory load for an entire computation (the sequence of states between the input and final output) is the maximum of the memory loads imposed at the various states of the computation.

The directed graph in Figure 1 represents a process for adding two n-digit numbers.

$$\begin{array}{r} a_n \ a_{n-1} \ \dots \ a_2 \ a_1 \\ + \ b_n \ b_{n-1} \ \dots \ b_2 \ b_1 \\ \hline S_n \ S_{n-1} \ \dots \ S_2 \ S_1 \end{array}$$

* It is possible to generalize directed graphs to include decision making capabilities as well (e.g., see Gorn, 1973), but we have not done so here.

TASK NCA

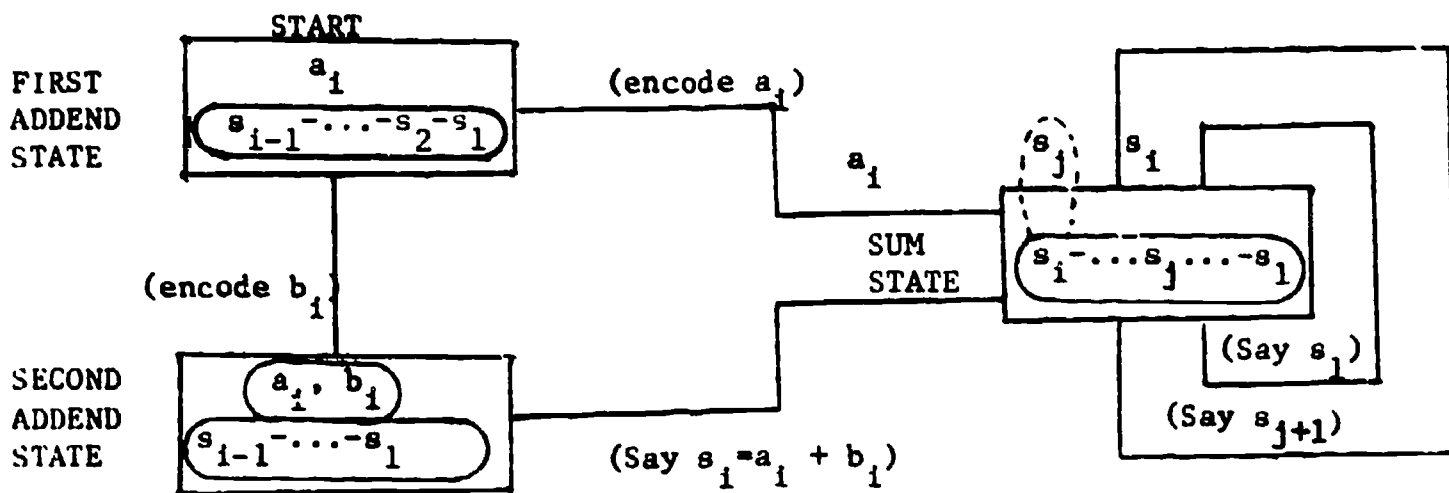


Figure 1. Node labelled directed graph for Task NCA showing states and transformations.

This process includes three nodes (states): **FIRST ADDEND STATE**, **SECOND ADDEND STATE**, **SUM STATE**. The first addend a_i is presented and encoded prior to the **FIRST ADDEND STATE** and b_i is presented and encoded prior to the **SECOND ADDEND STATE**. The sum of the last two addends (less than 10) is determined prior to the **SUM STATE**. At each state, the processor must retain the partial sum ($s_{j-1} \dots s_2 s_1$) obtained prior to encoding the new addends a_j and b_j .

In adding two three-digit numbers, for example, the maximum load of six occurs at the **SECOND ADDEND STATE** just before the third pair of digits (a_3, b_3) is added: two for the column sums s_2 and s_1 , one for the chunk (unit) ($s_2 s_1$), two for the new addends (a_3 and b_3), and one for those addends as a distinct unit. The rationalization for distinguishing (counting) ($s_2 s_1$) and (a_3, b_3) as substimuli over and above their constituent elements is that these units enter into subsequent operations, ($s_2 s_1$) during recall at the **SUM STATE** and (a_3, b_3) during addition between the **SECOND ADDEND STATE** and the **SUM STATE**. Notice that after the sum is generated, the load is four: three for the column sums s_3, s_2 , and s_1 and one for the unit $s_3 s_2 s_1$, which is subsequently repeated. The maximum load at the **FIRST ADDEND STATE** is also four: two for the column sums s_2 and s_1 , one for ($s_2 s_1$), and one for a_3 .

Test of the Analytic Model. The empirical viability of the proposed model was studied by Voorhies and Scandura (reported in Scandura, 1973; Voorhies, 1973). The method and results are summarized below.

(Method). Six tasks were used in the experiment: repeating digit lists (L), repeating digit lists and extra time (XL), saying "1" before repeating digit lists (1L), non-carry addition (NCA), carry addition (CA), and mixed addition (MA) with both carrying and non-carrying.

In the three list tasks, the digits were presented orally, one at a time. After each digit was presented, the subject was required to say the digit and then repeat every digit presented up to that point. After hearing the instruction "Repeat" in task XL, the subject was required to repeat the last string of digits an extra time. In task 1L, the subject had to say "1" before repeating the last string. The lists used in tasks XL and 1L were of predetermined length; in task L, new digits were added until the subject made a mistake.

The addition tasks were presented similarly. Each successive pair of digits corresponded to the digits in one column of a column addition problem. After each such pair, the subject was required to say the sum of the two digits. In the NCA task, the next pair of digits was presented immediately thereafter, and the process was repeated until the subject made a mistake. In the CA task, the subject verbally separated the tens and units digits of each column sum before continuing and added one to the first input digit from the next column. In the MA task, the tens digit of the sum was sometimes zero so that no carrying was involved. In all addition tasks, the subjects were required to repeat each partial sum after adding the digits in a given column. This repeated overt responding on the part of the subject made it possible to both monitor progress during processing and help insure that rehearsal was of a fixed, known variety.

The digits used in the individual problems were selected randomly subject to constraints designed to minimize undesired chunking. On each task, the subjects were also asked to evaluate the degree to which they followed the assigned processing procedure.

The subjects were six volunteer graduate mathematics students at the University of Pennsylvania who were tested individually. The subjects were thoroughly trained on the nature of the study and on the six tasks over a period of approximately three months prior to the experiment. Their cooperation in following each processing procedure to the best of their ability was strongly encouraged. In addition to extensive practice, a metronome was used to pace both the experimenter's presentation of digits and the subject's processing. The metronome speeds varied over subjects and were chosen so as to maximize the subject's "comfort." Preliminary training and practice continued until the experimenter was confident that the subjects were able to process the numbers automatically, under the experimental conditions and without errors or hesitations. During the experiment proper, each subject was tested within a ten day period in a room as free of distractions as possible.

Each task was given on a different day with the order of tasks counterbalanced over subjects. After a brief review of the procedure in question and being reminded of the need for cooperation, the subject was given five warm-up problems. Performance was evaluated on the 20 problems which followed. The main dependent variable was the percentage of criterion strings of a given length that were repeated correctly.

The analytic model described in the previous section was used to determine memory loads for each task at each stage of processing. Loads of 7 and 8, respectively, were predicted for lists of length 6 and 7, non-carry addition with 4 and 5 digits per addend, carry addition with 3 and 4 digits per addend, and mixed addition with 3 and 4 digits per addend.

(Results and Discussion). The data of the individual subjects was analyzed to determine sharp drops (over 40%) in percentage correct from one load level to the next. The average of the base load levels (just prior to such drops), rounded to the nearest integer, was used to estimate each subject's processing capacity (C). Two of the six subjects had estimated capacities of 7 and four had capacities of 8.

Figure 2 summarizes the experimental results adjusted for capacity and averaged over subjects for each task.




Figure 2. Performance at various capacity levels pooled over subjects, but adjusted for varying capacities.

Although the differences between C-1 and C ($t_{10} = 2.5, p < .02$), C and C+1 ($t_{10} = 5.35, p < .001$), and C+1 and C+2 ($t_{10} = 1.9, p < .05$) are all significant, the drop at capacity C from 65% correct to 27% correct at C+1 is more than twice that for the drops of 17% between the other successive loads.

In general, the analytic model seemed to provide reasonably good estimates of the memory loads involved in applying the assigned procedures to the various tasks. In the vast majority of cases (26 of 36) actual percentage drops in individual performance between successive loads on the various tasks were as predicted. That is, these drops tended to occur at the same analytically determined (predicted) load level regardless of the task in question. Because all subjects were using essentially the same processing rules, this result suggests that information processing capacity has a physiological base and is not itself subject to training (although task performance may be). We agree with Miller (1956); the "magic" of the number 7 (or 8) is one manifestation of being human.

Incidentally, rehearsal was observed to act just as any other processing procedure, and did not seem to improve performance. Although rehearsal may provide opportunity for recoding, where subjects inject rhythm or other process deviations into the procedure, rehearsal in and of itself does not seem to affect retention.

A Fixed Capacity, Stochastic Model of Human Information Processing. In spite of the heuristic value of the analytic, Scandura (1973) model, actual predictions were something less than perfect. For example, at load C-1, 18% of the items were missed, while at load C+1, 27% of the items were repeated correctly. In spite of our Herculean efforts to minimize idiosyncratic variation, by intensively training subjects on particular procedures, such variation still existed. Undoubtedly, there were lapses in attention, unwitting chunking, and other deviations from the ideal which

influenced performance.*

One obvious way to account for such deviations is to introduce random variation into the Scandura (1973) model. There are two basically different ways in which extraneous processing may influence performance. Such processing may increase memory load, for example, where new, irrelevant information is unintentionally activated (e.g., via self monitoring, thinking of things external to the algorithm). Such information may replace needed information and thereby hinder performance. On the other hand, extraneous processes may reduce memory load, for example, where two digits are recoded as one unit (chunked).

Letting α be the probability of an extraneous element entering the processor and displacing an element already present, and β be the probability of two critical elements being chunked into one unit, then α^n and β^n , respectively, are the probabilities of n extraneous elements entering the processor and $n+1$ elements being recoded as one. Accordingly, the probability of an error at the j^{th} stage of processing is

$$P(E_j) = \alpha^{C-L_j+1} \left(1 - \sum_{i=1}^{\infty} \beta^i\right) + \sum_{i=1}^{\infty} \beta^i \alpha^{C-L_j+1+i}$$

which is equivalent to

$$P(E_j) = \alpha^{C-L_j+1} \left(\frac{1-2\beta}{1-\beta} + \frac{\alpha\beta}{1-\alpha\beta}\right) \quad (1)$$

where C is the individual's capacity, L_j is the load at stage j , and $i+1$ is the number of critical elements reduced to one chunk. Assuming the subject never corrects him or herself once an error is made, the probability of making an error in an n -state computation is

$$P(E) = P(E_1) + (1-P(E_1))P(E_2) + \dots + (1-P(E_1)) \dots (1-P(E_{n-1}))P(E_n) \quad (2)$$

Implicit in equation (1) are the assumptions that there is a positive probability that no critical element is displaced $(1 - \sum_{i=1}^{\infty} \alpha^i) > 0$ and that no chunking occurs

$(1 - \sum_{i=1}^{\infty} \beta^i) > 0$. These relations lead to the boundary conditions $0 < \alpha < 1/2$ and

$0 < \beta < 1/2$.

* Although our training procedures were fairly intensive, our subjects were volunteers and hence not subject to the type of control that might have been possible had they been explicitly hired for purposes of the experiment, with attendant reward for "good" performance. To some extent, the algorithms taught could also have been somewhat ambiguous regarding memory load. For example, we did not attempt to formally represent distinctions stemming from whether items were to be stored or held in active memory.

Test One of the Stochastic Model. In view of the complexity of the above equations, parameters α and β cannot easily be estimated using standard maximum likelihood methods. Instead, a computer program was constructed which considered all possible (α, β) pairs, representative of intervals of size .01 (i.e., $(.01, .01)$, $(.01, .02)$, ..., $(.49, .49)$). The estimated values $(\hat{\alpha}, \hat{\beta})$ were the means of those (α, β) pairs for which the sums of the absolute differences between obtained and predicted data was less than some predetermined standard. (Details of the program may be found in Voorhies, 1973.)

The obtained $(\hat{\alpha}, \hat{\beta})$ estimates for each of the six subjects on tasks L, NCA, and CA are shown in Table 1. Figure 3 shows that these estimates provided a good fit of the data.

Table 1
 $\langle \hat{\alpha}, \hat{\beta} \rangle$ values
 at critical loads
 C-1, C, C+1

	L	NCA	CA	Mean
S _H	.06, .21	.17, .40	.46, .09	.22, .23
S _J	.66, .43	.42, .49	.46, .48	.45, .47
S _B	.33, .26	.41, .06	.67, .49	.40, .27
S _F	.14, .36	.26, .30	*	.20, .33
S _L	.30, .33	.30, .30	.37, .11	.32, .25
S _C	.38, .39	.26, .44	.16, .21	.27, .35
Mean	.28, .33	.30, .33	.38, .28	.31, .31

* None of the possible (α, β) pairs met the required standard suggesting that these two parameters were not adequate to account for the data.




Figure 3. Obtained and predicted performance based on the stochastic model. Note: obtained and predicted performance of L overlap.

In general, the estimates of $\hat{\alpha}$ and $\hat{\beta}$ were relatively high, $\hat{\alpha} = .32$, $\hat{\beta} = .31$. There were only five cases where $\hat{\alpha}$ or $\hat{\beta}$ was below .15. This result seems consistent with both the experimenters' informal observations and the subjects' evaluation of their own performance, and suggests that both extraneous processing and chunking were involved to a considerable degree.*

Study Two. In spite of precautions to the contrary, the subjects' processing in the first study frequently appeared to deviate from the training procedures. This was evidenced by the use of (auditory) rhythmic patterns, slurring numbers, and even verbal reports by the subjects themselves. To determine whether this effect could be eliminated (or modified) a second study was conducted in which the subject was encouraged to also bring visual and tactile processes into play.

(Method). Four of the six subjects in the first study volunteered again to serve as subjects. Three of the six tasks (L, NCA, and CA) were used and presented as in the first study with the following differences: (1) Subjects were asked to "bring a picture to mind" of each digit presented but to eliminate the picture before processing the next digit. The analogue of an automatic slide projector was used. (2) Subjects were required to move their arms in sharp, up and down movements as digits were spoken in cadence with the metronome. After each task instance (problem) the subjects were also required to respond to an explicit evaluation sheet concerned with the extent to which they felt the problems had been processed as planned. In addition, one-fourth of the 36 problems for each task were presented on each of the four experimental days, instead of one task per day as in study one. The reported results are based on that two-thirds of the data, where the processes used were rated highest by both the experimenter and the subject.

(Results and Discussion). Using the same criteria as study one to determine the individual capacities (C) of individual subjects, the obtained estimate for each of the four subjects was one higher than in the previous study. For example, S_j's average capacity in study one was 7 3/8 (rounded to 7), whereas in study two, it was 8 1/4 (rounded to 8). Because of the higher incidence of chunking reported by subjects in study two and to facilitate comparison, the lower estimates were used consistently.

The experimental results are summarized in Figure 4. The drops between C and C+1 ($p < .05$), C+1 and C+2 ($p < .05$), and C+2 and C+3 ($p < .025$) are all significant.

Overall, the obtained results differed in two ways from those in study one: (1) the subjects performed at a generally higher level and (2) the relatively sharp drop in performance which was previously found between C and C+1 was replaced by a more gradual decline, spread over several load levels.

Although it was impossible to specify precisely the source of the generally improved performance, it seems unlikely that the experimental manipulation (between studies) more sharply distinguished the digits during the processing. Indeed, the subjects' comments on the post hoc evaluation form suggested that forming images (cf. Paivio, 1971) of the digits may have increased chunking. For example, the subjects reported that "pictures" of the digits often tended to linger and become superimposed on one another in memory, enabling them to "read" the digits off during recall.

* Although the small number of cases precluded statistical tests, the parameters tended to vary more over subjects than over tasks. The respective means of $\hat{\alpha}$ and $\hat{\beta}$ over subjects ranged from .20 to .45 and from .23 to .47, while over tasks $\hat{\alpha}$ ranged from .28 to .38 and $\hat{\beta}$, from .28 to .33. Furthermore, with the exception of S_H the standard deviation of each subject's $\hat{\alpha}$'s and $\hat{\beta}$'s about his or her mean value was less than or equal to the standard deviation of the $\hat{\alpha}$ values of the task means. In addition, with the exception of data points (.41, .06), (.46, .09), and (.37, .11), there is an apparent linear relationship between $\hat{\alpha}$ and $\hat{\beta}$, with $\hat{\beta}$ roughly equal to $3/5 \hat{\alpha} + .16$. These observations could be merely happenstance, but they could also reflect the possibility that α and β correspond to fundamental human traits, which like information processing capacity, may vary over individuals but be relatively stable across tasks, at least in certain kinds of task situations.




Figure 4. Performance at various capacity levels pooled over subjects, but adjusted for varying capacities.

Application of the stochastic model tended to support this interpretation; the suggested increase in chunking is reflected in the estimated values of $\hat{\alpha}$ and $\hat{\beta}$. (With respect to the individual estimates in Table 2, χ^2 tests revealed no significant differences between obtained and predicted performance. In study one, χ^2 was significant in only 1 of 30 cases.)

Table 2

$\langle \hat{\alpha}, \hat{\beta} \rangle$ Values
at Critical Loads
C - 1, C, C + 1

	L	NCA	CA	Mean
S _J	.44, .49	.24, .45	.31, .44	.33, .46
S _B	.30, .43	.15, .43	.43, .48	.29, .45
S _F	.33, .45	NS	.43, .46	.38, .46
S _G	.12, .47	.21, .45	.33, .39	.22, .44
Mean	.30, .46	.20, .44	.38, .44	.30, .45

The obtained estimates of $\hat{\beta}$ averaged .14 above the corresponding mean in study one. In only two cases did $\hat{\beta}$ decrease, and then only by .04. Although the $\hat{\alpha}$ estimates for individual subjects and tasks, differed from study one, the average value ($\hat{\alpha} = .30$) remained at approximately the same level ($\hat{\alpha} = .32$). Furthermore, whereas the variability of $\hat{\alpha}$ was approximately the same in both studies ($\sigma = .13$ in study one and .11 in study two), the standard deviation for $\hat{\beta}$ was reduced from .14 in study one to .03 in study two. The cause of this effect was not clear. One possibility was that directing the subjects to process digits simultaneously in multiple modes reduced the individual variation open to the subjects in study one. Another (related?) possibility was that the result was an artifact reflecting a boundary condition on $\hat{\beta}$.

Concluding Comments. Basic theory and individual results aside, perhaps the most important contribution of this research is the new method it provides for studying a wide variety of memory and information processing phenomena. Rather than having to rely exclusively on highly prescribed experimental conditions to control subjects' behavior, the proposed method of representing procedures (including processing loads) provides a feasible way of also manipulating the processes individual subjects use in dealing with such tasks. This possibility would appear to be especially important in working with relatively complex tasks. (In this regard, we note parenthetically that if there is one way (procedure) for accomplishing a given task, then according to a well known result in computer science there is an infinite number of other ways that will also work.)

Research along these lines could have important practical as well as theoretical implications. In mathematics education, for example, considerable attention has been given to research concerned with the relative efficiency of various algorithms in arithmetic computation -- between, say, the borrowing and equal additions methods for subtraction. The results of such research have traditionally been ambiguous. The proposed model provides a way not only for the explicit analytic determination of memory loads in using the respective algorithms but also a feasible method for obtaining definitive, empirical information.

On the theoretical side, the proposed methodology could provide a useful way of dealing with such perennial experimental problems as ensuring that an item is in short term store (or long term memory) when that is what we want (cf. Waugh & Norman, 1965). In most such applications, of course, further development of the theory itself will almost certainly be required. It is an open question, for example, whether the parameters α and β are independent of processor load. Indeed, it is reasonable to argue that α and β depend on processor load in two diametrically opposed ways. (1) The likelihood of extraneous processing and/or chunking could be a direct function of the availability of space for "turning around" on oneself. (2) On the other hand, extraneous processing and/or chunking could vary (directly) with cognitive strain. Another open question concerns the relative advantages of assuming that C has an integral value rather than estimating C as a data parameter.

References

- Anderson, J. R. and Bower, G.H. Human Associative Memory. New York: Halstead Press, 1973.
- Bower, G. H. Chunks as interference units in free recall. Journal of Verbal Learning and Verbal Behavior, 1969, 8, 610-613.
- Broadbent, D. E. Perception and communication, New York: Pergamon Press, 1958.
- Brown, J. Some tests of the decay theory of immediate memory. Quarterly Journal of Experimental Psychology, 1958, 10, 12-21.
- Conrad, R. Serial order intrusions in immediate memory. British Journal of Psychology, 1960, 51, 45-48.
- Corballis, M. C. Rehearsal and decay in immediate recall of visually and aurally presented items. Canadian Journal of Psychology, 1966, 20, 42-51.

- Crowder, R. G. Prefix effects in immediate memory. Canadian Journal of Psychology, 1967, 21, 450-461.
- Dalrymple-Alford, E. C. Perception and immediate memory. British Journal of Psychology, 1967, 21, 63-67.
- Drachman, D. A. and Zaks, M. S. The "memory cliff" beyond span in immediate recall. Psychological Reports, 1967, 21, 105-112.
- Gorn, S. On conclusive validation of symbol manipulative processes. How do you know it has to work? Journal of the Franklin Institute, 1972, 499-518.
- Katz, L. Limited capacity hypothesis (?). Journal of Verbal Learning and Verbal Behavior, 1968, 7, 942-944.
- Meunier, G. F., Ritz, D. and Meunier, J. Rehearsal of individual items in short-term memory. Journal of Experimental Psychology, 1972, 95, 465-467.
- Miller, G. A. The magic number seven, plus or minus two: Some limits on our capacity for processing information. Psychological Review, 1956, 63, 81-97.
- Murdock, R. B. Test of the "limited capacity" hypothesis. Journal of Experimental Psychology, 1965, 69, 237-240.
- Paivio, A. Imagery and verbal processes. New York: Holt, Rinehart & Winston, 1971.
- Posner, M. I. Rate of presentation and order of recall in immediate memory. British Journal of Psychology, 1964, 55, 303-306.
- Posner, M. I. and Rossman, E. Effect of size and location of information transforms upon short-term retention. Journal of Experimental Psychology, 1965, 70, 496-505.
- Rumelhart, D. E., Lindsay, P. H., and Norman, D. A. A process model for long term memory. In E. Tulving and W. D. Donaldson (Eds.) Organization of memory. New York: Academic Press, 1972.
- Savin, H. B. On Conrad's prefix and grouping in short term memory. Quarterly Journal of Experimental Psychology, 1968, 20, 123-128.
- Scandura, J. M. Structural Learning I: Theory and research. New York: Gordon and Breach Science Publishers, 1973.
- Sitterly, T. J. Short-term retention of sequentially presented digits as a function of interdigit interval, digit duration, and series length. Journal of Experimental Psychology, 1968, 78, 174-178.
- Somers, R. H. "Limited capacity" hypothesis in incidental and intentional learning. Psychological Reports, 1967, 21, 545-548.
- Sternberg, S. The discovery of processing stages; Extensions of Donder's method. Acta Psychologica, 1969, 30, 276-315.
- Suppes, P. Some theoretical models for mathematics learning. Journal of Research and Development in Education, 1967, 1, 5-21
- Voorhies, D. Information processing capacity. University of Pennsylvania: Ph. D. dissertation, 1973.
- Waugh, N. C. and Norman, D. A. Primary memory. Psychological Review, 1965, 72, 89-104.
- Whimbey, A., Fischhof, V. and Silikowitz, R. Memory span: A forgotten capacity. Journal of Educational Psychology, 1969, 60, 56-58.
- Whimbey, A. and Lieblum, S. Individual differences in memory span with and without activity intervening between presentation and recall. Journal of Educational Psychology, 1967, 58, 311-314.
- Whimbey, A. and Ryan, S. F. Role of short term memory in solving reasoning problems mentally. Journal of Educational Psychology, 1969, 60, 361-364.
- Williams, J. R. and Fish, D. Effect of item length and number of different elements on immediate memory. Psychonomic Society, 1965, 3, 353-354.
- Woodworth, R. S. and Scholberg, H. Experimental Psychology. New York: Henry Holt & Co., 1954.

HOW CAN ONE CHARACTERIZE KNOWLEDGE STRUCTURES . . . ?

F. Schott

Technische Universität Braunschweig
West GermanyAbstract

An approach to the construction of the contents, objectives and sequences of teaching within a formalized descriptive system.

The formalized descriptive system here proposed has two objects:

1. The exact definition of the structure of educational contents, teaching objectives, teaching sequences and behavioral change in a single formalized descriptive system.
2. The construction of a terminology for "educational contents", "teaching objectives" etc., in which all the concepts can be referred to the same formal system.

The proposed formalized descriptive system has three types of constructional elements: "building blocks", situations and changes. The "building blocks" consist of elements and relations. From these "building blocks" situations can be constructed; these represent propositions. From situations, changes can be constructed; these may represent teaching objectives, behavior or behavioral change.

The descriptive system permits both a quantitative determination of the contents and of the similarity between contents, and also the presentation of the structure of contents. Moreover, it permits precise discrimination between the content aspects and the behavioral aspects of objectives. Propositions can be represented unequivocally in three dimensions. The method of presentation derives from the work of Newell, Shaw and Simon, and that of Dörner.

The usefulness of the proposed formal system of description in the empirical investigation of different educational problems will be demonstrated through examples.

TEACHER PLANNING: A FUNCTION OF MANAGEMENT

Yehoash S. Dworkin
The Ideas System

Nancy E. Dworkin
New York University

The system described below has been developed to facilitate teacher planning through the utilization of concepts and approaches drawn from information feedback and management. The basic assumption of the system is that each teacher can be viewed as a manager of materials, manpower, and production schedules, who is locked into a closed system of set demands and yet subjected to continuously shifting conditions requiring on-going decision making.

The closed system is the pre-set demands of a particular curriculum, bounded by the administrative strictures of school districts, demographic conditions, budget constraints, etc.; the manpower is the student population with which the teacher must interact on a daily basis; the materials are all of the non-human resources on which the teacher can draw in order to meet production standards, and the production standards are both the curricular and achievement averages which are posited as norms for any given grade or curriculum area.

Although the initial impetus for the system was a result of work done by the authors in the Program for Learning Studies at Children's Hospital National Medical Center under the direction of Dr. Mark N. Ozer, it is now utilized as a planning system for all classroom teachers in the elementary grades. Additionally, the initial target populations were learning disabled children who were identified as beyond the scope of the classroom teacher. At present, however, the system is viewed as generally applicable to any individual child, small group, or total class within a given learning structure.

System Principles

The system itself uses a demonstration-planning approach through which all of its major principles are exemplified. The administrative path through which the system travels is:

- 1) clinical interaction with selected learning disabled children and their teachers;
- 2) transfer of system-planning principles from the interface of single child and teacher to groupings within the class;
- 3) transfer of system-principles to planning for the class;
- 4) use of the planning system on a selective basis where standard classroom planning techniques prove inadequate.

The system addresses itself to the principles of Energy Conservation (on the part of the teacher); Differential Production Frameworks (on the part of the child); Information Feedback (between teacher and child). Each of these principles is exemplified in a specified segment of the Demonstration System described below.

The Demonstration System

The Demonstration System is composed of three major phases, each one illustrating a particular strategy framework through which information and problems can be presented to the child. Each of these frameworks illustrates an organizational principle relating to the information delivery, and a functional principle relating to problem solving. The organizational

principle refers to the teacher's role, and the functional principle to the child's role. These frameworks are:

- I Programming/Logic
- II Modality Variation/Patterning
- III Focus/Reasoning

Each of the strategy frameworks encompasses a series of tasks of graduated difficulty, with each task utilizing the problem-solving information developed through its predecessor task.

Cutting across all of the strategy frameworks is a cueing system which is operative for every task. The cueing system allows for the manipulation of information vis-a-vis the initial instruction, so that each child can be guaranteed successful solution and/or completion of the task. It addresses itself to elements of the initial problem and the range of possible solutions, and involves the following three steps:

I. Identification. This cue reformulates or restates the initial instruction, but maintains all of the complexity of the original problem and the original solution.

II. Isolation. This cue focuses on specific parts of the problem, and the objective is to reduce the probability of student error. Although the initial presentation of the problem is reduced, the complexity of the solution remains intact.

III. Manipulation. This cue allows for changes in any or all parts of the original problem. Further it may bring both teacher and child into physical contact with the materials of the problem and the original solution may be radically changed so long as a single element of each is retained. It must guarantee ultimate solution.

Surrounding both the strategy frameworks and the cueing system is a dual information feedback loop through which teacher and child are informed of progress at any given point in the system, and through which they are presented with the information necessary for decision making. This dual information feedback loop is triggered by a pre-arranged signal system in which the teacher and the child agree on a set of differential reinforcers. The first reinforcer is operative when the child has attempted solution but needs further help; the second reinforcer is operative when the child has achieved partial solution to a given problem; the third reinforcer is operative when the child has arrived at solution to a single problem; the fourth reinforcer is operative when the child has completed all his tasks within a single strategy framework, or set of problems; and the fifth reinforcer is operative at the completion of the entire demonstration.

1. Non-solution. The teacher uses a neutral phrase to indicate that further work is needed in order to complete the problem ("O.K. Let's work on it together").

2. Partial solution. The teacher identifies the correct parts of the solution and reintroduces the initial instruction in order to complete the problem.

3. Single solution. The teacher utilizes a pre-arranged signal to let the child know that he has successfully completed the problem. This signal is agreed upon prior to the demonstration by teacher and child, usually by the child selecting a word or phrase through which he would like to be informed of his own success.

4. Solution sets. The child identifies the ways in which he arrived at a solution to X number of tasks, following the teacher's positive reinforcement of his success.

5. Completion. The teacher summarizes all of the successes (hits) achieved by the child, and asks the child which tasks and solutions he found most helpful.

The differential production principle is carried by the strategy frameworks. One of the major issues confronting most teachers in planning the management of their class, beyond the choice of curricular units, is the production framework through which they would like the children to arrive at problem-solution. Unfortunately, most teachers have not been trained to differentiate varying levels of production complexity. Thus, although their selection of materials and curricular units may be entirely appropriate for the class or any individual children, the method of production may be well outside their [the children's] range. In the demonstration, the Focusing/Reasoning Framework illustrates the simplest form of production, where all of the information is given to the child, the child's task then is to identify, or in some way manipulate that information without either creating new information or changing the original elements.

The Programming/Logic section identifies the final solution, but provides only part of the information necessary in solving the problem. The Production Method, therefore, (an intermediate one), demands no change in the initial elements of the problem, but requires that the child move from the known to the unknown on a binary decision path.

The most complex production tasks demand that the child totally reconstitute given information in order to arrive at exclusively independent production. This is illustrated through the Modality Variation/Patterning Framework. It is quite evident in the demonstration that each production method allows a range of tasks from the most simple to the most complex. It is not the problem and the solution that constitute the greatest challenge to the teacher's planning capabilities, but, rather, her ability to identify appropriate production methods for each child or each group in the class.

The energy conservation principle is carried by the cueing frameworks. Each of the three cues demands a differential amount of energy on the part of the teacher, vis-a-vis helping the child arrive at successful solution. In any learning framework, it is desirable to transfer maximum amount of energy expenditure from the teacher to the learner. In simple mathematical terms it is quite evident that the expenditure of a single energy unit from child to teacher has a value of one (1), whereas the expenditure of a single energy unit from teacher to class has a maximum value of $n(1)$.^{*} Thus, each time any part of the class experiences difficulty with problem solution, the teacher must be concerned with her selection of cues, in effect diminishing her available supply of energy. The cueing system used throughout the demonstration has a geometric progression in terms of energy expenditure. Each additional cue doubles the amount of manipulation and thus doubles the amount of energy expenditure. The first cue simply changes the nature of the verbiage involved in the initial instruction. The second cue requires manipulation of at least one part of the solution and one part of the

* n = size of class

problem. If we establish a model problem in which a teacher addressed 40 students and had half of them arrive at successful solution with each additional cue, we would see the following differences in energy expenditure:

EU = ENERGY UNIT = CUE	
Problem I	Problem II
(EU1) 40 = 40	(EU4) 40 = 160
(EU2) 20 = 40	(EU2) 20 = 40
(EU4) 10 = 40	(EU1) 10 = 10
120	210

(CODE: Identification = EU1; Isolation = EU2; Manipulation = EU4)

What is illustrated in these two solutions is the difference in energy expenditure occasioned by differential ordering of cues.

The above problems do not mean to imply that teachers will invariably achieve 50 percent success at each cueing state. Rather, the first of the demonstration is to have the teacher recognize that starting with the simplest cues (in terms of EU), will ultimately allow her the maximum amount of time for those children experiencing maximum difficulty arriving at solution. Unfortunately, most teachers will start with the maximum EU cues, namely, manipulation of problem and solution, before they have tried the others.

Administrative Path

In the clinical application of the system, the teacher selects children with whom she is experiencing difficulty. Generally, these children are identified as learning disabled, but their problems may lie in either behavior or academic areas. The demonstration team works through the problems with each individual child in the presence of the teacher and any other school specialists who have had prior contact with the case. Following the demonstration, the teacher is asked to identify strategy and cueing combinations which [she felt] were most effective in helping the child. The teacher then selects a set of teaching objectives for the coming two-week period, and together with the demonstration team plans a series of lessons using the strategy-cueing combinations and the selected objectives. As part of the planning, they arrange a dual feedback schedule through which child and teacher will be informed of the move towards solution.

At the second level, from two weeks to a month after the original demonstration, the teacher and the demonstration team review the results of the initial planning, and select a new set of objectives. At this stage, however, the objectives are related to a grouping in the class, rather than towards an individual child. Once the strategy-cueing combinations are agreed upon, the demonstrators carry out an on-the-spot demonstration with the designated group of children.

The third level of the system is divided into two phases, again coming from two weeks to a month after the second level, with two weeks separating the two phases. Phase A is in the form of a Teacher Seminar in which all of the teachers involved in the system exchange information concerning effective techniques and reinforcement schedules. The demonstration team serves as a catalytic agent for this discussion. Following this exchange of information, each teacher selects classroom objectives and goes through a planning session with a member of the demonstration team.

During Phase B the teacher carries out a demonstration with the entire class. The demonstration team lends support, and the teacher and demonstrators review effective strategy-cueing combinations following the demonstration. At this point, the demonstrators are ready to withdraw from the system since the teacher has confirmed her ability to handle the management elements inherent in the planning system.

Summary

The Demonstration-Planning System rests on two cardinal assumptions: the first is that learning disabled children can function best if they are maintained in the regular classroom, and the second is that once teachers receive sufficient help in planning the management of their classrooms, they are fully capable of dealing even with the disabled learner and helping him move towards success. The principles of the system emphasize energy conservation, differential production and information feedback through the use of a demonstration model divided into strategy and cueing components. Ultimately, the thrust of the system is to diminish the necessity for outside consultants and specialists by enhancing the teacher's ability to organize and manage the classroom.

HETERARCHIES AND THEIR RELATIONSHIP TO
BEHAVIORAL HIERARCHIES FOR SEQUENCING CONTENT IN INSTRUCTION

M. David Merrill

Andrew S. Gibbons

Brigham Young University

As a background note it should be indicated that over the past several years we have been involved in attempts at content analysis working with subject matter experts in several disciplines as a necessary step in the design of instructional materials. We started in this process by a rather literal application of Robert Gagne's hierarchial analysis. More recently we have developed primitive heterarchies as opposed to hierarchies in statistics, algebra, accounting, physical fitness, and to a somewhat lesser extent in freshman composition. Another group under Steve Anderson (BYU-ITV) have identified primitive heterarchies in about a dozen areas including such diverse topics as child development and interview skills. This experience has convinced me of the following:

1. Adequate instructional development cannot proceed without some kind of careful specification of the content structure involved.
2. Instructional development proceeds more efficiently and effectively when guided by a previously analyzed content structure.
3. Content analysis can proceed quite independently of concerns for instructional strategies or delivery systems, or student behavior. It should be noted that it is the interface of concerns with all of these facets of the instructional process that results in instructional products or programs.

This paper is not a report of this practical experience but is much more abstract and theoretical.

It should be noted that the author is not a philosopher. The content of this paper however, is very much in the realm of philosophy. The author is not a practicing experimental psychologist. However, this paper is very much related to work in cognitive psychology. The author is not a linguist. However, this paper is very much concerned with language. What I have to say has probably already been said by scholars in these fields. As indicated my orientation is instruction and instructional development. These ideas have been useful in helping me to think about describing subject matter for the purpose of sequencing it for instruction. If the ideas are not new at least they are not widely known among those who are involved in the day-to-day analysis of subject matter for preparing instructional materials.

Summary

A heterarchy is a representation of the relationship between the components of subject matter. A hierarchy (especially as defined by

Robert Gagne (1968) and as widely applied in instructional development) is a representation of the relationship between the behaviors of students as they interact with subject matter. A premise of this presentation is that the components of subject matter and the behavior of students are independent phenomena. Hierarchies have frequently been used to sequence instructional interactions. Applying student behaviors to a heterarchy suggests that there are many alternative ways to sequence instructional interactions only one of which is hierarchical.

It is proposed that subject matter consists of the imposition of structure on the world of objects and events. Content analysis is the process of representing this structure by means of some symbolic metaphor. Instructional analysis is the process of sequencing components of subject matter and prescribing student interaction with these components in such a way that the student can remember and use all or part of this structure in explaining and manipulating his environment in predictable ways.

Subject Matter

Subject matter consists of referents which have been organized into concepts which are related by operations. Subject matter is represented by symbols which are combined into a language which forms a metaphor for the referents and their relationships that exist in the "real world." Let's explain this statement.

A REFERENT is an object or event which did exist, which does now exist, or which can at some future time exist "out there" in the "real world." By events we include behaviors of humans as well as natural events which occur without human intervention. In the "real world" there are innumerable such objects and events. However, they do not constitute components of subject matter until they are given some organization. Without organization they are raw facts -- matter unorganized.

In order to comprehend the world, man imposes order by arranging referents into categories which share common characteristics. The formation of selected sets of objects and events into classes constitutes the first step toward defining subject matter. Note that this organization can be imposed by a single individual and in fact, some esoteric organization is probably imposed by each of us as we experience the world. However, it is not until an individual persuades a significant group of other individuals to adopt his categorization that we have socially significant subject matter.

Objects and events are not static but interact with each other. One significant class of categories are those events that relate one object or event to another. Such events may be relational such as inclusion, exclusion, and order or they may involve processes such as transformation, composition, and decomposition where objects are changed by the interaction into members of a different class.

SYMBOLS represent a special class of objects which can be used to represent the members included in a given class of referents. The word referent suggests that given objects or events are those elements that are

referred to by a given symbol. Such symbols can be words, mathematical notation, musical notation, and a wide variety of other special notational systems which have been developed to represent the classes of referents that have been defined.

LANGUAGE consists of procedures for combining symbols such that the combinations form a metaphor for the interactions of referents. Languages also have the characteristics of having logical relationships of their own such that not every statement or even every symbol constituting a language system must have a metaphorical relationship to referents in the real world. Languages are usually not completely isomorphic with referents. A given language statement often adequately represents only some of the attributes of a given relationship between referents or particularly between classes or referents.

For purposes of content analysis for instruction, subject matter can be said to be comprised of only two components; concepts and operations. A CONCEPT has three aspects: a set of referents, a symbol-label which stands for the set, and a language statement which provides a representation of the critical characteristics of members of the set. Such a statement is usually called a definition.

An OPERATION also has three aspects: a set of event referents, a symbol-label which stands for the set, and a language statement which provides a representation of the critical characteristics of members of the set. When such a statement is expanded to include identification of the concepts affected by the operation it is called a rule.

Some additional elaboration of representation is probably desirable. A given referent can be represented in varying degrees of fidelity. An object can be represented in varying degrees of fidelity. An object can be represented by a model which may be almost completely isomorphic with the referent except in perhaps one dimension such as size; a picture is less isomorphic with relevant characteristics than a model but much more isomorphic than a verbal description. A similar argument can be made for events using such forms of representation as simulations, motion pictures, video tapes, etc.

Content Analysis

Given the above definition of subject matter, content analysis is, on the surface, a very simple process. First, identify the concepts that constitute a given discipline or portion of the discipline that is to be taught. Second, identify the relationships between these concepts. Third, using an appropriate symbol system represent this structure. Such a representation constitutes what has herein been called a heterarchy.

Heterarchy Representation--Problems and Challenges

As it turns out the adequate representation of heterarchies is somewhat more complex than at first appears. Listed very briefly without elaboration or review of possible solutions are some of the difficulties and challenges involved:

First: The first problem is the selection or invention of an adequate symbol system for metaphorically representing subject matter structures. It is quite clear that major developments in other areas of inquiry were greatly accelerated when an adequate notation system was developed. Michael McDonald-Ross (1974) clearly identified the challenges for developing such a notational system. His paper also reviews some of the systems which have been suggested by other investigators.

Second: Most subject matter can not be described in a single dimension or even in a 2-dimensional space, it is instead multifaceted. It is rather apparent in many subject matters that assuming a procedural orientation to the subject matter in one set of concepts and relations but taking a theoretical approach results in a different configuration of concepts and relations. To be sure, some of the concepts identified as elements of one facet enter into the relationships of other facets. This multifaceted nature of subject matter also becomes apparent when one takes a descriptive as opposed to an application approach to a given body of subject matter.

McDonald-Ross (1974) and Pask (1971) feel that in almost every content analysis situation there are task structures related to problems which the learner is learning to solve and knowledge structures related to the logical interrelationship of concepts and operations. When one considers the variety of task structures which are apparent for a given rule set in many subject matters one gets a glimpse of the multifaceted nature of organized content. This problem adds a challenge to the identification of a notational system for representing content structures in that such notation must be able to unambiguously represent this multifaceted nature of subject matter structure.

Third: Not only is subject matter multifaceted but the concepts and relations constituting the subject matter are not all at the same level of inclusiveness. That is, some concepts are more general than others and completely subsume a group of other concepts. Some relations are between concepts at one level of generality and other relations are between concepts at quite another level of generality. If these inclusion relationships were nice orderly taxonomic structures the problem would be somewhat simplified. Often however the inclusion relationships involved are only partial. Furthermore many of these inclusion relationships cut across facets of the subject matter in different directions. Another challenge for an adequate notational system is to capture the essence of these layered relationships.

Fourth: Evidence from cognitive psychology would convince one that none of us learn new concepts or relationships in isolation. We apparently build our mental representations of subject matter structures from that which we have already stored in our long term memories. This suggests that part of many subject matter structures are the classical analogies which can serve as a scaffolding metaphor for learning the new structure. A common example of such scaffolding content is the use of water flow as a metaphor for electron flow. Many other analogies exist or could be invented. An adequate notational system must represent such scaffolding and indicate the degree to which it is isomorphic with the concepts and relations constituting the target subject matter.

Student Behavior

In several other sources (Merrill, 1973; Merrill and Boutwell, 1973) it has been suggested that student behavior related to cognitive material can be classed into four types. (This is rather direct application of Robert Gagne's conditions of learning and is not new with the author.) These types are recall and recognition (recognizing the differences these are nevertheless treated as a single category), classification, rule using, and rule finding. It is assumed that these categories of behavioral events are independent of subject matter content. That is, any type of behavior might be demonstrated in relationship with any type of content. This suggests that several types of instructional interactions are possible:

1. The student can provide a label for a specific referent.
2. The student can state a definitional proposition.
3. The student can state a relational proposition.
4. The student can find and state a new proposition either definitional or relational.

Note that all of the above deal with the student's ability to handle the language metaphor which represents referents "out there." In addition the student can manipulate referents as follows:

5. The student can classify a referent as a member of a particular concept.
6. The student can use a given operation to manipulate referents.
7. The student can invent or discover a new operation and manipulate referents in a new way.

Note that these latter behaviors are usually necessary before one is willing to infer understanding or comprehension of the language metaphor.

The whole question of representation between symbolic language and referents has been omitted. Obviously between each of the extremes identified a student can experience all manner of simulation. For example, the manipulation of models rather than referents. The classification of pictures rather than referents, etc.

The possibility of inverse representation also exists. That is, the manipulation of objects which represent the language metaphor. In these cases language itself is the referent and the metaphor for language is objects. An example is the use of cuis rods in mathematics.

The Problem of Sequence

By definition a heterarchy has no implied sequence. That is, one should be able to start at any concept or operation in the heterarchy and by tracing through the network determine its relationship to all other concepts and operations in the heterarchy. It is possible to follow such a course in sequencing instructional events. Most of us, however, would probably subscribe to the position that some of the many possible sequences must be more efficient and effective for a particular learner than others.

After all, the evidence is pretty convincing that learners have some type of short-term memory which has a very limited capacity. Until the base of information stored in long-term memory has been increased sufficiently to provide meaning to the inputs, random patterns through a body of subject matter must be difficult for a student to handle.

Hierarchies

What is a Gagne type hierarchy related to the above description? Several possibilities exist. The following is one representative of these possibilities. Given a segment of a content structure consisting of an operation relating several concepts, the top box of a hierarchy would be the student's ability to manipulate given referents by using the operation. The next level of the hierarchy might consist of the student's ability to classify referents of each of the concepts which comprise the rule. The next level would consist of the student's ability to classify referents of attribute concepts which are used to define the main concept. An even lower level would be the student's ability to provide a label for a specific referent from each of the attribute and main concept classes. As Gagne (1970) indicated in the early descriptions of the hierarchy the ability to state the definition or the rule are probably not part of the hierarchy but may be included as parallel behaviors.

Most interpreters of the Gagne type hierarchies assume that the correct sequence of mastering such a hierarchy is to start at the bottom and to make sure the student masters each prerequisite item as he proceeds through the hierarchy. In the case cited, perhaps the first task for the student should be in the "top" box. By giving the student referents and clearly indicating that they do indeed belong to the component concept classes the student can then be taught the use of the operation without having the necessary "prerequisite" classification skills. Granted this may not be the "real world" performance required when he won't be given clearly identified referents. The sequence would proceed by moving out in the hierarchy from the operation. After he has learned to use the operation given referents he is then taught to identify given referents of the concepts when the attributes are pointed out to him. Then he learns to point out the attributes himself and so forth. This is a simple hierarchy but the same principle applies to more complex hierarchies. The student is initially given enough information to enable him to perform complex operations. Having learned the operation given prerequisite information he is then progressively required to provide more and more of the prerequisite information himself. Again considering current work in cognitive psychology this sequence might provide a much better orientation for the student as he works back down the hierarchy because he knows the intent of the instruction. When he must learn to classify referents of a concept before he knows why or in what way the concept enters into meaningful relationships it is probably difficult for him to represent such concepts in his long-term storage in a way that facilitates retrieval.

Space limitations prevent the exploration of sequences involving different levels of the heterarchy or sequences employing analogous concepts. In all cases the consideration of heterarchies and the application of student

behavior to such heterarchies makes clear that hierarchies represent only one approach to the sequencing of instruction.

Why is Content Analysis Necessary

Why isn't task or behavioral analysis sufficient? In specifying subject matter for instruction three questions need to be answered: (1) What is there to teach? (2) Of what there is to teach, what should be presented as part of a given instructional system? (3) In what sequence? Limiting our analysis to tasks or student behavior increases the probability of the following types of errors: first, teaching performances rather than meanings; second, omitting significant content; and third, failure to teach organized subject matter systems as systems.

A task analysis can identify the steps that a student must take to perform a given task or to produce a particular outcome. If we assume that the proper sequence for teaching is the same as the performance sequence we are likely to greatly complicate the learning process for many tasks. The performance of steps in a sequence on the job often implies knowledge of why the step is being performed. This suggests that an understanding of the concepts and operations underlying the performance will in many cases facilitate acquisition of the performance. Furthermore, in many cases a given algorithm or set of performances represent only one of the many possible ways that a given outcome can be produced. Limiting the instruction to the performance behaviors makes new learning necessary for the student to engage in alternative procedures. Providing a grasp of the underlying content structure not only enables the student to more easily grasp alternative procedures it may even facilitate his developing his own algorithm for performance.

Second, when we limit instruction to the behaviors identified as a result of the analysis of a given performance we are likely to omit understanding that may be necessary for performance in other situations where some of the conditions for performance may have been altered. If on the other hand all of the concepts and operations involved in the performance of a given task have been taught and the performance has been taught as a special application of this knowledge it is much less likely that significant concepts or operations will have been omitted from the instruction.

Third, there is merit and necessity for students to learn organized subject matter without ever applying it to specific performances. This is especially necessary for learners who will pursue creative careers which will push back the frontiers of knowledge in such areas. Limiting the analysis of subject matter to task or behavioral analysis makes the identification of organized bodies of subject matter difficult. Based only on task or behavioral hierarchies, it is not easy to determine a sequence for presenting such organized bodies of knowledge. A careful subject matter analysis and a resulting heterarchical representation of that subject matter greatly facilitates the organization of that content for instructional presentation as well as giving the student a representation of the subject matter which facilitates his own storage and retrieval of the content in memory.

Conclusion

A hierarchy is a structure that results when student behavior is applied to the components of concepts and relations of a part of a heterarchy. The assumption that instruction should follow a sequence from the bottom of a hierarchy to the top may not be justified based on work in cognitive psychology. By defining heterarchies the determination of many other sequences becomes more apparent.

In content analysis, the major tasks lying ahead involve the more detailed identification of the interrelationships of facets of subject matter structures and the identification and specifications of an adequate representational notation for representing this subject matter.

BIBLIOGRAPHY

- Gagne, Robert M. Learning Hierarchies. Educational Psychologist, 6, No. 1. 1968.
- Gagne, Robert M. The Conditions of Learning. New York: Holt Rinehart and Winston, 1970.
- McDonald-Ross, Michael. Glass Beads and Geometric Monsters: The Problem of Representing Knowledge. Paper presented at AECT Convention Atlantic City, N. J., March 1974.
- Merrill, M. David. Content and Instructional Analysis for Cognitive Transfer Tasks. AV Communication Review, 21, 1973, 109-125.
- Merrill, M. David and Boutwell, Richard C. Instructional Development: Methodology and Research. In Kerlinger, Fred (#d.), Review of Research in Education, vol. 1, Itasca, Illinois: F. E. Peacock (AERA), 1973.
- Pask, Gordon. A Cybernetic Experimental Method and It's Underlying Philosophy. International Journal Man Machine Studies, 1971, 3, 279-337.

A NEW TYPE OF TEST OF COGNITIVE FUNCTIONING:
AN INTELLIGENCE TEST BASED ON FINITE AFFINE GEOMETRY

Zoltan P. Dienes
University of Sherbrooke

The construction of this series of tests started with a mathematical analysis of Raven's Progressive Matrices. It seemed that some of the mathematical content was interesting, but that there was no well discernible system in building up mathematical hierarchies. Perhaps a test might be interesting in which such mathematical hierarchies were used, reaching more and more complex levels of construction and thinking. Finite vector spaces, and the corresponding projective geometries seemed a good subject matter, and experiments with individual children and with classes of children seemed encouraging. Reasonably normal curves were obtained on achievement in solving problems in this area by ten year old children. In order to provide room for generalization, a line was withdrawn from a finite projective geometry, thus creating an affine geometry. Such a geometry always has a square number of points and so would lend itself to a matrix type of layout.

The first experiments were done on a 9-point affine geometry, obtained from a 13-point projective geometry, itself obtained from a 27 element vector space, where the scalars were 0, 1, and 2. Pictures of boys and girls and dogs were used, so that no picture contained more than two of any kind of creature. The pictures were paired as pairs of additive inverses, meaning that the sum of two pictures in a pair, taking the creatures independently, always came to zero modulo 3, namely to 3 or to 6 in practice. The layout was as given in Figure 1.

The first test was to learn to recognize properties that were true of a certain number of pictures, and find in fact all the pictures which possessed that property. The properties increased in level of complexity as the test was done. There were, in the 9 by 9 matrix four types of property. (see also table 1).

- (1) The absence of a certain kind of creature.
- (2) The presence of two kinds of creatures in equal numbers.
- (3) The "sum" (modulo 3) of two kinds of creatures is zero (namely zero, 3, or 6).
- (4) The "sum" (modulo 3) of two kinds of creatures is equal to the number of the third kind of creature.

Subjects were taught one kind of property, then they had to find the remaining pictures with that property, then all the pictures possessing similar properties. The subjects were taught the next type of property, and so on, until they had gone through all four types of properties.

Next they were told that all the pictures of the kind they have been working on had not been drawn. They were to draw them in the four divided circles provided, or write their names if they preferred.

It was interesting to score on the following points:


- (i) Had the subjects noticed that there were never more than two of a kind in a picture?
- (ii) Had the subjects noticed that the sum modulo three of the upper half and of the lower half of each disc was zero?

Finally they were asked to find a common property to all their drawings. This was that the total sum was zero. This property had not been met in the drawn pictures in the matrix, and had to be "discovered" by the subjects.


Part B of the test was an inversion of the procedure. A list of twelve properties, pictorially and verbally presented, was provided. On each page subjects found a certain number of discs marked in black, and they had to identify the property possessed by all the pictures encircled.

In other words, the first part of the test required a passage from properties just learned to pictures possessing these, the second part from sets of pictures to the common property of all the pictures.

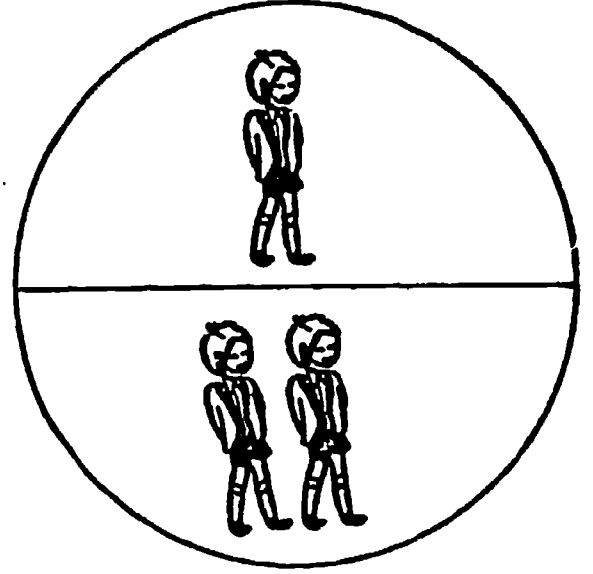
Figure 1




1




2



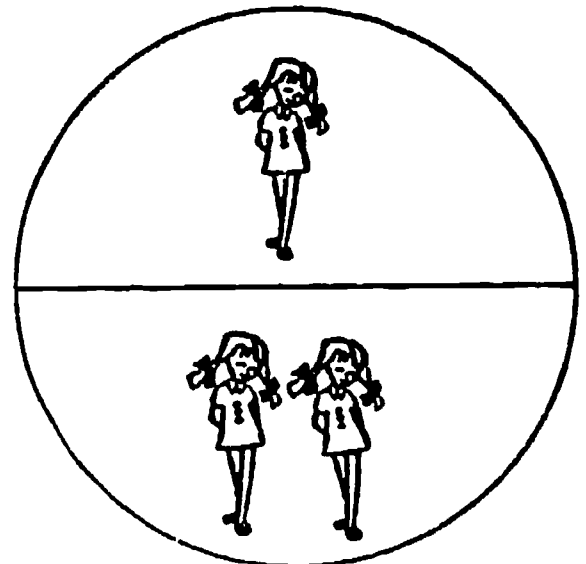
3




4



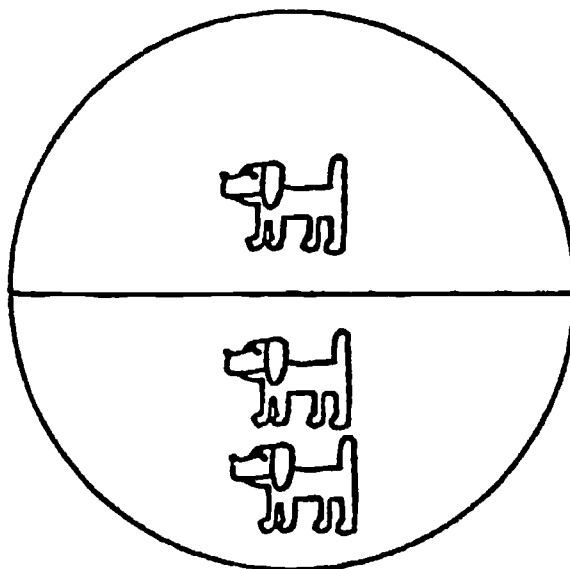
5




6



7



8



9

LIST OF PROPERTIES

<p>A</p> <p>NO DOGS</p>	<p>B</p> <p>NO GIRLS</p>	<p>C</p> <p>NO BOYS</p>
<p>D</p> <p>BOYS same as GIRLS</p>	<p>E</p> <p>DOGS same as BOYS</p>	<p>F</p> <p>GIRLS same as DOGS</p>
<p>G</p> <p>BOYS together with GIRLS are equal to ZERO</p>	<p>H</p> <p>BOYS together with DOGS are equal to ZERO</p>	<p>I</p> <p>GIRLS together with DOGS are equal to ZERO</p>
<p>J</p> <p>BOYS together with GIRLS are equal to DOGS</p>	<p>K</p> <p>DOGS together with BOYS are equal to GIRLS</p>	<p>L</p> <p>GIRLS together with DOGS are equal to BOYS</p>

Don't FORGET that when you ADD 3 counts as 0, and 4 counts as 1.

CRITERIA FOR COGNITIVE STRUCTURES: REFLECTIONS ON
 PIAGET'S THEORY OF GENETIC EPISTEMOLOGY AND INTROSPECTION¹

John Limber and Michael Chiseri

University of New Hampshire

It is obvious even to the casual reader of any of Jean Piaget's many volumes, that he is concerned with much more than child psychology. For over 50 years Piaget has developed and refined a theory of genetic epistemology (TGE), a theory about the growth of human knowledge. Among the many implications of that theory is a claim that scientific thought has its beginnings in the characteristic development of children and, furthermore, that the growth of knowledge in an individual regulates the species acquisition of knowledge. The final cognitive level attained by a mature individual is determined not only by his biological development but by the cognitive level of the culture of that individual. Knowledge is thus normative. A member of a pre-scientific culture would not be expected to attain a level of formal or scientific thought.

Not infrequently psychologists have neglected these admittedly complex yet basic objectives of Piaget's work in their analysis of its concepts and structures. As a consequence criticisms of TGE are frequently off the mark. One cannot evaluate any concepts, cognitive structures in TGE for example, without a clear understanding of the role which those concepts play in the theory.

What are cognitive structures?

Cognitive structures are elements of explanatory theories of cognitive processes. Structures, concepts, or constructs are postulated to account for various psychological phenomena just as are constructs in physics or biology postulated to account for certain physical or biological phenomena.

It may be useful here to review a number of commonly accepted (e.g. Kaplan, 1964; Fodor 1970) features of such theoretical entities. (1) Structures are not to be confused with the behaviors they are presumed to explain. While the structures are inferred from these behaviors they are categorically distinct from observational measures like verbal reports, reaction times, or pointing responses. (2) The significance or meaning of any particular structure accrues in a number of ways including the multiple indices used to relate those concepts to observables and the network of relationships among the various concepts within a given theoretical system. As a result, one cannot simply evaluate the validity of one element of a theory without careful consideration of the overall role of that element within

1. Draft of a paper presented at the 5th Annual Interdisciplinary Structural Learning Meeting, April, 1974.

the theory. (3) While the ultimate validity of a theory--and the elements within the theory--rests on their relationship with observables, it is doubtful that we can demand or specify any interesting logical (i.e. necessary and/or sufficient) relationship between theoretical entities and those observables. In other words it is unrealistic to expect to define a set of observable conditions that will define a theoretical entity, e.g. a child has a certain operational structure if (and only if) she solves correctly at least two conservation tasks. At best we can expect to draw such inferences with a certain probability of being correct. There is the inevitable error.

Criteria for cognitive structures

At this point we can turn to the primary subject in our paper--the issue of criteria or indices of cognitive structures in Piaget's theory of genetic epistemology (TGE). The broad substantive question here is--What are the empirical consequences of the various cognitive structures within TGE in order to account for the means by which the human mind goes from a state of less sufficient knowledge to a state of higher knowledge (Piaget, 1970, p.12).

A long-standing controversy in the literature on concrete operations concerns the indices of operational structures. Piaget and his associates typically require successful performance and an adequate verbal justification of that performance before concluding that a child is utilizing the structures postulated as emerging in the operational periods of development. Critics, from the earliest publication of Piaget's works, have attacked the use of verbal justification from a number of perspectives. For example, one recent critic (Brainerd, 1973) concluded that TGE fails to justify the use of verbal justification as a criterion and that judgments in themselves, that is a child's response to a dichotomous question about some quantitative relation between objects (p.173) appears to be the "minimum necessary" evidence for inferences as to the presence or absence of a particular logical structure.

Against the background of these general introductory remarks we would like to focus on two specific propositions:

- (1) that judgments are insufficient criteria to assess cognitive structures of logical thought.

(2) that a verbal justification criteria is clearly justified within TGE especially in the context of the construction of logical necessity and the consequent development of formal thought both within an individual and the culture.

On the insufficiencies of judgments

The basic point here is simply that reliance on a judgment criterion can lead to unacceptable and sometimes unnecessary errors as to the mode of thinking, or structure utilized on a given task. There is simply no one to one correspondence between judgments and structures. A child may get the judgment right for a number of reasons including multiple structures within TGE capable of solving the problem and a host of irrelevant hypotheses, subtle experimenter cues, etc. On the other hand the child may get it wrong because the appropriate structure is unavailable, an inappropriate structure was used though a suitable structure was available, memory deficits, inadequate experience, momentary distractions, boredom, etc.

Clearly the entire range of inferential errors open to the experimenter precludes any absolute certainty as to the structures giving rise to a particular judgment. The experimenter as always must attempt to minimize these errors using what the theory predicts and whatever other information he has about the situation.

It is useful here to distinguish two classes of inferential errors. Those that are specifically predictable from within TGE, that is intrinsic to it; and those errors about which TGE has nothing specific to say, but are part of the common lore of experimental psychology. Of these latter errors we have nothing much to say other than to point out that "fortuitous conservation" as Chiseri (in preparation) terms it, is much more of a problem with dichotomous judgments than with verbal justifications where the probability of a child coming up with an acceptable justification irrelevantly is very small in comparison with the probability of the child's making the correct judgment by chance.

The intrinsic errors, however, are much more important to use at this time. Within TGE it is not unusual that alternative solutions to a given problem are available; that is there exist two or more structures which in certain cases may be adequate to generate the identical judgment response. When this situation obtains, inferences from the judgment response back to the causal structure are subject to intrinsic error (a false positive inference). There is no mandate in TGE, insofar as we know, that individuals must use the most sophisticated mode of thought available to them.

Consider an extremely simple task. A child is presented with two different length and different color sticks. He is asked to say which of the two is bigger, that is, make a relational size judgment. In addition we can vary the means of presenting the sticks to the child: we can let him hold the sticks; we can hold the sticks ourselves out of his grasp; and we can present the sticks in sequence--not simultaneously--to the child. In all cases the child's task is the same--which is bigger? It would be no surprise to find that older children (5-6) solve all three tasks equally well, while

younger children (1-3) have increasing difficulty with the increasingly abstract, i.e. removed from direct action, tasks. Within TGE one might distinguish three modes of thought--three distinct structures--operating in the three versions of the task. These are a direct action mode, a perceptual action mode, and an image mode. In each case a comparison is to be made; but the manner of comparison differs.

The point of this somewhat obvious example is that particularly for the older children, using the judgment criterion alone, we would be unable to infer which mode of thought was employed in giving a correct judgment.

This fundamental point has been made time and time again by Piaget and his associates; it has also been made quite clearly by Smedslund (1969). Indeed the Piagetian framework is notable or notorious for the many alternative constructions that can be placed on a given class of behavioral events. This is perhaps the most serious difficulty with TGE. As another example, consider Piaget's (1968) comments on T.G. Bower's work on object constancy with young infants. Piaget argues that those experiments only show that recognition of a previously seen object can occur very early in life; positive performance on a recognition task alone does not imply conservation.

...it does not tell us whether for the baby the hidden object continues to exist when it is placed behind another object, or whether the object has momentarily ceased to exist, like an image that can disappear and reappear and then be recognized as having been seen already. The hypothesis of conservation would suppose a substantiation of reality and an organization of space; recognition by itself does not indicate whether this organization is present or not, or even whether such organization is possible or not at this age. (p.979, fn.7).

Another example concerns the so-called transitivity or seriation tasks. These are much more germane to our interest as they concern the distinction between pre-logical and logical modes of thought. According to Piaget, transitivity arises from the development of the seriation operation, approximately at 7 years. As always it is necessary to understand just what this notion implies within the Piagetian framework:

The child's understanding of transitivity may be tested by having subjects compare A and B and then B and C perceptually (from a series $A \ll B \ll C$) but then they must deduce A's relation to C, without comparing the two, which preoperatory subjects refuse to do. (1968 ; p. 102, emphasis added).

The operative word here is "deduced" in contrast to other words Piaget might have used, for example "compute" or "correctly judge". Merely eliciting the correct judgment that $A \ll C$ is insufficient evidence to warrant the inference that a logical solution was employed. It is the mode of solution that is important not merely

whether or not the child solved the problem. Interesting experiments, such as those reported by Brayant and Trabasso (1971) that do not report assessing the mode of solution in transitivity tasks are just irrelevant to the validity of TGE. Brayant and Trabasso argued that memory or attentional deficits in children might well be responsible for younger children's failure to make correct "transitivity" judgments. They suggested that perhaps the younger children forget the pairwise relations, $A < B$, $B < C$, etc. and as a result are unable to compute that $A < C$. They convincingly demonstrated that if the younger children were trained on the pairwise comparisons until all ages reached a comparable level of recall, then all age groups made correct judgments at nearly the same high level in a five series $A < B < C < D < E$ judgment task. This general argument, buttressed with their data, would be a powerful alternative to the Piagetian theory if the judgment criterion along were relevant to inferences about the transitivity mediating structures. That this is not the case is clear from Piaget's remarks cited above and elsewhere (Piaget, 1967; p.50) where he explicitly cautions against permitting imaged-based solutions to substitute for logical solutions in such tasks. Although the matter deserves careful experimental inquiry, the Brayant and Trabasso procedure seems likely to maximize the use of imagery.

In order to infer that a logical structure was employed in a task, one must have evidence that the child see the logical necessity of his judgment as a function of the given premises $A < B$ and $B < C$ in contrast to more empirical, pre-logical, action, perceptual, or image-based solutions independently postulated within TGE. Alternately--or preferably in concert--the experimenter might seek a task where the effectiveness of various prelogical structures is minimized, or even better, would lead to an incorrect judgment. This is more nearly the case in a properly conducted conservation of liquid paradigm where the child's non-logical empirical strategy, e.g. "tall is more", comes into conflict with his intuitive tendency to conserve. It is the resolution of the disequilibrium brought about by such conflicts which ultimately brings about logical modes of thought.

While the judgment criterion is generally inadequate to discriminate among alternative structures within TGE, this in itself is not justification for the important role verbal criteria have in the investigation of logical structures. To do this we must turn to the extremely important and complex role that language and the notion of logical necessity play within TGE.

LOGICAL NECESSITY AND INTROSPECTION

Our argument is essentially that logical necessity is an intrinsic conscious correlate of logical thought and that evidence for the presence of necessity is to be found in the child's verbal reflections or introspections about his causal mental operations.

In addition, we want to emphasize that neither necessity nor its verbal correlates are merely incidental indicators of logical structures, comparable say to other indices like GSR or pupillary dialation. Instead these verbal reflections--a product of one stage of development--become the raw material for the transition into the stage of formal, analytical, thought processes within an individual and the emergence of science in a culture.

Logical necessity -- a state of consciousness.

It is no coincidence that critics of TGE who tend to disregard the verbal justification criteria for logical structures are also likely to completely misconstrue the notion of logical necessity. One such critic (Brainerd, 1973, p.175) in the course of dismissing Smedslund's arguments relating verbal justification to necessity remarks:

Many Piagetian tasks involve only approximate physical properties of objects about which there is no logical necessity at all. This is especially true of conservation problems...whose "logic" is disproven in modern relativity theory.

Necessity, of course, has absolutely nothing to do with physical properties of the world or physics other than as a construct in a theory, TGE, that seeks to explain the development of physics as a product of the human mind. Logical necessity in TGE is a state of mind or consciousness that emerges into awareness upon construction of logical structures in the operational child. Piaget (1973, p.26) says:

...following the formation of a structure one observes on its completion some modifications in the subject's behavior which are difficult to explain otherwise than by that completion itself, in other words by the "closure" of the structure. These are fundamental facts which are translated into consciousness of the subject by feelings of obligation or of normative necessity and in his behavior by obedience to rules...The feeling of logical necessity difficult to evaluate like all states of consciousness, will be translated in behavior by the use and recognition of transitivity.²

2. In this connection, recall Braine's (1959) observation that while many children solved his problem, his older children (6-7 years) displayed a different attitude or "set" than did the younger children who solved the problem (4-5 years). Braine says of the younger group: they seemed to lack the attitude of the older Ss expressed in the words, "I know one's bigger, but they look the same." Braine notes others have characterized abstract thinking as an attitude rather than as an ability.

Before proceeding, it is imperative for us to clear away a basic confusion related to the misunderstanding of necessity. The basic experimental paradigms in TGE designed to assess logical operations are surely not meant to determine if the child can perceive minute differences in length or amount. They are designed to ascertain whether or not a child can take certain premises and then logically derive a conclusion--the judgment response. For the most part the truth or falsity of the premises and/or the conclusion is a secondary consideration (cf. Flavell, 1963, p.160). Indeed critics (e.g. Hall and Kingsley, 1968; Brainerd, 1973) are quite right--but irrelevantly so--when they point out that rarely if ever are the amounts precisely the same after a transformation in a conservation task. Surely one does not want to be in the position of accusing a physicist who says "not same" in a conservation paradigm of utilizing pre-logical thought structures while the confident operational child who boldly asserts "of course they're the same", oblivious to effects of sublimation and surface tension is declared to be using logical thought structures. We do not anticipate anyone working with subjects drawn from a mixed population of young children and physicists. The point is that the use of judgments in themselves do not permit a theoretically motivated alternative interpretation of the same judgment or the same interpretation of alternative judgments. This is not the case within TGE where the mode of thought, not the outcome, is the focus of interest.

To summarize, logical necessity does not refer to states in the physical world; it refers instead to a conscious attitude that develops in the mind of an operational child as a function of the emergence of certain logical operations internally and certain external, normative factors. It still remains however to explicate the empirical consequences of this attitude in the behavior of the child.

Introspection and consciousness. The Piagetian answer to this perplexing problem has from the beginning relied on a child's verbal explanation of his judgment response. Our reconstruction of the role of language in TGE, while vague in places, is essentially as follows. (1) Developing mental structures are "translated" into consciousness as described above. (2) Via language one has access to consciousness although introspective reports are by no means veridical reflections of ongoing causal mental processes. They can however reflect the logical attitude or feelings of necessity engendered by the emergence of logical operations at the causal level. (Piaget often speaks of the logical and causal levels proceeding in parallel, e.g. 1973.)

While language cannot explain thought, language becomes increasingly necessary as thought becomes more abstract. Piaget (1967; p.98) says specifically that language is a necessary but not sufficient condition for the construction of logical operations.

It is necessary because without the system of symbolic expression which constitutes language, the operations would remain at the stage of successive actions without ever being integrated into simultaneous systems or simultaneously encompassing a set of interdependent transformations. Without language the operations would remain personal and would consequently not be regulated by interpersonal exchange and cooperation.

In brief, language serves to symbolically condense or generalize operations and to socialize thought. As a result of this socialization, knowledge becomes explicit and communicable in contrast to intuitive, tacit, or personal knowledge; in large part this is what is implied by describing both science and formal thought as formal.

Recall that a primary objective of the theory is to account for the achievement of the level of formal thought in certain (but not all) cultures. Formal thinking is characterized as analytic, hypothetico-deductive, propositional, and of course formal. It is no accident that these terms also apply to generally accepted characterizations of science. Nor is it, we believe, a coincidence that these descriptions are also commonly ascribed to the functions of the dominant (generally left) cerebral hemisphere--the locus of propositional language (and by some accounts consciousness itself, Jaynes, 1973).

Time does not permit further exploration of these ideas here. However we would like to conclude with a summary of several basic points.

1. TGE is the most comprehensive theory of its kind. It encompasses phenomena that need to be related within one theory but generally are not.
2. Its most troublesome aspect is a certain vagueness and richness of structures permitting it to assimilate all phenomena to it, rendering it impervious to refutation.
3. While we are sympathetic to the problems in using verbal criteria as indices of mental structures, these are pragmatic problems, e.g. a concern with reliability and experimenter induced demand characteristics rather than with the validity per se of verbal justification. The use of judgments alone is certainly insufficient.
4. Non verbal means of investigating logical structures should be encouraged, e.g. Smedslund (1969) and particularly Miller (1973). Braine's (1959) observation should be followed up. In addition we raise the possibility that the shift from pre-logical to logical thought may be related to a shift in relative reliance on processes characteristic of the (generally) right and left hemispheres, respectively.

References

- Brainerd, C. J.. Judgements and explanations as criteria for the presence of cognitive structures. Psychological Bulletin, 1973, 79, 172-179.
- Bryant, P., and T. Trabasso, Transitive inference and memory in young children. Nature, 232, 1971.
- Chiseri, M. J. The identity-equivalence Decolage in Piagets Conservation Problems: Paradigms and Implications for the Theory of Genetic Epistemology. Dissertation in Preparation, University of New Hampshire.
- Flavell, J. H. The Developmental Psychology of Jean Piaget, Princeton, N.J.: Van Nostrand, 1963.
- Fodor, J. Psychological Explanation. New York: Random House, 1968.
- Jaynes, J. The origin of consciousness in D. Krech and H. Levin (eds.), A Symposium in Honor of Robert MacLeod, Ithaca, privately printed. 1973.
- Kaplan, A. The Conduct of Inquiry, San Francisco: Chandler Publishing Company, 1964.
- Miller, S. A. Contradiction, surprise and cognitive change: the effects of disconfirmation of belief on conservers and nonconservers. Journal of Experimental Child Psychology, 1973, 15, 47-67.
- Piaget, J. Six Psychological Studies. New York: Random House, 1967.
- Piaget, J. Genetic Epistemology, New York: Columbia University Press, 1970a.
- Piaget, J. Structuralism. New York: Basic Books, 1970b.
- Piaget, J. Main trends in inter-disciplinary research. New York: Harper Torchbooks, 1973.
- Pinard, A., and M. Laurendeau, "Stage" in Piaget's Cognitive-Developmental Theory: Exegesis of a concept. in D. Elkind and J. H. Flavell(eds.) Studies in Cognitive Development, New York: Oxford University Press, 1969.
- Smedslund, J. Psychological Diagnostics, Psychological Bulletin, 1969, 71, 234-248.

THE STAGE BEYOND FORMAL OPERATIONS

John M. Kennedy

University of Toronto

It is the proposal of this paper that there may be a "stage", in the customary Piagetian sense of "stage", beyond the formal operations stage that is usually understood by Piagetians to be the capstone of intellectual development. It is the purpose of this paper to outline characteristics of this proposed new stage, and to show a basis for calling this proposed new stage a continuation of the kind of intellectual development Piaget has described.

Criteria for a Piagetian stage

Merely showing empirically that new responses are learned after formal operations settle stably in the intellectual functioning of the growing person is not a justification for saying there is a new stage beyond Piaget's formal operations. That is, new skills may be learned by an adult who developed formal operations as an adolescent. New social and emotional skills, on-the-job skills, artistic and educational skills might be learned, but none of these skills need be taken as further progress along the route marked by sensori-motor, preconceptual, concrete operational and formal operational thought. To show that a new set of skills is a part of the Piagetian course of development one must define the directions taken by that route, the modus-operandi and modus-vivendi captured by Piagetian theory. Then one must show that any proposed further stage fits the direction followed in the Piagetian map of the child's development. It should not merely be an amplification of a previous stage or a refinement of the previous stage. It should introduce new problems for intellectual growth, while following the pattern of growth evident across all the previous stages. Accordingly, this theoretical essay will attempt to extract from the Piagetian corpus essential features of the Piagetian child's intellectual development, and show what to expect in any further development along the same lines.

Since this is a theoretical paper, no presumption will be made that Piaget's theory is right or wrong. Piagetian theory will be considered "in abstract", the aim being to determine whether the set of Piagetian postulates inexorably leads to a stage beyond formal operations. While hypothetical examples and empirically-obtained examples of the types of thinking found in research in Switzerland, Britain and Canada will be offered in this paper, these are intended as illustrations of possible stages, not as proofs that large-scale coherent stages exist.

To document the appropriate references that have been influential in the genesis of this paper, line by line, would be very repetitious. It is hoped it may be acceptable to record that the major Geneva works included several by Piaget alone (1950, 1970a, 1970b, 1970c, 1971), ones by Piaget and Inhelder (1958, 1969), and among commentators Elkind (1970) was particularly useful on egocentrism, Wallach (1969) on conservation, and Langer (1969) on equilibrium. The general text by Flavell (1963) and the remarkably good brief introduction by Ginsburg and Oppen (1969) were also helpful.

Piagetian Theory

The theory developed by Piaget includes a set of stages, called the sensori-motor, preconceptual, concrete operational, and formal operational stages. At times, Piaget describes these stages in highly abstract terms, using nomenclature borrowed from modern mathematics. His purpose is to argue in explicit terms that the operations and activities referred to by this abstract and formidable terminology are interrelated in certain precise ways - for example, "lattice" or "group" relationships are sometimes actualized in the child's thinking. In particular, problems of "conservation" and "decentration" (notably in the form of "egocentrism") arise on each stage, and interrelated solutions to the problems of conservation and egocentrism, on each stage, are essential to the child's progress from one stage to the next.

Piaget has a dynamic for development as well as a description of development's stages. Thus, development results from "disequilibrium", a conflict between two conceptions, which may be a conflict between the child's deductions and his observations, or a conflict between two deductions, or a conflict between two beliefs. Disequilibrium sets the stage for development, and if the child progresses from disequilibrium to equilibrium then he moves upwards on the Piaget scale of development. Equilibrium is present when deductions are in harmony with each other, when observations are in harmony with each other, when action and observations are in harmony, etc.

To show that a proposed new stage is a part of Piagetian progress it is necessary that it involve problems not captured in the foregoing list of stages, and it should involve inter-relationships between its problems. These inter-relationships need not necessarily involve "group" or "lattice" structures, but at least they should involve a set of relationships that bring issues like conservation and egocentrism into question, and full development in the stage should result in inter-related solutions to the problems of conservation and egocentrism. Also, the genesis of the stage should involve disequilibrium and full adoption of the stage's characteristics should result in equilibrium. Thus

major criteria for a Piagetian stage of development include (a) solutions to problems of egocentrism (b) solutions to problems of conservation (c) the solutions (a) and (b) should be inter-related and (d) the stage should be marked by the acquisition of a new equilibrium.

A cautionary note is in order at this point. One should not expect exact parallels between the patterns of intellectual problems within particular stages. Thus, while sensori-motor intelligence is supposed to fall into six sub-stages (roughly (1) reflex use (2) primary circular reaction (3) secondary circular reaction (4) coordination of schema (5) tertiary circular reaction and (6) invention of means) one need not expect each stage to fall neatly into equivalent sub-stages. With this disclaimer noted, the analysis may proceed, using for its basis the major criteria espoused above.

To continue the analysis, it is now necessary to ask about the characteristics of development that link one stage to the next. In this vein, returning to the set of stages outlined by Piaget, it is important to look beyond the particular stages, to both the relationships between the stages and the limitations inherent in each of the stages. In particular, what domain is explored in any one stage, and what then occurs as the next stage dawns?

The sensori-motor stage is important for the exploration of the domain of action. The child is mastering the organization of action on objects - learning to distinguish actions from objects, learning to detach an action learned in commerce with one object and apply it to new objects, and learning to fit actions together (to repeat an action, or to follow an action with another action in a sequence never before used). Ultimately, the child learns to invent actions. The limitations on sensori-motor intelligence expressly include an inability to represent or image an action and manipulate the internal representation of the action to solve a problem. While forms of imitation do appear quite early, these cannot be manipulated by super-ordinate intellectual operations to solve problems. As the child learns to generate and manipulate an internal representation, he passes from the purely sensori-motor to the next stage on the Piagetian chain, the preconceptual phase.

The cross-over from the sensori-motor stage to the preconceptual stage is suggested by an observation on Lucienne. Piaget puts a watch chain inside an empty box (Piaget, 1936; English translation, 1952). There is a very small slit at the opening of the box. Lucienne looks attentively at the slit and then, several times in succession, opens and closes her mouth, increasingly wider and wider each time. Then, unhesitatingly, Lucienne inserts her finger in the slit and pulls at the box to enlarge the opening, allowing her access to the chain. Then she grasps the chain.

In this observation, Lucienne seemed to represent the opening by a part of her body, and the representation was used to investigate the problem of obtaining access to the chain. The beginning of an ability to work with problems on an representational level of thought is indicated.

In the preconceptual phase, the child is fully capable of imagining invisible actions, and combining the imagined or internally represented events in new ways. He can explore his mental symbols and can invent mental symbols. However, he is typically unable to follow more than one dimension at a time; thus, he cannot successfully compare "yellow flowers" and "all flowers". He is unable to engage in intellectual comparisons that require comparing two dimensions, one of which is a reciprocal of the other, that is, a comparison wherein one dimension alters to compensate for change in the other. He lacks skills needed for comparing states and events; the child cannot work on or check his logic - he cannot assert logical necessity, only empirical necessity. There is no assurance that he will ever be able to retrace logical steps, e.g. to return from a conclusion to the original evidence.

The next stage beyond the preconceptual child, is the age of concrete operations, a stage when logical relations between concepts are being mastered. In this age, the child can assert a kind of necessity beyond empirical necessity. He can charge that some things are logically entailed. He is capable of manipulating the logical relationships between concepts. He can assert that change in one dimension compensates for change in another dimension, and so he can predict unobserved events, deduce conclusions, and indicate how to return to original evidence or states by retracing a sequence, or by using another sequence, never seen before. The child can check his logic and analyze intellectual procedures to see whether they are correct. However, he cannot organize a total set of dimensions into one system, and therefore he cannot proceed systematically through all possible combinations of the variables in a problem. The more he learns to organize combinations of variables the closer he comes to the next stage, where problems are conceptualized in terms of all possible combinations of variables.

In the next stage - formal operations - the child arranges a fully systematic process for coping with the dimensions and variables of a problem. Presented with a problem like determining the factors involved in a pendulum's action the subject can arrange a testing procedure that will examine all the factors in a comprehensive scheme. Once observations have been made, the subject can isolate the relevant factors. Once all possibilities have been conceptualized and the relevant factors have been empirically assessed, the subject can deduce what will happen in other possible (i.e. theoretical or hypothetical) circumstances. Thus, having determined that the roughness of the surface and other friction forces control the movement

of a rolling ball, the subject can deduce what would happen in the hypothetical circumstance where all the friction forces are reduced to zero - the ball would roll indefinitely, an observation one may not be able to make in practice.

It is evident from the preceding description of stages that each successive stage overcomes limitations on the preceding stage. But that is not the whole story, or indeed the key to the story. If the growth of the child is put succinctly, the pattern in the growth becomes more evident, as follows: The sensori-motor stage is the stage of action. In the next phase, the preconceptual child can operate on a representational level -- representing actions. In the concrete operational stage, the child can be logical about the relationships between his imaginal representations. Finally, the formal child can organize the logical relationships he clarified as a concrete child, and as a consequence he can deal with the hypothetical.

It is important to note that the actions of the first stage are the alimant or material or substrate of the second stage. Also, the representations of the second stage are the materials inter-related systematically in the succeeding stage. The individual logical inter-relationships used by the concrete child are re-organized into a comprehensive structure at the formal stage. Action, representation, logic and the hypothetical - what should the succeeding stage be? The substrate of the next stage should be what was developed in the formal operations stage. If the material of the formal operations stages is to be organized afresh in the next stage, via "reflective abstraction" (Piaget, 1971), what will the processes of that stage be concerned with? Is there anything above and beyond what is actual and what is hypothetical - the real and the imaginary - the distinctions made by the formal operational child?

The formal operational stage is concerned with determining the rules for exploring causal variables; it should follow that the succeeding stage is concerned with examination of those rules and criteria for using those rules. Formal operations clarify the rules of explanation and verification, the laws governing seeking, finding, and applying explanations. As formal operations are the ground rules of empirical research, it is the rules governing these ground rules - a meta-theory, about those rules - that should develop in the succeeding stage. That is, the next stage should uncover procedures for determining what is and what is not a question that can be explored and answered. In place of the hypothetical, the post-formal-operations stage is a stage of meta-theory, in which the impossible or anomalous or insoluble becomes of concern.

Where the sensori-motor stage dealt with action and object, the preconceptual stage dealt with representations. The representations of the preconceptual stage are separated into "logically consistent" and "logically inconsistent" at the next stage. The representations of the concrete operational stage are reorganized at the formal operational stage into actual and hypothetical classes. In a meta-theoretical stage, presumably these classes are further subdivided into empirical or syntactical-concerned with physical variables or concerned with rules guiding description and exploration of theoretical statements. Can a distinction between empirical and syntactical matters form the kernel of a Piagetian stage? Perhaps this question can be examined in the context of problems of egocentrism and conservation.

Egocentrism and Conservation

The problem of egocentrism for the sensori-motor child was to distinguish his own actions from those of objects. His experiences had to be distinguished into those pertaining to his body, and those pertaining to his environment. Egocentrism for the preconceptual child involved distinguishing his own symbols or images or internal representations (or dreams) from public symbols, or socially-agreed conventions of reference. As it were, the child does not distinguish his symbols into socially-transmitted and privately, individually invented. Egocentrism for the concrete operational child involves the growth of a point of view. He realizes that some of his symbols are his own idiosyncratic inventions, and he also slowly comes to the realization that he can take many points of view, and at the end of the growth of concrete operations, he realizes that he is the only person at a particular physical viewpoint at a particular time, and he can successfully predict what is visible from another observer's point of view. Further, as concrete operations give way to formal operations, he realizes when his opinions have been tested, and when they are as yet untested - that is, he distinguishes his conception of what is possible from his concept of what is fact. He no longer considers his idiosyncratic assumptions to be facts, and he will only report an event supporting his beliefs if he actually did make the observation.

In a meta-theoretical stage his assumptions would become divided into those that are testable against fact, and those which are principles for conduct. The conduct may be moral, flowing from ethical principles. The conduct may be investigative, flowing from a decision that a statement is testable. Or the conduct may be "formal" or "syntactical" or "autoclitic" in the sense of amending and polishing and manipulating concepts and logic in a theory, to make it more parsimonious, elegant, clearer, or internally consistent, and, of course, none of these features of a theory are directly related to empirical verification of propositions of the theory. Consequently, the meta-theoretical stage offers the child the opportunity to question questions, to ask whether questions can be answered empirically.

The problems of metatheory make it necessary to organize questions (not physical dimensions) systematically, into different kinds like "circular" or "unanswerable" or "leading toward infinite regress." Where some queries like "what is the difference between a duck?" are merely ungrammatical and are probably recognized as such quite early in development, questions, like "are you real?" or "do you exist?" are more sophisticated and are grammatically acceptable. They may be deemed unanswerable because they violate the rules of a metatheory about empirical statements, not simply a theory about statements - in-general, ie. the grammar about all types of sentences.

The particular metatheory involved here can be described in alternative terms and related to egocentrism, as follows. In any metatheory, "levels" of different questions are distinguished. In the particular case that is of concern here, one level is concerned with empirical variables one can experience. Another level is the person who is investigating or experiencing. The levels are confounded (in different ways) by asserting things like "I always tell lies" (which requires that the statement itself is a lie) or "I am experiencing Nirvana, ie. nothing" or "I feel strangely as though I were dead" (which means that there is an experience of not-experiencing). For a more complex example, consider some issues centering on the concept of "free will".

It is one thing to assert that none of the organisms (including people) studied by an investigator has "free will". It is quite another thing for the investigator to assert that he, the person who studies the organisms, does not have free will. For some purposes, it might be useful - and consistent in a metatheory - to assert that he, the experimenter, freely follows verification rules, while determinism holds for the subject of the investigation. In particular, the metatheory might specifically indicate that there is no necessary conflict between the experimenter's experience that he makes free choices and his belief that his object of study is deterministic.

The belief can be an a priori assumption or a deduction from his evidence, but in neither case would it be necessary for him to conclude that he does not genuinely experience "freedom of choice" and therefore "free will" is not real. He can treat his own experience as one kind of evidence, and the results of experiments as another kind of evidence and simply assert that the two need not be considered as possible sources of contradiction. In this sense, he can legitimately and consistently want people to "feel free" while believing that individuals are governed by deterministic laws. (This solution to the free-will/determinism controversy may be a little weak, but our point here is simply that a decision about what sorts of evidence are useful to what issues can be taken at a metatheoretical level, to avoid some possible conflicts in evidence). If the experimenter offers these distinctions he is operating with metatheoretical

skills. Similarly, "the sentence I am uttering is false" might be described by a subject as invoking problems of reference in which useful distinctions between the sentences and its referents are not maintained. If this line of argument is invoked, the subject is acting on a metatheoretical stage.

The capacity to treat some facts or claims or even other people as objects of empirical analysis, and to hold oneself and one's theory in a different logical status is perhaps one of the key products of a post-formal-operations methodology for thinking. Notice that one is not simply asserting "I am different", which may simply mean I am more intelligent, or smaller, more careful or systematic, etc. The assertion is about logical status, not merely about different empirically-testable characteristics. The claim is about logical egocentrism not a political or social value-laden egocentrism. The claim is simply that questions like "Are Gandhi, Batman and Big Bird all alive today?" are different from questions like "Am I alive today?" It is perfectly acceptable for someone else to question whether you live and breathe, or whether you are a figment of someone's imagination, so some questions about "your reality" are sound, and empirical. The problem for the growing organism is to realize that some ways he proceeds to test the reality of his heroes can be applied by his friends to him, but he cannot logically apply them to himself. Thus the present analysis makes the purely logical suggestion that the formal operational child knows how to investigate - to use some methodology - but he is not clearly aware of the differences between questions that can and cannot be answered with his methodology, particularly questions about himself.

In this light consider the problem "prove that you exist": The formal operational subject would want to isolate all the relevant variables and obtain data from which to draw inferences. Consider how one Canadian undergraduate subject responded.

31. Existence is based on measurement. Therefore if I can be measured, I exist (scientifically speaking). I have height, temperature, etc. and therefore "exist".

The empirical bent is shown in this response. Notice how it avoids or begs the difficulty of identifying who is doing the measuring.

Given this capacity to suggest hypothetical circumstances, the formal operational child would be able to entertain the hypothetical idea that he did not exist. Thus, in part of one response is found:

S2. Carry out various experiments using the senses as apparatus to discover and test the place in space man occupies ... compile information and obtain results which would either validate or falsify the proposition.

The last phrase shows the thorough going consistency of the true sceptic, willing to put his own existence in doubt! Although consistent, and empirical, the response is limited in important ways. It may be more sophisticated than a flat undefended assertion like this problem is "funny" or "odd" or "meaningless". Yet its sophistication only opens the subject to troublesome predicaments that he cannot reject out of hand, like a less sophisticated subject. The sceptic's line of debate can echo and re-echo "prove it" even down to the level of questions like "prove you exist" and can entangle this kind of subject who treats all charges and questions as alike. In principle the formal operational child would not be able to respond to "Do you exist?" by taking the systematic view that the question is not open to empirical exploration in the same sense that "Does friction retard the movement of a rolling ball?" is open to investigation, proof and disproof. That is, the formal operational subject would not be able to aver that problems like "prove you exist" introduce special problems of definition and circularity and the logic of verification so that empirical methods are often irrelevant. In contrast, if a hypothetical subject responded as follows, his answer would have post formal operations characteristics:

S3. (Hypothetical): Either the question about my existence can be answered in ways that make reference to experience per se, or it does not. If it is to be answered with a reference to an experience, then the question presupposes an experiencing organism. Thus, it may be necessary to explore any one or any set of variables. Showing that a capacity for experience is assumed will suffice. If the question does not entail a reference to experience, then it cannot be answered at all. It would be, by definition of the relation between questions and experience, an insoluble problem. In like vein "prove that you are now dreaming" introduces logical conundrums.

As one undergraduate subject put it:

S4. I am not sure that this proposition can be evaluated successfully. Such an evaluation would deal with the assumed concepts of Dualism - the existence of both a spirit and a body entity. One must assume one or both of these entities exists.

Another put it, succinctly:

S5. What is existence? One must exist in order to prove one exists.

The problems of egocentrism involve failing to distinguish one's own place in the methodology for developing and representing knowledge. As a younger child thinks that his private symbols are public symbols - in the sense that he cannot distinguish the two - so a formal **operational** child has problems of egocentrism, because he treats certain kinds of insoluble and empirical questions as all of one type, missing the fact that some questions are merely questions of definition and he has a particular status vis-a-vis some questions. He egocentrically fails to distinguish his capacity for experience (as one level of analysis, or one type of evidence) from an object of study - which is just a part of his experience, or himself as viewed or considered by another person.

At this point in the analysis, a connection between egocentric problems and a conservation problem can be stated. Where above, the discussion centered on the Ego who asked questions, it is now necessary to focus on the parts of the questions that are asked, the terms of a discussion.

The non-metatheoretical subject may confuse levels used in analysis. This may have the result that a term introduced at one level or in one way is treated similarly to terms belonging to other levels of analysis or introduced in other ways. In other words, the non-metatheoretical subject would find it difficult to retain the fact that some terms are purely definitions and he will treat a term-by-definition like a term-with-empirically-determined-reference in the course of an argument or exploration. In other words, in the course of a train of thought, he cannot conserve the fact that some terms are introduced into the argument purely as definitions to aid the discussion. Where a younger child confuses fact and assumption - reality and imagined possibility - the older child may confuse definition and empirical finding.

Thus, if the above analysis is sound, the metatheoretical stage is concerned with ontology (levels of existence and experience), fiats (or terms established by definition, and, presumably, cognate rules governing internal consistency in the manufacture of theory) and verification procedures (rules governing the points in a discussion where one can and cannot introduce empirical exploration).

The preceding analysis of a possible link between conservation issues and post-formal problems can be most clearly seen by way of examples. The general point is straightforward. The post-formal operations stage is one where the subject recognizes how definitions offered to permit terms to have utility in a discussion will necessarily lead to statements which are true, by definition of the terms used but which create a misleading impression that something empirical has been established. To take a hypothetical example:

S6. (Hypothetical) "The child beatings of today's society are due to the fact that infants do not have strong negative sanctions to apply to caretakers."

Is S6 making a falsifiable empirical claim? He will not be if one way to assess whether children have strong negative sanctions is to determine whether adult caretakers beat them. A commentator (who may be S6 himself) who fails to conserve the way in which the term "sanction" was introduced by S6 may be misled into thinking S6's assertion is a new piece of factual information and not a definition's consequences.

The metatheoretical subject recognizes how different definitions for the same term will lead to superficially different or even contradictory statements, which because of the different root definitions are in fact compatible. This is the implication of the capacity to hold on to - to conserve - the nature of definitions of terms, including different definitions or points of view bearing the same term. For example, the claim "science fiction is a glimpse of tomorrow" can be empirical or not empirical as follows: It is empirical if science fiction is defined as fiction whose kernel idea is an ostensibly or avowedly scientific hypothesis. It is not empirical if it merely implies that the way to determine whether a piece of fiction is science fiction is to see whether it ostensibly or avowedly is an attempt at describing what a future world might be like. Notice three different levels of sophistication shown in excerpts from evaluations of this assertion by Canadian undergraduate subjects.

S7: Science fiction is as, its name implies, purely speculation and imagination with no basis in fact.

This statement is at the concrete operations level, in which assertions are often made without regard for their basis in observable facts. The statement preceding this one in S7's protocol was: "no one can say what is going to happen or what is not even in the slightest", a clearly counter-factual statement, if predictions about simple cyclical events like nightfall and sunrise are taken into common consideration. Taken as a whole, this subject's response indicates an opinionated form of thinking that avoids contact with an empirical ground.

S8. Consider several pieces of science fiction in great detail and decide whether (the) predictions are plausible or just someone's imagination at work. This might have to be done to many pieces of science fiction before a reliable conclusion could be made.

This response is more subtle than the previous one, in that the claim is related to a mode of testing, not flatly asserted. However, the testing procedure involves merely confronting the material, and making judgements in which the subject and the materials are the only items in question. The idea of introducing and considering or obtaining evidence about the referent of the predictions - the predicted events - is absent. The response is rendered somewhat sophisticated by a clear recognition of the problem of selecting an adequate quantity of material to be assessed. Nevertheless, a much better response is the next one, S9.

S9. Read up some old science fiction books written some time ago and see if the things written then are of any truth today.

This respondent went on to consider whether if old science fiction made accurate predictions, contemporary science fiction can be presumed to be predictive. The response is clearly empirical, and entails an investigator and two things to be related - contemporary events and old science fiction. Thus the subject is as empirical as S8, but more soundly so. In a class of 50 undergraduates given this question, an S9 type answer predominated. Yet perhaps it fails to grasp an important nettle. How are science fiction and the future related? Consider:

S10. In my opinion science fiction is a person's imagination presented for the enjoyment of others. These works could in fact be glimpses of tomorrow because it is people's ideas or imaginations that create the new ideas which make today different from yesterday.

In this answer, there is an initial care over definition, and a consideration of the necessary relationship between what has been defined and potential observations. The subject does not point out the particular difficulty mentioned above by the present author, that recognizing what is science fiction may entail noticing what is offered as predictive fiction, but his discussion has a comparable form in that important relationships between imagination and the future are involved in the assertion. Thus, the subject shows aspects of post-formal-operations thinking; in particular he showed how his way of defining a term would inevitably lead him towards other statements which could partly be outcomes of definitions and not completely fresh, testable and empirically verifiable assertions.

Consider a useful case from physics that shows how verification is at least in part a metatheoretical problem. At times, light has been considered as acting like waves. At other times, light has been considered as acting like particles. That is, light is sometimes X and sometimes Y, where X and Y are inconsistent. Inconsistencies in a theory are usually considered indications that the theory is faulty. Arduous attempts have failed to remove this inconsistency from the theory of light. Accordingly, it has been decided to allow the inconsistency to stand, and assert that in this case the inconsistency is acceptable. Thus, the rules governing the creation of a theory have been altered in this special case, pointing up the fact that these rules are arbitrary. Physicists prefer internal consistency but will accept inconsistency if need be.

Another example that throws light on verification and the aims of a theory can be drawn from Greek philosophy. As Baker (1947) puts it, Plato once defined man as a "two legged animal without feathers". Diogenes quickly appeared at the Academy with a plucked rooster and announced: "This is Plato's man". As a result Plato's definition was amended by addition of the phrase "with broad flat nails."

The story shows that enormous assumptions are made tacitly when a "featherless biped" description is offered. For example, this kind of definition is useful when one has to identify one class of animals from another in a natural or circumscribed setting. Definitions are usually useless when the setting is not clearly circumscribed, and Diogenes violated Plato's unspoken assumption that artificial intervention was not permissible. Diogenes could have repeated his disproof by taking a sledgehammer to the claws of the rooster, and "artificially" engineering broad flat nails. The lesson is that the general question "What is man?" can often be answered as a relative question: "What is the difference between man and X in certain settings?" Metatheoretical thinking deals in the organized understanding of these kinds of assumptions underlying such superficially simple and apparently purely-empirical questions like "What is man?".

The post-formal operations or metatheory stage is one that describes not merely the act of getting to know, but establishes what can and cannot be known, on logical grounds. It is a form of thinking that is important to all sciences, philosophy, and it has social ramifications too. That is, where the formal-operational child may assert that "all men are created equal" or "justice is an all-important principle" or "life is sacred" the meta-theoretical child would understand the belief that these statements are unverifiable. They are, he might assert, guidelines or significant presumptions rather than empirical assertions. The formal operational child may want to check religious beliefs for their accuracy. In contrast, the metatheoretical child may offer the view that claims like "There is a God" are not of the empirical type.

He might argue that definitions of God, Love, Justice, Beauty, the Rights of Man and related terms are not subject to worldly exploration. The basis of his argument would be the contention that some terms can be assessed for their consistency, but they cannot be proven or rejected on empirical grounds. They are inherently untestable as empirical matters. The definitions are bases for a discussion, for criteria for labelling events as "ethical" or "just" for example, not testable claims.

Within the logical system of a metatheory, the growing person comes to understand what is empirical, and what is not. He understands what is circular, and what makes empirical reference. When an argument is offered he may assert that the argument is about a definition, not about a testable proposition. For example, if A argues vehemently that Justice entails responding with an "eye for an eye" morality, and B contradicts A, demanding that A accept a view that Justice must be tempered with mercy, C, the metatheoretician will note that A and B are arguing about the definition of the word Justice, and their personal preferences.

The crux of logical puzzles or "conundrums" can be discovered at this stage of development, if the metatheoretician retains a grasp on levels of description. Consider the following conundrum:

No horse has two tails.
 Every horse has one more tail than no horse.
 Therefore, every horse has three tails.

A very interesting response - one that is critical for the present theory, one that has been given by several subjects in pilot testing - is the suggestion that ~~one should~~ take a sample of horses and count their tails. This is the epitome of the purely empirical response. Indeed, it is unduly empirical and it would be particularly unhelpful if it had happened that several pieces of totally inconsistent logic had stumbled onto a correct conclusion, as in:

$2 + 2 = \text{a rich man} = \text{a "square"} = \text{four-sided} = 4$

Therefore $2 + 2 = 4$

In the horse conundrum the effect depends on a play on levels of description, where "no horse" is made first a term in parallel to empirical horses, and then later, subordinate to empirical horses, as though it were one kind of empirical horse.

Hudson (1966) found one extremely high I.Q. highschool boy who responded to the horse conundrum:

S11. "This is an interesting point. Why shouldn't a horse have three tails. Live and let live. If horses had three tails they would not need to use one another's tails, which all goes to prove that the horse, not the dog, is a man's best friend. He can lend us a tail in time of need."

Piaget would probably say S11 was engaging in pure "assimilation", playing with the conundrum, and Piaget might say that this kind of assimilation would most probably occur when the conundrum presented no difficulty to the subject's intellectual structure. Accordingly, it is interesting that this boy indicated considerable sophistication on other questions.

This boy - Hudson calls him Tarry - used incisive meta theoretical thinking to respond to general statements, as follows:

Hudson: Truth, in matters of religion, is simply the opinion that has survived.

S12 (Tarry): Why bring in the bit about religion?

Hudson: Happy the man with an interest to pursue in his spare time.

S12 (Tarry): Therefore he has no spare time.

Tarry seized on the definition of "spare time" and played with it, just as he realized that "opinions that survive" is one definition of "truth".

The manner in which the metatheoretician can call a camouflaged assumption into the light is reflected in another Tarry comment.

Hudson: A man ought to read just as inclination leads him; for what he reads as a task will do him little good.

S12 (Tarry): Some men's inclination leads to things other than reading.

Probably every psychologist is familiar with his own metatheoretical knots, especially those that arise in teaching. One of the most notable is the heredity-environment issue! This issue, phrased in certain ways, may be empirically answerable. But phrased in some ways, it is most certainly not an empirical question. Thus, to ask which influence - heredity or environment - controlled the development of (say) Mao Tse Tung, may be to ask a nonsense question. There would be no Mao without heredity, and no Mao without an environment. Phrased differently - could one produce a Mao with a different heredity? - it is still not subject to empirical proof if one's model of heredity and environmental factors assumes as a matter of principle (or by default) that one can compensate for loss of one inherited factor by boosting appropriate tough artificial environmental factors (special education, medical operations, etc.).

Only when a heredity/environment issue is phrased in certain ways can it be answered empirically. For example, one can determine a variance on I.Q. tests for a population. Then one might determine a variance in heredity factors (with some assumptions about error factors) and determine a variance in environmental factors (again with some assumptions about errors). One might then ask whether the contribution of heredity to variance in I.Q. scores is stronger than the contribution of environmental factors. Does the variability in heredity factors co vary with I.Q. scores? Does the variability in environmental factors covary with I.Q. scores? These are empirical matters.

Equilibrium

Are there some disequilibrium conditions that are entailed by formal operations such that metatheoretical operations will result in equilibrium? The answer may be contained in some of the above examples and responses, notably the horse conundrum. In some responses, the role of empirical observations was inter-related, perhaps too closely, with the role of logic. Where all things about statements, including their logic, is assumed to be empirically testable, then an examination of an argument's conclusions is supposed to support or disconfirm the logic as much as anything else about the argument. This may put the cart before the horse in some crucial ways. That is, for some purposes it may be essential to assert that the only way to check the logic of an argument is by a further application of the operations espoused by the logic, not by an empirical test to determine whether the conclusion of the argument is correct. Otherwise, the logic would be deemed acceptable on some occasions, and unacceptable on other occasions, because a conclusion was untrue on some occasions and true on others. This, by the definition of disequilibrium, is an instance of a disequilibrium condition.

Once again, the analysis can be seen clearly in its application to examples.

Above, it was proposed that a hypothetical subject might confront an argument of the following kind:

$2 + 2 = \text{a rich man} = \text{a "square"} = \text{four-sided} = 4.$
Therefore $2 + 2 = 4$

The conclusion may be true, though the argument entails some unjustifiable links, illogical steps. In another instance, a comparable argument would lead to an untrue assertion:

2 + 2 = a rich man = a "square" = rectangular = made of right angles

Therefore 2 + 2 = made of right angles.

Should the logic be rejected if it leads to incorrect assertions and accepted if it leads to true assertions? The answer is clearly no, because it is the business of logic and logical deductions from a theory to determine whether a theory is correct or incorrect. It would be impossible to distinguish the theory from which deductions were drawn from the logic for making the deductions if both were to be tested with the same verification procedure. Therefore, it is essential that one set of tests determine whether assertions are true and another set of tests determine whether the logic is sound.

It is necessary to check the soundness of the logic in an argument before the fact that a deduction is untrue can be supposed to falsify a theory, and in this sense testing logic against the empirical correctness or incorrectness of a deduction, is putting the cart before the horse. Some physical examples may help to clarify this point. Consider, if A argues that rain falls from the sky, and therefore something is pulling the rain towards the ground, then the logic is sound and the assertion is true. If A asserts rain falls from the sky and therefore something is holding the sky up, his logic is unsound and his assertion is false. If A asserts rain falls from the sky and therefore something is pulling the rain up into the sky his logic is faulty, although (bearing extra-terrestrial matter in mind) his conclusion is true. If the above analysis is correct, then logic is self-justifying (like a definition) and a particular logical chain can be checked for its soundness only by reference to the rules of the logic that was supposed to have been used in creating the chain.

Hence, the metatheoretical subject who has attained equilibrium will not try to check the logic of an argument by reference to the truth and falsity of deductions. Hence too, the pre-metatheoretical subject will become involved in disequilibrium or inconsistency when he relies on empirical tests to check the logic of arguments, for the same logic will sometimes seem sound, sometimes unsound. The subject will be unable to distinguish a false theory (rain falls from the ground) from unsound logic.

Some of the significance of the empirically-oriented responses to the "no horse" conundrum is evident in the above analysis of disequilibrium. Consider some Canadian undergraduates responses to this conundrum.

S13. Logically the existence of a horse with any number of tails other than three will disprove this argument ... Observe horses and when (you) find one with any number of tails other than three (you) have disproved the proposition.

Notice in this response that "proposition" and "argument" are used interchangeably as though disproof of an assertion renders the argument unsound. If the subject was now confronted with the argument:

No horse has no tails.

Every horse has one more tail than no horse.

Therefore, every horse has one tail. The subject would be forced to conclude that the truth of the final assertion now justifies the logical connections between the sentences. He would then have contradicted himself, and be in disequilibrium through both rejecting and affirming the same logical chain.

It is useful to compare S13 to an inferior and a superior analysis, for a Piagetian would anticipate that the form of thinking shown here should have emerged from a set of less refined operations, and proceed, if the present paper is correct, towards a clearer separation of aspects of the puzzle.

S14. Great perceptual illusion! This effect is done by the last sentence. Nevertheless, dissecting it literally, it is stupid and doesn't make sense.

It is noteworthy that this subject probably was responding to the puzzle quite seriously, despite his exclamation. For example, immediately prior to the horse conundrum he was asked to evaluate the statement "Human nature is unchangeable" and he responded:

S14. Highly incorrect. Human nature changes with physical metabolism and societal environment ... together with other factors (these) assist in the change in human nature. Human nature is continuously changing.

And immediately after the horse conundrum S14 was asked to comment on the assertion that "Aeroplanes and sports cars - not paintings and music - are the twentieth century's work of art". His reply was:

S14. Refutable. Neither are 20th Century works of art, but (they are the) development of long engineered study of machines, etc. Yet the ideas of their shape and design are of 20th Century works of art. (sic)

Evidently, this subject is intelligent and thoughtful and intellectually alert - his scholastic performance at his college shows this, too - but his skills seem to be manifested in interesting opinions rather than an understanding of relationships between opinions and a data-base and logic. Both S13 and S14 contrast markedly with a superior subject, S15.

S15. In order to evaluate this (horse conundrum) the student must study the laws of logic. The student must decide whether "no horse" is a noun, making the conjecture valid, or if "no horse" is the negation of horse, making the statement invalid.

Here there is no **reference** to an empirical test, and a form of logic that would make the argument valid is described alongside a form that would be invalid.

Discussion

It is the claim of the above discussion that Piagetian developmental theory leads to a presumed final stage - formal operations - which, in principle, requires a further stage if cognitive equilibrium is to be maintained by the organism. This further stage, it has been argued, involves metatheoretical thinking, that is thinking about kinds of queries that might be investigated. To support the proposal that this metatheoretical stage is indeed Piagetian, it has been suggested that problems of egocentrism and conservation arise in the stage, from common roots, so that the solutions to egocentrism and conservation problems are inter-related.

Once again, it should be noted that the examples of thinking discussed above are offered by way of illustration, rather than by way of proof that a comprehensive and extensive stage of cognitive development does exist. It would be unwise to diagnose intellectual levels of subjects from a few responses to a specially selected set of questions. Nevertheless, the examples, taken together with the theory, do present prima facie indications that the problems to be solved in a metatheoretical fashion are important to scientific thinking, and are involved in a normal subject's evaluations of a wide range of propositions. Therefore, perhaps it may be reasonable to conjecture about some of the sub-stages that a subject may show when analyzing such propositions. That is, perhaps a consistent sequence which leads the subject to a metatheoretical level may be worth proposing.

Possible stage 1: Faced with a proposition, the subject indicates agreement or disagreement. Here, the subject shows no consistent method for testing propositions, no method that another person could use quite independently of the subject. Instead, the subject's own intuitions or opinions are the basis for evaluation. He will claim that a statement is nonsense, but not defend his view by reference to evidence gathered with an explicit method. He may simply answer a claim with a counter-claim that he suggests is the truth, while the first claim is suggested to be untrue. Asked to defend his assertion he may assert that he simply knows what is true, or repeat the essence of his statement as though he has defended his assertion. In a variation of this theme, the subject may assert that everything is a matter of opinion anyway, failing to offer the possibility of being shown to be correct or incorrect by procedures that he personally would accept as reasonable tests of his opinions.

Possible stage 2: This stage is characterized by the introduction of a method for settling disputes, but it fails to contain distinctions between disputes over definitions and disputes over fact. Where the previous stage held all beliefs to be opinions, this one holds all opinions, logic and definitions to be empirically testable. Presumably, there could be an in-between stage when the subject vacillates between the assertion that opinions are inherently not testable and the view that everything is testable, like a child evaluating quantities now by length, now by height, before settling on one criterion. An early part of this stage may be indicated by attempts to test propositions but with a procedure that only involves confronting the relevant matter, with no clear instructions about how to isolate the relevant characteristics of the material. In the later, fuller development of this stage, the subject would indicate precisely what features to select, and how to evaluate them and compare them. Yet there is no clear distinction between testing a theory for its validity and testing a logical argument for its soundness.

Possible stage 3: In this stage the subject notes how different conceptions of the key terms in a question will lead to different kinds of evaluations. He notes which parts of a theory are a priori terms, which putative operations are purely logical rules, and which operations are supposed to be referring to verifiable empirical processes. He shows how to test the logic of a deductive process, and how to test the truth of its conclusions, without confusing the two.

In sum, where Piaget posits a sequence of development coursing through action, internal representation, making logical connections and, then, understanding systematic verification, the present paper has proposed that this line of development requires a stage for examining the appropriateness of systematic verification - a metatheoretical stage.

References

- Baker, H. The image of man. Harvard, Boston, Mass., 1947.
- Elkind, D. Children and adolescents. Oxford, Toronto, 1970.
- Flavell, J.H. The developmental psychology of Jean Piaget. Van Nostrand: Princeton, N.J., 1963.
- Ginsburg, H. and Opper, S. Piaget's theory of intellectual development. Prentice-Hall: Englewood Cliffs, N.J., 1969.
- Inhelder, B. and Piaget, J. The growth of logical thinking from childhood to adolescence. trans. by A. Parsons and S. Milgram, Basic Books, New York, 1958.
- Langer, J. Disequilibrium as a source of development. In Mussen, P., Langer, J. and Covington, M. (Eds.). Trends and issues in developmental psychology, Holt, Rinehart, and Winston: New York, 1969.
- Piaget, J. The psychology of intelligence, trans. by M. Percy and D.E. Berlyne, Routledge and Kegan Paul: London, 1950.
- Piaget, J. Structuralism, trans. by C. Maschler, Basic Books, New York, 1970.
- Piaget, J. Genetic epistemology, trans. by F. Duckworth, Columbia: New York, 1970.
- Piaget, J. Science of education and the psychology of the child, trans. by D. Coltman, Orion: New York: 1970.
- Piaget, J. Psychology and epistemology, trans. by A. Rosin, Grossman, New York, 1971.
- Piaget, J., and Inhelder, B. The psychology of the child, trans. by H. Weaver, Basic Books, New York, 1969.
- Wallach, L. On the basis of conservation. In Elkind, D. and Flavell, J.H. (Eds.) Studies in Cognitive development, Oxford: New York, 1969.

HIGHER ORDER RULE CHARACTERIZATION OF HEURISTICS
FOR COMPASS AND STRAIGHT EDGE CONSTRUCTIONS IN GEOMETRY¹

Joseph M. Scandura, John H. Durnin,² and Wallace H. Wulfeck II
University of Pennsylvania

According to Polya (1962), perhaps the greatest value to be gained from the study of mathematics is the ability to solve problems. In spite of its importance, however, relatively little is known about how to teach people to solve problems, or how to program computers to do so. Specifically, one of the great mysteries of our time is why some problem solvers (human or computer) succeed on problems for which they have all of the necessary component skills (operators) whereas others fail.

In dealing with this question most research in AI has been concerned with the construction of powerful computer programs which can solve more or less diverse classes of complex problems. In computer simulation an attempt is made to also parallel human performance on such problems. In general, such systems (e.g., Newell & Simon, 1972; Minsky & Papert, 1972) have been comprehensive in scope; they have been concerned with problem definition (the construction of subgoals), memory, the derivation of solution procedures, and the use of such procedures.

The present research has adopted a somewhat different strategy. It seeks understanding by dealing separately with the various aspects of problem solving (e.g., the derivation of solution procedures). In particular, this research is concerned with the specification and testing of general, potentially useful heuristics for constructing procedures for solving compass and straightedge construction problems in geometry. The research also was concerned with developing and determining the feasibility of a general method by which heuristics may be identified in arbitrary problem domains.

One general point of departure was Polya's (1962) work on heuristics for geometry construction problems. These heuristics are purposely cast in a form designed to parallel human thought processes in much the same way as are such general heuristics as means-ends analysis (e.g., Newell & Simon, 1972). Human processing presumably is highly efficient in many situations, and the importance of paralleling human processing in AI, as well in computer simulation, has become increasingly well recognized as a means of significantly reducing processing time. Winston (1972), for example, has noted how constraining syntactic procedures to reflect underlying semantics in the recognition of block scenarios can drastically reduce the number of possibilities that must be considered.

In spite of the broad acclaim for Polya's work generally, however, and the intrinsic support for his notion of heuristics specifically, it sometimes has been difficult to capitalize on these ideas as fully as might be desired. Although often useful, his heuristics frequently are little more than general hints, and leave much to be desired insofar as pinpointing what a human or computer must know in order to solve specific kinds of problems. In order to lend themselves to technological treatment, heuristics must be transformed or incorporated into strictly mechanical procedures that can be more or less readily implemented on computers. Ideally, one might desire reduction of heuristics to algorithms; witness the alpha-beta "heuristic" (e.g., Nilsson, 1971).

Since heuristics tend to be (problem) domain specific, the potential value of more or less general and systematic methods for specifying heuristics in arbitrary problem domains seems fairly clear. Our approach to this problem was designed to be compatible with Scandura's (1973) theory of structural learning, and is an extension of a method used earlier by Ehrenpreis and Scandura (1972). That portion of the theory with which this research is most concerned has been shown empirically to reflect the behavior of individual subjects in particular situations where problem definition and

1. This research was supported by National Science Foundation Grant GW 6796 to the first author. An unabridged version of this paper is available from Joseph M. Scandura, Walnut Street, University of Pennsylvania, Philadelphia, Pa. 19174.

memory do not play an important role. Furthermore, this idealized theory has been shown with some empirical verification (Scandura, 1973a; Scandura, 1973b) to be extendable to situations involving memory and, apparently, also problem definition (Scandura, 1973a, p. 348) and perception (Ch. 5), without essential change. The structure of the theory must be enriched in these cases but without affecting its basic character (i.e., underlying behavior mechanism). This research is based directly on one part of the idealized theory, in particular that part which is concerned with competence--the specification of rule sets which account for classes of problems. In this theory, a rule set is said to account for a class of problems, roughly speaking, if for each problem in the class (1) there is a solution rule (operator) in the rule set which has the problem in its domain and whose range contains the solution to the problem or (2) there is a higher order rule in the rule set which applies to rules in the set and generates a solution rule. In such a rule set, higher order rules correspond to heuristics. (For a more general and formal formulation, which allows for any number of levels of derivation and in which the rules are not in a fixed hierarchy, see Scandura, 1973a, Ch. 5; 1973b.)

It seems unlikely, of course, that algorithmic methods can be found for devising nontrivial rule sets or heuristics. Indeed, as Chomsky (1968) has argued in the case of linguistics, no such method exists for dealing with observables as complex as language. Work in automatic programming, on the other hand, while it is quite far at present from a satisfactory solution, is proceeding as the authors³ understand it, on the assumption that significant progress in this direction can be made.

In the present research, the task of specifying heuristics is made simpler in at least two ways. First, and most important, the type of competence theory proposed imposes important constraints on the nature of allowable rule sets, and in turn on the form of the heuristics (higher order rules). In particular, higher order rules are assumed to operate on component (lower order) rules to generate integrated problem solution rules (procedures). These rules may simply compose component rules but may also modify them, for example, by generalization or restriction rules (Scandura, 1973a).

Second, restricting the level of analysis to that of flow diagrams, rather than computer programs, makes it natural to represent the constituent operations and decision making capabilities at whatever level seems to most adequately reflect human knowledge rather than at a level predetermined by some programming language. (We do not mean to minimize the importance of devising working programs. In fact, parts of this analysis have been implemented by one of the authors.) While no general assurance can be given with regard to any particular method, it would seem that a method which results in heuristics (and simple operators) that appear consistent with human thought would have a reasonable chance of having general value.

METHOD OF ANALYSIS

Our method of analysis went something as follows. First, we attempted to set some reasonably explicit bounds on the class of geometry construction problems to be considered. In particular, we considered only those problems in or like those of Chapter 1 of Polya (1962).

Our next step was to classify these problems on heuristic-intuitive grounds. Our aim was to place similar problems in the same categories, in accordance with the general form of their solutions. We were one step up in this regard, since Polya had already done part of the categorization for us. All of his problems can be solved according to some variant or combination of the three general heuristics he describes: (1) the pattern of two loci, (2) the pattern of similar figures, and (3) the pattern of auxiliary figures.

After the various tasks had been classified, we made sure that the domains and ranges of each task were fairly explicit. Then we identified explicit procedures for solving each type of task. Care was taken to insure that these procedures reflected our

3. The authors do not profess to be experts in AI or in computer simulation as such but rather in the adjacent and we think complementary domain of structural psychology which is, in our view, considerably broader than contemporary cognitive and information processing psychology.

intuitions as to how intelligent high school students might go about solving the problems. In some cases it was possible at this point to subclassify some of the tasks.

The most critical step was to identify general parallels among the procedures developed for the sampled problems within each of the various classifications, and even more important to devise higher order rules (operator combination methods) which realized these parallels as relatively formal, but still general, procedures. The higher order rules so identified (together with the component lower order rules on which they act) provided a general basis for constructing solution rules for the sampled problems.

Then we attempted to refine the resulting higher order rules with regard to specific sampled problems. This was done systematically; where a higher order rule failed to yield an adequate solution rule for a sampled problem, appropriate modifications in the higher order rule were made. A serious attempt also was made to insure that the higher order rules were compatible with human knowledge.⁴

PATTERN OF TWO LOCI

Our first step was to select a broad sampling of two-loci problems and to devise procedures for solving each. For example, consider the problem: "Given a line and a point not on the line, and a radius R , construct a circle of radius R which is tangent to the given line and which passes through the given point." This problem can be solved according to the following procedure: "Construct the locus of points at distance R from the given point; construct the locus of points at distance R from the given line; construct a circle using the intersection point of the two loci as center, and the distance R as radius."

This solution rule clearly involves the pattern of two loci. In this case, as with all of the problems in Polya's first category, the tasks may be characterized according to the form of their solution procedures: two loci are determined one after the other; the point of intersection of these loci in turn makes it possible to construct the goal figure.

Further analysis of the class of two-loci problems, however, revealed differences in the ways problems are solved. In many solution rules, for example, like the example above, the two loci can be found independently, in either order. Furthermore, at no point in the course of applying the solution rule is it necessary to measure a distance. Some form of distance measurement, however, is required with other tasks. Some of the sampled tasks require measurement in order to construct the goal figure; the solution rule for another problem involves measurement before the second locus can be found. In still another task, one of the loci is actually given, or equivalently, can be thought of as obtained by applying an identity rule. The goal figure in still another task is simply the point of intersection of the two loci.

An initial characterization

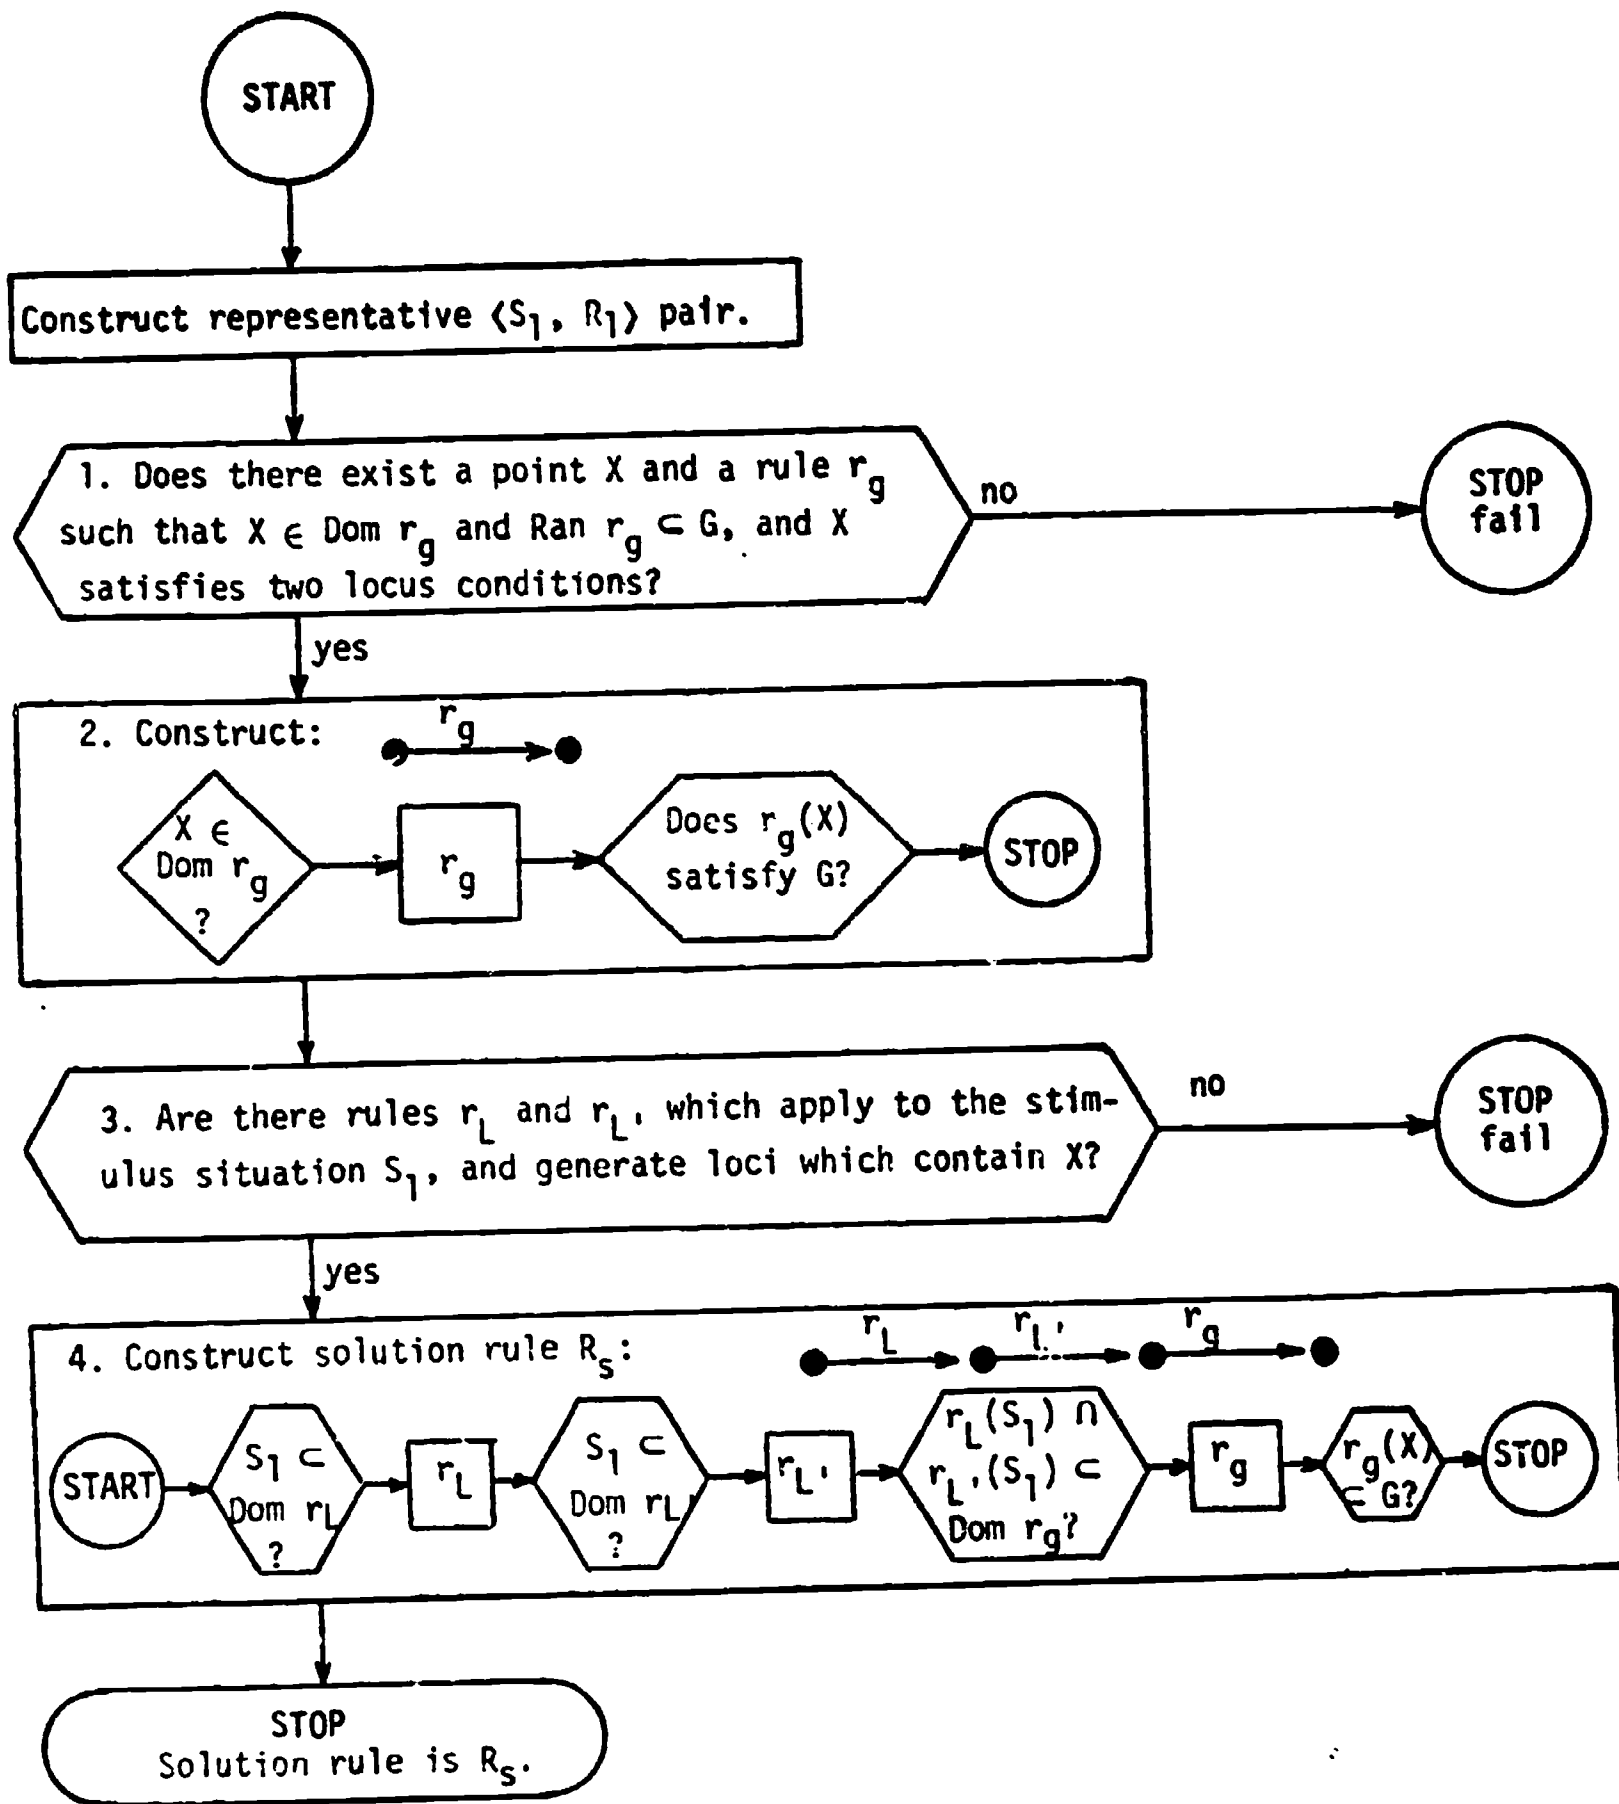
As a first step in characterizing a two loci higher order rule, we systematically went through the various solution rules for the pattern of two-loci tasks and identified all of the different component rules that appeared in our sample problems either (1) in constructing one of the loci, or (2) in constructing a goal figure. The lower order rules we identified were mostly common constructions (e.g., perpendicular bisector, circle, parallel line). Some of the lower order component rules were used to construct a needed locus, others were involved in constructing goal figures, and some served both functions.⁵

The higher order rule in Figure 1 shows schematically how the various solution rules may be constructed from the component rules.

4. See Appendix A.

5. Lists of the component rules involved in our analyses are available in the unabridged report. See footnote 1.


Figure 1



The higher order rule in Figure 1 applies to the problem (i.e., the stimulus situation, S_0) and to the goal (G) itself, as well as to the lower order component rules.⁶

First, an arbitrary representation $\langle S_1, R_1 \rangle$ analogous to the solved problem is constructed. In our illustrative task, a sketch like Figure 2 would serve this purpose.

Figure 2



Note that constructing such a representation is not the same either as solving the problem, or as constructing a solution rule for the problem. The sketch in Figure 2, for example, can easily be generated by first drawing an arbitrary circle, then drawing an arbitrary line tangent to it, and placing an arbitrary point on it. More generally, an arbitrary representation (R_1) of the goal figure (R_0) is constructed first. Only then is a representation (S_1) of the information given in the stimulus situation (S_0) constructed in relation to the representation of the goal figure. In effect, the first operation on the higher order rule amounts to representing geometrically the meanings of goal situations⁷ (i.e., goals plus stimulus situations) by a "sketch," or some equivalent representation.

The second step is the question: "Is there a point X in $\langle S_1, R_1 \rangle$ which satisfies two locus conditions - and, if so, is there a goal constructing rule (r_g) such that point X is contained in the domain of r_g ($\text{Dom } r_g$) and such that the range of r_g ($\text{Ran } r_g$) is contained in the goal, G ?"

As shown in Scandura (1973a), decision making capabilities can be characterized as partitions on a class of input situations; in the present case, each representation $\langle S_1, R_1 \rangle$ either contains a point X which satisfies two locus conditions or it does not. If it does satisfy two such conditions, then the next operation involves forming the rule consisting of (1) a decision which asks whether there is a point X in the domain of r_g which satisfies two locus conditions, (2) the rule r_g , and (3) stop.

Next, the available component rules are tested to see whether there are two of them which apply to the represented stimulus (S_1) and generate loci which contain the point X . Given that such locus rules exist, the next operation constructs the solution rule R_s in which first one locus rule r_L is applied (after testing to see whether the stimulus situation is in its domain), when the other r_L , and finally the goal construction rule r_g .⁸

A more rigorous analysis

This level of description is sufficient to give one an intuitive feeling for how the higher order rule operates. But the rule is ambiguous, especially for computer implementation purposes. In the first decision making capability, for example, it is not clear just what constitutes a locus condition. Similarly, in the second decision making

6. See Appendix B.

7. See Appendix C.

8. See Appendix D.

capability the notion of a rule applying to a stimulus situation is something less than precise.

Closer perusal of the individual tasks made it possible to overcome these ambiguities. In many cases, the desired point X is a given distance from one or two given points and/or lines. In the example above, for instance, the point X is a distance R from the given point and from the given line. This suggested the following more rigorous characterization of the first decision making capability:

(1) Does there exist a point X in $\langle S_1, R_1 \rangle$ and a rule r_g such that (X, E) is contained in the domain of r_g where E is a given distance, and the range of r_g is contained in the goal ($\text{Ran } r_g \subset G$) such that X is a given distance from one or two given points and/or lines?

A similar analysis suggested reformulating the second decision making capability as:

(2) Is there a rule r_L such that a pair consisting of given points, lines, and/or distances in S_1 is in the domain of r_L ($\text{Dom } r_L$) and such that X is a member of L (i.e., a point on L) where L is contained in the range of r_L ($X \in L \in \text{Ran } r_L$)?

A similar characterization is required for r_L .

A higher order rule incorporating these refinements can be used to generate solution rules for many two-loci problems. For example, in the illustrative problem there is certainly a point X in the representation $\langle S_1, R_1 \rangle$ which is at the given distance R from a given point and from a given line in S_1 . It is also true that there is an r_g rule which applies to the pair consisting of the point X and the given distance, and whose range consists of circles and is thereby contained in the goal.

Unfortunately, as it stands, the modified higher order rule does not provide an adequate means for characterizing solution rules for other sampled two-loci tasks. In certain tasks, for example, no distance is given. The important requirement in such cases is often that the point X be equidistant from a given pair of elements, points and/or lines, in two different instances (i.e., for two given pairs of elements). Thus, in the tasks, "Inscribe a circle in a given triangle," the desired point X is equidistant simultaneously from two different pairs of sides of the triangle, or equivalently, the point X is equidistant from the three sides.

Still other tasks involve the (lower order) rule for constructing the locus of vertices of an angle of given measure subtending a given line segment. The task, "Given side a of a triangle, the median M_a , and the measure of angle A opposite side a , construct the triangle," is of this type. The locus of vertices, in this case, is an arc but the points on it are not at a fixed distance from any point on the given segment. Nor are the points of the locus equidistant from any two particular points on the line segment.

In order to take these possibilities into account, the decision making capability was generalized so that the point X could be equidistant from a pair of points or lines, or could serve as a vertex of an angle of given measure whose sides subtend (i.e., pass through the end points of) a given segment. Decision making capability (3) was also enriched so that pairs consisting of angle measures and/or segments could be in the domain of a locus rule.

Further, in the problem, "Given three intersecting lines, not all intersecting at a common point, construct a circle which is tangent to two of the lines and whose center is on the third," we have a situation where one of the loci, the line containing the point X , is already given. To handle this possibility we simply assume an "identity" lower order rule, one which identifies a given line as a required locus.

With these modifications, the higher order rule handled almost all of the pattern of two loci tasks we had sampled. We ran into difficulty, however, with another task: "Given two parallel lines and a point between them, construct a circle which is tangent to the two lines and passes through the point." This difficulty involved the second decision making capability (3). There is a pair of lines in the domain or one of the

locus rules - one which constructs the locus of points equidistant from the two given parallel lines. The second locus rule, however, requires that we first measure a distance between two parallel lines, one of which is not present in the stimulus S_0 until after the first locus rule is applied. That is, we need to determine the distance between one of the parallel lines and the locus of points equidistant from the two given parallel lines. This distance serves as the desired radius.

Application of the higher order rule in this case results in failure at decision making capability (3). Fortunately, it is easy to modify the higher order rule to take this possibility into account. Furthermore, as we shall see, this modification serves an important purpose in dealing with the larger class of construction problems solvable either by the pattern of two loci or by the pattern of similar figures.

Instead of stopping when the second decision fails, we simply add another group of tests (A-C). (A) and (B) duplicate (1) and (2) except that X must satisfy only one specific condition. (C) asks: "Is there one component rule such that a pair of given points and/or lines is in the domain of that rule and is there a locus L such that the point X is part of L and L is contained in the range of r_L ?" If the answer to this is no, we stop, but if the answer is yes, we can ask whether there is another locus rule r_L such that the represented stimulus situation S_1 , together with the preceding locus $r_L(S_1)$, contains a pair of given points and/or lines that are in the domain of r_L .

A revised higher order rule which incorporates all of these modifications is shown in Figure 3, found on the following page.

In checking this higher order rule we found it to provide an adequate account not only of all of the pattern of two loci problems sampled, but others as well. For example, consider Task A: "Given sides a, b, and c of a triangle, construct the triangle." In this case, application of the higher order rule generates the solution rule. This solution rule involves: (1) application of the rule, "Construct the locus of points at a given distance from a given point," to the end point of one line segment using another side as distance, followed by (2) another application of the rule to the other end point using the remaining side as radius. Then, the triangle rule, "from a point not on a given segment, draw segments to the end-points of the given segment," is applied to the intersection of these two loci to obtain the desired goal figure.

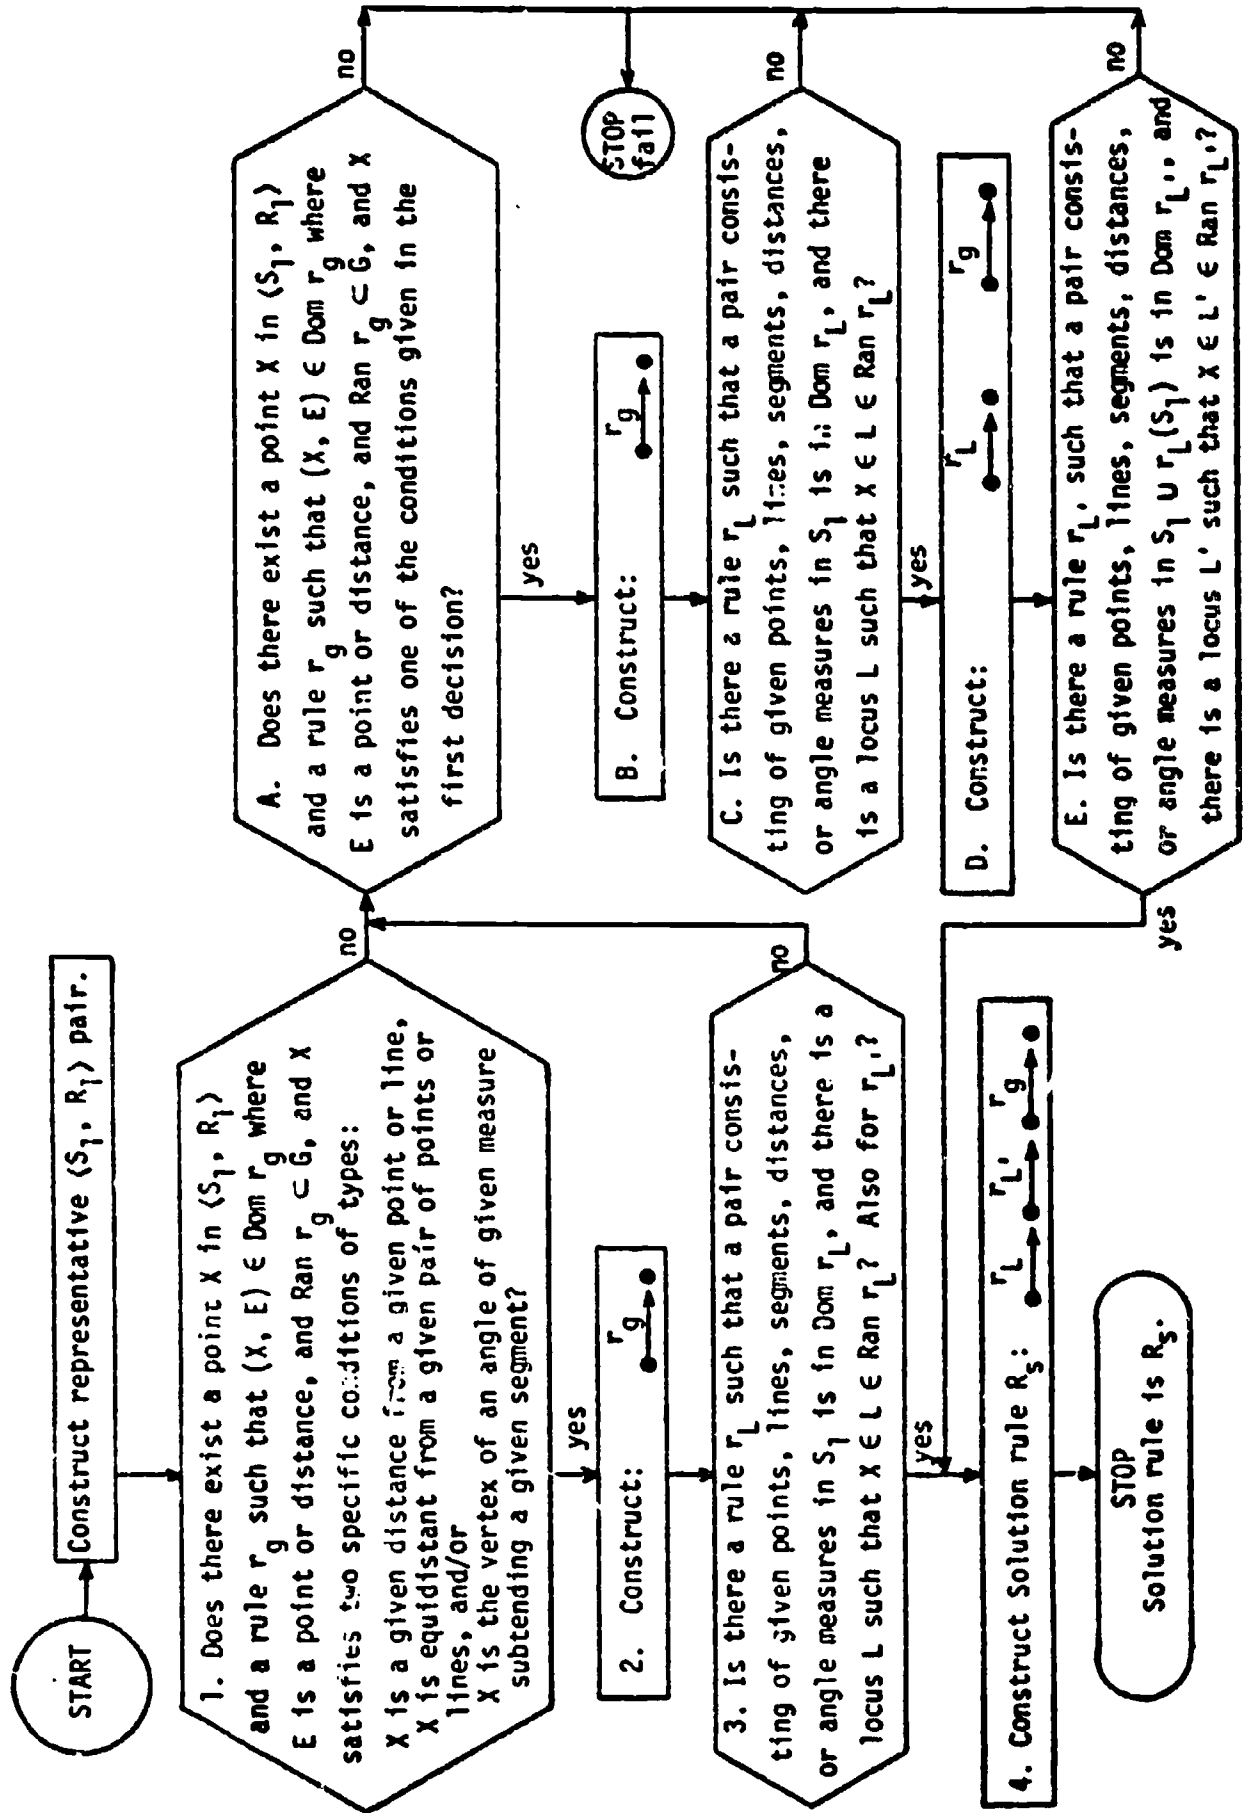
In some cases, of course, different lower order (component) rules were involved. For example, consider task B, "Given two intersecting lines and a point of tangency on one of the lines, construct a circle which is tangent to the two lines and which passes through the given point of tangency." In this case, the locus rule for constructing perpendiculars to lines through points on the given lines had not been required with any of the sampled problems.

Discussion

Aside from the possibility that new two-loci problems may require additional lower order rules, the higher order rule appears adequate. In particular, the higher order rule not only generates solution rules for each of the sampled two-loci problems, but also seems compatible with human knowledge.

As the form of the higher order rule suggests, the component decision making capabilities play a crucial role in deriving solution procedures. These decision making capabilities are designed to reflect the underlying semantics of the problem situations by referring directly to figural representations of semantic information implicit in the problem descriptions. In general, parts of a figural representation (S_1, R_1) will represent the meaning of a task statement and reflect the relation between the given stimulus (S_0) and the goal figure (R_0). Notice that while the relation between S_1 and R_1 will be the same as between S_0 and R_0 , S_1 and R_1 will not in general be the same as S_0 and R_0 , respectively.

Figure 3



For purposes of our analysis, the decision making capabilities were viewed as atomic although they can also be analyzed into more basic components. The first decision making capability in the second two loci higher order rule, for example, involves both a conjunction and disjunction of a number of simpler conditions. This decision making capability could be subdivided, for instance, into the following two decisions: (A) Is there a point X that is a given distance from a given point and/or line? (B) is there a point X equidistant from a pair of given points or lines? Instead of having one decision making capability involving conditions A and B, then, we could have one decision making capability involving A, and a subsequent one, B.^{9, 10}

In addition to its purported compatibility with human knowledge, the higher order rule is also sufficiently precise to be mechanizable. One of the authors (Wulfeck) has recently written a program in SNOBOL 4 which uses an intermediate version of the two loci higher order rule (see the unabridged report) to generate solution procedures for many of the problems we sampled. A naming system replaced the figural representation described above (see footnote 7). Routines corresponding to many of the lower order rules (see the appendices of the unabridged report for lists of the component rules) were also written.

Granting the adequacy of the higher order rule for purposes of our analysis, we wish to comment briefly on some limitations in regard to the compatibility of the lower order rules with human knowledge, though the specification of component rules is not our central concern. These limitations are all variants on a common theme: The lower order rules we have identified can be constructed from more basic components. This fact is reflected in at least three ways.

First, many of the simple rules have components in common. Several rules, for example, all involve constructing a locus of points (circle) at some distance from some point. The differences lie in whether or not the distance and/or center points are given directly or must be determined first. The construction rules needed to determine these distances and/or center points are quite basic and are apt to be useful in a wide variety of construction situations. Any reasonable account, designed to deal with a wider variety of problem situations, would undoubtedly include these construction rules directly in the rule set.

Second, certain of the identified lower order rules, particularly the rule for constructing the locus of vertices of an angle of a given measure subtending a given line segment, are complex in themselves and cannot automatically be assumed to be available to many problem solvers.

A third limitation is closely related to the first and was mentioned earlier: the lower order rules are to some degree specific to the tasks we have identified. To some extent this may be unavoidable because there are always certain problems which require "trick" solutions. It would be desirable, of course, to keep this to a minimum. In this regard, it should be emphasized that the simpler the lower order rules the greater the problem solving flexibility.

One way to modify our characterization to handle these limitations would be to "reduce" the lower order rules into their components and, correspondingly, to "enrich" the higher order rule by adding sub-routines for constructing the needed locus, r_L , and goal, r_g , rules.¹¹ Such rules would correspond to the type of knowledge that a person just having been taught the basic construction rules would need to have in order to generate solution rules directly.

For example, consider the rule: "Determine the distance between a given point and a given line and then construct the locus of points at the obtained distance from the given point." This rule can be divided into two subrules: (1) "Determine the distance between a point and a line," and (2) "Construct the locus of points at a given distance from a given point." To compensate for the reduction in the latter case, the higher order rule could be "enriched" so that more complex r_L and r_g rules can be

9. For a discussion of how new decision making capabilities are learned from simpler ones, see Scandura (1973a).

10. Such refinement may be useful in the assessment of behavior potential (Durnin & Scandura, 1973), specifically in increasing the precision of diagnostic testing.

generated where needed. Specifically, instead of selecting a composite rule directly when it meets certain prescribed conditions, as we have done so far, we include in the higher order rule a simple sub-routine for combining component lower order rules. Such a sub-routine for example, might select (sub)rules until one is found whose domain includes a pair consisting of a point and a line (e.g., the distance measuring rule (1), and another (e.g., the circle rule (2)) such that its range consists of circles (loci). To make the search more efficient, it is natural to add the requirement that the range of the former be contained in the domain of the latter. After the component rules have been identified, the sub-routine would form the composite of these rules, and finally, would test the composite against the condition in the initial higher order rule.

As attractive as this possibility might appear at first, a little thought suggests its implausibility as a way of modeling human knowledge. This can be seen by noting that all geometric constructions with straightedge and compass are generated by just three basic operations: (a) using a straightedge (e.g., to draw a line, ray, or segment through two given points, or through one point, or intersecting a line, etc.) (b) drawing an arc given a compass set at some fixed radius, and (c) given two points, setting a compass to the distance between those points.

As we have seen, many of the lower order rules are really quite complex. Requiring a higher order rule, designed to reflect human knowledge, to generate such rules from elemental components is unrealistic. It is unlikely that a subject who is only able to perform the three indicated operations above would also have at his command a rather complex and sophisticated higher order rule. The acquisition of such complex capabilities by naive subjects, whether of a higher or lower order, would almost certainly have to come about gradually through learning, presumably by interacting with problems in the environment.¹²

PATTERN OF SIMILAR FIGURES

Three classes of similar figures problems

The pattern of similar figures problems were analyzed in similar fashion. Again, we began with a broad sampling of problems from Polya (1962). One of the problems identified was, "Given a triangle, inscribe a square in it such that one side of the square is contained in one side of the triangle and the two other opposite vertices of the square lie on the other two sides of the triangle." The second step was to identify a solution rule for each of the problems. For the problem above the solution rule was, "Construct a square of arbitrary size such that one side is contained in the side of the triangle which is to contain the side of the goal square, and such that one vertex is on another side of the triangle. Draw a line through the point of intersection of those two sides of the triangle and through the fourth vertex of the arbitrary square. From the intersection of this line and the third side of the triangle (which is the fourth vertex of the goal square) construct a segment perpendicular to the side of the triangle which is to contain a side of the goal square. Complete the goal square using the length of the perpendicular segment as the length of the sides."

Similar figures problems, like the example task above, may be characterized as those whose solution procedures involve a similarity mapping process: from some center (point) of similarity a figure or set of points is mapped onto another. Further, the solution procedures always involve constructions according to geometric invariants under similarity mappings, either parallel lines, since parallelism is preserved, or equivalently, "copying" angles, since similarity maps are conformal.

Further analysis of the similar figures problems revealed three relatively distinct classes of solution rules. In the sample problems above, and in other problems in the same class, the solution rules all involve first constructing a square of arbitrary size which is in the same orientation as the desired goal square, and which meets as many of the task conditions as possible. (Rules of this type for constructing similar

12. See the section on "future directions".

figures are denoted by r_{gs} .) The second step in each solution rule uses two pairs of corresponding points in the goal and similar figures (i.e., in $\langle S_1, R_1 \rangle$ superimposed with the similar figure) to determine the point of similarity (P_s), and then, constructs a line through the point of similarity and a point on the similar figure which corresponds to a needed point of the goal figure. (Point of similarity rules are denoted r_{ps} .) Finally, the obtained point on the goal figure is used as a basis for constructing the goal square.

The second class is well represented by the problem, "Given angles B and C of a triangle, and the median M_a to side a, construct the triangle." The corresponding solution rules begin similarly by applying a similar figures rule (r_{gs}) to two given angles to construct an arbitrary sized triangle similar to the goal triangle, with medians, altitudes, etc., as required. Then a modified point of similarity rule (r_{ps}) is used to determine the point of similarity (P_s , the vertex of the non-given angle), and to construct the given segment (e.g., M_a), such that one endpoint of the segment is the point of similarity, and such that the segment coincides with the corresponding segment in the similar triangle. Finally, a line is constructed, through the other endpoint of the constructed segment parallel to the side of the similar triangle that is opposite to the point of similarity. The remaining sides of the goal triangle are obtained by extending two sides of the similar triangle to intersect the constructed parallel line.

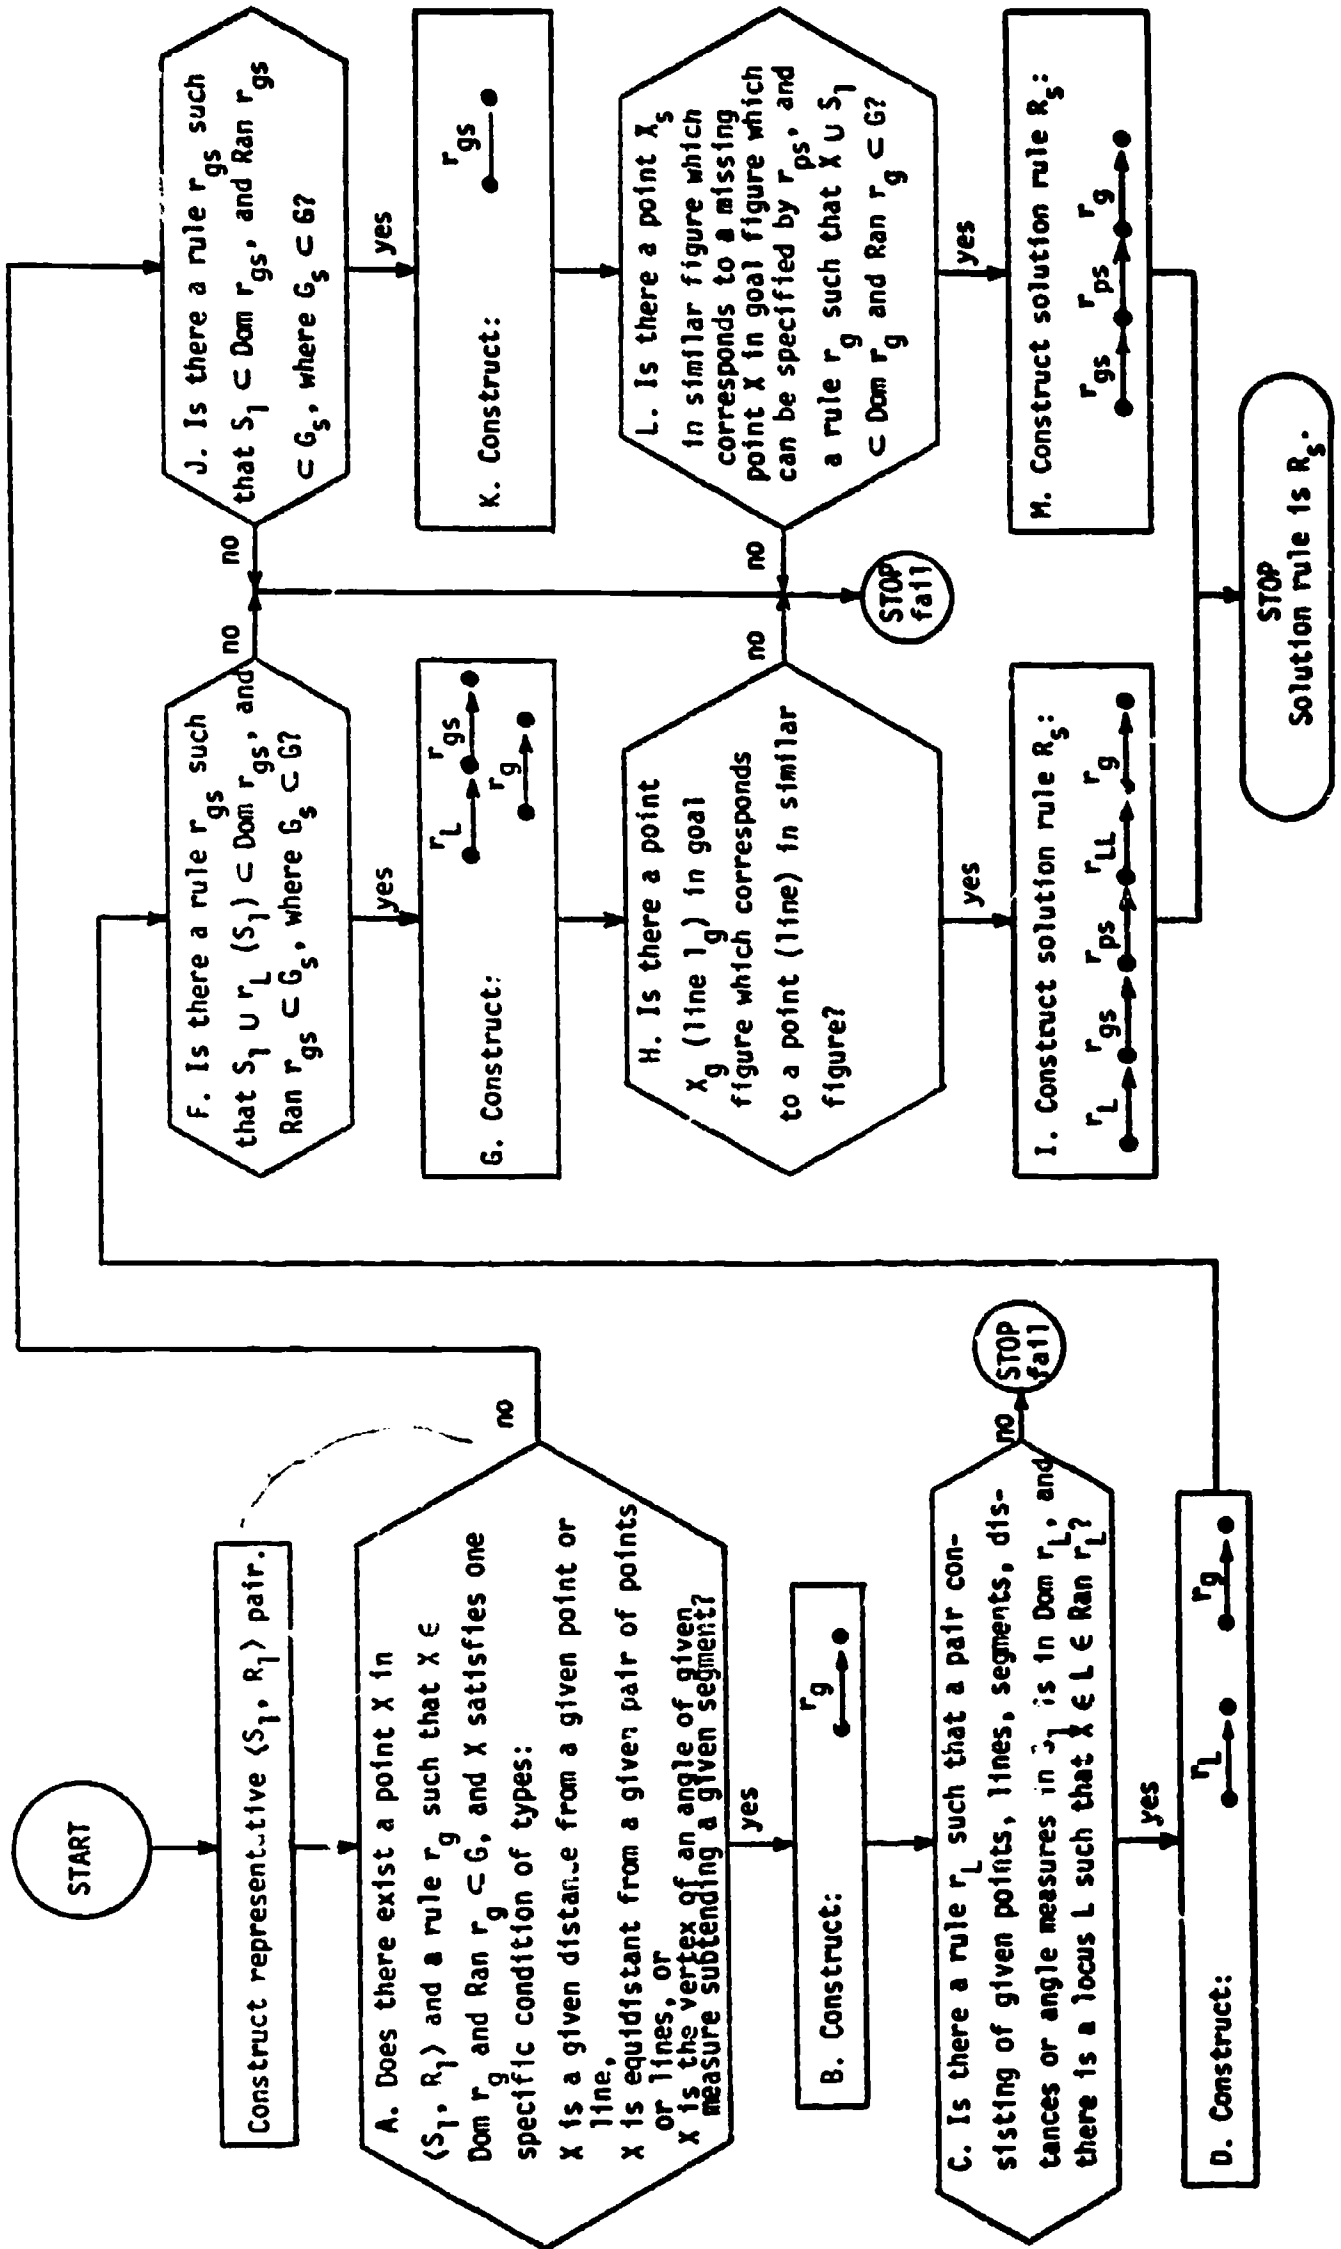
The solution rules for the third class of problems differ in that the first step in each is to use an r_L rule to construct a locus of points which contains a critical point, specifically the center of the goal circle. In the problem, "Given a line and two points (A and B) on the same side of the line, construct a circle tangent to the line which passes through the two given points," for example, the locus of points (L) equidistant from the two given points contains the center of the goal circle. Also, the point of similarity is the intersection of the locus and the given line. The second step is to construct a similar figure (circle, C_1), which satisfies part of the goal condition. In our example, a circle is constructed with center on the constructed locus and tangent to the given line. Next, another version of the point of similarity rule is applied; this time the point of similarity (P_s) and a given point on the goal figure (e.g., B) are used to determine a corresponding point (B') on the similar circle. Then, parallel lines involving corresponding points are constructed to determine the center of the goal circle. Finally, the goal circle is actually constructed.

The similar figures rule

The higher order rule shown in Figure 4 (together with a set of applicable lower order rules) provides a sufficient basis for solving all of the sampled pattern of similar figures problems. Furthermore, the higher order rule appears to reflect the underlying semantics (Figure 4 found on the following page). For example, let us see how a solution rule for the first illustrative problem above (inscribing a square in a triangle) can be generated by application of the higher order rule. The first decision making capability (A) asks essentially whether a point X is needed to serve as the center for a goal circle. As the goal figure is a square, the answer is obviously no. Decision making capability J then asks if there is a goal similar figure rule (r_{gs}) which applies to representing stimulus S_1 and generates squares that satisfy part (G_s) of the goal condition (i.e., the range of r_{gs} is contained in G_s which in turn contains G - equivalently, anything which satisfies G_s satisfies G , but not necessarily conversely). The lower order rule, "Construct a square in a triangle with one side coincident with one side of the triangle and one vertex on another side of the triangle," satisfies these conditions so the rule is retained as indicated in operation K.

Decision making capability L asks two things: (1) Is there a point X_s which corresponds to a missing point X in the goal square? (2) Is there a rule r_g such that

Figure 4



the stimulus S_1 , supplemented with the point X (i.e., $X \cup S_1$), is in the domain of r_g , and r_g generates a goal-like figure ($\text{ran } r_g \subset G$)? In short, is there a point X_s in the similar square which corresponds to a point X from which the goal square may be constructed? Clearly, there is such a point X_g and the rule, "Determine the distance from a point to a given line segment and construct a square with sides of that length" satisfies the necessary conditions. Operation M forms the solution rule consisting of the two rules above with the point of similarity rule between them.

To see how the higher order rule works with the second class of problems, consider the second illustrative problem above (constructing a triangle, given two angles and a median). In this case, the answers to decision making capabilities A and J are again "no" and "yes," respectively. Here, r_{gs} is, "Construct a triangle of arbitrary size using two given angles and add parts corresponding to given segments." At decision making capability L there is a point X in the goal figure, the endpoint of median M_a , which can be specified by r_{ps} . Operation M again forms the solution rule.

Notice that the first two classes of problems involve the same path in the higher order rule. Each solution rule requires a goal similar figure rule (r_{gs}), the point of similarity rule (r_{ps}), and a goal constructing rule (r_g). The only difference is whether the goal and similar figures are squares or triangles, with all that implies for the particular r_{gs} and r_g rules required. In short, this example illustrates how what may appear initially to be basically different kinds of problems may turn out to have a common genesis.

The third problem (constructing a circle tangent to a given line and passing through two given points) illustrates the other path through the higher order rule. In this case, if we knew the center (X) of the desired circle, we could solve the task. Furthermore, this missing point X is on a locus, namely the locus of points equidistant from the two given points. Hence, decision making capabilities A and C are satisfied, and we retain the circle constructing rule (r_g) and the perpendicular bisector rule. Decision making capability F asks if there is a rule (r_{gs}) which applies to the stimulus S_1 as modified by the output of the locus rule (i.e., $S_1 \cup r_{pb}(S_1)$). Condition F is satisfied by a rule that generates circles with centers on a given line (the locus) and tangent to another given line. The answer to the decision making capability H is also "yes." The two given points on the goal figure obviously correspond to two points on the similar circle. By operation I , the solution rule follows directly: "Construct the locus of points equidistant from the two given points; construct a circle with center on that locus tangent to the given line; apply the point of similarity rule, and then the parallel line rule to determine the center of the goal circle; construct the goal circle using this center and the distance between it and a given point of radius."

It should be noted that in one of the sampled tasks the "locus" is given. The easiest way to handle this special case is to simply add an identity locus constructing rule as before. It would also be a simple matter to modify the higher order rule to take this possibility into account by asking, prior to or at decision making capability C , whether there is a line in S_1 which contains X .

Combined rule for two-loci and similar figures problems

It would appear from our analysis that the two higher order rules, together with the necessary lower order rules, would provide an adequate basis for solving the sampled two loci and similar figures problems and others like them. Indeed, there are two possible modes of solution in the case of one of the sampled similar figures tasks: "Inscribe a square in a right triangle so that two sides of the square lie on legs of the triangle, and one vertex of the square lies on the hypotenuse." Instead of using the pattern of similar figures, as illustrated in our first example, the pattern of two

loci rule can be used to construct the bisector of the right angle. The intersection of this locus with the hypotenuse (the other locus) is the "missing point" X and provides a sufficient basis for constructing the goal square.

Although it is not always critical to distinguish between different modes of problem solving, any complete account designed to reflect human behavior must specify why one mode of solution is to be preferred over another (cf. Scandura, 1973a, Ch. 8). In the present case, there are two possible ways of handling this. First, we can add a higher order selection rule to the rule set which says simply, if both higher order rules apply, select the pattern of two loci. The rationale is that the pattern of two loci rule will generally yield a simpler method of solution.

A second way to handle the problem is to devise a single higher order rule which combines the advantages of both higher order rules. The higher order rules in Figures 3 and 4 can be combined to yield the higher order rule depicted in Figure 5. The path in this higher order rule designated $\langle 1,2,3,4 \rangle$ corresponds to that path of the two loci higher order rule which deals with those cases where the two loci may be found in either order. The path $\langle 1,2,3,A,B,C,D,E,4 \rangle$ deals with those two-loci problems where one locus must be found before the other. The other two paths correspond to the similar figures higher order rule.

PATTERN OF AUXILIARY FIGURES

Not all compass and straightedge problems can be solved via the pattern of two loci or the pattern of similar figures. In this section, we describe a higher order rule for dealing with the third class of problems identified by Polya (1962), the pattern of auxiliary figures. We also show how the combined higher order rule (above) may be extended to account for essentially all of the construction problems identified by Polya (1962).

Auxiliary figures higher order rule

Our initial analysis was based on a sample of five diverse auxiliary figures problems. One of the problems used was, "Given the three medians of a triangle, construct the triangle."

The analysis proceeded as before. First, we identified a procedure for solving each problem. Then, we looked for similarities among the solution rules and identified the component rules involved. In general, the required goal figures were not constructable via either the two loci or similar figures higher order rules. However, in each case the goal figure could be obtained from an (auxiliary) figure that was constructable from the given information. In the problem above, for example, a triangle can be constructed from segments one-third the lengths of the given medians. The goal figure is obtained by extending two of the sides of this auxiliary triangle to the respective median lengths and drawing lines through the resulting endpoints.

The analysis resulted in the auxiliary figures higher order rule shown in Figure 6. This higher order rule generates a solution rule for the illustrative task above as follows. First, an arbitrary representation for the solved problem $\langle S_1, R_1 \rangle$ is constructed. In this case, an arbitrary triangle is "sketched," and its medians are represented on it. The first decision asks whether there is (1) an auxiliary figure, and (2) a rule r_g which operates on the auxiliary figure and generates the goal figure. In this task, there is such an auxiliary figure, a triangle having sides one-third the lengths of the given medians.¹³ In addition, the rule, "Extend the constructed segments to their given lengths and draw lines through their endpoints," satisfies condition (2). The next decision (III) asks whether or not a point is needed, in addition to the auxiliary figure, to construct the goal. Here, the answer is "no"; no other point is needed. Finally, decision IV asks if there is an auxiliary figure construction rule (r_a) available whose domain contains S_1 ($S_1 \in \text{Dom } r_a$)

13. See Appendix F.

Figure 5

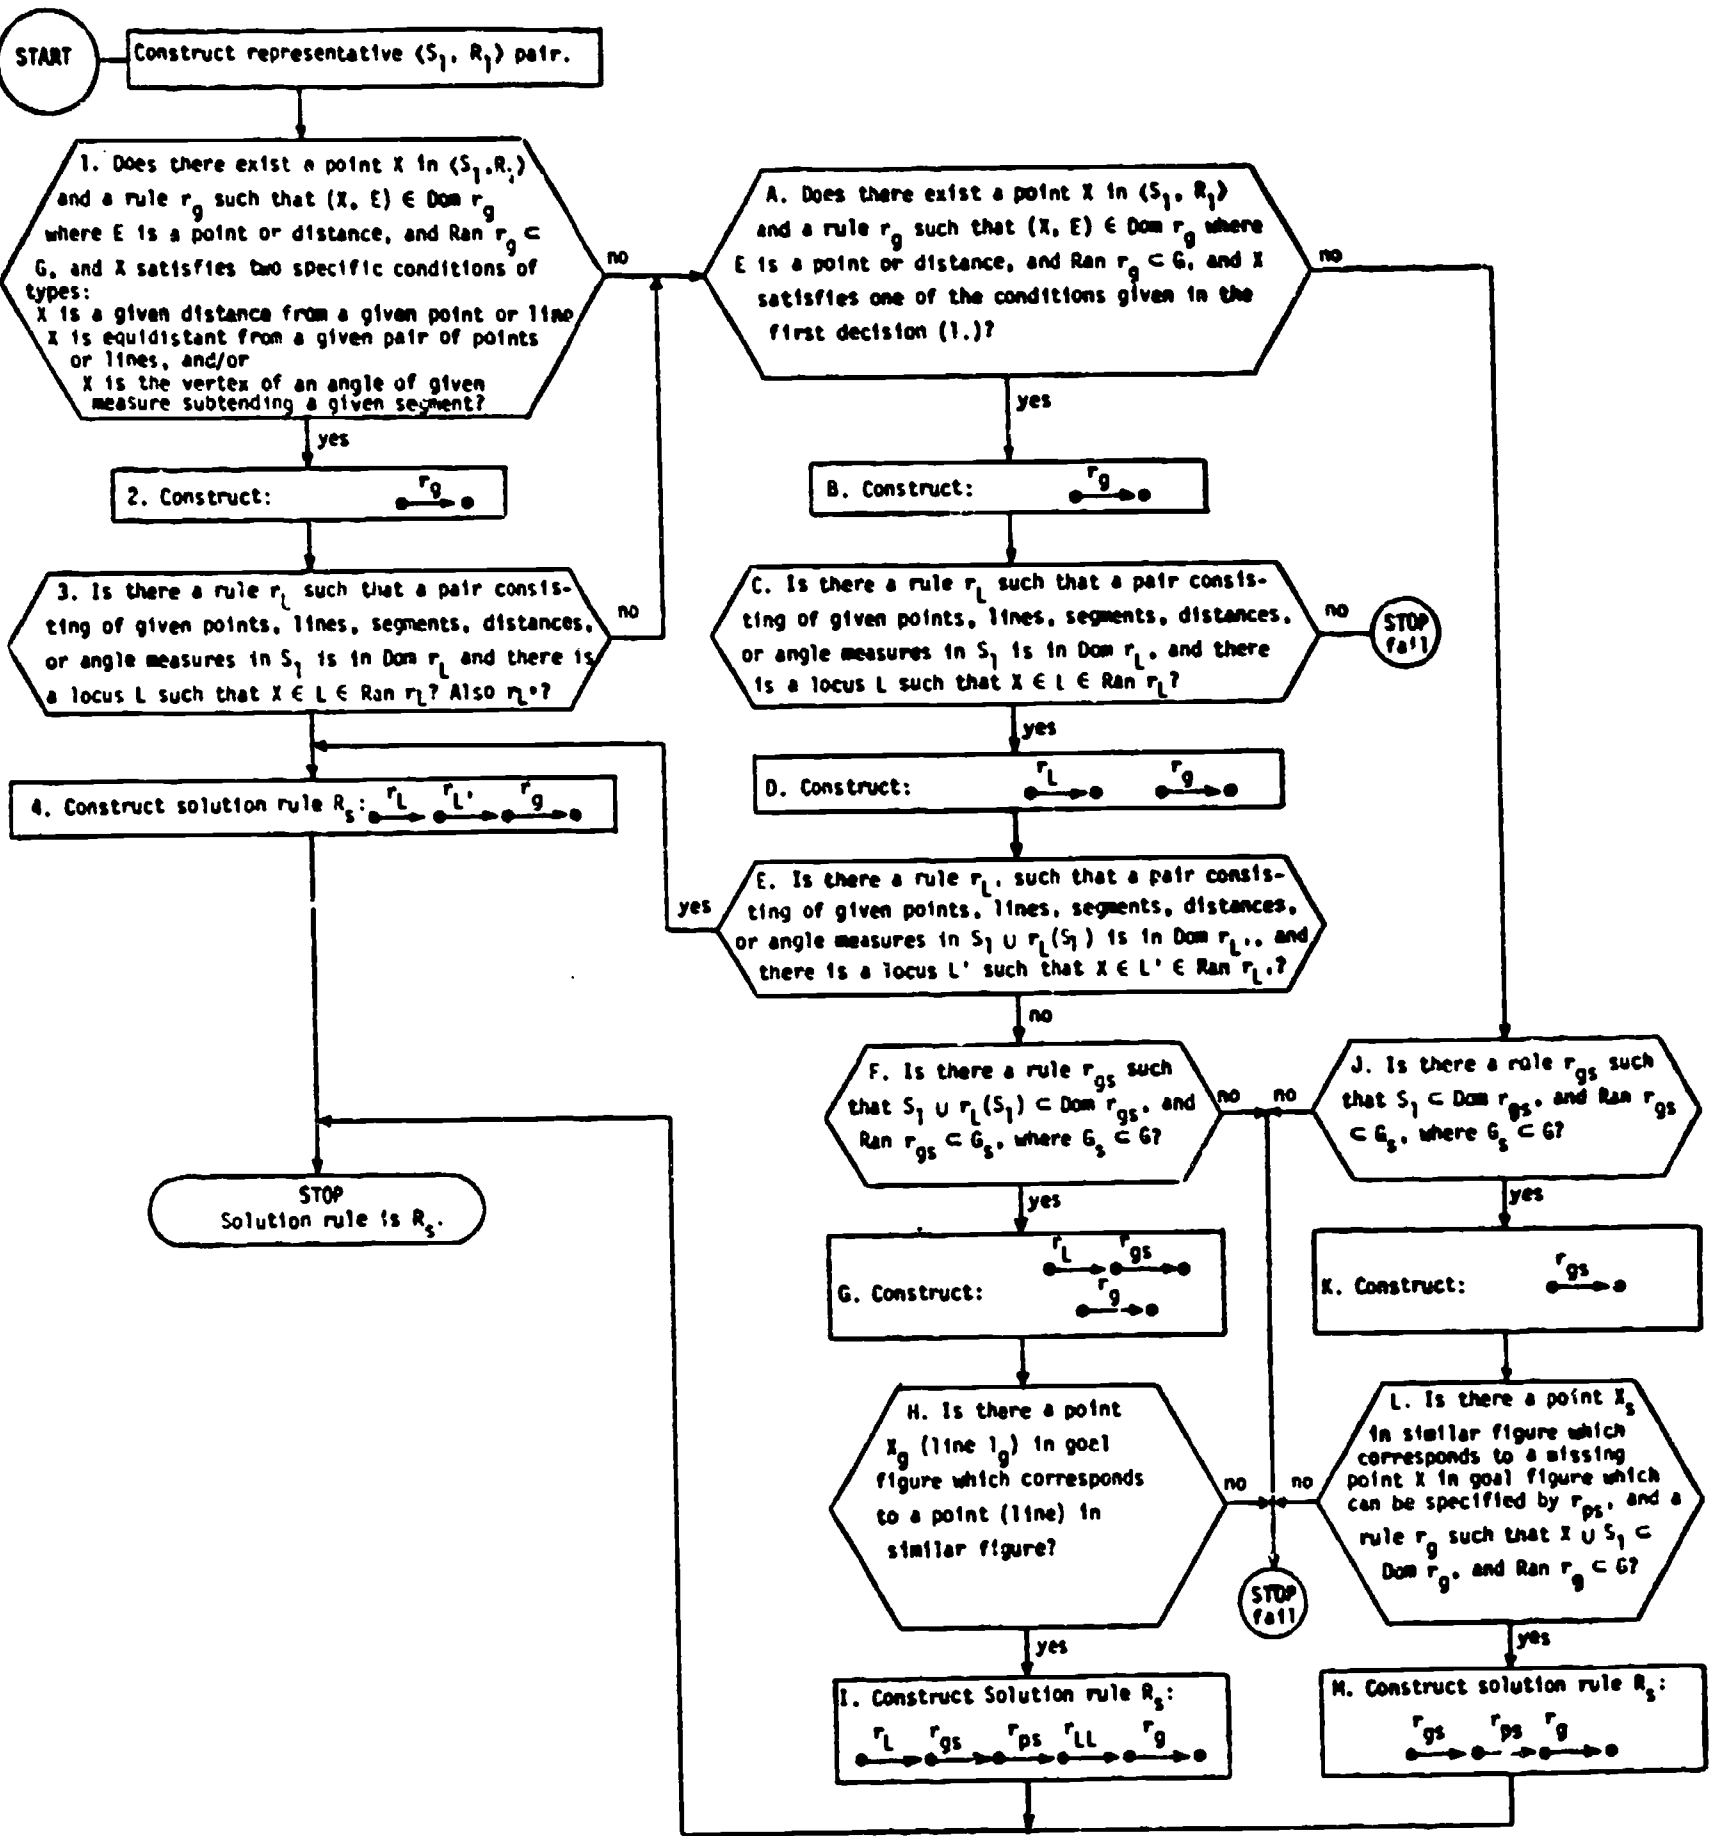
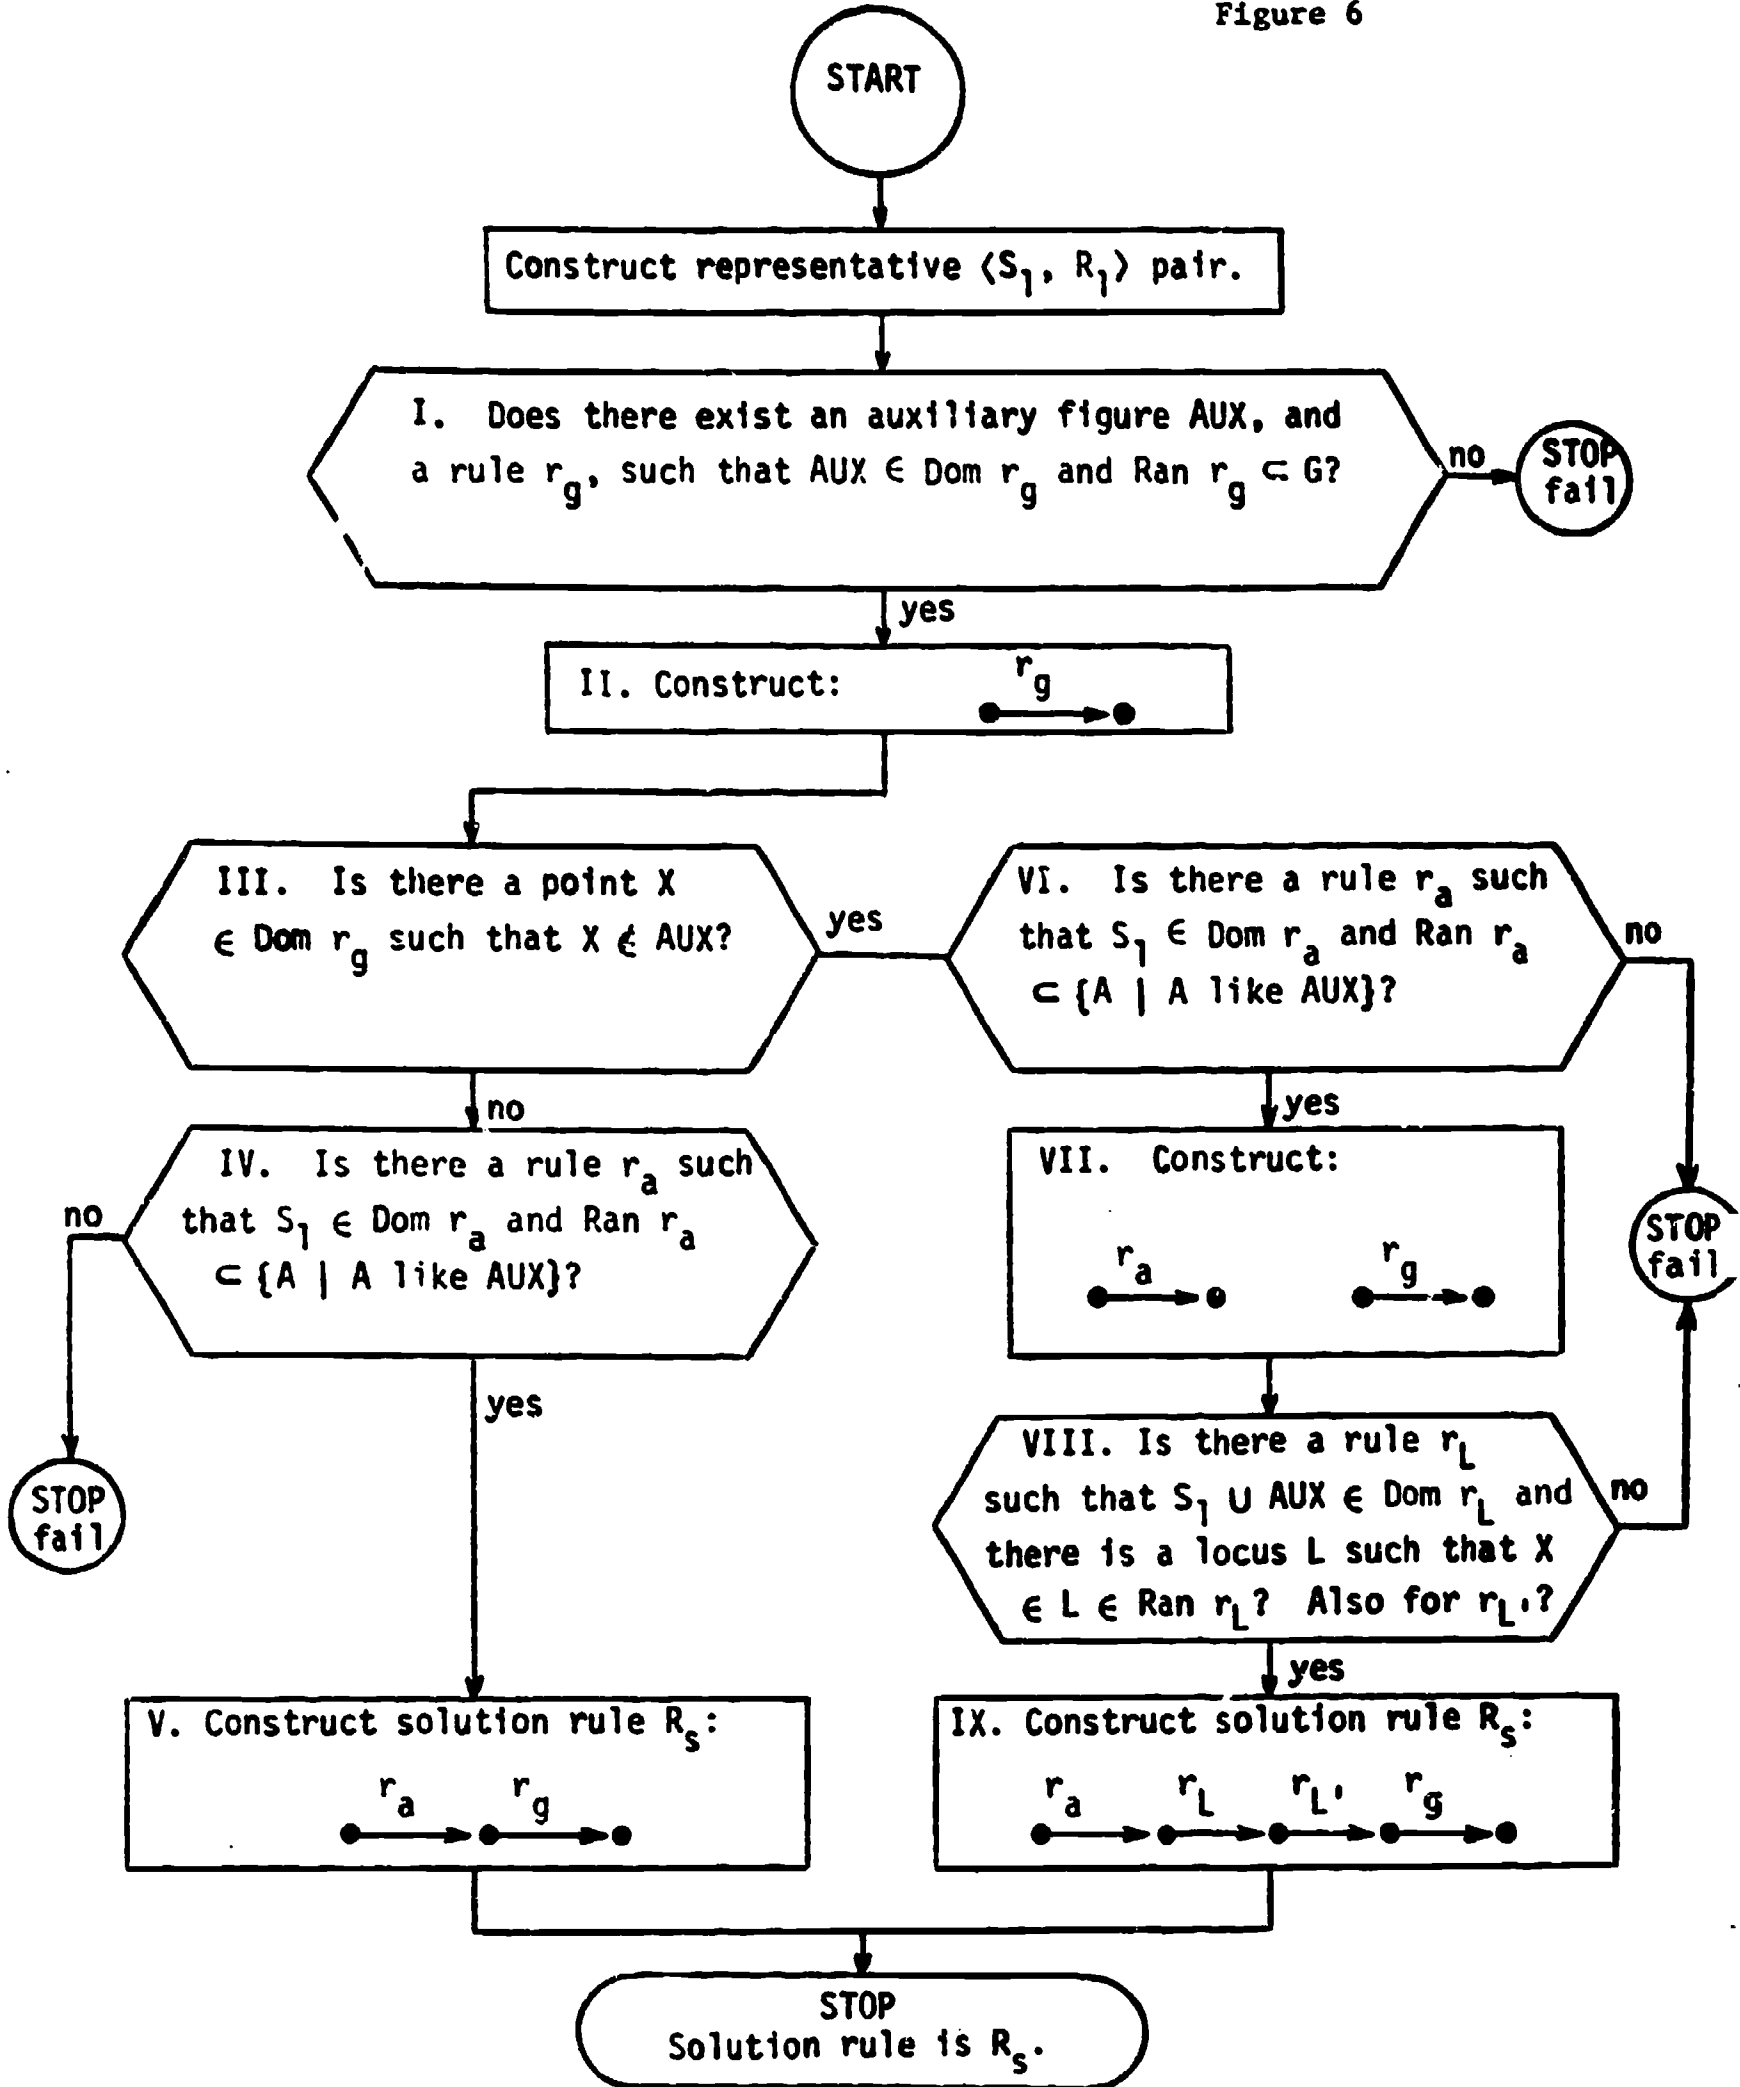


Figure 6



and whose range contains the auxiliary figure (i.e., $\text{Ran } r_a \subset \{ A | A \text{ is like AUX} \}$). In this case, the rule, "Construct a triangle from segments one-third the lengths of three given segments (medians)" satisfies these conditions and operation V constructs the solution rule, "Construct a triangle having sides one-third the length of the given medians; extend two segments of the constructed triangle to the respective median lengths, and draw lines through the endpoints of the medians to construct the goal triangle."

The other path through the higher order rule may be illustrated using the task, "Given the four sides a, b, c, d of a trapezoid ($a < c$), construct the trapezoid." Again, the answer to decision I is "yes." (Where the answer is "no," the higher order rule fails.) The triangle with $c-a, b, d$ as sides, serves as the auxiliary goal figure and the goal rule, "Through corner points of an auxiliary figure and through another point not in the auxiliary figure, draw segments to complete the goal," is selected. Unlike the first path, however, the answer to decision III is "yes" since the goal rule (r_g) acts on pairs ($X \cup \text{AUX}$) consisting of an auxiliary figure and a critical point X . The next decision (IV) asks if there is a rule r_a that constructs the auxiliary figure from given information. This condition is satisfied by the r_a rule which constructs the auxiliary triangle from the sides of a trapezoid. Decision VIII asks whether there are two locus rules (r_L and $r_{L'}$) which apply to the auxiliary figure and/or other given information (S_1) and whose ranges contain X . The circle rule (r_C), applied to different portions of $S_1 \cup \text{AUX}$, plays the role of both locus rules. The solution rule (Operation IX) is a concatenation of the component rules.

Combined two-loci, similar and auxiliary figures higher order rule

Taken collectively, the three higher order rules described above can be used to construct solution procedures for a wide range of geometry construction problems. Furthermore, they appear compatible both with human behavior and with the heuristics originally identified by Polya (1962).

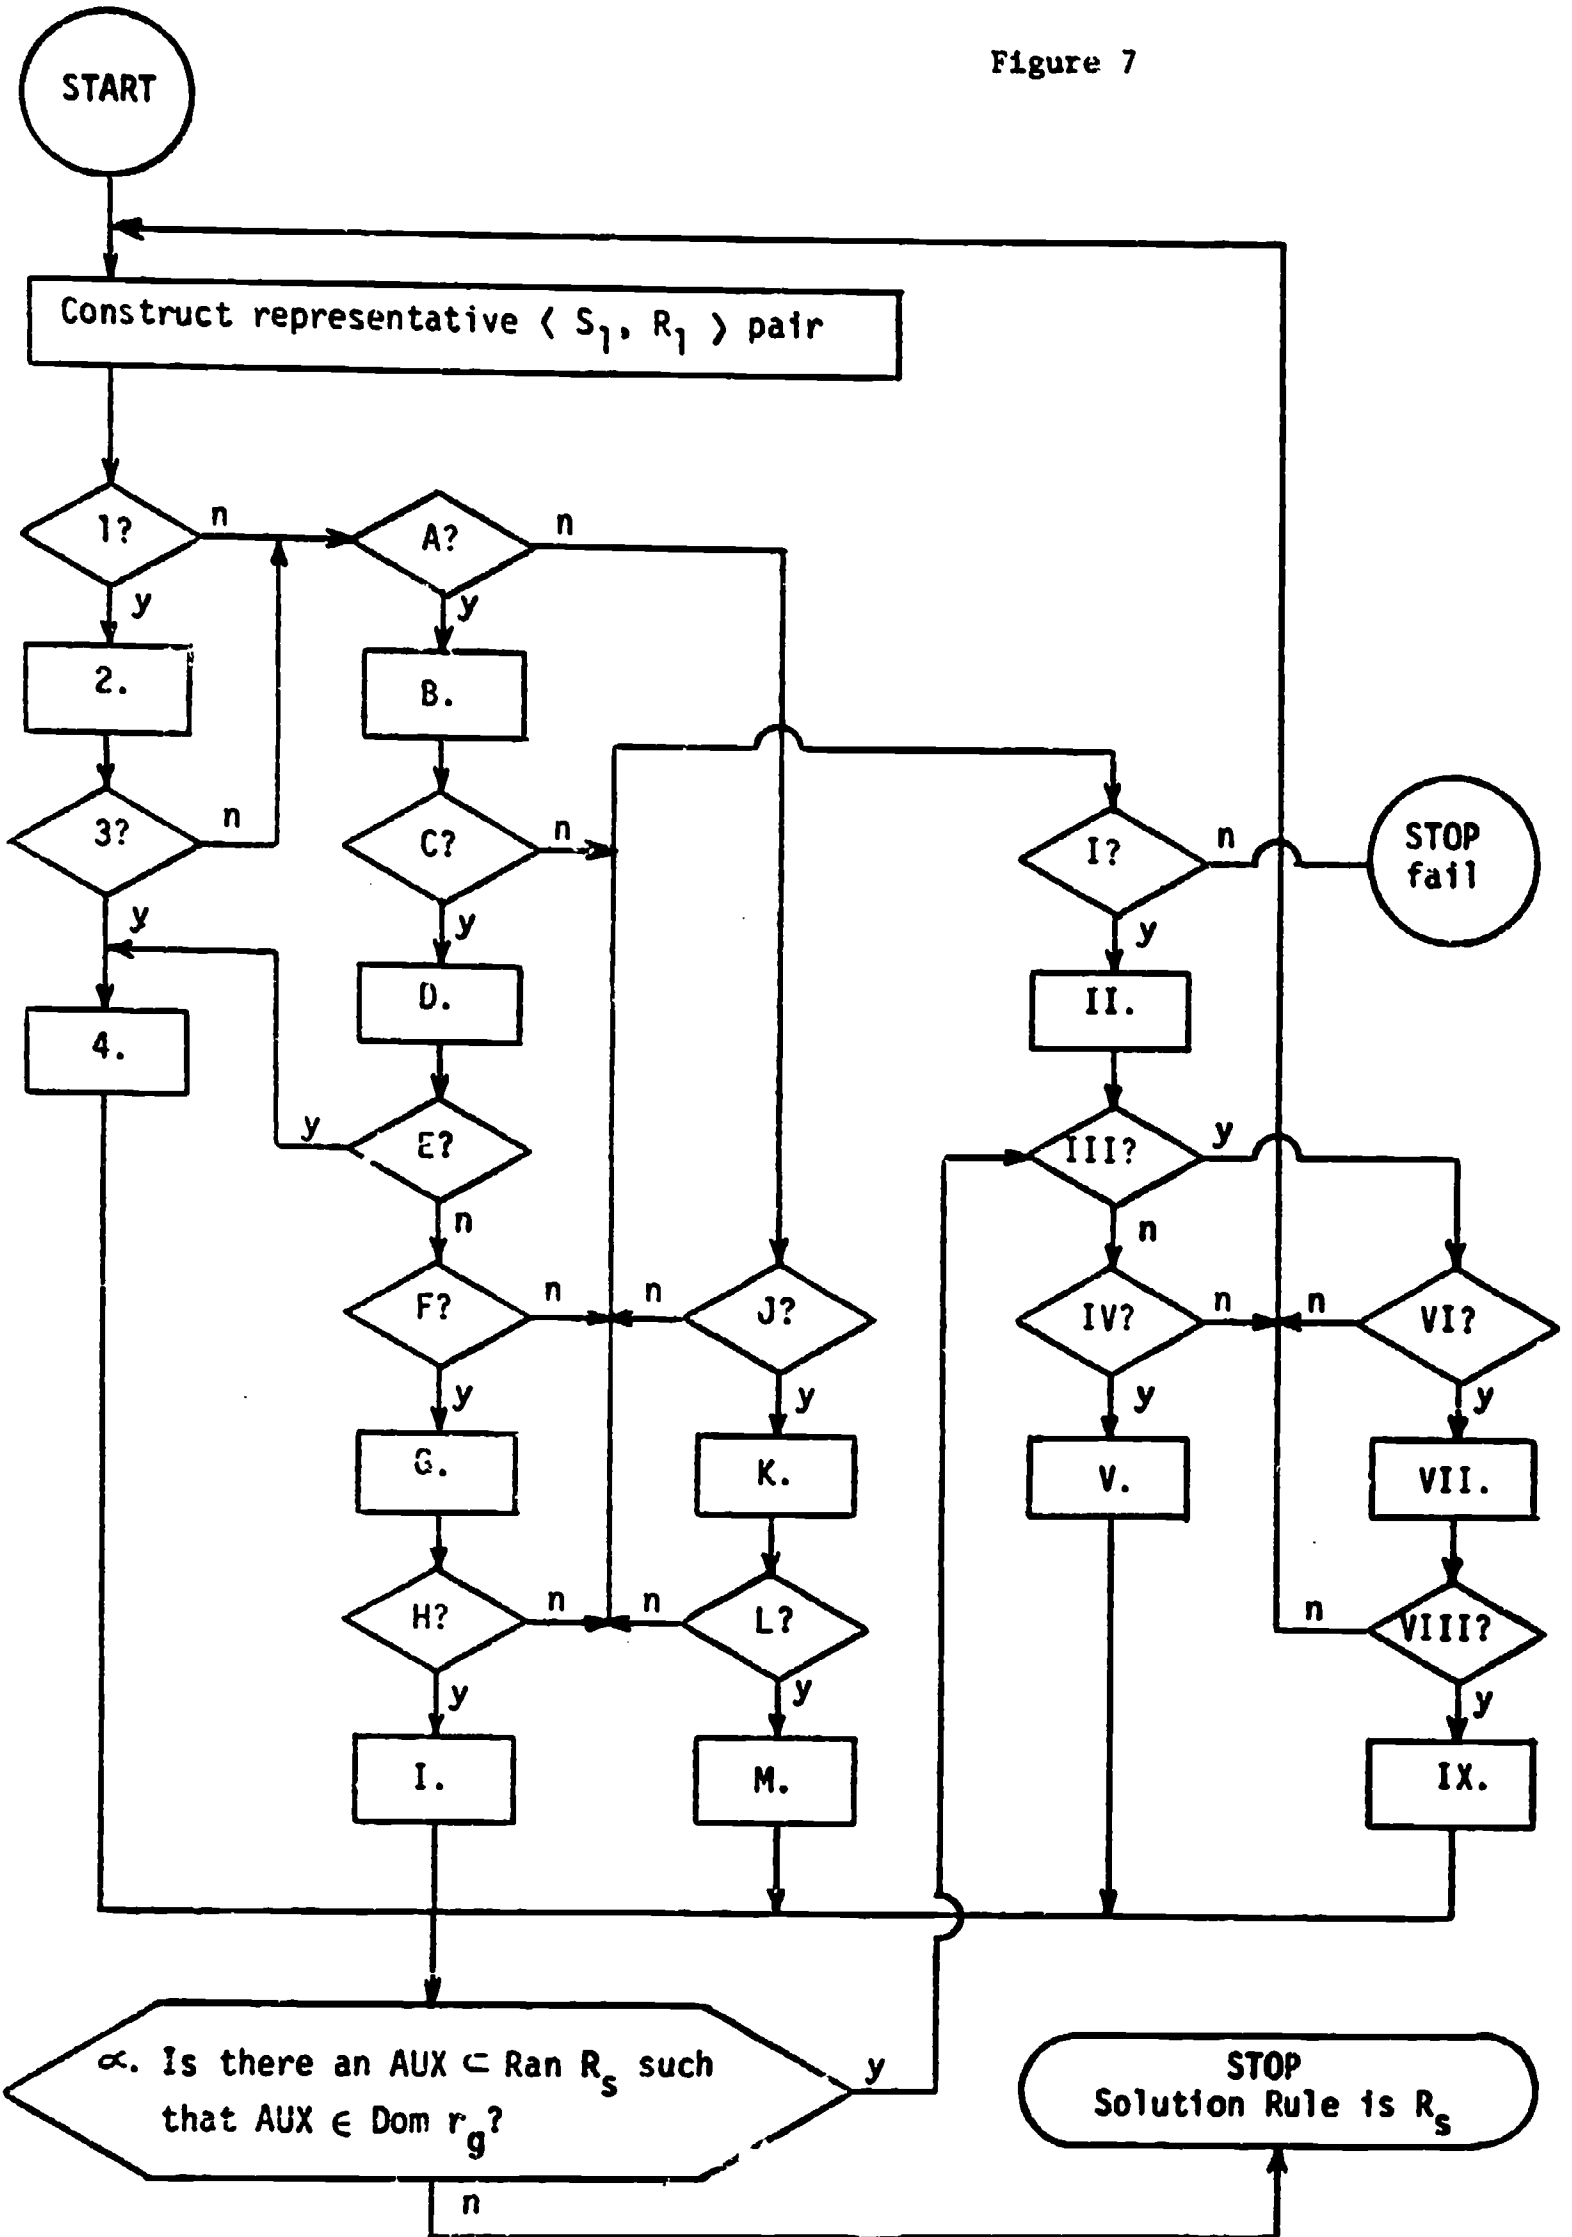
This is not meant to imply, however, that the three higher order rules are unrelated to one another. Both the needed point X in the pattern of two loci, and the similar figure in the pattern of similar figures can be regarded as special auxiliary figures. Indeed, one could modify the auxiliary figure higher order rule so that it, together with the relevant lower order rules, would account for all three classes of problems. In addition, the similar and auxiliary figures higher order rules may be viewed as progressive generalizations of the two-loci higher order rule. It is not difficult to conceive of third level higher order generalization rules which have the two loci higher order rule and a similar or auxiliary figure as inputs, and a more general higher order rule in which a similar or auxiliary figure is substituted for the missing point X , as the corresponding output.

Alternatively, the combined two-loci, similar figures higher order rule (Figure 4) can be extended to include auxiliary figures. In addition, the extended higher order rule depicted in Figure 7 allows recursion on the higher order rules.

To see this, notice that the higher order rule shown in Figure 6 can terminate at several points without finding a solution rule. In some problems this is unavoidable; there may not be an auxiliary figure from which the goal figure can be constructed. Sometimes, however, there is an auxiliary figure, but one which is not directly constructable from the given information. Such auxiliary figures can often be constructed via the pattern of two-loci, the pattern of similar figures, or the pattern of auxiliary figures itself. In those cases where such an auxiliary figure exists, we allow for this possibility by returning control to the start of the combined higher order rule in order to derive an r_a rule for constructing the auxiliary figure. Once an auxiliary figure (r_a) rule has been derived, the original procedure resumes.

To see how this higher order rule works, consider the following task, "Construct a trapezoid given the shorter base a , the base angles A and D , and the altitude H_t ." As


Figure 7



in the trapezoid example given earlier, the needed auxiliary figure is the triangle having sides $c-a$, b , and d . But, this triangle is not directly constructable from the given information. None of the assumed lower order rules is adequate, so the higher order rule breaks down at step VI. The flow of control therefore returns to step I with the aim of constructing the auxiliary figure.¹⁴ Beginning here, the problem of constructing this auxiliary figure is a straightforward similar figures task, one in fact which we had sampled.

The higher order rule of Figure 7 also generates solution rules for even more complex problems, provided we assume the necessary component rules. For example, consider the problem, "Given three noncollinear points A , B , and C , construct a line XY which intersects segment \overline{AC} in the point X and segment \overline{BC} in the point Y , such that segments \overline{AX} , \overline{XY} , and \overline{YB} are all of the same length."

Figure 8



The reader may wish to derive the solution rule for this more difficult problem himself. (Hint: Several recursions are required. For details see the unabridged report.)

DISCUSSION

Summary

In summary, a quasi-systematic method for characterizing heuristics involved in problem solving was proposed and illustrated with compass and straightedge constructions in geometry. Higher order rules, together with corresponding sets of lower order rules, were constructed for the two-loci, similar figures and auxiliary figures problems identified by Polya (1962). First, the two-loci heuristic of Polya was made precise. We saw how decision making capabilities (decisions), and particularly the conditions used to define decisions, play a central role in higher order rules. The similar figures and auxiliary figures heuristics were similarly formulated. We also showed how the two-loci and similar figures higher order rules could be combined to form one higher order rule, which (together with appropriate lower order rules) provides a basis for solving both

14. This involves memory and is not indicated in the flow diagram.

kinds of problems. Finally, a combined two-loci, similar and/or auxiliary figures higher order rule was constructed. This higher order rule allows recursive returns to components of the higher order rule, corresponding to the individual higher order rules, and was considerably more powerful than the others. Its use on some complex problems was illustrated.

Overall, the analyses demonstrated the viability of the analytic method. The higher order rules identified were precise, compatible with the heuristics identified by Polya, and intuitively seemed to reflect the kinds of relevant knowledge that successful problem solvers might have.

The central role played by semantics in the analysis should be emphasized. The meaning of each task was represented by a goal figure $\langle S_1, R_1 \rangle$ representing the given goal situation $\langle S_0, R_0 \rangle$. The relations among, and properties of, the elements of these figures, together with the domains and ranges of individual rules, were reflected directly in the higher order rules. Although little attention was given to the formal representation of semantic features, the goal figures clearly placed powerful constraints on the rules selected at each stage in applying the higher order rules. Representation in terms of some arbitrary (e.g., random) syntax, unconstrained by goal figures, would have necessitated backup capabilities and, in principle, could easily increase the number of possible construction rules at each stage beyond any reasonable computational capability. That is, without the constraints imposed by the goal figures, the number of possible points, arcs, and lines that might be constructed could be almost unlimited. The effect of using goal figures is very much the same as that referred to by Winston (1972) in a recent paper on vision. He argued that although the number of combinatorially possible arrangements of vertex types (Guzman, 1968) is very large, the number of types that yield real figures is much smaller.

Limitations

Nonetheless, the present study has certain limitations which, in principle, could be overcome. First, as in existing state space formulations, all of the higher order operations were limited to compositions of rules. In future research, more attention should be given to other kinds of operations. Generalization, restriction, and selection rules (e.g., Scandura, 1973a), for example, might well be expected to play an important role in problem solving.

There are a variety of ways in which such rules might enter. (a) In discussing the two-loci higher order rule, we have already seen how the scope of a decision (making capability) may be generalized to generate solution rules for a broader range of problems. In particular, we saw how the first decision, which was initially restricted to situations where the desired point X was a given distance from two given points, could be generalized, for example, to allow the point to be the same distance from two given points. It is not hard to envisage a generalization rule by which such shifts might be made. The relationships observed previously between the missing points X , and the similar and auxiliary figures, suggest another kind of generalization involving the identified higher order rules.

(b) There are a wide variety of construction problems which might require the independent derivation of more than one missing point X , similar figure, or auxiliary figure. As a simple example, consider the task of constructing two circles, one of which is to be inscribed in a given triangle and the other, to pass through its vertices (i.e., to circumscribe the triangle). In this case, the problem can be solved by applying the two-loci higher order rule twice. The higher order derivation rule here can be thought of as a generalization of the two-loci rule in which two or more applications (i.e., recursions) may be allowed. One can easily conceive of a simple higher order generalization rule which operates on rules and generates corresponding rules which are recursive. The combined two-loci, similar and auxiliary figures higher order rule is one possible consequence of apply some such higher order rule.

(c) If we had allowed unsolvable variants of the problems considered, truly viable solution rules would have to be appropriately restricted. The solution rule for "constructing a triangle with sides of predetermined length," for example, works only when the sum of each pair of sides of the triangle is greater than the third. A completely adequate solution rule would have to test this possibility. It is possible to conceive of higher order rules, which operate on rules of various kinds together with special restrictions (e.g., the triangle inequality) to generate correspondingly restricted rules.

(d) It is also possible to conceive of three dimensional analogues of compass and straightedge constructions. In this case, the higher order rules would operate on the usual two dimensional construction rules and would generate their three dimensional analogues. For example, a rule for constructing the locus of points equidistant from a given line (i.e., a pair of lines) corresponds to a three dimensional rule which constructs a cylinder about the line.¹⁵

A second limitation is that nowhere did deduction play a role in our analysis. In solving constructions, real people frequently attempt to justify logically the various constructions they make. Constructing a triangle given its three medians, for example, requires that a person know or deduce the fact that the medians intersect at a point two-thirds of the way from each vertex to the opposite midpoint (see footnote 13). To this extent, our analysis is limited and may not adequately reflect human knowledge. Our rules reflect semantics, but not inference. Extension of the proposed analysis to deduction should be a first order of business. It is likely that existing geometry theorem proving systems (e.g., Gelernter, 1959) may be useful in this regard.

A third major limitation of this research is that cumulative effects of learning were not considered: each problem in our analysis was considered as *de novo*. If one wishes to characterize solutions to problems in a given class (e.g., the two-loci tasks) relative to a fixed, self-sufficient set of rules, some fairly complex rules (e.g., the angle vertices rule) must be included. Furthermore, and in many ways more important, such characterizations, at any particular level of analysis in a task domain tend to lack flexibility. The atomic elements are so large, relatively speaking, that there are many intermediate level problems that cannot readily be solved using such rule sets exclusively. Also important from the standpoint of behavioral analysis, it is doubtful that such lower order rules would adequately reflect the knowledge had by most subjects assumed to know the identified higher order rules. Such subjects would almost certainly also know a wide variety of simpler construction rules, even though we might not explicitly include them in a rule set determined by sampling complex problems of the sort we used. Future work is planned which is designed to meet many of these objections.

Future directions

The method of analysis used in the present research is based on Scandura's (1973a) theory of structural learning, more particularly on those aspects of it which deal with competence. The aim of the latter is to specify (hopefully mechanizable) procedures which characterize the knowledge underlying given classes of behaviors (e.g., problem solutions) that one might wish to attribute to an idealized knower. As noted, our approach to this problem involves the invention of finite sets of rules (including higher order rules which may operate on other rules as well as on data elements) which can be applied as indicated for example, to generate problem solutions.

This level of theory, of course, applies only at an analytic level in the sense generative grammars account for language behavior. The relevance of the theory to actual human behavior, or, for that matter, to the design of artificial intelligence systems, depends fundamentally on our ability to specify mechanisms by which such rules are to interact in specific situations, and what effect if any such interaction has on the nature of the rule set itself.

5. Implicit in the above examples is another limitation to which we have indirectly referred previously. Our original analyses were limited almost exclusively to single higher order rules. In no case did we attempt to identify rules which may operate on higher order rules, although our examples make it clear that we could have done so. The rules involved in accomplishing this would be practical rather than theoretical.

The structural learning theory (Scandura, 1973a) is partly concerned with the specification of such mechanisms. The theory rests on the fundamental and widely held assumption that in problem solving people are attempting to achieve some goal. In the simplified version of the theory considered here, the basis mechanism which governs the use of available rules is as follows: (A) The subject tests his available rules (r) to see if one (or more) of them satisfies the given goal situation (i.e., If $S_0 \in \text{Dom } r$ and $\text{Ran } r \subset \text{Goal}$). If so, the subject will apply it. (B) If a subject does not have a rule available for achieving a given goal, then control automatically shifts to the higher level goal of deriving a procedure which will satisfy the original goal. (C) If a higher level goal has been satisfied (that is, if some new rule has been derived which contains the stimulus situation in its domain and whose outputs satisfy the original goal criterion), the derived rule is added to the set of available rules and control reverts back to the previous goal. The third hypothesis allows control to return to lower level goals once a higher level goal has been satisfied. (For more general and rigorously formulated sets of hypotheses see Scandura, 1973a.)

Putting all this together, we see that if an appropriate higher order rule is available when control shifts to a higher level goal, then the higher order rule will be applied and control will automatically revert to the original goal. The subject will then apply the newly derived rule and solve the problem. If the subject does not have a higher order rule available for deriving a procedure that works, then control is presumed to move to still higher levels (e.g., deriving a rule for deriving a rule that works). Although this process is assumed to go on indefinitely in the idealized theory, memory places strict limits in actual applications.

Even this simple assumption provides an adequate basis for generating predictions in a wide variety of problem solving situations. Consider the problem of converting a given number of yards into inches. There are two possible ways in which a subject might solve the problem. The first is to simply know, and have available, a rule for converting yards directly into inches: "Multiply the number of yards by 36." In this case, the subject need only apply the rule according to hypothesis (A). The other way is more interesting, and involves the entire mechanism as described above. Here, we assume that the subject has mastered one rule for converting yards into feet, and another for converting feet into inches. The subject is also assumed to have mastered a higher order composition rule.

In the second situation the subject does not have an applicable rule which is immediately available, and, hence, according to hypothesis (B), he automatically adopts the higher level goal of deriving such a procedure. Then, according to the simple performance hypothesis (A), the subject applies the higher order composition rule to the rules for converting yards into feet and feet into inches. This yields a new composite rule for converting yards into inches. Next, control reverts to the original goal by hypothesis (C) and, finally, the subject applies the newly derived composite rule by hypothesis (A) to generate the desired response.

Moreover, this mechanism provides a basis for an efficient characterization of learning, since, according to hypothesis C, newly derived rules are added to the knowledge base (rule set). Such (additional) rules are in no way distinguished from any others in the rule set; for example, they may serve as component rules in new higher order rule applications. (Also, it should be noted that derived rules may themselves be of higher order and may, thus, be used to satisfy future higher level goals.)

To see how knowledge may cumulate according to this mechanism, let us assume that the learner initially knows rules for converting miles into yards, yards to feet, feet to inches, and the higher order composition rule above. Suppose also that the learner is first presented with the problem of converting miles to inches. In this situation, the learner will fail to solve the problem, since the composition rule we specified above applies only to pairs of rules. (We assume that it does not apply to itself.) However, if the problem of converting yards to inches is presented first, the

subject will solve it as before, and derive a yards to inches rule in the process. Further, if the miles-to-inches problem is then presented, it can be solved using the derived yards to inches rule and the miles to yards rule as components. Although this example is obviously very simple, it does illustrate the potential importance of problem sequence in a growing (learning) system.

Although other investigators have made use of similar notions in varying degrees, the type of mechanism proposed appears to make more general use of rule and higher order rule constructs. Frequently, for example, procedures which are allowed to operate on procedures are not themselves part of the knowledge base; they are viewed as control processes. (In the present case, only the learning mechanism itself acts as a control process.) Nor are newly derived solution procedures often added to the set of available procedures. Newell & Simon (1972, p. 135), for example, allow the Logic Theorist to add proved theorems to an initial set of axioms, but this is essentially at the level of data, upon which proof generation procedures operate, and not at the level of the procedures themselves. Viewing learning as "debugging" (e.g., Minsky & Papert, 1972) or as "means-ends" analysis (Newell & Simon, 1972) is essentially analogous to the introduction of higher order rules except that in these cases implicit restrictions are imposed on the allowable higher order rules.

In any case, most investigations in artificial intelligence have involved some kind of state space representation (e.g., Nilsson, 1971), with problem solving involving some type of search. No generally agreed upon way of representing learning seems to have emerged, however. Sometimes, learning is treated as the modification of parameters in evaluation functions which select 'promising' nodes for expansion (e.g., Samuel, 1959). In other cases, learning systems have been devised to reflect stimulus-response principles in psychology (e.g., Feigenbaum, 1961, Bower, 1972). Where considered by information processing psychologists who have adopted this point of view (e.g., Rumelhart, Lindsay, & Norman, 1972), learning involves the transformation of one state space to another (Scandura, 1973b).

Though the proposed representation may be formally equivalent, it is our belief, based on a variety of studies with human subjects (e.g., Scandura, 1973a), that it is not psychologically equivalent. For one thing, our search for basic psychological mechanisms (e.g., of learning), which reflect commonalities in human behavior, differs in important ways from that in computer simulation, where the essential goal is to parallel overt human behavior in complex instances of problem solving and where the basic mechanisms (e.g., means-ends analysis), therefore, are often judged on more immediately pragmatic grounds.

Irrespective of one's opinion on the issue, the laws which govern the interactions among individual rules are assumed to be fixed once and for all and have potentially important implications for computer implementation. In particular, the fixed mode of interaction would make it possible in principle to modify and/or extend an artificial intelligence system rule by rule, without having to worry about the effects of these changes on other parts. (This latter property appears to some extent to be shared by Newell and Simon's (1972) production systems.)

One of the major complications in current artificial intelligence research is that even minor changes in one part of a system may have unpredictable effects which may require compensating changes elsewhere. The switch to heterarchical systems (e.g., Minsky and Papert, 1972) in which control may shift among individual programs in some predetermined manner, does not appear to alleviate this problem.¹⁶ In contrast to the above mechanism, the mode of control in heterarchical systems may vary from system to system, and worse, from the standpoint of debugging, may interact with the individual programs themselves. In short, the important point for artificial intelligence research is the possible advantage for implementation of a fixed mode of interaction.

Whether or not the mode of interaction is restricted to that proposed here is not the most crucial point. To the extent that artificial intelligence research may

16. See Appendix G.

benefit by taking account of such mechanisms, psychological research aimed at discovering what these mechanisms are would appear to be a first order of business for those interested in human thought. (For a "richer" theoretical mechanism which incorporates memory, see Scandura, 1973a, Ch. 10.)

With the foregoing in mind, an alternative which we are now pursuing is to begin initially with rule sets composed of simpler rules, and to allow these rule sets to grow gradually by interacting with a problem environment.¹⁷ In the present case, only three atomic operators (lower order rules) will be introduced initially: (a) setting a compass to a given radius, (b) drawing a straight line (segment), and (c) using a set compass to make a circle. It is not immediately clear what the higher order rules should be but, presumably, any reasonably satisfactory rule set would include some types of simple composition, conjunction, and generalization higher order rules, together, possibly, with variants of the two loci and other higher order rules identified above. It should be emphasized in this regard that the initial selection of rules would not in itself be sufficient; the choice and sequencing of to-be-solved problems may also be expected to have important effects on both the rate and type of knowledge acquisition. For obvious reasons, computer implementation seems almost essential in this research and is the course we are pursuing.

IMPLICATIONS

Artificial Intelligence

The present research appears to have three general implications for work in simulation and artificial intelligence.

First, the rules we have identified may be implemented relatively easily (some have already been). As such, they would be useful either directly in systems concerned with geometric figures and constructions, or indirectly in research having more encompassing aims as described above.

Second, the results are suggestive of how the construction of at least certain artificial intelligence systems might be partially systematized. In this regard, the topic of compass and straightedge constructions is not nearly as important as is the fact that the analysis serves as a prototype for the proposed method of analysis. At the present time this method is being used to analyze the proofs contained in an experimental algebra I high school text based on axiomatics.

Third, our use of flow diagramming as a mode of representation of individual rules suggests that perhaps such representation might play a somewhat larger role in the exposition of future artificial intelligence research. The routine use of a large number of different and highly technical programming languages is often enough to turn away outsiders (such as ourselves) who might otherwise be interested.¹⁸ The limitations of flow diagrams with regard to memory considerations may be a small price to pay for a more neutral and familiar form of representation. Furthermore, flow diagrams have a flexibility as to level of representation which is not shared by particular programming languages. This makes it possible to more readily represent basic components at a level of atomicity tailored to immediate needs, and to psychological reality (cf. Scandura, 1973a), rather than to basic components determined by some programming language. These comments, of course, apply only to psychological and expository considerations and say nothing of the more strictly technical problems of representation which must be dealt with in computer implementations.

Education

The results of this study also have both long range and immediate implications for education. The promising nature of the results attests to the practicability of the proposed approach as a means of identifying the knowledge underlying reasonably complex kinds of problem solving. In addition to serving as a prototype, the identified rules themselves could be helpful in teaching high school students how to solve compass and

17. See Appendix H.

18. See Appendix I.

straightedge construction problems.

By identifying precisely what it is that students must know (i.e., one possible knowledge base), these rules provide an explicit basis for both diagnosis and instruction. In particular, the methods of analysis formalized by Scandura (1973a) and developed empirically by Scandura and Durnin (1973) and Durnin and Scandura (1973) can be applied directly to assess the behavior potential of individual subjects on the individual rules, including the higher order ones. Operationalizing the knowledge of individual subjects in this way, and comparing this knowledge with the initial competence theory (i.e., set of rules), provides an explicit basis for remedial instruction (Durnin & Scandura, 1973). In effect, each subject can be taught precisely those portions of each competence rule which testing indicates he has not mastered.

Care was taken to help insure that the higher order rules reflect the kinds of ability individual subjects might have, or use. To the extent that the identified higher order rules are unknown to high school students, instruction in these rules ought to facilitate problem solving performance. The diagnostic and instructional efficacy of these higher order rules has been demonstrated in a recent field test (Scandura, Wulfeck, Durnin, & Ehrenpreis, 1974).

The above discussion of how knowledge is acquired through interaction of the learner with a problem environment also has educational relevance. Specifically, by assigning values to various objectives and costs to particular kinds of instruction (or rules), it should be possible to study the problem of instructional sequencing and optimization in a way which is both precise and relevant to meaningful education. We view this as a critically important problem for future research.

REFERENCES

- Amarel, S. On representations of problems of reasoning about actions. In D. Michie (Ed.), Machine intelligence 3. New York: American Elsevier Pub., 1968, pp. 131-171.
- Bower, G. A selective review of organizational factors in memory. In E. Tulving and W. Donaldson (Eds.). Organization of Memory. N.Y.: Academic Press, 1972.
- Chomsky, N. Syntactic structure. The Hague: Mouton Press, 1957.
- Durnin, J. H. & Scandura, J. M. An algorithmic approach to assessing behavior potential: Comparison with item forms and hierarchical technologies. Journal of Educational Psychology, 1973, 65, 262-272.
- Ehrenpreis, W. and Scandura, J. M. An algorithmic base for curriculum construction: A field test in Mathematics. Journal of Educational Psychology, 1974, in press.
- Feigenbaum, E. A. The simulation of verbal learning behavior. Proceedings of the Western Joint Computer Conference, 1961, 121-132.
- Gelernter, H. Realization of a Geometry Theorem Proving Machine. Proc. Intern. Conf. Inform. Proc. Paris: UNESCO House, 1959. Reprinted in E. Feigenbaum & J. Feldman (Eds.) Computers and Thought. New York: McGraw-Hill Book Co., 1963.
- Guzman, A. Decomposition of a visual scene into three-dimensional bodies. Proceedings of the Fall Joint Computer Conference, 1968, pp. 291-304.

- Minsky, M. and Papert, S. Research at the laboratory in vision, language and other problems of intelligence. Artificial Intelligence Progress Report, Memo 252
Cambridge, Mass.: M.I.T. Artificial Intelligence Laboratory, 1972.
- Newell A., & Simon H. A. Human Problem Solving. Englewood Cliffs, N.J.: Prentice Hall 1972.
- Nilsson, N.J. Problem-solving methods in artificial intelligence. New York: McGraw-Hill, 1971.
- Polya, G. Mathematical discovery, Vol. 1. New York: Wiley, 1962.
- Rumelhart, D. E., Lindsay, P. H., & Norman, D. A. A process model for long term memory. In E. Tulving and W. Donaldson (Eds.). Organization of Memory. New York: Academic Press, 1972.
- Samuel, A. Some studies in machine learning using the game of checkers, IBM Journal of Research and Development. 1959, 3 (3), 211-229.
- Scandura, J. M. Deterministic theorizing in structural learning: Three levels of empiricism. Journal of Structural Learning, 1971, 3(1), 21-53.
- Scandura, J. M. Structural learning I: Theory and research. New York: Gordon & Breach, 1973. (a)
- Scandura, J. M. The structure of memory: Fixed or flexible? University of Pennsylvania Structural Learning Series, Report 73, May, 1973. (b)
- Scandura, J. M. & Durnin, J. H. Assessing behavior potential: Adequacy of basic theoretical assumptions. Journal of Structural Learning, 1973, in press.
- Scandura, J. M., Wulfeck, W. H. II, Durnin, J. H., & Ehrenpreis, W. Diagnosis and instruction of higher order rules for solving geometry construction problems. University of Pennsylvania Structural Learning Series, Report 80, 1974.
- Winston, P. H. Summary of selected vision topics. M.I.T. Artificial Intelligence Laboratory, Flash 30, July, 1972.

APPENDICES

Appendix A.--The only really adequate way of determining whether a rule is compatible with human behavior is to effect a behavioral test; that is, to see whether a rule provides an adequate basis for assessing the behavior potential of individual subjects, thereby making it possible to predict the behavior of individual subjects on new instances (of the rule). (The theoretical foundations for such tests have been worked out and tested empirically [Scandura, 1971, 1973a, Scandura & Durnin, 1973; Durnin & Scandura, 1973].) The basic idea is to determine each subject's behavior potential with respect to each rule in an identified rule set, and then to use the theory as a basis for making predictions concerning performance on problems which require interactions among the rules. The closeness of fit between the predictions and observed behavior would provide a direct test of the adequacy of the rule set. A study reported in Scandura (1973a) on rule generalization was of this type. Since this was impractical in the present study, we adopted the weaker and less rigorous criterion of requiring that the rule sets be compatible with our intuition (cf. Chomsky, 1957).

Appendix B.--Strictly speaking, human subjects are presented with statements of problems as stimuli. Throughout this and our subsequent analyses we assume that the subject's initial subgoal is to interpret the goal statement (i.e., determine its meaning). The second subgoal is to solve the problem. In effect, the initial goal is divided into a pair of subgoals to be achieved in order. Our analysis is limited to the second part of this task, and then only on the assumption that there is no further division of the problem into subgoals. We also assume that the given problem statements can be uniformly and correctly interpreted.

Although we do not pursue the question here, we have reason to believe that forming subgoals is closely related to the question of (problem) representation (cf. Amarel, 1968).

Appendix C.--Other representations would probably be more efficient for computer implementation, since graphic systems are relatively complex to implement. For example, some sort of naming system for points, lines, etc. could be devised together with appropriate interpretive routines to identify relations of interest among elements. In fact, the naming system for triangles in common use evolved for just this purpose; names for sides, vertices, medians, etc., if correctly interpreted, carry much information about relative position, intersections, etc.

Appendix D.--In the structural learning theory (Scandura, 1973a), it is assumed that the problem solver automatically tests the solution rule R_S to see if it satisfies a higher level goal condition. That is, is $S_0 \in \text{Dom } R_S$ and $\text{Ran } R_S \subset G$? If the higher level goal is satisfied, control is assumed to revert to the original goal so that R_S will be applied.

Appendix E.--In evaluating alternative rule-based accounts for a given class of tasks, decisions must always be made concerning exactly how the computational load should be apportioned to the higher and lower order rules. Any number of alternatives exist; at one extreme, the lower order rules may do all of the computation, in which case a separate rule would be needed for each type of problem, and, at the other extreme, the component lower order rules may be of minimal complexity with the higher order rule assuming most of the computational burden. The requirement of compatibility with human knowledge, of course, substantially reduces the number of plausible characterizations.

Appendix F.--We do not attempt to spell out the procedures necessary for finding auxiliary figures. However, in all of the sampled auxiliary figures problems, it was necessary to construct a line parallel to some "distinguished" line through some "distinguished" point not on that line. Such procedures also frequently require special knowledge -- for example, that medians intersect at a common point that is $2/3$ of the distance from the respective vertices to the midpoint of opposite sides. Such knowledge is frequently logically deducible, but for our purposes, may be represented in terms of simple "associations" for example, between triangles with their medians and the common intersection property.

Appendix G.--Scandura's (1973b) comments regarding relationships between the structural learning mechanism, and the notion of heterarchical control in systems of artificial intelligence (Minsky & Papert, 1972) may be relevant here.

"For a time artificial intelligence systems were viewed as wholes, as frequent complex programs. As work in the area progressed, the difficulties of building upon earlier work became increasingly clear because of the close interrelationships among various parts of such systems. To overcome this limitation, heterarchical, or modular planning has been used (e.g., Minsky & Papert, 1972). Heterarchical systems consist of sets of programs (modules) pertaining to syntax, semantics, line detection, and so on, together with a heterarchical executive which switches control among these "modules in accordance with a predetermined plan.

"Modules in heterarchical systems correspond essentially to rules in the structural learning theory; the executive control structure corresponds to the basic mechanism. There is, however, an important difference between the two. In heterarchical systems, the basic goal is pragmatic. Such systems make it easier to modify and build upon previous work. No one seriously means to imply that heterarchical control reflects the way people perform, although in developing artificial intelligence systems intuitive judgements are sometimes made with this in mind.

"In contrast, the structural learning mechanism is assumed to be built into people (presumably from birth); it is not learned and need not be taught. While the rules a person knows may increase from time to time, the mechanism is assumed to remain constant.

"This is a strong claim, something which no responsible person would make concerning executive systems currently used in heterarchical systems. Among other things, it is very unlikely that an existing control system would be useful in systems other than the one for which it was designed. It is my contention that benefits might accrue in artificial intelligence and, of course, in simulation if structural learning like control structures were used [pp. 42-43]."

Appendix H.--Such rule sets have been called innate bases (Scandura, 1973a, Ch. 5). In general, innate bases lack the immediate, direct computing power of comparable rule sets composed of more complex rules but, theoretically at least, can grow to become more powerful.

Appendix I.--We realize, of course, that some computer specialists may not take our suggestion very seriously. We, however, find the work in simulation and AI highly suggestive for our own studies and hope in the interest of interdisciplinary communication that some readers may be moved more in this direction.

DIAGNOSIS AND INSTRUCTION OF HIGHER ORDER RULES
FOR SOLVING GEOMETRY CONSTRUCTION PROBLEMS

Joseph M. Scandura, Wallace H. Wulfeck II, John H. Durnin, and Walter Ehrenpreis
University of Pennsylvania

The viability of systematic analyses of real problem domains in terms of rules and higher order rules has been demonstrated by Scandura, Durnin, and Wulfeck (1974) and Durnin and Scandura (1973). The practical importance of such analyses in such areas as artificial intelligence and education, however, is still an open question. Although an attempt was made to insure in those analyses that the rules identified reflect human knowledge, it was not demonstrated there that they do. Rigorous tests of this thesis require experimental data. Furthermore, even if the rules do turn out to be compatible with what human subjects are likely to know, it is not clear whether, and to what extent, instruction in the higher order rules will result in improved problem solving performance.

The research reported in this paper deals with these questions in education with respect to the geometry construction analyses in Scandura et al. (1974). Specifically, the purpose of this research was to determine the extent to which: (1) the basic higher order rules identified in the analyses are compatible with the knowledge had by a group of average ability teenagers, (2) instruction in the higher order rules facilitates performance on geometry construction problems and (3) instruction in some higher order rules influences (i.e., facilitates or hinders the learning and/or use of) subsequent ones.

A total of four paths of the two loci, similar figures, and auxiliary figures higher order rules were considered in the study. According to the analysis given in Durnin and Scandura (1973), the central question in determining behavioral compatibility is whether the paths of the higher order rules act in (near) atomic fashion. The major task, here, is to determine whether the ability, or lack thereof, to appropriately combine available lower order rules in one problem situation is reflected in other problem situations of the same type. This ability can be determined either directly in higher order task situations, where the subject is required to derive solution procedures for given problems, or indirectly, as we have done below, by asking the subject to actually solve problems (i.e., derive solution procedures and then use them).

In theory, when paths of a rule act in atomic fashion with a given population of subjects, inadequacies determined through testing can be overcome through direct instruction on the paths involved. If a person can solve problems whose solution rules require one path of the similar figures higher order rule, but not problems involving the other path, for example, then instruction presumably would be required only on the latter path.

Previous research provides more or less definitive guide lines on how to proceed and what to expect with regard to the first two questions. Thus, the research reported in Durnin and Scandura (1973) suggests that introspection as to how one actually solves a class of problems often results in the identification of procedures which appropriately partition the class (into equivalence classes). Although we know of no empirical research in the literature which bears on the second goal, members of the MERG group have developed materials for diagnostic testing and remediation in the arithmetical skills which are based directly on these ideas (for details, see Scandura, 1972). Formal data concerned with instruction have not been obtained but informal tryouts attest to the effectiveness of the remedial materials.

With regard to the third goal, very little can be said on the basis of available evidence. The fact that various paths of the higher order rules share many steps in common suggests that there might be positive transfer from one path to another. Thus, having learned one path, there is apt to be less to learn on subsequent ones so that learning them will require less time. On the other hand, one could argue that

similarities among the higher order paths could result in interference. In attempting to generate solution procedures to given problems, the subjects might use the wrong paths.

The present study was designed both to provide answers to the first two questions and to determine transfer and relative learning efficiency resulting from prior training.

METHOD

Tasks and Materials

Four paths of three higher order rules were considered in the study. Figure 1 depicts the two loci rule, and Figure 2 depicts the two paths of the similar figures rule in Scandura et al. (1972). Path 1 (restricted similar figures) involves steps J, K, L, and M, and path 2 involves steps A, B, C, D, F, G, H, and I. The fourth path involves the auxiliary figures rule (Scandura et al., 1974). For instructional purposes, the decisions and operations of each path were written, respectively, as simple lists of questions and imperative statements.

FIGURE 1

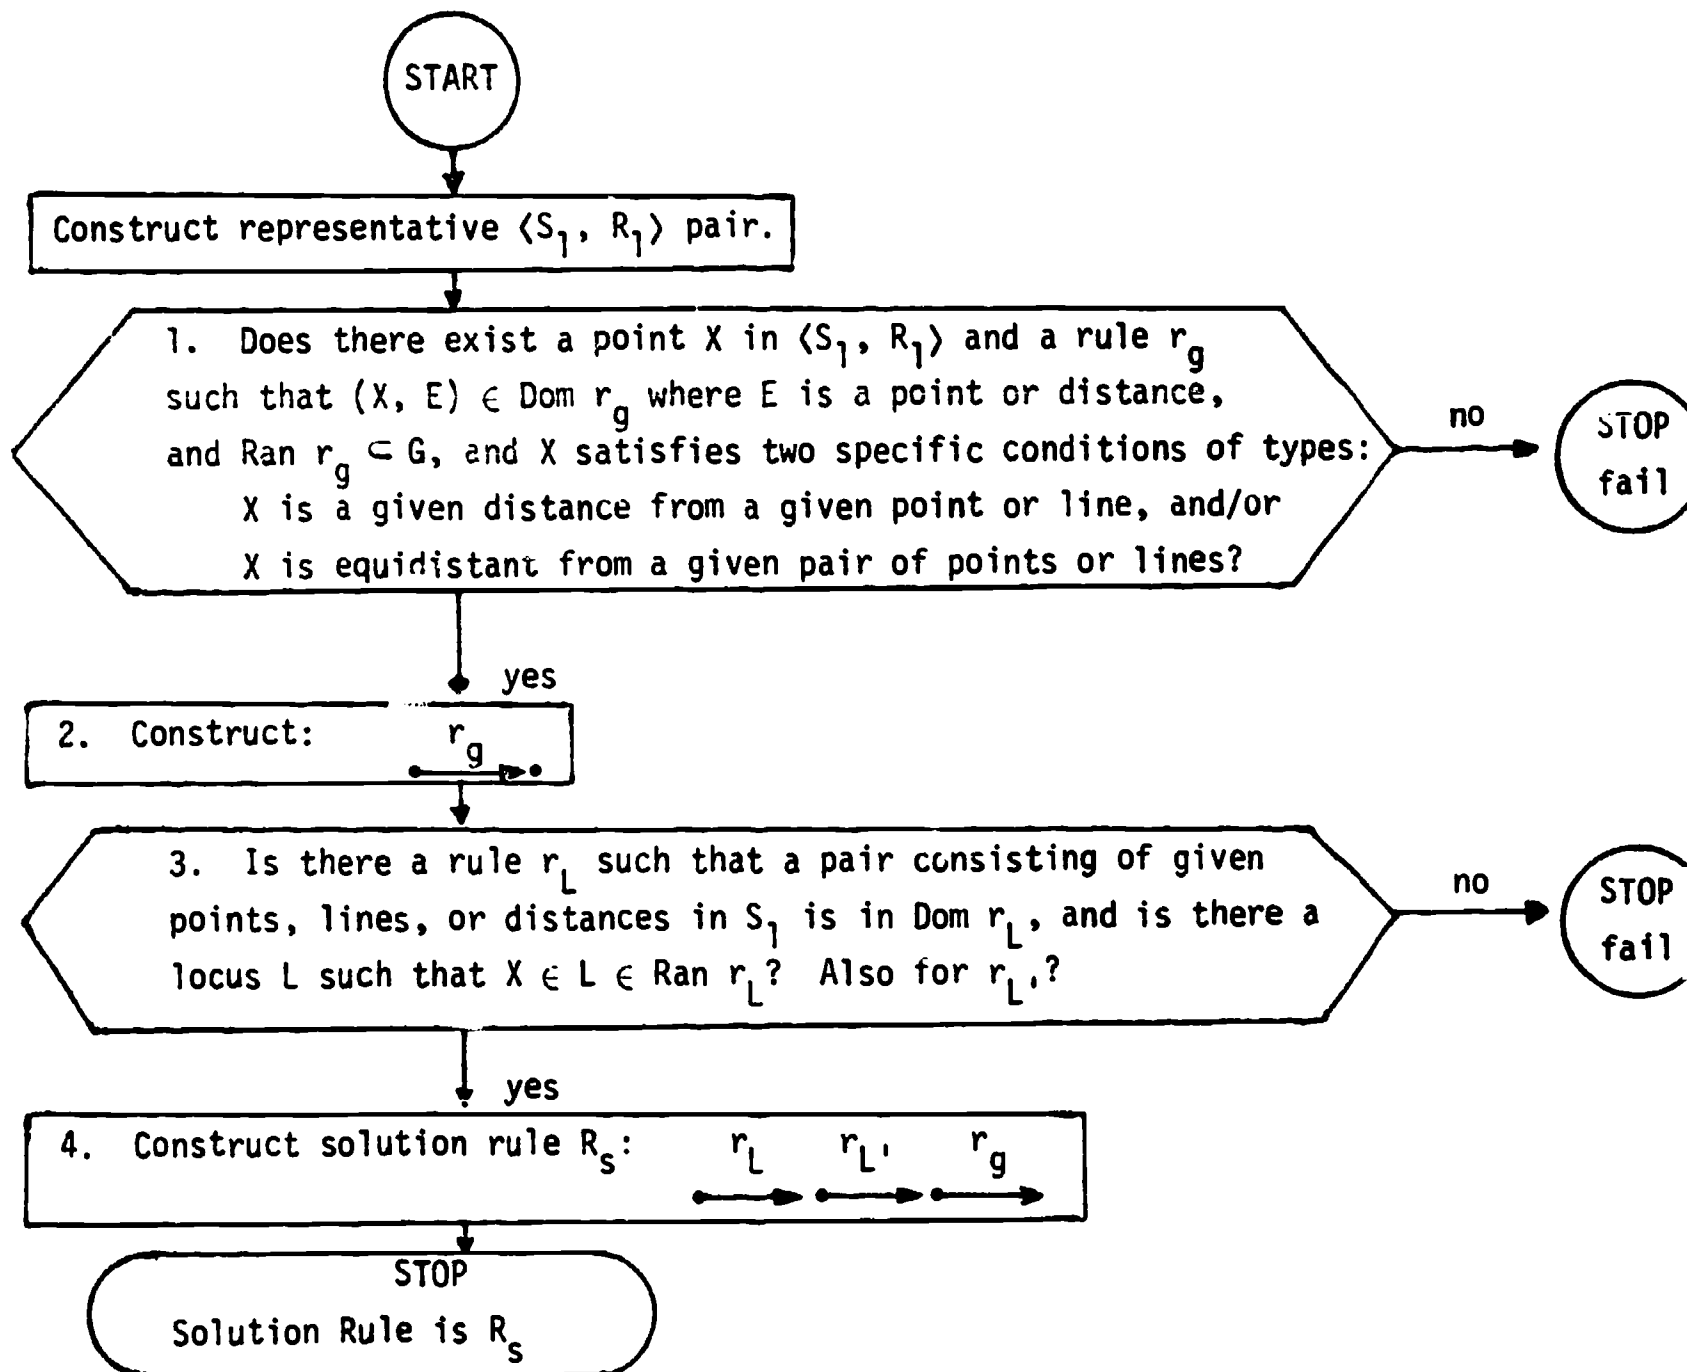
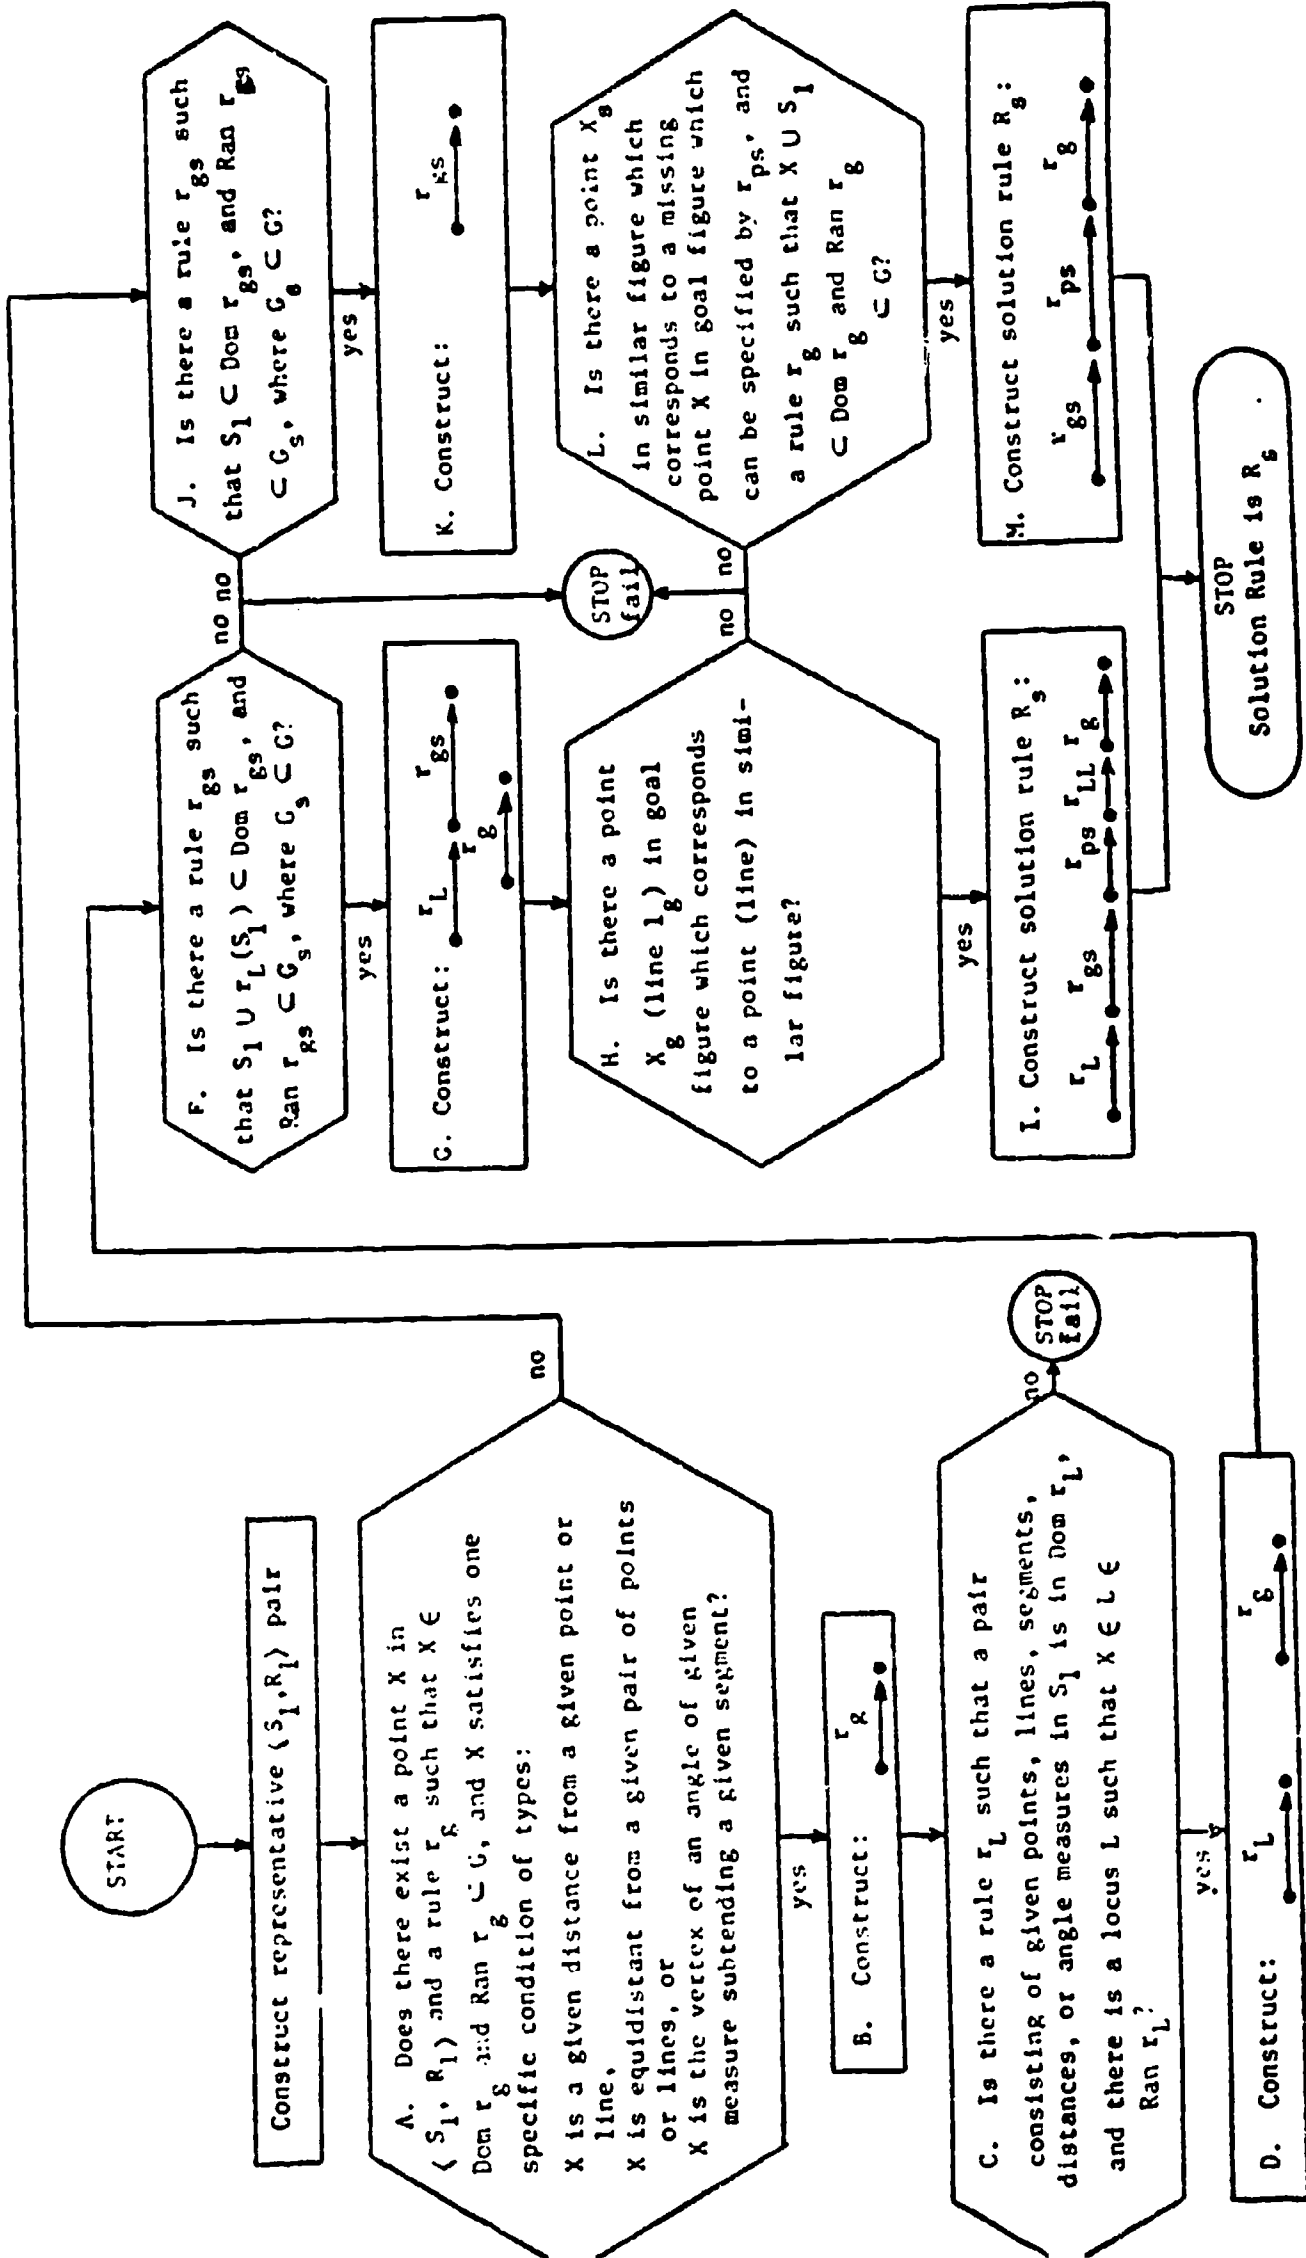


FIGURE 2



The experimental tasks were 13 geometry construction problems (see Table 1) taken from Scandura et al. (1974). These problems may be categorized according to which of the four paths of the higher order rules may be used to generate an appropriate solution rule. The two loci higher order rule constituted one path; the similar figures rule yielded two paths (restricted similar figures and similar figures); the auxiliary figures rule yielded one.

TABLE 1
Geometry Construction Tasks

Test	Problem Number	Solution Type	Problem Statement
Pretest I	1	Two loci	Given a line and a point not on the line, and a radius R , find a circle having the given radius R , which is tangent to the line, and passes through the given point.
Pretest I	2	Path one similar figures	Given angles B and C and the altitude H_a , construct the triangle.
Pretest I	3	Two loci	Given side a , and the median M_a , and the height H_a , construct the triangle.
Pretest I	4	Path one similar figures	Given angles B and C and the angle bisector D_a , construct the triangle.
Posttest I	5	Two loci	Given sides a and b and the median M_a , construct the triangle.
Posttest I	6	Path one similar figures	Given angles B and C and side b opposite angle B , construct the triangle.
Posttest II	7	Two loci	Given two intersecting lines and a radius R , construct a circle with radius R tangent to the two given lines.
Posttest II	8	Path one similar figures	Given angles B and C and the median M_a , construct the triangle.
Choice Test	9	Two loci and Path one similar figures	Given right triangle ABC with right angle at B , inscribe a square in it such that two sides of the square lie on the legs (AB and BC) of the triangle and the fourth vertex of the square (the intersection of the other two sides) is on AC .
Pretest II	10	Path two similar figures	Given two intersecting lines m and n and a point A not on either line, construct a circle tangent to lines m and n which passes through point A .
Pretest II	11	Path two similar figures	Given two intersecting lines m and n and a point P on line m , construct a circle whose center is on line m , which passes through point P and is tangent to line n .
Posttest III	12	Path two similar figures	Given line m and points A and B on the same side of line m , construct a circle tangent to line m which passes through points A and B .
Posttest III	13	Auxiliary figures	Given sides a and c and the altitude H_b , construct the triangle.

The lower order rules needed (in addition to the higher order rules) for solving the experimental problems are shown in Table 2. For purposes of instruction, these rules were refined to the level of actual compass settings and placements. Each consisted of a sequence of operations to be performed. Accompanying sketches illustrated the result of each operation.

TABLE 2

Lower Order Rules

1. <u>Circle rule</u>	Construct the locus of points at a given distance from a given point.
2. <u>Median locus circle rule</u>	Construct the locus of points at a given distance from the midpoint of a given segment.
3. <u>Point-line circle rule</u>	Determine the distance between a given point and a given line and then construct the locus of points at the obtained distance from the given point.
4. <u>Parallel line rule</u>	Construct the locus of points at a given distance from a given line.
5. <u>Angle bisector rule</u>	Construct the locus of points equidistant from two given intersecting lines.
6. <u>Triangle rule</u>	From a point not on a given line segment, draw segments to the endpoints of the given segment (i.e., construct a triangle given a side and an opposite vertex).
7. <u>Perpendicular bisector rule</u>	Construct the locus of points equidistant from two given points.
8. <u>Similar triangle rule</u>	Construct an arbitrary triangle from a pair of given angles, and construct on it parts corresponding to other given segments.
9. <u>Goal triangle rule</u>	Construct a triangle having some part a given length similar to a given triangle with a corresponding part.
10. <u>Point of similarity rule</u>	Select a point of intersection of two lines through corresponding points of goal and similar figures as the point of similarity, then construct a line through the point of similarity and a point on the similar figure, to intersect the goal figure at a corresponding point, from which the goal figure may be constructed.
11. <u>Similar square rule</u>	Construct an arbitrary square in a right triangle with two of its sides contained in the legs of the triangle.
12. <u>Goal square rule</u>	Given a right triangle and a point on its hypotenuse, construct a square with that point as one vertex, such that its two opposite sides are contained in the legs of the triangle.
13. <u>Similar circle rule</u>	Construct an arbitrary circle with its center on one line and tangent to another line.

*The fact that there were 13 tasks and 13 lower order rules is strictly happenstance. The only connection is that the senior author was married on August 13 and is in the 13th year of marriage.

Descriptions of all problems and rules were reproduced on 21.59 cm. x 27.94 cm. (8 1/2" x 11") paper. Each problem appeared on a separate page so that constructions could be done on that page. The 13 problems were arranged into six separate tests as shown in Table 1.

The instructional materials were arranged into seven training booklets. Booklet 1 contained lower order rules 1-10 and a sample task for each. Booklet 2 contained 10 review tasks for the rules of Booklet 1. Booklet 3 contained path 1 of the two loci higher order rule, with the two Pretest I, two loci problems as practice. In parallel fashion, Booklet 4 contained path 1 of the similar figures higher order rule with the two Pretest I, similar figures problems as practice. Booklet 5 contained lower order rules 11 and 12 of Table 2 with corresponding practice tasks and Booklet 6 contained Rule 13 from Table 2 with a practice task. Booklet 7 contained the second path of the similar figures higher order rule along with the two problems from Pretest II as practice.

Pencils, compasses, and straightedges were available where subjects did not provide their own.

Subjects, Design, and Procedure

The subjects were 30 Trenton State College students enrolled in an undergraduate college geometry class.

A repeated measures design was used. The first phase involved lower order rule training (Booklets 1 and 2) and Pretest I. Its main purpose was to obtain information regarding the adequacy of the two loci higher order rule and path one of the similar figures rule as a basis for assessing the (higher order) behavior potential of subjects. A secondary purpose was to obtain success or failure profiles, so that the subjects could be stratified before assignment to experimental groups.

The first meeting with the subjects occurred during a regularly scheduled 75-minute class period. One instructor and two experimental assistants were available to help the subjects and to evaluate their work. They were given Booklet 1 and instruction on lower order rules 1-10 contained in it. The steps of each rule were read aloud and the corresponding constructions were performed on the blackboard. Each subject then completed the corresponding practice problem.

During a second regular class meeting, the subjects were given the practice problems in Booklet 2 and were required to perform at least one correct construction for each of rules 1-10. Achievement of this criterion level was verified by one of the experimenters. As soon as they reached criterion, individual subjects were given Pretest I. All subjects were instructed to attempt all problems in Pretest I before the period ended; no subject "ran out of time." Pretest problems were scored "passed" if a correct solution figure was constructed. Minor deviations ("compass errors") were allowed. Each pretest was scored individually by three experimenters; there was no disagreement. The pretest results were used to stratify the subjects as shown in Table 3.

TABLE 3
Pretest Results

Similar Figures Problems	Two Loci Problems			Σ
	Passed Both	Passed One	Passed None	
Passed Both	6	1	0	7
Passed One	2	1	2	5
Passed None	2	6	10	18
Σ	10	8	12	30

On the basis of the pretest results, the subjects were randomly divided into two groups of 15 each, the two-loci-then-similar-figures (TS) group and the similar-figures-then-two-loci (ST) group, with the constraint that each of the cells in Table 3 was split evenly. (The two "singleton" subjects were placed in different groups.) Individual or small group sessions were arranged with each subject for all subsequent training and testing. Throughout the experiment, each subject retained all instructional materials, but not completed tests.

During the remainder of the study, instruction was provided on three of the four higher order paths and performance was measured on both within and extrascope problems. At the third meeting, the TS subjects were given Booklet 3, path one of the two loci higher order rule, and instruction in how to apply the rule. Specifically, they were shown how to determine whether particular lower order rules from Booklet 1 (which was available) were relevant to solution, and if they were, how to combine them so as to generate solution rules for the Pretest I, two loci problems (1 and 3). No actual constructions were performed. (One TS subject failed to attend this or any other instructional session and was dropped from the study.) The subjects in group ST received Booklet 4, path one of the similar figures higher order rule, and instruction on the application of that rule using problems 2 and 4 from Pretest I. After three subjects had been trained, the instruction was modified slightly so that additional emphasis was given to the stopping decisions (i.e., to conditions where the rule did not apply).

Immediately following instruction each subject was given Posttest I. The subjects in both treatment groups received exactly the same problems. Booklet 1 containing statements of rules 1-10 was available throughout. Also Booklets 3 and 4 containing the higher order rules were available to the subjects in groups TS and ST, respectively. Following Posttest I, one subject in the ST group became ill (no causal relationship implied) and had to be dropped from the experiment. (The two subjects dropped from the study had both failed all pretest problems, and had been assigned to different training groups.)

At the fourth meeting, those subjects who had received the two loci training received path one, similar figures training and vice versa. Instruction was given exactly as before. Posttest II paralleled Posttest I and followed immediately after training. Booklets 1, 3 and 4 were available to the subjects throughout the testing.

At each subject's fifth meeting, he was given Booklet 5 containing two new lower order rules (11 and 12), and training proceeded as with Booklet 1. With all previous training booklets available, the subjects then took the Choice test problem. (This problem could be solved by either of the two higher order rules on which the subjects had been trained.)

Next, at the sixth meeting, subjects were trained on the lower order rule (13) in Booklet 6. With this rule and all previously learned rules also available, the subjects then took Pretest II. The purposes of Pretest II were similar to those of the first Pretest, but dealt with the second path of the similar figures higher order rule. (At this point, two additional subjects who were failing the course, dropped out of the study. The remaining 26 subjects completed the experiments, 13 in each group.)

Finally, at the seventh meeting, each subject was trained as before on Booklet 7, the second path of the similar figures higher order rule, using the problems in Pretest II. After training, the subjects were given Posttest III. One problem of Posttest III was within the scope of the second similar figures path; the other was an auxiliary figures problem not solvable by using any of the three higher order paths on which instruction was provided.

Approximate times required by each subject were recorded for each session of the experiment.

RESULTS AND DISCUSSION

Assessment Results

Pretests I and II contained a total of six problems grouped on an a priori basis according to their solvability via the three higher order paths on which training was provided. To test the behavioral atomicity of the identified higher order rules, contingencies among within-class (path) problems were examined. Table 4 presents the Pretest I and II results on the three classes of problems. On Pretest I, the subjects' performance on the first two loci problem (problem 1 from Table 1) was significantly correlated with performance on the second two loci problem (problem 3) (Fisher's exact probability = .00485; one tailed), and similarly for the path 1, similar figures problems (2 and 4) (exact probability = .00165; one tailed) and, on Pretest II, for the path 2 problems (10 and 11) (exact probability = .00794; one tailed).

TABLE 4
Results of Pretests I and II

	Pretest I				Pretest II	
	Two loci		Path 1 of similar figures		Path 2 of similar figures	
Problem 1	Problem 2		Problem 2		Problem 2	
	Pass	Fail	Pass	Fail	Pass	Fail
Pass	10	1	7	3	2	1
Fail	7	12	2	18	0	25

These results strongly suggest that the identified paths, both collectively and individually, acted in atomic fashion for the experimental subjects. Some of the deviant cases, furthermore, are due to two particular subjects who initially were obviously uncooperative but later applied themselves. Nonetheless, the relatively large number (5) of remaining "fail-pass" cases on the two-loci problems requires some discussion. In particular, this result suggests the possible desirability of further refinement of the two-loci higher order rule into a larger number of distinct paths. This would require analysis of the atomic operators and paths in terms of sub-operators and sub-decisions and, thereby, substitution of a number of paths with more limited domains for the original path. Because the various decisions of this rule involve disjunctions of properties, a basis for such refinement follows directly. The second decision making capability of the two loci higher order rule, for example, refers to a disjunction (A or B or C) of properties, any one of which, if satisfied, is sufficient to direct a computation to a particular sub-operator. With some subjects at least, it is certainly possible that the ability to decide, say, on whether there is a rule containing a point and a line in its domain is independent of the ability to decide on whether a domain contains segments or angle measures. In this case, the ability to solve the problem involving property A would say nothing as regards the ability to solve a problem involving property B, as was the case with problems 1 and 3. In effect, such refinement would not only be consistent with our results but would follow directly from our analysis.

Instructional Effectiveness

Over the entire experiment there were 51 cases where subjects failed all pretest problems from a given class prior to training on the path corresponding to that class. On posttests immediately following such training, new problems from the same classes were solved in 45 (88.2%) of the 51 cases. (The 95% confidence interval for this percentage is 79.3% to 97.1%.) In addition, there were 14 cases where subjects had solved only one pretest problem in a given class. The (new) posttest problems were solved in 13 of those cases.

Table 5 summarizes the results from Posttests I, II, and III on problems for which training immediately preceded testing, arranged according to the number of pretest problems passed.

TABLE 5
Results on Posttest Problems Within Scope of
Immediately Preceding Training

Number of pretest problems within scope of training passed prior to training	N	Number of Ss passing within scope problem after training
Posttest I		
0	14	11
1	6	6
2	9	9
Posttest II		
0	13	12
1	7	6
2	8	8
Posttest III		
0	24	22
1	1	1
2	1	1
Total Cases		
0	51	45
1	14	13
2	18	18

After training on either the two loci higher order rule or path one of the similar figures rule, 11 of the 14 subjects who failed both Pretest I problems in a given class succeeded on the corresponding Posttest I problem (binomial, $p < .05$). Not surprisingly, all six subjects who solved one Pretest I problem and all nine subjects who had solved both Pretest I problems also solved the Posttest I problem.

On Posttest II, 12 of 13 subjects who failed both pretest problems in a given class succeeded after training on the higher order rule for that class. Of seven subjects who solved one pretest problem, six solved the corresponding Posttest II problem after training. All eight subjects who solved both pretest problems succeeded following training. On Posttest III, 22 of 24 subjects who failed both Pretest II problems succeeded following training. Two other subjects, who had solved one or both pretest problems, also succeeded on the posttest problem.

Overall, there were seven cases of posttest failure out of 83 cases where success was expected. Five of the seven discrepancies occurred on Posttests I and II. After failure, these five subjects were retrained on the respective higher order paths, and retested. All five succeeded on the second trial.

Inspection of the seven individual problem attempts which resulted in failure showed that, in six of those cases, mistakes occurred at points corresponding to disjunctive decision points in the two loci rule or in path 2 of the similar figures rule. As noted earlier, disjunctive decisions may be broken up (refined) to form separate paths. This suggests that more explicit attention to the molar nature of such (disjunctive) decisions might possibly have reduced even the small number of inconsistencies noted.

Instruction in the higher order rules was not only effective but also was relatively efficient. Subjects were able to solve relatively complex construction problems, once they knew the component rules involved, after only about 75 minutes of higher order rule training. Instruction on the lower order rules took an average of about 100 minutes. In all, less than three hours of actual instruction was required.

Sequence Effects

In addition to the positive assessment and instruction results, a number of interesting sequence effects were found. On Posttest I, positive transfer to problems, for which training had not been provided, occurred in approximately 40 per cent of the cases. Of 13 subjects who had failed both Pretest I problems in the class for which no training was given, five solved the corresponding Posttest I problem. (Two subjects solved the two loci problem; three solved the path one, similar figures problem.) Three of seven subjects who had solved one of the untrained problems on Pretest I solved the Posttest I problem. (All three solved the two loci problem after path one, similar figures training.) On the other hand, one subject who had previously solved both Pretest I two loci problems failed the Posttest I two loci problem after being trained on the path one, similar figures rule. Apparently, this was due to his misunderstanding of the posttest instructions; the subject thought he was required to use the trained higher order rule (which was inadequate). After this misapprehension was corrected, the subject was retested and passed the problem. (It was at this point that more stress in the instruction was placed on when a higher order rule would not work -- i.e., when to stop.)

The results on Posttest II, restricted to the problem for which training had been given prior to Posttest I, suggests that training on the second higher order rule did not interfere with earlier training. Twenty-seven of the 28 subjects who completed Posttest II solved the problem corresponding to the first-trained rule, after second training. The remaining subject passed the problem on a second trial.

This lack of interference was also reflected in performance on the Choice problem, which could be solved using either the two loci or path one similar figures rules. All 28 subjects solved the problem. There were 17 two loci solutions and 11 path one similar figures solutions; this difference was not significant. Furthermore, there were no significant differences as to solution preference due to order of training. Group ST had eight and six, respectively. This suggests that when two (or more) rules are available at the time of problem solving (as they were), selections do not depend on when the rules were originally learned.

It is also of interest to note that the higher order rule selections after learning were in approximately the same 3:2 ratio as observed success on Pretest I, where 18 and 12 subjects, respectively, solved one or both of the two loci and similar figures problems. This observation suggests that subject preference after learning is somehow related to subject likelihood of having learned two loci-like higher order rules prior to training. Our data are inadequate to determine why this is so but it could have something to do with the involvement of similar, previously learned selection rules (cf. Scandura, 1974).

Positive transfer was also found from training on one higher order rule to the next. When given first, the two loci training required an average of about 32 minutes (25-60 min. range) and the similar figures training, 31 minutes (20-65 min. range). When the training came second, the corresponding times were 21 minutes (15-45 min. range) and 28 minutes (18-60 min. range). The third training session on path two of similar figures higher order rule required only about 18 minutes (10-30 min. range), even though the relatively large number of failures on the corresponding Pretest I problem suggests that this path was more difficult than the others. Overall, then, ignoring the particular training involved, first training took an average of 32 minutes; second training, 25 minutes; and third training, 18 minutes. Differences among these means were highly reliable ($F_{2, 47} = 34.70, p < .005$). Individual comparisons of first and second training and second and third were also reliable ($t_{25} = 4.85, P < .001$ and $t_{25} = 4.14, P < .001$, respectively).

Rather surprisingly, performance on the auxiliary figures problem, which could not be solved using any of the three trained higher order rules, depended on the sequence in which the higher order rules were learned ($\chi^2 = 7.58, df = 1, p < .01$). Eleven of the 13 TS subjects solved it while only three of the 13 ST subjects succeeded. No other relation was observed between the results on this problem and on any of the previous ones.

It is impossible to say with any certainty the source of this rather striking sequence effect. One possibility is that, in attacking the problem, subjects may have tended to select the first higher order rule on which they were trained. In this case, the TS subjects could have had an advantage because the auxiliary figures problem may be solved by repeated application of a variant of the two loci rule. Equivalently, it is possible that the subjects combined the higher order rules into more encompassing rules (see the combined two loci-similar figures rule in Scandura et al., 1974) as they were learned. If so, those subjects who tried the two loci path of the combined rule first, most likely the TS subjects, would again have an advantage, especially if the effects of limited memory are taken into account (cf. Voorhies & Scandura, 1973).

Both of these explanations, unfortunately, imply differential solution type preferences on the choice test problem. Since no such effect occurred, some alternative accounting seems necessary. One plausible explanation stems from the fact that the similar figures and auxiliary figures higher order rules may be regarded as progressive generalizations of the two loci rule. That is, all of the higher order rules begin by identifying constructable elementary figures upon which further operations may act to generate a goal figure. In the two loci rule, the elementary figure is "the missing point X." In the similar figures rule, the elementary figure is more general; it is no longer a "degenerate" point, but is still constrained by similarity. Finally, in the auxiliary figures rule, the elementary figure is arbitrary. (For details, see Scandura et al., 1974.)

Because the TS subjects were taught the procedures in a "natural" order of generalization, while the ST subjects were not, the former may have been more likely to have "induced" a generalization procedure. More specifically, the TS subjects may have learned a "higher, higher order rule" for making generalizations. Such a rule could have been used to derive some form of auxiliary figures higher order rule, which in turn would have allowed derivation of an adequate solution rule.

CONCLUSIONS AND IMPLICATIONS FOR MATHEMATICS EDUCATION

In view of the clarity of these results, it would appear that the identified higher order rules can be used effectively and efficiently both to diagnose difficulties subjects are having with geometry construction problems and to provide instruction in how to solve such problems. Furthermore, training on prior higher order rules seemed to facilitate the learning of later ones.

This is not the first time that beneficial effects have been found for instruction in heuristics. Ennis et al (1969), for example, found that training subjects in general heuristics such as means-ends analysis and planning improved mathematical problem solving. At the present time, Hatfield* also has a study under way in which he is trying to facilitate performance in mathematical problem solving by the informal (clinical) introduction of heuristics.

Although the present study is based on more rigorous and exacting analyses of heuristics in terms of higher order rules (as well as lower order ones), it must not be thought that these varying kinds of studies are incompatible. The present research is completely neutral as regards how information is to be imparted. Motivating the child to learn and the actual mode of presentation in the classroom is up to the teacher's judgement. However, because our higher order rules have been so explicitly characterized, instruction in whatever form (including diagnosis of individual sources of difficulty) is potentially more efficient and feasible than with loosely formulated heuristics.

Although a considerable degree of transfer was evident from training on one higher order rule to another, it is still an open question as to whether explicit instruction in higher order rules (by whatever means) also helps the learner develop new "heuristics" on his own. Earlier, of course, Koughhead and Scandura (1968) found that "what is learned" in making simple discoveries can be presented in expository form with equivalent results. In the present case, however, the task of identifying "what is (to be) learned" is far from trivial and, initially, may require more informal, inductive methods (cf. Lowerre and Scandura, 1974). Whatever the answer, there is certainly no reason why the teacher might not encourage discovery (of higher order rules) in addition to whatever explicit training is provided. Indeed, one good teaching strategy would appear to be to present a variety of situations where learners are required to discover higher order rules. Even in the present study, the higher order rules were not taught explicitly as formal (some would say "rote") procedures. Representation of the rules as flow diagrams simply made the experimenters more aware of exactly what it was that was to be taught.

In spite of the positive nature of these results, it should not be forgotten that they deal primarily with the question of how subjects perform in particular problem solving situations given what (rules) they know on entering into the situation. Any complete prescription for problem solving instruction must deal in detail with the course of solving whole classes of problems. Our findings concerning the sequential effects of instruction on higher order rules has demonstrated the importance of such study, and is one step in this direction, but it is a small one indeed. Consider the complications introduced in considering a continually changing set of lower order rules (as learning progresses), not to mention the difficulties in attempting to explicate precisely the source of the sequential effects we observed with the higher order rules. Nonetheless, we are optimistic concerning the progress that might be made in this direction, and considering the obvious implications for mathematics education, to use a time-worn phrase, "we had better begin."

*Personal communication.

References

- Durnin, J. H. and Scandura, J. M. An algorithmic approach to assessing behavioral potential: Comparison with item forms and hierarchical technologies. Journal of Educational Psychology, 1973, 65, No. 2, 262-272.
- Ennis, R. H., Finkelstein, M. R., Smith, E. R., and Wilson, N. H. Conditional logic and children. Ithaca, N.Y.: Cornell Critical Thinking Project, 1969.
- Lowerre, G. and Scandura, J. M. Conceptually based development of individualized materials for critical thinking based on logical inference. Reading Research Quarterly, 1973-74, IX, 186-205.
- Roughead, W. G. and Scandura, J. M. "What is learned" in mathematical discovery. Journal of Educational Psychology, 1968, 59, 283-289.
- Scandura, J. M. Rule selection, problem definition (subgoal formation), and retrieval in problem solving. University of Pennsylvania Structural Learning Series Report 79, 1974.
- Scandura, J. M. Structural learning I: The theory and empirical research. New York: Gordon and Breach Science Publishers, 1973.
- Scandura, J. M. A plan for the development of a conceptually based mathematics curriculum. Instructional Science, 1972, 247-263.
- Scandura, J. M., Durnin, J. H., Wulfeck, W. H. II. Higher order rule characterization of heuristics for compass and straightedge constructions in geometry. Artificial Intelligence, 1974, in press.
- Voorhies, D. and Scandura, J. M. Determination of memory load in information processing, 1973.

INDIVIDUALIZATION OF INSTRUCTION AND MULTIDISCIPLINE STRUCTURAL LEARNING

Robert J. Starr
University of Missouri

The force generated by student, public and governmental sources in calling for systematic approaches to teaching-learning, accountability and relevance have resulted in the structuring of educational environments for individualized instruction. Unfortunately such "innovations" are often produced as a stop-gap measure to placate negative publicity; hence, are void of theoretical or research foundations.

Individualized instruction is seen by many as innately superior to conventional teaching. Some view individualization as a philosophy rather than a set of techniques; thus, individualization occurs only when an instructor embraces the philosophy and actively meets pupil differences within the classroom (Kinsella and Kinsella, 1972). The result of having a philosophy of individualization is the development/collection of many teaching-learning alternatives for singular concepts. When we individualize, we recognize a person's unique qualities and we build on strengths and the potential found within each pupil; thus, to individualize is to humanize learning (Bagby, 1974).

Philosophically then, the basis for individualization appears to fall within the realm of idealism. One cannot realistically expect to move toward complete individualized instruction because of current educational constraints and yet the goal of individualization remains before us. Instructors seeking to move toward this beacon must plan teaching-learning alternatives which build on the individual and add to his/her self worth. This focus on the person seeks to free potential which has been bound up due to the lock-step nature of conventional (traditional) schooling.

In the last few years a rather uncomplicated model has appeared and is being widely used for categorizing individualized instruction (Edling, 1971). This model, see Figure 1., serves as a basis for matching goals of the school with the goals and strengths of various systems of individualized instruction.

		Objectives	
		School Determined	Learner Selected
Media	School Determined	Individually Diagnosed and Prescribed	Personalized
	Learner Selected	Self Directed	Independent Study

FIGURE 1. Individualized Instruction

It is readily seen that the above model compares school versus pupil selection of learning objectives (what is to be learned) and school versus pupil media for achieving the learning (how the objectives are to be reached). According to one author (Hull, 1972) when the school selects both the learning objectives and the media for attainment, the category is termed individually diagnosed and prescribed learning. When the school determines what is to be learned but allows the learner freedom to determine how he will attain the objectives, the category is termed self directed learning. In situations where the learner selects the objective but the media are determined by the school, the category is termed personalized learning. Finally, if the student selects both what is to be learned and how to learn it, the category is termed independent study.

With this rather neat categorization of individualization one may turn his attention to a comparison of this format and structural learning. It has been suggested that improvement in the quality of education requires more than rhetoric and/or superficial proposals, that we need a deeper understanding of the teaching-learning process, itself (Scandura, 1973). A new way of viewing teaching-learning is that of structural learning which is defined on three levels of theory. Level one consists of a first step of competence which accounts for potentially observable behavior of interest to an observer; a second level deals with the behavior of humans under certain idealized conditions and the third or unrestricted step is concerned with memory and the limited capacity of subjects to process information.

It appears then that level one when considered in the context of individualized instruction is concerned with the match up of learning and individual interest and/or ability rather than the definition of laws which govern the ways that rules may interact in accounting for behavior. We are simply restating the "fact" that there are a variety of ways to learn any knowledge or skill and the individualized format must actively seek the mating of learning style and proper learning environment. Viewing all learning as problem solving permits a more systematic approach to experience and accretion of learning. Although many would claim that problem solving is inherent in individualized instruction students may not learn or even be able to build on prior learning without a categorization system. This essence of the structural learning system has been part of Brunerian theories of teaching-learning (Bruner, 1966).

Basically then, individualized instruction and structural learning differ in the exposition of rules. Structural learning seeks the rules, the finite testing procedure, to measure each individual's knowledge relative to the rules in a given competence theory. Individualization, however, is not as concerned with the rules of learning as it is with the individual. Individualized instruction can occur without rules of learning as seen in the independent study cell of Figure 1. The only requirement for individualized instruction seems to be the philosophy of a teacher which permits self pacing.

The real contribution of structural learning and individualization is seen in the desire for quality teaching-learning. Individualization then, cannot contribute to the problems in learning per se. The philosophy of individualized instruction which is transformed into reality within the classroom in today's schools is a concept which seeks to quiet reactions to the teaching-learning environment. Ultimately individualization will gather together the rules for learning and fit these to individual learning styles.

As each discipline begins to specify level one and two of the structural learning format the rules may be compared and consolidated to arrive at multidiscipline structural learning, see Figure 2.

	Math	Poli Sci	Phys Sci	Bio Sci
Rules for Competence				

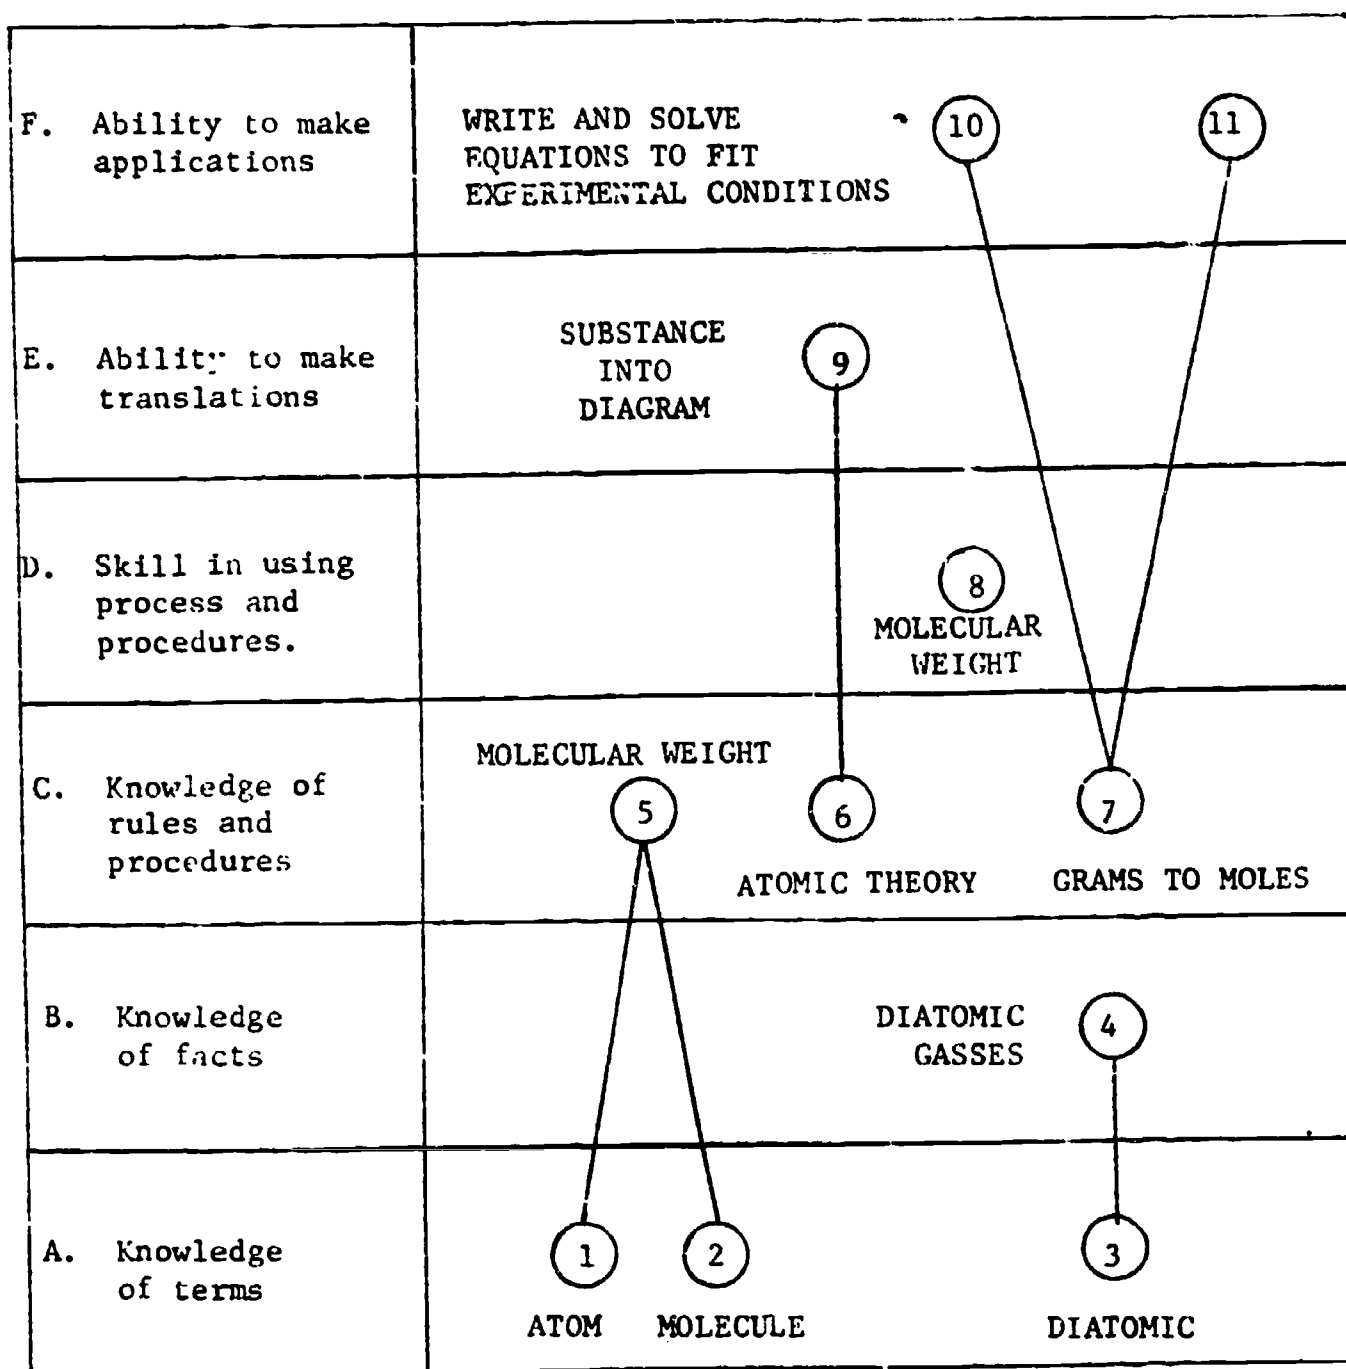
FIGURE 2.
Defining Level One for Multidiscipline Structural Learning

Thus similar and sometimes identical rules can be used in expository teaching-learning based on a variety of disciplines. As one moves toward the various types of individualization then, the rules can be seen as truly interrelated. The ultimate in individualized instruction can be reached as the levels of structural learning are achieved in the various subject areas and are then amalgamated into multidiscipline learning. The use of a table of specifications, see Figure 3., may be one way to better understand the need for rules for competence within each discipline (Bloom, Hastings and Madaus, 1971).

We note that the development of this type of a table gives us the behavior and content which is associated in a particular unit hence, we are moving toward the fruition of the structural learning concept.

In conclusion, it would appear that individualized instruction is a nebulous term. Individualization though, is humanizing the classroom and this in itself is to be encouraged and applauded. Structural learning requires that instructors begin looking for rules which hold true under specific conditions thus permitting the instructor to structure learning for faster and quicker elicitation from the computer like cells of the brain.

FIGURE 3. Model of a table of Specifications



The future of multidiscipline structural learning must await publication and amalgamation of rules within the disciplines. This systematic approach to teaching-learning appears bright if . . . research precedes implementation. The educational movement into the specification of objectives should lead to the derivation of rules to meet specific circumstances. In addition, the production of behavioral objectives should permit finite testing.

As one considers the possibilities of individualized study he/she must be aware of the need for human interaction. No person is going to live in this world alone, the recent shortages of beef, gasoline, etc.,

emphasizes our dependence on others. During the formative years of learning the youth of our country must not be isolated into complete programs of individualized instruction. Each pupil must learn to participate in the give and take activities of our world, he/she must know that this is a necessity for the survival of Homo sapiens. Armed with these thoughts the future of structural learning within an individualized format seems quite bright.

REFERENCES

- Bagby, G. "Precollegiate Education," Danforth News and Notes, 9 (2), 1-5, February, 1974.
- Bloom, B.S., et.al. Handbook on Formative and Summative Evaluation of Student Learning. St. Louis: McGraw Hill, 1971.
- Edling, J.V. "Individualized Instruction," ERIC Prep Kit #16, ED 041 185, 1971.
- Hull, R.E. "Selecting an Approach to Individualized Education," Phi Delta Kappan, 55, 169-173, November, 1973.
- Kinsella, M.T. and B.W. Kinsella, Impact on Instructional Improvement. New York State Association of Supervision and Curriculum Development, 7 (4), 1972.
- Scandura, J.M. "Structural Learning and the Design of Educational Materials," Educational Technology, 13, 7-12, August, 1973.

HOW MUCH STRUCTURE IS THERE IN MEMORY?*

T. K. Landauer
Bell Laboratories

The purpose of this talk is to see if you will be persuaded that the title constitutes an interesting question.

I have been exploring a model for the basic mechanisms of storage and retrieval in human memory which makes minimal assumptions about organization. In fact, it assumes that data are entered at random places and are searched for in a completely unsystematic, undirected manner. This rather structureless memory system turns out to exhibit a surprising variety of human-like behavior.

What I will do here is present a description of the model, pause to see whether you have questions about how it works, and then give as many illustrations of its applications as possible in the time allotted.

The storage space consists of a large three-dimensional array containing a very large number of relatively small storage loci. These storage loci may be thought of as similar to the storage registers in a computer's core memory. They may have some maximum size, as yet not specified, but their content is assumed to be quite flexible and variable. Precisely what is entered into these loci, and when, is presumably under the control of parts of the system other than the memory (by the way, it is not denied that the system has other components than memory; e.g., perceptual pre-processors, parsers, etc., and an executive to organize it all). However, where in the memory a particular datum is entered is a random event. The place of storage is not under the control of the rest of the system. Data are entered at the tip of a wandering pointer. This pointer meanders through the three-dimensional memory space in a slow and inexorable random walk. The random walk is simply a function of time; the pointer moves with equal probability, -1, 0, +1 steps in each of three directions in each successive small epoch of time. The pointer does not move towards any particular location or area as a function of the nature of the data it is currently processing. Any kind of data can go anywhere.

When the system wants to retrieve information, it institutes a blind search. The search is accomplished by sending out a signal from the current pointer location. This signal spreads from the pointer tip like the spreading of

*This talk is based partly on a paper entitled "Memory without organization, Properties of a model with random storage and undirected retrieval" presented at the Psychonomic Society annual meeting, St. Louis, November 1972. A longer working paper of the same title, giving many additional details of the formulations and results summarized here is available on request. I thank D. E. Meyer, N. Nilsson, S. Sternberg and L. Streeter for help and comments.

a wave around a rock thrown into a pond. The rate of spread of the signal is constant and its maximum distance is fixed. If this signal encounters a register containing appropriate matching information, then the total contents of the register are returned to the pointer. In other words, the contents of registers are accessed by a sufficient match with the signal. The signal may consist of the total contents of a register, or only some part of it. An analogy might be finding the name of the owner of a car with a certain license plate by announcing the number in a large audience and waiting for a return from an individual who recognizes the signal. (The time for the acoustic signal to reach the individual and for his answer to travel back is the analog of the search-time.)

It needs to be emphasized that the pointer does not move in making the search, it sends out a signal that constitutes the search process. The position of the pointer is not influenced by the occurrence of the search, or by its outcome. In particular, the pointer does not move to the location where particular information is found. The signal is assumed to spread to different maximum distances depending on various circumstances. However, almost never can it spread to the whole memory (more typically to perhaps one-third of it). Therefore, the search will often fail.

That is the model. Now, I will describe some applications to psychological phenomena.

I will start with some cases of simple human verbal learning--cases in which quantitative relations are well established.

The system learns, using this memory, according to a classical negatively accelerated learning curve. This is because each time a particular fact is encountered, it is stored at a new, random place. Since the required information can be retrieved by finding any one replica of the information, the more of them there are, the more likely it is that retrieval will be successful. The model also yields strikingly similar functions to those found in human data on forgetting. Forgetting occurs in the short run by the pointer walking away from the storage location, and over longer periods by the recording of contrary information or by the relatively rare occurrence of write-over of previous data.

The memory shows effects of spacing of trials of several kinds. Basically, spacing is beneficial because it causes the several entries of the same information to be more widely distributed in the memory space. As a result, there are more pointer locations from which a search can reach the desired information. Several previously unexplained complex spacing phenomena are also predicted from the model.

Because data about the same conceptual topic can be stored in different, independent loci, sometimes not accessible at the same time, certain other interesting phenomena of human memory are implicit in the model. For example, the tip-of-the-tongue phenomenon, in which one can remember something but not everything about, say, a name, is explained simply by assuming that one stores only part of the information about it on one occasion, and the rest on another.

Perhaps more interesting than these applications to simple learning phenomena, are some predictions concerning the effects of usage frequency on the retrieval of information about words. One reason that these predictions may be more interesting is that they are more novel; there are no previous theories to account for these relations in a detailed way.

The time required for memory access depends primarily on how long it takes the match signal to reach an appropriate register. This is assumed to be a linear function of how far away the closest locus containing the desired information is from the pointer at the time of search. The larger the number, n , of different loci in which the same information is stored, the closer to the pointer will be the nearest one of these, on the average. Therefore, the more times something has been stored, the quicker will be retrieval. For n large, the mathematical relation is approximated by the function

$$RT = a + bn^{-1/3},$$

where n is the number of replicas stored, and a and b are constants. This general function was compared to reaction time data for retrieving various facts about words of different frequencies. One set of data were observations by Oldfield and Wingfield (1965) of the time required to name line drawings of common objects. Another was of times required to say that a word like "CHAIR" was not the name of an ANIMAL. (In this case it is assumed that the subject must first look up the name of the visually presented word, CHAIR, in his memory, then, in a succeeding stage, determine that it does not belong to the category.) A third set of data came from an experiment in which categories were artificially learned in the laboratory. In all three cases, the model gives a very close quantitative fit to human data.

As a side-light, these kind of data may be used to fit a more general form of the model in which the dimensionality of the storage space is a free parameter whose value is estimated from the data. In all three sets of data, the best fit of the function is obtained with three dimensions.

The model can also be applied to the recognition of words under suboptimal perceptual conditions. In this application, it is assumed that the search signal does not provide a sufficient match to determine a unique word, but rather the item "recognized" is the first one found whose stored representation is sufficiently like the degraded stimulus. In this case, the probability of finding the correct word is determined by the number of places in which it is stored relative to the total number of places in which it and other things that can be confused with it have been stored. The model, in this instance, reduces to a special form of the so-called "urn-model" previously proposed by Pollack, Rubenstein, and Decker (1959). Simulations were performed, based on the assumption that degradation of words results in the loss of information about just one letter, and using a published sample of English reading material as the assumed contents of memory. These found a close resemblance between the function for word recognition predicted by the model and those found in human recognition experiments.

A final set of applications involves some recent findings about "semantic memory" (Rubenstein, Lewis, & Rubenstein, 1971). When people judge that a string of letters is a word, they are faster if the word has more than one meaning than if it is an equally common word with only one meaning. The explanation follows from that of spaced practice. It is assumed that polysemous words are encountered at more widely dispersed times, and, consequently, their storage loci are more widely distributed. This means that, on the average, the pointer will be nearer to one of the replicas of such a word than to a word with a single meaning.

The model has been applied to a number of other relations, phenomena and situations that cannot be described for lack of time. However, I have presented a sampling that gives essentially the flavor of the kind of applications in which it has been successful.

This demonstration that a model relying on random processes can account for many significant aspects of human memory may contain an important lesson. In my opinion the lesson is that the potentialities of such simple mechanisms ought to be explored much more thoroughly before complex and sophisticated self-organizing schemes are invoked for the explanation of human knowledge and performance.

References

Oldfield, R. C., & Wingfield, A. Response latencies in naming objects. Quarterly Journal of Experimental Psychology, 1965, 17, 273-281.

Pollack, I., Rubenstein, H., & Decker, C. Intelligibility of known and unknown message sets. Journal of the Acoustical Society of America, 1959, 31, 273-279.

Rubenstein, H., Lewis, S. S., & Rubenstein, M. I. Homographic entries in the internal lexicon: Effect of systematicity and relative frequency of meanings. Journal of Verbal Learning and Verbal Behavior, 1971, 10, 57-62.

AN OUTLINE OF CONVERSATIONAL DOMAINS AND THEIR STRUCTURE

Gordon Pask
Systems Research Ltd.

1. Introduction

The knowledge structures described in this paper have been chiefly applied in the field of education; to represent subject matter from such areas as probability theory, science and sociological disciplines. The structures are called conversational domains. They are generated and also used in the context of a tutorial conversation between a student and a real life teacher or between a student and a machine instrumented heuristic (not, in the usual sense, a teaching machine or CAI program.) Acting in its most restrictive mode, the heuristic instruments a teaching policy for the subject matter in the domain; in its least restrictive mode, (which is of immediate interest) it constrains otherwise free transactions so that (a) the student directs his attention to some topic in the domain and (b) the dialogue of the conversation is segmented into occasions upon which cognitive processes called understandings (a technical term) take place; by this means, the conversation may be decomposed into a series of non arbitrary beginnings and endings. If (a) and (b) are satisfied, for a domain R, the conversation is called a strict conversation, anchored on R.

Most of the experimental results which underpin the structural postulates of this paper regarding cognition, memory, and the like, have been obtained in a facility for instrumenting and regulating strict conversations. This facility is known as the "Course Assembly System and Tutorial Environment" (henceforward, called by its acronym "CASTE"; Pask, Kellikouris, and Scott 1972, 1973 a, 1973 b; Pask 1971, 1972)

The conversation, in this case, is non verbal and takes place in a stratified, mechanical, command and question language, designated $L = L^1, L^0$. In CASTE transactions, one participant (A) is a human being and takes the role (initially, at least) of a student. The other participant is the CASTE heuristic (B) and, unless specifically constrained, it is as neutral as possible. If desired, constraints may be added to convert B into a tutorial heuristic.

The paper outlines the structure of conversational domains in which strict and tutorial conversations take place with certain intuitively obvious requirements upon the learnability and memorability of topics in the subject matter; it indicates how conversational domains are generated from an expert's or source's thesis about a subject matter and gives an overview of the operation of a program, EXTEND, through which the student in CASTE can play the role of source and develop extensions or extrapolations of the conversational domain. The only serious proviso is that the resulting structure satisfies all the constraints upon form (not upon the content or even the magnitude of an innovation) that are imposed for the sake of learnability and memorability, upon the original.

2. Theoretical Underpinning for Conversational Domains

As a preface to the discussion, it is essential to outline enough of a theory of learning and teaching, called conversation theory (Pask, Kellikouris, and Scott 1974 a, Pask 1974 b) to make sense of the term "understanding" in its technical usage and to motivate the (otherwise quite arbitrary) restrictions imposed to ensure learnability and memorability upon a conversational domain. The interconnection of a cognitive theory and the structure of a conversational domain is not accidental. It appears impossible to speak cogently of knowledge without some theory, albeit oversimplified, of knowing. Conversely, it is impossible to furnish a coherent account of cognition in the absence of a theory about what may be known.

According to conversation theory there exist L processors, typically, though not necessarily, brains, able to interpret and compile L procedures. Here, L is the conversational command and question language ($L = L^1, L^0$ of Section 1). Regarded, for this purpose, as a programming language. " L procedure" is a general rubric intended to encompass serial programs and non deterministic (parallel execution) programs as special cases, but usually, to designate a fuzzy program or a class of programs.

If asked what a program does, one reply is to say it computes a relation (or, in case it is a regulatory or control program) that it stabilizes some relation. The same comment applies to a fuzzy program and thus, in our jargon, to a procedure. It is pedantic, but consistent, to remark that a procedure or a program, for that matter, is characterised as reproducing or producing a relation (though the exact meaning given to the word "relation" will only be evident when conversational domains have been discussed). Suppose that we only intended "relation" in a syntactic sense: namely, given as a mathematical entity in extension, as a listing of ordered subsets of elements in some coordinate space or product of sets. If so, and given some important caveats regarding termination any pair or more of programs operating upon the same field (i.e., the coordinate space, or product set, in question) are extensionally equivalent if they produce or reproduce the same relation. In general, fuzzy programs compute fuzzy relations. Hence, if a pair or more of procedures act upon the serial fuzzy field to produce or reproduce the same fuzzy relation they are tolerance equivalent.

Of course a program or procedure only does anything if it is executed. In order to be executed, an L procedure must be compiled in an L processor. We call the actual or potential execution of an L procedure, in fact its compilation, a concept: if the L procedure produces or reproduces a relation R_i , then it is a concept of i . Similarly, the compilation of a class of tolerance equivalent procedures all of which produce or reproduce i is a concept of i . However, in deference to a fact (which soon becomes evident) that we mean more by "relation" than "relation in extension" there may be two or more compilations of tolerance equivalent L procedures that reproduce or produce two syntactically identical but distinctly interpreted relations R_i and R_j . These figure as two or more distinct concepts, namely concept i and concept j , and the two relations R_i and R_j are isomorphic but not identical.

Under what circumstances do the compilations of L procedures exist, as stable entities, in L processor. Such questions are usually glossed in connection with standard processors like digital computers (given L as a programming language and a machine architecture able to accommodate L expressions and sufficient storage medium). The notion that questions of this kind must be asked of concurrent processors, notably those that are able to execute fuzzy programs without numerical resolution, may strike a chord with computer theorists. Undoubtedly, the question is realistic in psychology and, insofar as the L processors are brains, it is mandatory.

The tenets of conversation theory are as follows and are compatible with discursive as well as experimental evidence; namely

- (a) a compilation of a procedure (a concept i) is stable if and only if there exists a further procedure able to operate upon a description of concept i (alias the procedure and a tag for its compilation) in order to produce or reproduce this concept (notice, a tolerance equivalent procedure, duly compiled, is a reproduction).
- (b) Such a procedure, if compiled and undergoing actual or potential execution, is a memory.
- (c) As a convention, introduced to avoid vicious circularity, the concept is called an L^0 procedure and the memory an L^1 procedure.
- (d) An understanding of R_i is the production of a concept i in the context of a memory i that reproduces it.

(e) Evidence for an understanding of R_1 is obtainable from the dialogue going on in a strict conversation between participants (A and B) provided the conversation is anchored upon a given conversational domain. The evidence consists in an explanation of R_1 (which may, without prejudice to the argument, be a non verbal explanation elicited as a modelling or programming operation) and a derivation of this explanation (usually manifest as part of a learning strategy).

(f) If the verbal or non verbal explanation is called an " L^0 explanation" (in $L = L^1, L^0$) then "derivation" figures as a synonym for " L^1 explanation."

A regulatory heuristic such as the conversational heuristic employed in the CASTE facility of Section 1 is primarily a device for ensuring that all or some of the topic designating relations in a conversational domain are understood by the participants, in such a way that each understanding ends an occasion and that occasions are strictly ordered and fully observable by an experimenter. To secure these conditions the experimenter designs or instructs a participating agent (the CASTE heuristic, B, for example; equally in principle, a Piagetian interviewer) who interacts with the student or subject (A, the other participant). If any topic is to be understandable, some care must also be exercised in specifying the form of conversational domain and (insofar as evolution is permitted) the means by which it can be enlarged. The essential requirements are that the relations designating the topics can be explained, that separate topic relations are properly distinguished, and that (given certain assumptions about the innate cognitive repertoire of a student) any topic relation can be derived, without loss of specificity, from others in the same conversational domain. As an entirety, these requirements guarantee the productibility and the reproducibility of topic relations and are succinctly expressed as conditions for learnability and memorability.

Systematic growth of a conversational domain, either from scratch or by way of learnable and memorable additions, can, in large measure, be mechanized. One method consists in a discussion, not to be confused with a strict conversation about the resulting domain, between a subject matter expert and someone who follows a series of instructions. Most of the following commentary is couched in terms of this method because the discourse and the constraints that govern it are quite easily comprehended. It should be emphasized, however, that the subject matter expert may equally well interact with a mechanically executed heuristic, the EXTEND program, so that all apparently sloppy regulations can be stated as precisely as desired. As noted in Section 1, EXTEND can also be used in the context of a strict conversation, provided the participating student changes role. Finally, EXTEND has recently been augmented by a further heuristic, THOUGHTSTICKER, which realizes a different, but still systematic type of evolution.

3. Some Epistemological and Ontological Preliminaries

A conversational domain represents the thesis, entertained by a source or subject matter expert. In calling it a knowledge structure, (of chemistry, for example) we mean "what may be known and done according to a (specified) source's thesis about chemistry" and certainly do not mean "all that can be known and done in chemistry, unqualified."

When a source entertains a thesis, it will be assumed that the source's purpose is to enter into discourse with respect to his thesis; in particular, the thesis is used as a base for tutorial dialogue. This fixes, albeit loosely, the pragmatics of a situation in which the source delineates a thesis to form a conversational domain; the pragmatics of a tutorial situation (for example, a strict conversation, between participants A and B, as in Section 1) are firmed up by the conditions under which the A, B dialogue is anchored upon the conversational domain.

What does a source, with such a broadly specified purpose, actually do when asked to express his views on a subject matter (chemistry, say). The least he can do is to mention certain topics by name; to specify, if interrogated, the meaning of each topic and to specify the meaning he attaches to their relations, one to the other, as constituents of a thesis. A source might do much more than that; he might, for example, express a coherent theory of chemistry without prompting. But most sources are unable to do so and report, as their introspective experience, a flux of more or less perspicuous concepts, amongst which some are closely knit, some vague; some clearly associated with the name "chemistry", some tenuously connected. The process of specifying a thesis, so that it can be represented in a conversational domain, involves winking out a picture of chemistry that assumes shape, retrospectively. In practice, we aid the source in this task by providing an analyst/interrogator (Section 2) who also ensures that the thesis is not only coherent, according to the source's personal criteria, but, in addition, satisfies the special learnability/memorability conditions imposed upon a conversational domain. It turns out that these conditions impose a form upon the syntactic components (i.e., the interpretation-free-components) of the thesis and amount to the requirement that the thesis has a Gestalt structure.

With these comments in mind, let us dissect the source productions (a series of topics and mooted relations between topics) into their logical parts, in order to characterize a thesis (in contrast to an amorphous, freely associated, melange of ideas). The perspective to be described does not necessarily tally with the source's view. But it does correspond to the view of an analyst/interrogator/cooperating with the source, and these points of view are, in principle, translatable. A topic (which the source calls T, perhaps) contains a nameable topic relation which we (or the interrogator/analyst) call R_i , using the index, i , as a name for the topic relation.

R_i has an extension (a possibly finite, possible indefinite listing as a subset of values of certain coordinates) and an intension or rule (roughly a rule in Scandura's sense). Though the extension of relation R_i may, in principle, be specified as a listing, it may also be represented by means of relational operators (Codd 1966) which produce R_i , in extenso, from a specific transformation of other relations (R_p, R_q) cited by the source. Such a representation is dubbed a derivation of R_i and there may be numerous different but extensionally equivalent derivations: for example, R_i is derived as the Join (one relational operator) of R_p and R_q on a common coordinate or, another derivation, as the projection of R_s onto its first and fifth coordinates composed with R_j (Projection and Composition are other relational operations.) Two major advantages are obtained by the derivational construction of relations in extenso: (a) The representation is uncommitted regarding the nature of the coordinates which may be and usually are complex (i.e. value sets of other relations) rather than simple (values of "1-adic-relations", alias "properties") so that a relation can be specified without, at this point, dogmatizing about what "really" is "simple"; (b) The representation is uniform and natural. To see this, notice that relational operations are rules (like intensions, except that they act upon relations rather than coordinate values) and, roughly speaking, serve as higher level rules in Scandura's sense.

There is no suggestion that the source either does or should see a topic T as a relation only: he may very often have an interpretation in mind, as well. But surely, if T belongs to a thesis, the relation in register with T (namely, R_i) has an extension and an intension. These are syntactic entities, insofar as they can both be imaged by grammatically legal expressions in a language. Without prejudice to the source's language (same compartment of natural language, presumably) the interrogator/analyst will express these syntactic entities in a metalanguage designated L^* . In L^* parlance the extension of R_i is represented as one or more derivations of R_i (clearly, the representation only becomes possible after several topics have been cited): the name of the topic (T) is the index, i , of R_i ; the intension of R_i is represented as a program graph called $PG(i)$ such that if any program in $PG(i)$ were interpreted and executed in a suitable processor, it would bring about or satisfy the relation R_i .

If the outpourings of a source are to count as a thesis we do require that

the source can cite (on being told by the interrogator/analyst/that his name for topic T is i) the intension of R_i as $PG(i)$ and one or more derivations that delineate the extension of R_i . The source may output both the intension and the extension of R_i idiosyncratically, in his own language; typically, he is required to furnish an explanation of R_i as the intension, and an explanation of how this explanation is obtained (that is, a derivation) as the extension of R_i . Quite commonly, the explanation and derivation are stated informally. But we do require that the source statement and the L^* statement are translatable; that is, if the interrogator/analyst/submits several L^* intensions, $PG(i)$, of R_i and several L^* derivations, as the intension of R_i then (after arbitrary repetition) the source agrees that the L^* intension and the L^* extension of R_i are, in all ways pertinent to his thesis, equivalent to his own explanation of[#] and derivation of T.

In particular, there is no supposition that the source's cognitive repertoire contains the relational operators that are used to represent derivations in the metalanguage L^* . Since sources commonly employ words like "generalisation", or "abstraction", or "cohesion" as the means by which the topic relations are derived, such a supposition is generally untenable. On the other hand we do assume (in requiring "translatability") that any cognitive operation employed in the source's derivations can be represented, in L^* , as some relational operator or some collection of relational operators applied in series.

To be explicit, the relational operators we use are listed below:

- Natural Join (on specified coordinates of relations).
- Projection (onto coordinates specified).
- Permutation (of coordinates of a relation).
- Composition (on a common coordinate of relations).
- Partial Complement (with respect of a coordinate space).
- Union and Intersection (under constraints to avoid vacuity).
- Restriction (of one relation by another relation).
- Isomorphism (between relations).

The operator isomorphism has special significance. Any conceivable relation between topic relations is a morphism of some kind and might be expressed as such. The variegated relational operators serve as a device for eliciting the morphisms of a thesis directionally and in a manner that reflects the systematic application of principles, axioms, and generative rules. Isomorphism is peculiar insofar as it stands for a relational identity (this does not mean that an isomorphism is an identity: definitely it is not an identity of values). In the conversational domain, however, the isomorphisms are, because of this, ascribed priority.

Two important consequences stem from effecting and instrumenting a dissection of topics into relations R_i with extension and intension stripped of a semantic interpretation. (A) It is possible to build up a relational network of L^* derivations which can be mechanically handled and manipulated and given the fact of a $PG(i)$ for each R_i , (or, as in the footnote, belief that it can be constructed) to ensure that the resulting structure satisfies the learnability and memorability conditions of a conversational domain. For these conditions, as noted previously, are only imposed upon syntactic entities (the L^* extension and the L^* intension of R_i in a thesis). (B) It was hinted, in Section 1, that the command and question language, L , used for dialogue in a strict conversation, is stratified. The stratifications of L into levels L^1 and L^0 secures a distinction between "problem solving" or "explaining" (manifest at level L^0) and "learning" manifest at level

[#]In practice, it is usual to defer eliciting the explanation of R_i until later: to accept the existence of an explanation for each topic on trust (which is reasonable in view of the source's pragmatic orientation). Under these circumstances, the source's explanations are often obtained directly as programs, written in an agreed sublanguage of L^* .

L¹). Exactly this distinction is retained, in the thesis embodied in the conversational domain, as a distinction between "explanation" on the one hand and "derivation" on the other. We stress that the stratification of language L, in common with the stratification implicit in the structure of the conversational domain, are imposed as a matter of observational convenience. Though important for manipulative purposes such as regulating a strict conversation, the strata are artificial and have no pretensions to psychological reality. For example, we do not presume that a student necessarily distinguishes between learning and problem solving: he may consider learning to be just "problem solving of problems posed for problem solving." By the same token, a source may regard derivation and explanation as essentially the same kind of process. Such identifications are perfectly legitimate. However, they lead to non vicious circularities which though humanly intelligible, are barely tolerable when it comes to mechanical processing. One circularity of this kind is evident below, on examining the condition of cyclicity imposed upon a conversational domain.

4. The Conditions In Learnability and Memorability.

Figure 1 shows the type of relational network produced by translating and recording source dialogue. If the thesis, represented syntactically in this network and its attachments, is to count as a conversational domain, the structure must be cyclic and consistent (recall, these are the conditions for learnability and memorability).

Cyclicity is stated in the source idiom as the following, humanly pellucid but mechanically intractable, requirement. "If you derived the topic called R_1 from topics called R_j and R_k thus explaining R_1 in terms of R_j and R_k you must also be able to derive the explanation of R_j and R_k from your explanation of R_1 ; and so on for all topics named R_1, R_j, R_k in your thesis." From a "formal" or mechanical processing stance, cyclicity is a condition prohibiting loss of specificity in a network considered as a whole; this requirement imposed to guarantee the possibility of reconstructing relations in extension from other relations in extension and, (provided a minimal number of intensions are specified) to guarantee the possibility of constructing a canonical intension for all other relations (i.e. a program which, if interpreted and executed, is extensionally equivalent to the programs in the program graph of the relation in question).

Consistency is a condition of one to one correspondence between topic names and the index values (i) that name relations in L^* . It is, incidentally, a prerequisite for cyclicity and the condition is secured continually as a thesis is produced.

Consistency, though superficially simple minded, has fairly profound implications.

The first part, "You shall not give the same name to two or more different relations" is straightforward. If the source calls a topic "T" and the interrogator/analyst inscribes the relation it contains as R_1 , then, on a subsequent occasion, the source may not use T to stand for a relation R_j where $i \neq j$; in particular, if a topic is derived from T it must be called S, where $S \neq T$, though it is quite possible that T is derived from topic S. The expedient of checking the network to ensure that this does not happen prevents the source, perversely or accidentally, vitiating distinctions he has already drawn, and it has practical significance (words like "Institution", in sociology, and "Distribution", in statistics, are often used ambiguously, in a way that obscures the distinctions inmanent in the thesis).

The next part of consistency, "You shall not give two or more names to the same relation", is more subtle. If the relations were really stripped of a semantic interpretation there would be no difficulty. A check could be made on a structure such as Figure 1 to ensure that R_1 is not called T on one occasion and S on another occasion, disallowing the locution, if it occurred. But, as

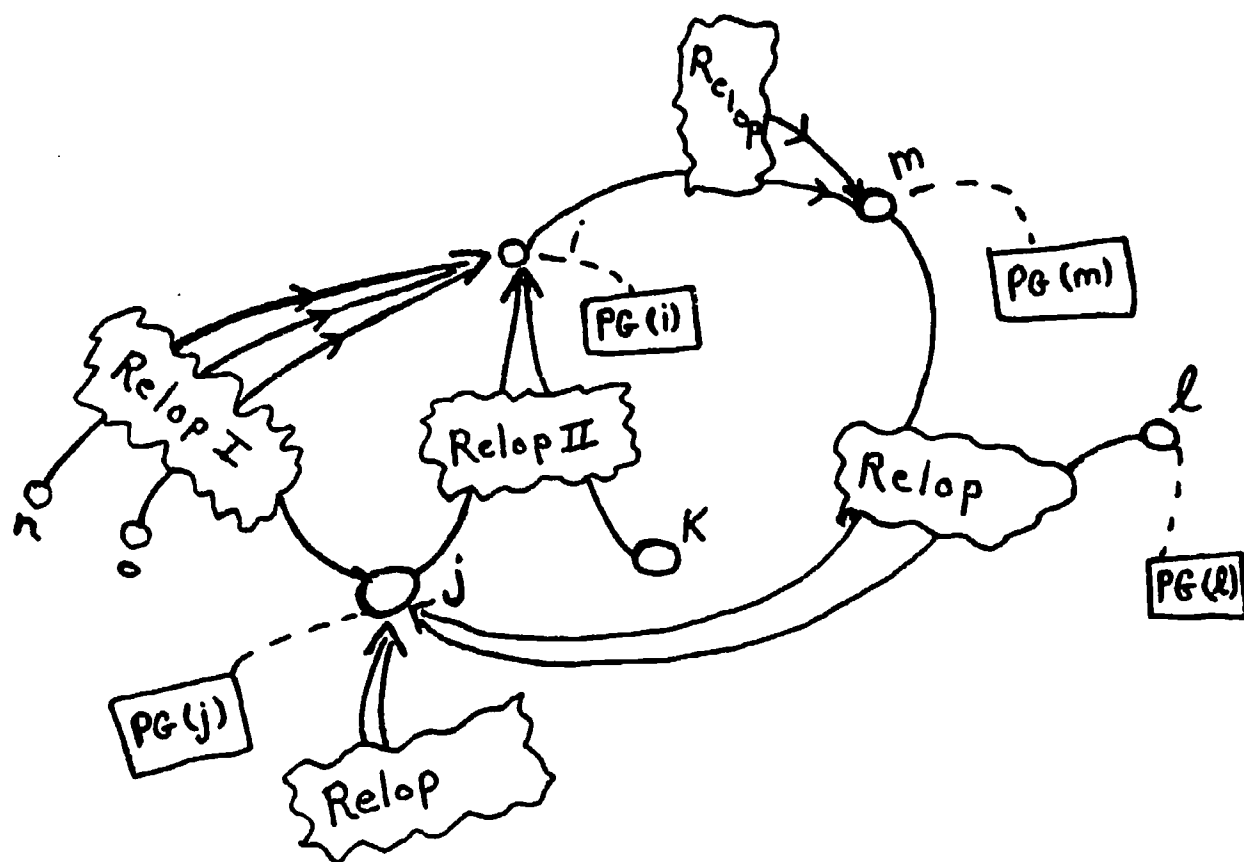


FIGURE 1. Nodes i, j, k, l, m, n, o stand as placeholders for relations in extenso R_1, R_2, \dots constructed by applications of relational operators (Relop) and for intensions (procedures) represented by $PG(i), PG(j), \dots$ connected through data links (not arcs in the graph of Relop derivations).

already emphasized, the syntactic dissection is a trick: a gambit for obtaining coherence, rather than a statement about what relations really are (namely, what the source, in a thoughtful mood, refers to as topics). We do not wish to obfuscate any distinction the source might make in his thesis. So, for example, we must distinguish between the case in which R_i is determined by two or more paths (that is, the relation derived by one path is extensionally equivalent to the relation derived by another path and is called R_i) and the case in which two or more relations are isomorphic (in one to one correspondence but not identical).

Consider, for example, the subject matter of elementary physics and the construction of a conversational domain headed by the topic "oscillator." A source versed in electrical theory might submit a thesis in which topic relations such as "movement of charge" and "inductance" and "resistance" and "capacitance" were combined to yield the topic relation "(electrical) oscillator". Even in the department of electricity this is not a unique derivation; for instance, the same thesis might furnish a derivation of the same topic relation using the notion "negative resistance". The result would be a two part derivation and, as a matter of fact, the head topic might be further abstracted to express the equations of an oscillator.

In contrast, suppose the source thesis spans the departments of mechanics and of electricity. The electrical derivation is, perhaps, of the kind already suggested. But the topic relation "mechanical oscillator" is derived within the cross departmental thesis from "displacement" and "mass" and "friction". Many other derivations are conceivable without the slightest difficulty but are of no immediate concern. The fact is, whatever derivation paths are used, a "mechanical oscillator" is isomorphic to but not identical with an "electrical oscillator" and both topic relations are isomorphic to an oscillator regarded as an abstract, mathematically interpreted, system. The derivation of "mechanical oscillator" and of "electrical oscillator" in the same thesis exemplifies derivation of two topic relations that are distinguished, in this thesis, on semantic or interpretative grounds.

In general, if a source insists that topic T (mechanical oscillator) is distinct from topic S (electrical oscillator) even though S and T are explained (in $PG(S)$ and $PG(T)$) by the same equations and even though their one or many derivations can be placed in register, the interrogator/analyst/ must instate T at a node i and instate S at a node j. By convention, the relations (even in extension) of distinct nodes are allocated distinct names: here, the names R_i and R_j . But as a result, the source will be required to delineate an analogy (we deal seriously, in this paper, only with isomorphic analogies,). The requirement impels the source to specify a relation, R_0 , between R_i and R_j , which is a one to one correspondence under some permutation of the coordinates of R_i and R_j (any permutation, including the null permutation). Further, (and this leads directly to semantic considerations) he must assert a distinction between the universes of interpretation of R_i and R_j or (if no distinction is made) he must state this fact explicitly[#]. Thus it may be that T is a mechanical oscillator and S is an electrical oscillator when R_0 , in addition to the permutation and the operator isomorphism, involves a distinction between "electrical" and "mechanical". Or it may be the case that S and T are isomorphic mathematical topic relations, when no distinction is required. But, whatever else may be the case, the universe of interpretation of R_0 is distinct from the universe of interpretation of R_i and R_j .

Though we used to consider isomorphisms, mechanically speaking, on a par with any other relation between topic relations (see, for example, Pask, Kallikourdis,

[#] Depending upon the detailed formulation, the distinction may be regarded as part of the intension of a concept or may appear in an interpretation function from the intension to a universe of interpretation.

and Scott, 1974) we currently credit them with special status and (mechanically speaking) assign them priority. That is, all isomorphisms are collected so that they apply between nodes at the same level in a superordinate/subordinate description scheme to be described in Section 5. If an isomorphic analogy is asserted at a given level and could be made to satisfy the validity requirements only if the surrounding mesh is modified (as it stands, for instance, it relates nodes at different superordinate/subordinate levels) then the disposition of other topic relations is modified in order to render the isomorphic analogy legitimate. One justification for this expedient has been hinted already (that an isomorphism stands for a relational identity). The other justification will become obvious in the next Section.

5. Retrieving the Semantic Interpretation of a Conversational Domain. Models and Descriptors.

Turn, now, to the questions of restoring the semantic interpretation immanent in a topic to the chiefly syntactic structure so far extracted from the source's thesis. According to the procedure employed and realized by the interrogator/analyst/ the question of giving a semantic interpretation is intrinsically bound up with the issue of describing the relational network and its appendages (the GP(1) for each node 1), in terms that are comprehensible to a student using the conversational language $L = L^1, L^0$ of Section 1. (rather than the metalanguage L^*). Other restoration methods are possible, in principle at any rate, but this is the restoration method employed in practice.

At some point in his exhibition of a thesis, the source is prone to say what he is anxious to teach; a heading for his course material. He might say this at the outset; for example, by saying "chemistry". But generally, such an assertion would be ill considered. The head topic (or a class or isomorphic head topics) are only manifest in retrospect, as the thesis unfolds. Definitive statements are seldom available until an appreciable chunk of the network is displayed for scrutiny by the source. So, at some point (at any point he chooses) the source is allowed to declare a head topic relation; given the proviso that he also specifies a depth, signifying the scope of his thesis.

On submitting a head and depth, the relational network is subject to an operation known as pruning, it being assumed that if the PG(1) are absent from any nodes 1 in the network, then they can certainly be supplied, or devised by the source. Pruning yields a quasi hierarchical structure, with the head located at the top and having nodes extending to the stipulated depth. The pruning process works if and only if the network is cyclic and consistent under the specified head and depth: failing that, it is rejected as inadmissible/prohibited and dialogue between the source and the interrogator/analyst/ continues (the source being furnished with advice regarding nodes at which the structure is defective). The entire heuristic, pruning and cyclicity/consistency/ testing included, is described in detail in Pask, Kallikourdis, and Scott, 1974; Pask, 1974.

Assume that a pruned structure has proved legitimate. If so, it resembles an and/or/ tree apart from the fact that analogical relations between topic relations that form part of the structure are retained as cyclic sub-structures; a typical mesh being shown in Figure 2. Derivation paths, distinct in the original network, are retained in the pruned structure, but the relational operators which provide an L^* account of the derivations are consigned to an operator data base, which is accessible only to a regulatory or developmental heuristic (either CASTE conversational or EXTEND of Section 1). The derived connecting arcs are called entailments; an omnibus term encompassing all L^* distinct means of derivation. Since the means of derivation represent (by hypothesis) those cognitive and constructive operations available, alike, to the source and any student, the entailments represent permissions to make a discovery (the corresponding omnibus term in psychology) and the distinct entailment paths

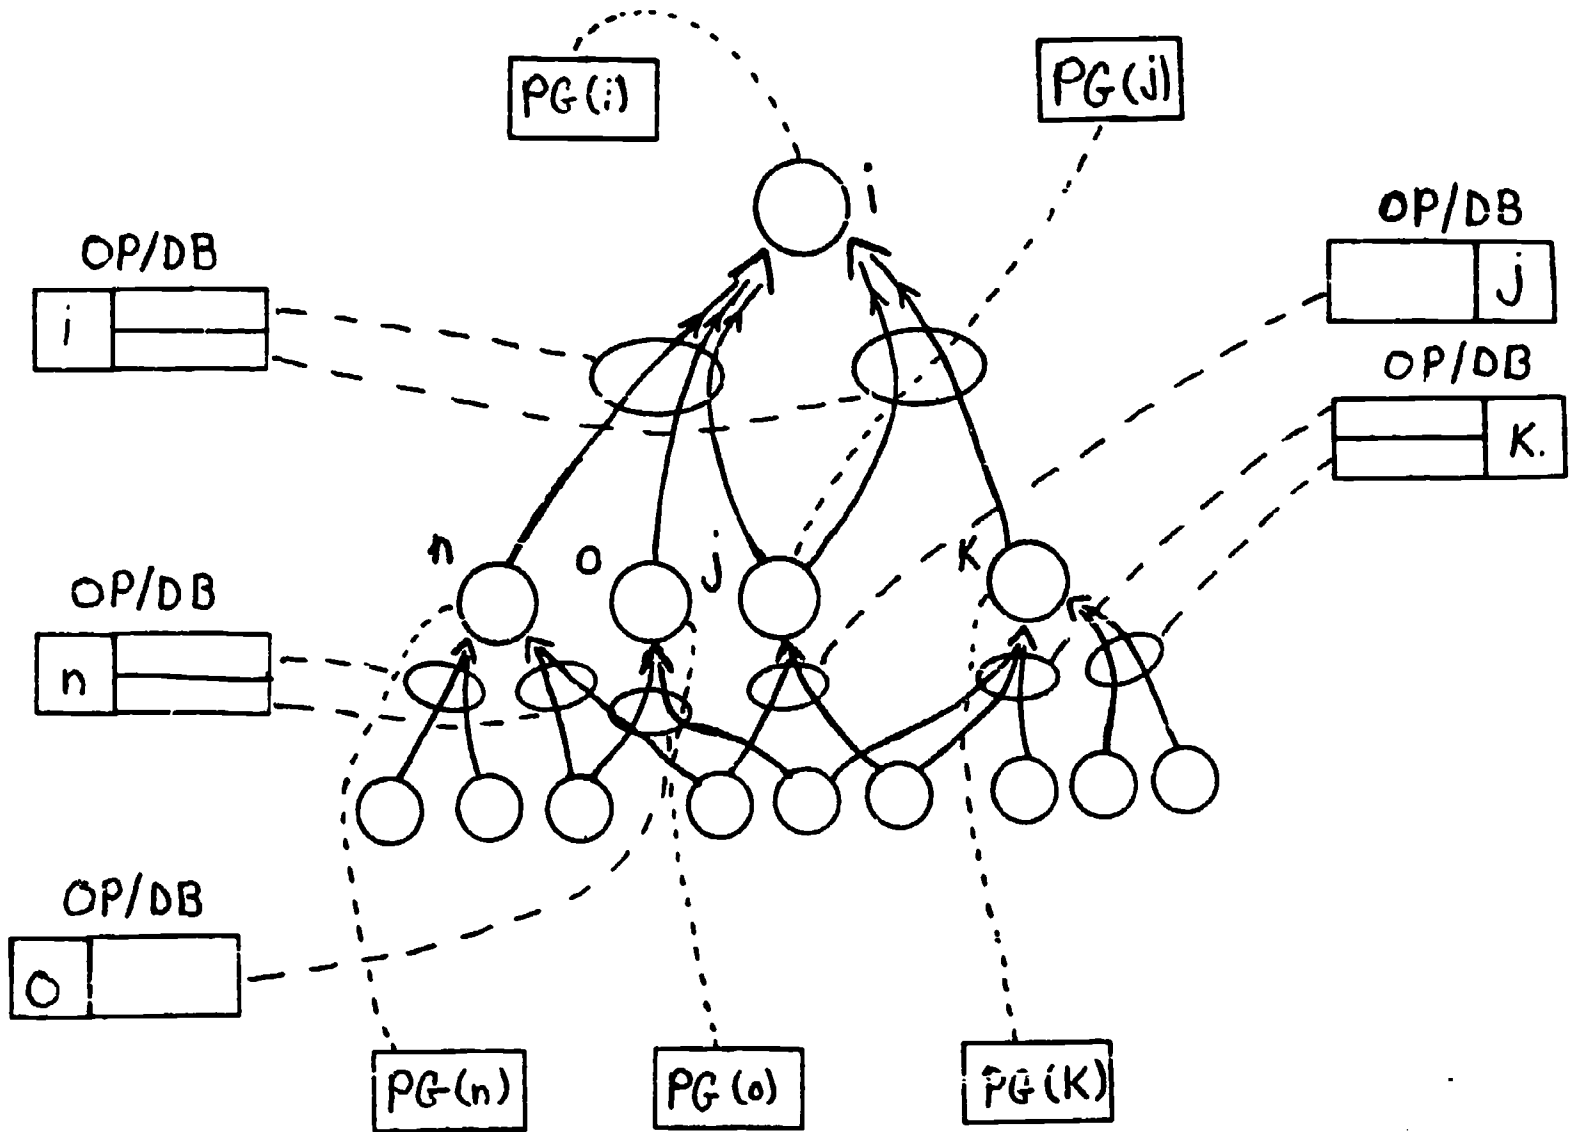


FIGURE 2. Pruned cyclic and consistent mesh with $PG(i)$ and operator data base OP/DB for each node.

constitute permissions to know as laid down by a thesis having the requirements mandatory for inclusion in a conversational domain; cyclicity and consistency. By the same token, the program graphs attached, or capable of attachment, to each node i (the $PG(i)$) represent permissions to do modelling operations in such a way that the resulting model (on execution) brings about the topic relation R_i . More succinctly, they are permissions to explain.

Now the source's choice of a head and a depth is presumably based upon semantic criteria; his own interpretation of the thesis he has expanded. In general, many choices of head and depth are acceptable (as yielding legitimate, pruned, entailment structures) for exactly the same relational network. For example, in a thesis on probability theory, one source chose the head topic relation "statistical inference" and another source, working in concert, chose "probabilistic automaton theory": the underlying structures being, of course, identical. But the choice also gives syntactic information which is formally or mechanically handled. The choice, essentially of a noun that names the entire structure, induces a description of the structure in terms of arc distance, between nodes representing the topic relations (Figure 2). Since each node is named by a noun (it stands for several things, most of them verb like) the arc distance from head is aptly christened the "superordinate/subordinate/descriptor", or the "syntactic descriptor", of the superficially taxonomy like, but, in fact only quasi hierarchical pruned structure. This superordinate/subordinate or syntactic descriptor has a special significance, as follows. It layers out the isomorphisms in the conversational domain so that any isomorphism is between nodes located at the same superordinate/subordinate level. #

It should be noted that any node in the structure with the exception of certain low level nodes, that protrude in Figure 2 as primitive or elementary, could be identified by other syntactic descriptions. For example, it is possible to elaborate the notation scheme of Figure 3 by numbering the nested subsets of arcs from left to right, or as a less arbitrary expedient, it would be easy enough to retrieve the relational operators from the operator data base.

These possibilities are deliberately set on one side, and the source is required to choose a set of semantic descriptors of the nodes in the mesh, which satisfy the following requirements.


(a) The descriptors may be regarded as unary but many valued predicates of the set of nodes, denoted $i \in I$ so conceived, they stand for node properties and have values, true or false, in expressions $Pred_{ji}(i) = Value_{jr}$ which is true if and only if Node $i \in$ the r th value of Pred j : that is to say, Node i belongs to a disjoint subset S_{jr} of a subset S_j of I having the property designated.

(b) Any descriptor has the possible dummy value "*" on Node i where "*" denotes inapplicable to or irrelevant to this node.

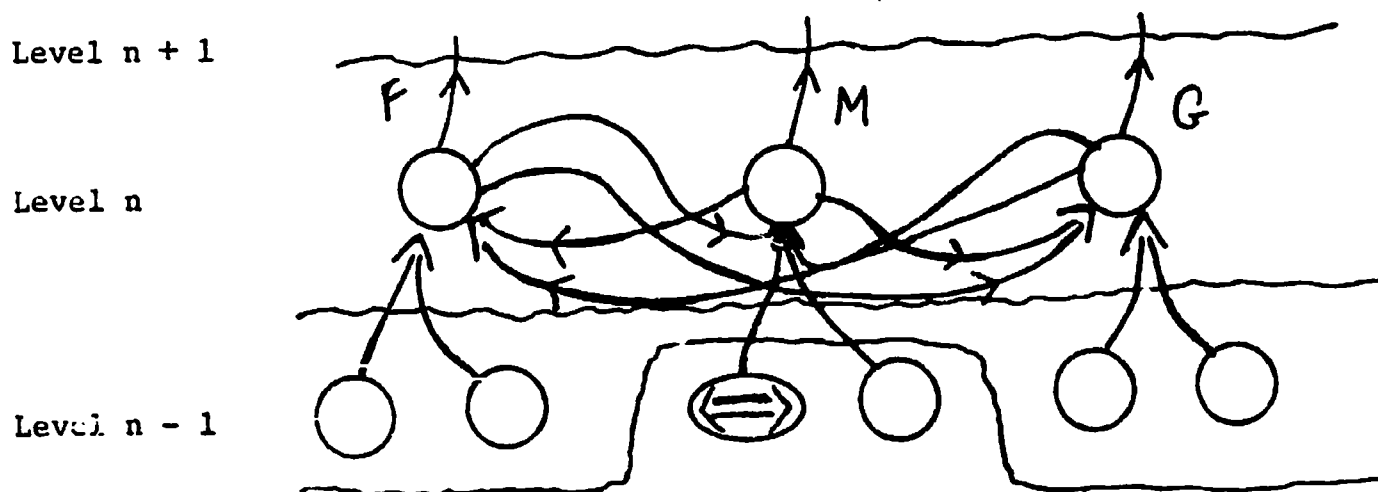
(c) A family of descriptors is a fine structure family in I , in Benerji's 1970 sense; that is, minimal conjunctive statements in the Pred j of this family delineate unit sets containing, at most, one node. Hence, a family acts as a unique indexing scheme on the nodes.

(d) The descriptors may be redundant (for example, there may be more than one family of them).

The reader has probably gathered from the gist of the argument (though the point was only mentioned obliquely) that all relations between topic relations are morphisms of some kind: (homomorphisms, for example) so that isomorphisms have special status. The relational operators are no more nor less than an expedient for extracting morphisms directionally, as derivations that follow some axiomatic principle.



a Non-analogical structure with arc distances.



b Arc distance assignment over isomorphic analogy
Relation M between F and G. Shorthand Convention
is shown below-F and G are
known as "Terms" or Relata" of analogy relation.




FIGURE 3. Syntactic description of entailment mesh by arc distance.

(e) At least one family (perhaps there is only one) is understood by any student using the conversational domain. The word understood is used in its technical sense to imply (as later) that all descriptors in the family, whatever they may actually be (they are only "regarded as" being unary but many valued predicates) can be treated as properties by a student.

(f) Any family "is not a full family"; that is, minimal conjunctive expressions in its predicates designate some unit sets that are not currently occupied by nodes (the motivation for this condition appears later, on considering the extrapolation of the conversational domain).

On delving deeper, the descriptors assume a different complexion. What, really, is the source doing when he classifies and indexes the nodes?? The nodes are placeholders for R_i and $PG(i)$ and it happens that a universe (I) of placeholders is momentarily delineated. Unless the source classifies nodes as having syntactic properties (i.e., connections to other placeholders in the mesh which is discouraged with the exception of the superordinate/subordinate descriptor) he must be classifying the $PG(i)$.

There are several ways in which the source could do so, depending upon the meaning ascribed to an intension. Here, we develop one meaning, which, though idiosyncratic, pervades the use and generation of all conversational domains employed with a facility like CASTE of Section 1.

For mechanical and observational convenience, we elicit one part of the evidence for an understanding: namely, the explanation of Section 2, through a non verbal explanation or modelling operation.

A modelling operation is precisely the construction of an interpreted program (or several of them) that is executed in a modelling facility. For example, Papert's programming language, LOGO, combined with a suitable processor and "toys" such as the "Turtle" is a modelling facility. A LOGO program (for instance, in Computational Geometry) is the syntactic or intensional part of a model; it represents a concept in "computational geometry". The program is executed in the LOGO processor and is interpreted in respect of "Turtle" motions. If the interpreted program works, to bring about a certain relation (say to delineate stellate figures, by "Turtle" movements) then it represents a concept for this relation. We say it represents (rather than is) a concept as there is a fundamental restriction upon LOGO (and other modelling facility) processors; the student can only contrive one model at once (though interrupts are possible) and execution can only occur in the absence of model building operations, and vice versa. The modelling facility (STATLAB) employed in our own studies of statistics and probability theory learning is, in some ways, less versatile but it allows for several models to be executed simultaneously and independently if they are interpreted in distinct (that is, independent) universes of interpretation (for example, sets of real or temporal events, sets of elements of atemporal abstract sets).

In view of these comments upon modelling operations, the meaning given to the semantic descriptions is as follows. They classify the models (or interpreted program graphs $PG(i)$) that are attached to each of the nodes i . In particular, their values are subsets of one or more universes of interpretation, namely, one or more parts of a modelling facility (for example, STATLAB). Several universes of interpretation exist in the modelling facility (or there are distinct modelling facilities) insofar as certain parts are independent (so that independent program execution is possible).

The following points are worth stressing as they are fundamental in the construction and use of a conversational domain.

(A) Strictly, the modelling facility is described (thus also designed) only after the descriptor values have been elicited. Waivers to this order of precedence have been noted; in particular, in the caveat about the fictional character of the fashion in which the source and interrogator/analyst dialogue is conducted. But, even if the source has a modelling facility in mind as he exteriorises his thesis, this modelling facility must be regarded as tentatively specified until the thesis is fully expressed.

(B) Although independence between universes of interpretation for the PG(i) is determined, if a conversational domain is given, the elements of these sets (i.e., the universes of interpretation) are not exhibited except in a generative manner. That is, the universes of interpretation consist in sets with elements that satisfy the properties that are stipulated by the descriptor values.

(C) The student participating in a conversation anchored on this conversational domain understands certain topic relations. This understanding of a topic relation implies (amongst other things) that he is able to treat this topic relation as though it is a property (hence, as though it is a set) and may test elements for membership in this set or generate elements that are members of this set. It follows that the student can, in principle, make an indefinite number of models (interpreted programs) by way of non verbal explanation, even though the ossature of the universe of interpretation is firmly established.

6. Methods of Eliciting the Semantic Descriptors.

One advantage of giving priority in the mesh to isomorphic analogies is that it allows for a neatly and readily instrumented method of eliciting the semantic descriptors. The source is presented with triples of nodes consisting in an isomorphic analogy and its terms (Fig 3) or relata. The first triple belongs to the first level of the superordinate/subordinate descriptor. The triple in question is regarded as a triple of objects in the repertory grid technique (Kelly, 1956; Bannister and Mair, 1967) and the source is required, as in the repertory grid technique, to cite a construct (alias, a descriptor). However, the current method deviates from the standard technique in several particulars. The terms of an isomorphic analogy may have the same or different values on a descriptor; but it is easy to see that the relator node standing for the isomorphic analogy, has no "real" value (it has the "*" value), though it most certainly has a value on some other descriptor (that is, if the descriptor is regarded as a unary but many valued predicate and the relator is a relation between topic relations).

Call the descriptor having values on the terms of the isomorphic analogy (but * value on the relator or analogy node itself) the term descriptor and the descriptor with some value on the relator node (but * value on the term descriptor) the residual descriptor. Both the term descriptor and the residual descriptor are obtained as source inventions. The source is next required to give values (real or "*") on the term descriptor to all isomorphic analogy terms at the same superordinate/subordinate level and values (real or "*") on the residual descriptor to all relata nodes at the same superordinate/subordinate level.

In case there is a superordinate/subordinate level, perhaps the first, that contains no isomorphic analogy, descriptors and their values (real or *) are elicited over any triples of nodes. In case there are several isomorphic analogies at the same superordinate/subordinate level, they are dealt with in an sequence. In either case, the process of evaluating descriptors is iterated for all other-than-relator nodes in the conversational domain and the process of eliciting semantic descriptors is repeated, layer by layer, for all isomorphic

analogies in the conversational domain (repetition is obligatory until the descriptor values form a unique indexing system; after that, it is optional). Typical semantic descriptors are those cited already; for example, descriptors having values designating broad partitions of a subject matter and the departments of a discipline ("mechanics" and "electricity" in physics).

Given prior specification of the mesh, certain inferences are possible: essentially based on two ideas. First, that an isomorphic analogy between R_i and R_j exists if and only if $PG(i)$ is tolerance equivalent (Section 2) to $PG(j)$ but that $PG(i)$ and $PG(j)$ may be independently executed, either in the same or different universes of interpretation. Next, that any node (apart from the lower most nodes) has at least one fully conjunctive substructure: a kernel. Generally, a node has several kernels (Figure 4). The meaning of a kernel is as follows. Any model attached to node i (an interpretation of $PG(i)$ or some part of it) is the composition of models that are attached to nodes in a fully conjunctive kernel of node i; these kernel models being interpreted in the same universe of interpretation. For example, an electrical oscillator model is the composition of other electrical models. This assertion does not mean that an electrical oscillator cannot be derived from mechanical models, (PG interpreted in a mechanical rather than an electrical universe of interpretation). But if such a derivation is part of a thesis (and it usually would be, in practice) the electrical to mechanical derivation is manifest at some point as an isomorphic analogy, bridging electrical and mechanical domains.

As a result of these two lines of inference, certain semantic descriptor values can be filled in automatically as soon as some semantic descriptors and their values have been announced. The automatic valuation process must search through the kernels of all nodes given values and apply the following rules.

- (a) If a semantic term descriptor has a real value (other than *) on node i, then it has the same value on all members of at least one kernel of node i.
- (b) If a semantic descriptor has a real (other than *) value on node i, it has some real (other than *) value on all nodes in all kernels of node i.
- (c) If a term descriptor has the same value (x, say) on two terms in an isomorphic analogy, the modelling facility must contain two independent parts, of the same kind (to model the product $X \times X$).
- (d) If a term descriptor has different values (X and Y) on two terms in an isomorphic analogy, the modelling facility is divided into two universes of interpretation, X and Y, capable of modelling the product X and Y and differing in kind.
- (e) If so, the X, Y distinction, written Dist (X,Y) is the residual descriptor.

These rules are applied systematically by searching the kernels of all nodes for which descriptors have been elicited and given values. Some of the more important contingencies are illustrated in Figure 5 using TD to stand for "term descriptor" and RD to stand for "residual descriptor"; using "real" as any real (other than *) value and "+" and "-" for particular "real" values. For example, if TD has real values on A and B in Figure 5 (1) and RD has a real value on M we know that TD has * values on M and RD has * values on A and B. In Figure 5 (2) TD is given + value on A and - value on B then TD has a real value (+) on at least one kernel of A and a real value (-) on at least one kernel of B. Given the evaluation of TD on any node in some kernel of A (Figure 5 (3) the value of TD on the remaining nodes in this kernel is specified: similarly (Figure 5 (4)) for some kernel of B and similarly (Figure 5 (5)) for any other kernel (if it exists) of A or B. Finally, the value of RD must (if TD has distinct values X and Y on the terms of the isomorphic analogy) distinguish the universes of interpretation underlying X and Y. This distinction is shown in Figure 5 (6) as the special descriptor Dist (X, Y).


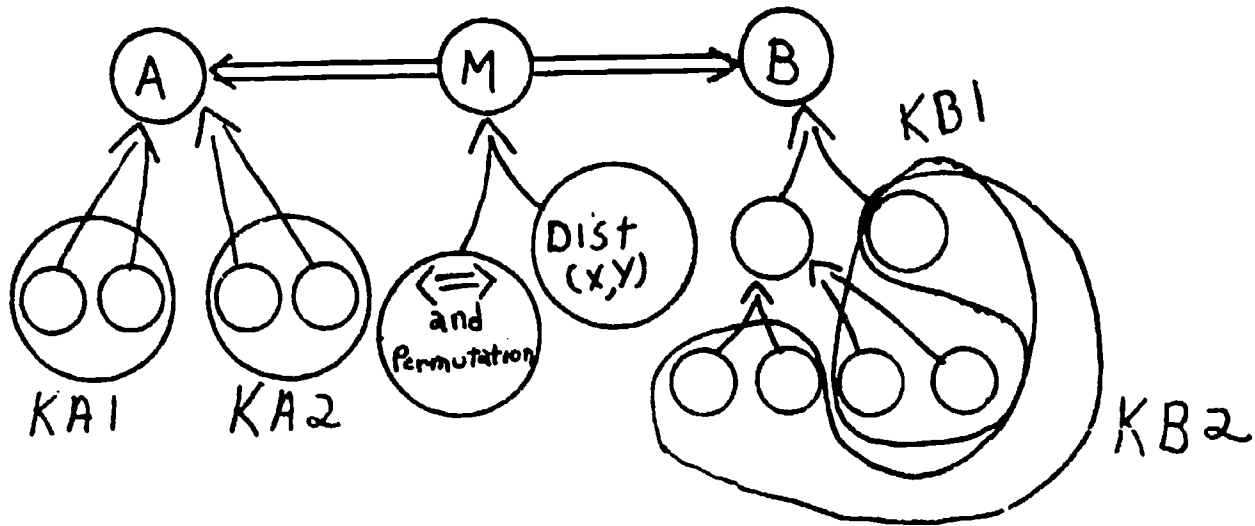


FIGURE 4. A node and its conjunctive kernels at arc distance depth 2.



Tabulation	A	B	M	KA1	KA2	KB1	KB2
1) TD	Re	Re	*	Re	Re	Re	Re
RD	*	*	Re	*	*	*	*
2) TD	+	-	*	+ if KA2 is -	+ if KA1 is -	- if KB2 is +	- if KB1 is +
3) TD	+	-	*	+	Re	- if KB2 is +	- if KB1 is +
4) TD	+	-	*	+	Re	-	Re
5) TD	X	Y	*	X	X or Y (Re)	-	X or Y (Re)
RD	*	*	Dist(X,Y)	*	*	*	*

FIGURE 5. Isomorphic analogy relation and table to concisely express major inferences. Each node A and B has a pair of conjunctive kernels; namely kernels (K) 1 and 2.

7. The Universe of Nodes or Placeholders.

We have stated that a modelling facility is (one or more) universe(s) of interpretation for the program graphs containing an intension for each topic relation (the $PG(i)$ attached to nodes i) and we have reviewed the "upwards down" description of this (or these) universe(s) of interpretation by contriving and evaluating semantic descriptions over the nodes $i \in I$. All this leaves a feeling of unreality about the nodes: they are a universe of placeholders which serves as a universe of interpretation for something: the question is, for what??

The question is properly answered by describing the transactions of a strict conversation (Section 1) and is glossed by the following (truthful but still rather enigmatic) comment. The nodes are a universe of interpretation for exteriorised learning strategies: Phrased differently, just as a modelling facility is a universe of interpretation for explanations (Section 2), albeit non verbal explanations, evident as model-building-operations, so the nodes are a universe of interpretation for derivations (Section 2). It will be recalled that an occasion in a strict conversation, is bounded by an understanding of some topic relation in the conversational domain and that a derivation, as well as an explanation is required to furnish evidence (Section 2) for an understanding.

This slightly esoteric notion is illuminated by sketching out the tangible arrangements used to exteriorise a derivation (alias, a learning strategy) as a strict conversation takes place with respect to the domain in question (Fig. 6, p. 19). The derivation is inscribed or registered by placing markers on the nodes, to indicate the completion of transactions and by retaining some markers in position. One transaction, for example, is an understanding: hence, under the proper circumstances, (explanation and derivation of a topic relation) a node is marked as being understood. Other transactions, such as aim (alias a focus of attention) and goal (one or several nodes that are currently being learned about, with a view to explaining the topic relations they denote) are only meaningful in the context of the conversation itself. In order to place such markers (in practice, they are often illuminated signal lamps) each node must be associated, as in Fig. 6, with marker value storage. A learning strategy (or equisignificantly a derivation) appears as a series of distributions of marker tokens and is visible to the conversational participants, as well as the outside observer.

8. The Evolution of the Conversational Domain.

The constructive operations, discussed in terms of a dialogue between the source and an interrogator/analyst have been partially mechanised. The source is still required to predicate or make distinctions but all the chores of the interrogation/analyst (and most of his advice giving) is equally well handled by a program (EXTEND, of Section 2 or THOUGHTSTICKER, of Section 2). Whether the process is handled mechanically or not, the conversational domain embodying the source thesis is static and used as the basis for a quite distinct strict conversation between participants (A and B of Section, 1, where B may be the CASTE heuristic, A a real student).

During the strict conversation the domain is frozen; a body of ossified knowables. Surely, the source could assess the domain afresh and modify his thesis; as a result, he might enlarge the domain without limit. But A, whilst he keeps the role of student, is unable to do so. The real strength of the mechanised procedures (EXTEND or THOUGHTSTICKER) is that they can be called as routines by the CASTE heuristic (acting as B) provided that A opts out of the student role and becomes a source or innovator. Under these circumstances, the conversational domain can evolve systematically, but, under user control (the user being, in this case, either A as student or A as innovator).

In this context, the EXTEND program operates as follows:

A, as student, sees a description of the conversational domain: Fig. 7 is a parody of such a thing, likening it to the grid lines on a map; each indexed cell being a conjunction of descriptor values. Some cells are occupied by nodes, some (by edict) are not. All nodes are ostended, in A, B transactions, by means of descriptor values.

At any instant, A directs his attention to one overall topic relation which he can describe: this topic relation is instated (by the CASTE heuristic) as his aim node. Consequently, A has addressed some cell as aim and, if A is in the role of student, any legitimate aim points to a cell that does contain a node. The act of changing role (student A, into innovator A) is instigated by A addressing an unfilled cell and informing B that he intends to fill it. If this event takes place, B calls for EXTEND with which A enters into a dialogue of the type already described (A in the capacity of source and EXTEND in the capacity of interrogator/analyst). Under these conditions just one unfilled cell may be filled at once (an extrapolation of the CASTE restriction to one aim node at once; or one focus of attention at once). But the cell may be located anywhere, neither necessarily or usually adjacent to the topic relations that A was learning about in the role of student. Hence, in order to permanently impress his systematic innovation upon the conversational domain, A is returned to the student role, and must demonstrate an ability to learn his freshly constructed topic relation on the basis of whatever topic relations he currently understands.

"One aim node at once" innovation is highly restrictive but, even so, non trivial. To exhibit this point it is necessary to recall the status of the descriptors: in particular, that at least one family of descriptors is understood by the student and that an aim node, even if the student is unable to learn about it, is meaningfully described. Fig. 8 illustrates the proper connotation of this remark for a conversational domain with only one family of descriptors. Arcs from all descriptor values enter each node of the mesh conjoined with other ingoing arcs; in other words, descriptors are simply relations that the student is guaranteed to understand and which he may thus treat as though they are properties. The relations dubbed descriptors and regarded as unary predicates of the nodes differ from topic relations which, in this sense, they glue together. Moreover, since they are described but unoccupied cells, the descriptors glue a thesis expressed in the conversational domain to other theses, that are not represented, but which may be invented. It is conceptually easy to generalize Fig. 8 to the usual case of a redundant description system with several families of descriptors, only one of which need be understood though several may be understood. The difficulty is graphical; arcs for nodes representing redundant descriptors ramify throughout the mesh and gain entry to the nodes of topic relations by disjunctive as well as conjunctive substructures.

Hence, the "one aim at once" operation of EXTEND is a non trivial device for sallying into a field of relations that may be described and assimilating them into the corpus of what may be known. The chief limitation of EXTEND is the one aim at once clause, due to which it is impossible (directly or explicitly) to introduce an isomorphic analogy, unless the terms of the analogy exist to begin with in the conversational domain. In this respect, EXTEND is more restricted than an actual dialogue between a source and an interrogator/analyst: since it is quite possible that the interrogator/analyst overrides the serial construction process.

THOUGHTSTICKER is a program designed to repair this deficiency of EXTEND: the heuristic it realizes is a mechanism for assimilating isomorphic analogies as a whole (terms and relata taken all together) in a format based upon "conditional problematic" statements such as: "If there was a topic U like topic V




FIGURE 6. Addition of marker storage locations to other connections of node i .

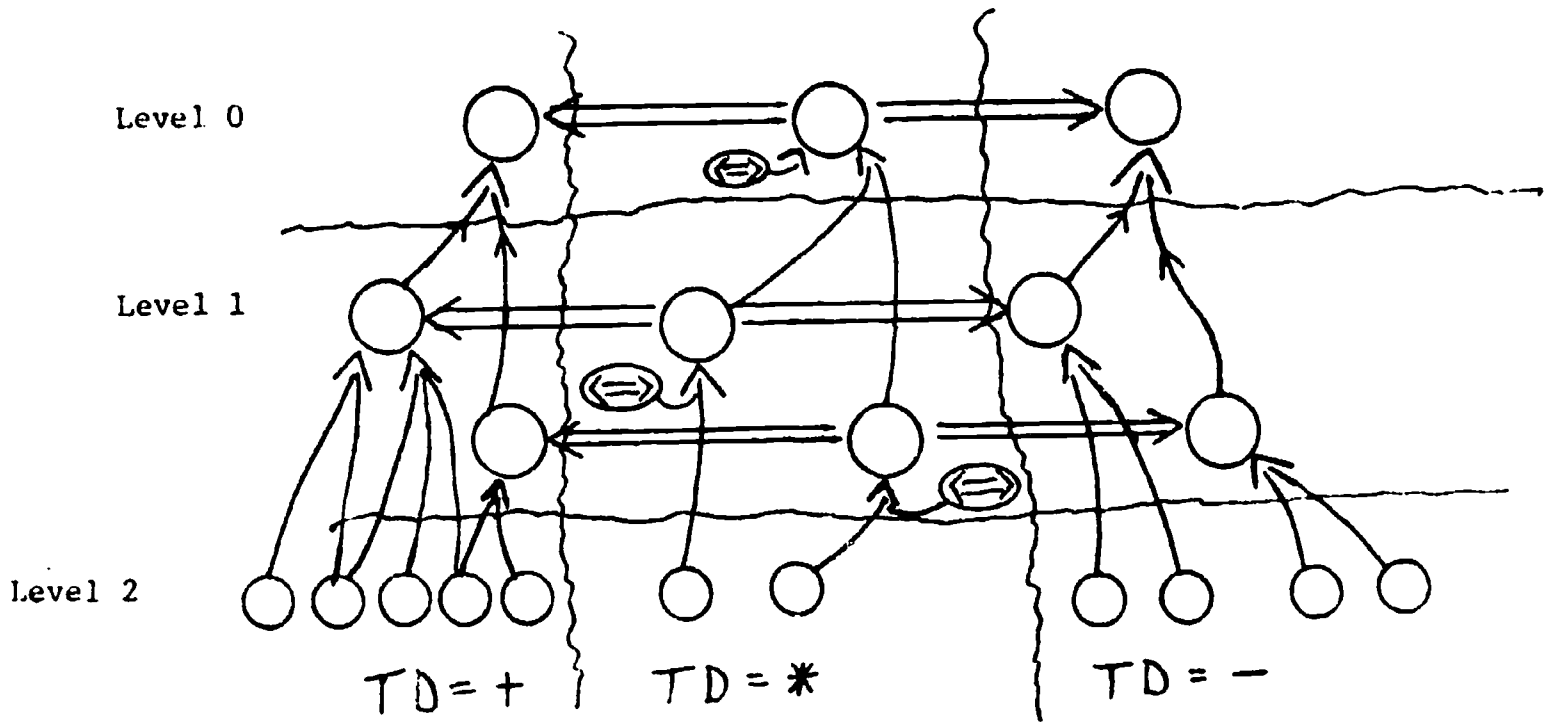


FIGURE 7. Descriptors arc distance and TD regarded as grid lines.




FIGURE 8. Alternative construction for FIGURE 7.

and they were isomorphically analagous...then certain requirements must hold relating U's coordinates and V's coordinates...if they do (which is mechanically checked) U and V are isomorphically analogous provided there is a distinction Dist X, Y that can be made and leads to the execution of PG(U) as well as PG(V) if the former is modelled in a universe of interpretation X and the latter in a universe of interpretation called Y." Pilot experiments indicate that THOUGHTSTICKER is a very powerful tool; for studying and encouraging "large" innovations. From a mechanical point of view, it is beset by the following difficulty. The THOUGHTSTICKER heuristic acts on the assumption that the innovator can have several simultaneous foci of attention; hence, any realization must embody several aim nodes. It follows that THOUGHTSTICKER cannot act as a routine of the one aim version of CASTE and the many aim versions developed to accommodate it are complex and currently very inefficient. However, the results are promising enough to outweigh these technical considerations, for many purposes. Further, the augmented system opens up possibilities beyond the compass of this paper. It is, perhaps, obvious that an organization able to accommodate several foci of attention on the part of one user is equally well able to accommodate several users with one aim (at least) for each user. This arrangement has been implemented and the innovative behaviors of groups appear to parallel, in a very interesting fashion, the behaviors accompanying the event called "large" innovation, on the part of only one user.

REFERENCES

- Banerji, R. B. Theory of problem solving. New York: American Elsevier, 1970.
- Bannistir, D. and Mair, J. The evaluation of personal constructs. New York: Academic Press, 1968.
- Kelly, G. A. The psychology of personal constructs, Vol. I & II. New York: Norton, 1955.
- Pask, G. Interaction between individuals: its stability and style. Math Biosciences, 11, pp. 59-84.
- Pask, G. Scott, B. C. E., & Kallikourdis, D. "Conversational domains." Submitted to A. I., 1973a.
- Pask, G. Scott, S. C. E., & Kallikourdis, D. Entailment and task structures for educational subject matter. Final Report, SSRC Project HR 1876/1, Systems Research Ltd., 1973b.

THE STUDY OF SYNTAX AND COGNITION: AN IDEA WHOSE TIME HAS COME

Barbara A. Hutson

State University of New York, Albany

It is sometimes useful for students of developmental psycholinguistics to view the development of their field as they would the development of a child--a child whose development reveals both his complex heritage and its interaction with patterns of experience. Like a child, the field progresses by coordinating diverse strands into an increasingly well-integrated whole. At times, new information from the environment is imperfectly assimilated, yielding rules which are initially over-generalized.

A child's overall cognitive organization may be characterized as in a given developmental stage, even though some of his responses echo an earlier period or foreshadow a later period. Analogously, a field of study may usefully be characterized as in a given state, even though some theorists anticipate a later period.

Earlier Assumptions about Syntactic Development

Several strands in recent research and theoretical development support the need for study of the relationship of syntax and cognition during their development. To place these trends in perspective, it is useful to review the assumptions accepted in the mid-sixties, assumptions which seem to rule out the possibility of any strong relationship between syntactic comprehension and general cognitive development.¹

1. The ages of acquisition were believed to be widely separated. McNeill in 1966 stated that a child's comprehension of syntax developed rapidly after age $1\frac{1}{2}$ and was essentially completed by age $3\frac{1}{2}$, in contrast to cognitive development, which matured later and more gradually.

2. Early and rapid development of syntax would seem to rule out a role for cognitive development in facilitating acquisition of linguistic competence, and by default lend some measure of credence to the hypothesis of innate knowledge of syntactic form.

3. The abstractness of a child's knowledge of syntactic form was contrasted with his decidedly concrete knowledge of other aspects of his environment at age 2-4 (McNeill, 1966). Again, accepting this assumption led quite easily to statements that because a child could not possibly have learned such abstract concepts, he must be born knowing them.

¹Although any of these comments would apply to syntactic production, the relationship of comprehension and cognitive development seems fundamental.

4. The adult speech serving as input to the child's language acquisition (with the exception of adult expansions of a child's sentences) was described as a "completely random, haphazard sample, in no way contrived to instruct a child on grammar" (McNeill, 1966). Hearing sentences which were rambling, broken, full of false starts and mazes, the child could only discover the local expression of linguistic universals if he were equipped with a set of "templates" (innate knowledge of syntactic form) against which to compare the speech he happened to hear from his parents.

5. Syntactic interpretation was assumed to occur prior to semantic interpretation and thus to be largely dependent upon syntactic markers and linguistic competence rather than semantic or extralinguistic considerations. In comprehending sentences, the surface structures was believed to be analyzed into structural features. Analysis of the syntactic deep structure then guided semantic interpretation of the lexical items within the underlying structures (Chomsky, 1965, p. 141). Since semantic interpretation was seen as secondary and subordinate, emphasis was placed on analyzing the child's grasp of abstract syntactic form, with minimal reference to meaning.

6. Chomsky's (1959) telling arguments against associationist learning theory as a credible explanation of a child's acquisition of syntactic comprehension were expanded by Fodor (1965). Imitation, reinforcement and generalization were shown to be inadequate to explain acquisition of the complexities of deep structure the child was believed to comprehend.

7. Guidelines for understanding the relationship between competence and performance were lacking. Chomsky's theory of transformational grammar, claimed as a theoretical base by many investigators of children's grammar in the early and mid-sixties, attempts to describe competence, the implicit knowledge of language structure possessed by a literate adult. Even for adults, there is a gap between competence and actual performance, which reflects competence, but is limited by performance factors such as memory, attention, motivation, etc. For children, such performance factors may be especially significant, and basing judgment of their competence on performance is hazardous (Chomsky, 1964).

For those who saw fully developed competence as innate, anything less than full adult competence could be attributed to performance factors. Others had the option of using a theory of competence as though it was a theory of performance or using performance as though it were an index of competence, either ignoring the inconsistencies or accepting the imperfect match as the best available. In any case, conditions did not facilitate studying the possibility that competence as well as performance develops over a period of time.

Recent Trends

Recent trends, however, bring some of these assumptions into question and suggest that the role of cognitive development in the acquisition of syntactic comprehension deserves careful study.

1. The ages of acquisition of syntax and cognition appear to overlap considerably. Piaget (1950) has traced the development of cognition from birth through adolescence. In the area of syntactic development, comprehension of simple sentences is present even before production of sentences. Bloom (1970) and Macnamara (1972) also point out that even before age $1\frac{1}{2}$, the child uses single-word utterances to express a variety of semantic relationships, and these relationships are reflected in the syntactic structures of early two-word sentences. Carol Chomsky (1969) and Palermo and Molfese (1972) point out significant development throughout the elementary school years. Kramer, Koff and Luria (1972) and Hutson and Shub (1974) note development of syntactic concepts which appears to extend to adulthood. The span of development of syntax is apparently long; it begins earlier and continues longer than McNeil had stated in 1966.

2. Evidence of such overlap in time of acquisition of syntax and cognition reopens the question of possible influence of development in one area upon development in the other. Not only is it possible that cognitive development provides a foundation for early acquisition of linguistic concepts, but, if development in both areas continues over a considerable period, it is also possible that at some points in development, cognitive development may provide support for development of syntactic comprehension, while at other points, maturity of certain syntactic concepts may facilitate further cognitive development (Jenkins, 1969).

3. For both syntax and cognition, the child's concepts appear to develop from concrete to abstract. Evidence has not been lacking that young children could display, when using concrete or familiar materials, capacity for performing some logical operations (Ricciuti, 1965; Piaget, 1952), and that their capacities gradually increase at least through adolescence. Many concepts in language also develop over several years (Brown, 1973; DeVilliers and DeVilliers, 1972). There appears to be a period of development spanning several years from the time when a concept is available under conditions of concrete support and minimal conflict, to the time when the concept is understood at such an abstract level that it is almost invulnerable to situational factors (Flavell and Hill, 1969; Hutson, 1973). In reviewing early statements concerning a child's abstract knowledge of syntax, it is important to question whether the criteria were demanding enough to warrant the conclusion that the child grasped the concept abstractly, or whether he was able to find in the context or in his experience sufficient support for a germinal concrete concept.

4. While it may be true that the speech heard at a convention of psycholinguists can be described as "haphazard," the speech input which is most directly relevant to a child's language acquisition appears to be more suitable for instruction. First, such speech is often embedded in the context of ongoing activity and familiar topics, providing a redundancy of cues for the interpretation of a sentence. Second, the speech directed toward young children by adults (Nelson, 1973) and by

other children (Shatz and Gelman, 1973) tends to be phrased in shorter, simpler sentences than is speech directed toward other listeners. These sentences often use forms just a little more advanced than the child's present usage, forms which will in a few months appear in the child's spontaneous utterances (Buium, 1973).

5. Two lines of evidence call into question the advisability of considering syntax in isolation from semantic and extralinguistic factors. Several theorists have disputed Chomsky's interpretation of the role of deep structure, suggesting either that the concept of deep structure is not needed or that the function of deep structure is as much semantic interpretation as syntactic analysis (Lakoff and Ross, 1967; Fillmore, 1968; McCawley, 1968; Olson, 1970).

From studies of children's language acquisition, there is abundant evidence of the utility of admitting semantic information early and explicitly in the analysis of children's sentences, and crediting children with the same tendency. Context is used by parents in interpreting children's utterances and by children in interpreting adults' sentences (Bloom, 1970; Carter, 1973; Rommetveit, 1968). Expectations based on a child's non-linguistic experience of the world may also affect his interpretation of syntactic structures (Slobin, 1971; Gowie, 1973; Hutson and Powers, 1974; Golinkoff, 1973). These trends toward richer interpretation and greater emphasis on semantic considerations, skillfully interpreted by Brown (1973) and by Olson (1970b), encourage viewing the acquisition of syntactic competence not as an isolated set of abstract concepts relatively little influenced by other areas of a child's development, but as a set of concepts intimately related to other areas of development, concepts which initially require considerable support from semantic and nonlinguistic features.

6. Strong currents in the area of learning insist that learning and association are not synonymous; while Chomsky's comments about the inadequacy of associationism to explain the acquisition of complex concepts may be valid, this would not rule out the richer formulations of learning set forth by theorists such as Gagne' (1966), Piaget (1950), or Jenkins (1973). While imitation and reinforcement may account for only limited aspects of language learning, a learning theory which views the child as actively constructing simple rules and hypotheses, based on all available sources of information, may have sufficient scope to deal with the rule-governed behavior required for comprehension of syntax.

There are, however, some provocative claims that transformational rules can be induced through modeling or imitation (Zimmerman and Rosenthal, 1974). It would be useful to test these methods against the framework developed from the many training studies of Piagetian logical concepts.

In the Piagetian training studies, it is common to find that although a response can be trained, an induced concept differs from a spontaneously developed concept in several ways. The induced concept (some would call it simply a situation-specific response) is less likely

to transfer to other materials or to related tasks, and is more easily extinguished through trick procedures (Smedslund, 1961). The prior learning or readiness of subjects is also an important factor. Subjects who initially show some signs of understanding the concept profit much more from training than do those who initially show no understanding. Our understanding of language development would be enhanced by examining whether there are similar possibilities and limitations in teaching linguistic concepts.

7. The relationship of competence and performance in syntax is still unresolved, despite laudable efforts by Cazden (1967), Fodor and Garrett (1966) and Flavell and Wohwill (1969). Borrowing liberally from these authors and extending their positions in directions suggested by the well-documented findings in cognitive development, the following synthesis is offered. It is possible that in syntax as well as in cognitive development, competence as well as performance develops gradually. A child is born with a few germinal capacities basic to language acquisition, and the genetically-endowed capacity to refine his initially gross hypothesis (Slobin, 1966). His competence increases as a result of his total experience of linguistic and nonlinguistic aspects of his environment, neural maturation, and his developing ability to organize his perceptions in increasingly complex fashion. Performance usually shadows competence, though the discrepancy is likely to be less after a concept has matured. Neither competence nor performance is monolithic; for both, the level achieved will differ for various concepts at any point in development.

Diverse strands, then, have converged, creating a climate which encourages exploration of the relationship of cognition and semantics to syntax during development. Semantics has been more broadly defined by Olson (1971), Schlesinger (1971) and Slobin (1970) than by Chomskian transformational grammarians such as Katz and Postal (1964). This aspect of knowledge appears to be quite closely related to overall cognitive growth, quite likely more closely related than is syntax. Semantics, thus broadly defined, may serve as a major component in syntactic acquisition.

Appropriate perspectives for viewing the relationship of semantics and general cognition are not yet well specified. For the moment, however, we leave this challenge to others, suggesting only that semantics and cognition may at first be largely undifferentiated but gradually semantics takes on a more specific character, while retaining close ties with general cognitive development. For purposes of the present discussion, semantics will be treated as an unspecified subset of general cognitive development.

Approaches to the Study of Syntax and Cognition

Even during the mid-sixties, although the general climate discouraged consideration of the possibility of an important relationship between

language and cognition, there were a few who expressed such a need. Sinclair de-Zwart (1969) stated that "linguistic universals exist precisely because thought structures are universal." Slobin (1966) suggested that a child was born, not with a set of linguistic categories, but with "an ability to learn certain types of semantic or conceptual categories" and "the knowledge that learnable semantic criteria can be the basis for grammatical categories."

In the past few years there have been several intriguing approaches to analysis of this issue. Huttenlocher's (1971) study of young children's manipulation of words provides insights as to similar sources of difficulty. A number of studies have attempted to analyze "the syntax of behavior," deriving "grammars" of rule-governed behavior in such diverse areas as cup-nesting (Greenfield, Nelson, and Saltzman, 1972), design-copying (Goodnow and Levine, 1973) and imitation of placements on a Piagetian landscape board (Pufall, 1966). These studies may lead to recognition of common knowledge structures in operations across diverse content areas while preserving the differences due to unequal difficulty of content and to individual differences.

Slobin (1970) sets forth bold, testable hypotheses as to probable cognitive correlates and foundations of syntactic operations, particularly in the early years. His thinking will surely stimulate much-needed analysis and empirical study, though results are not yet widely available. Perfetti's (1972) tough-minded analyses have stimulated spirited forays such as Golinkoff's (1973) innovative study of prelinguistic concepts of actor and agent. Brown (1973) sets forth initial postulates of non-linguistic concepts underlying early syntactic forms such as negation and questions, primarily in the sensorimotor period.

Parallels Between Syntactic and Logical Transformations

During the early school years, there is a rapid growth in cognitive development. It has become obvious that important growth in syntax also occurs during this period. Olson (1970a) points out that in both areas the child must overcome difficulties in dealing with transformations.

Both Chomsky and Piaget assign great importance to transformation. Without attempting here to deal in depth with either theory, it is sufficient for our present purpose to note certain parallels which suggest some common knowledge structures underlying behavior in these diverse content areas.

Transformations play a pivotal role in Chomsky's theory of transformational grammar, and on those modifications which insist on a greater role for the semantic component. The deep structure or relational structure of elements in a message is related to the surface structure, or actual expression of that message, by transformational rules which specify how that concatenation of elements may be expressed. The underlying structure of a message in an active, declarative sentence such as "John shut the door" can be transformed into a question "Did John shut

the door?" by placing the auxiliary verb do with its tense marker, in the initial position. The same underlying structure of elements can be expressed in the negative "John did not shut the door" or in the passive "The door was shut by John" by employing other transformational rules.

Children's initial attempts at expressing these relations do not reveal perfect knowledge. For example, the negative may at first be expressed as "Shut the door no" or "John no shut door," a question may be expressed as "How the toy works?" or as an active statement with rising intonation. If young children have fully developed competence, it is well disguised.

Children often comprehend a syntactic structure before they produce that structure accurately (Fraser, Bellugi and Brown, 1963). In naturalistic observation to date, however, it is important to consider whether cues other than syntactic comprehension could lead to the response from which comprehension is inferred.

In mastering their grammar, children learn the use and implications of various operations such as addition, deletion, and reordering of elements. The central concept acquired is the understanding of the means by which surface structure is related to its underlying relational structure by transformational rules. This is most difficult when previous experience and a new rule conflict, as in the passive voice, or in the exceptional structures studies by Chomsky (1969).

Under transformation, the structural relationship of sentence elements to one another is conserved. If the underlying structural relationship contains the proposition that A operates on B, the proposition is constant regardless of the specific transformational rule applied. That is, for the active sentence "Jane hit the ball," one might conceive of the negative "Jane did not hit the ball" as implying "It is untrue that Jane hit the ball." The question "Did Jane hit the ball?" implies that "It is uncertain that Jane hit the ball."^{1, 2}

Transformations are also pivotal in Piaget's theory of cognitive development. The transformation in logical conservation tasks that parallels linguistic transformation is not the physical change in appearance of the object, but the application of mental operations to various static configurations.

¹While these concepts can perhaps be conveyed more precisely by means of propositional calculus or linguistic notation, simple sentences are used here in order to communicate most readily.

²The relationship of elements within the proposition is unchanged. The specific transformational rules applied may, however, have implications for the relationship of the sentence to its social context, emphasizing various aspects of the proposition, indicating evaluative judgments about the proposition, or in other ways contributing to the total message.

In acquiring logical conservation, children gradually come to understand the system of relationships by which various dimensions are related. The base structural component, only gradually derived, is the system of interrelationships of various physical dimensions. The surface component consists of the actual object and the physical operations (such as addition, subtraction or rearrangement) performed upon it. The underlying relational structure and the surface structure are related by transformational rules such as reversibility.

If a child has not yet developed an understanding of the system of relationships of various dimensions, he tends to interpret the situation in terms of his general experience with physical objects - taller glasses usually contain more to drink than do smaller glasses, wider objects are usually heavier than narrower objects.

By attending to the dynamic process of transformation rather than to the static configurations of the object, the child comes to recognize the equivalence of dissimilar objects.

For example, when a ball of clay is changed into a sausage shape, the sausage no longer resembles the ball. Tracing the object through its various states, however, makes it possible to recognize that the process by which the change occurs can be reversed, changing the sausage back into a ball. This is possible only because the changes in appearance of the object do not change the underlying structural relationships. Changes in width, when compensated for by changes in height, do not change the weight or substance.

As noted earlier, it has become apparent that the time span of development of syntax and cognition overlaps considerably, from the sensorimotor stage through adulthood. There are indications that in neither area do all adults attain the same level of competence (Piaget, 1972; Kramer, Koff and Luria, 1972). In both areas, children appear to learn at first in very concrete ways and only very gradually do they achieve predominantly abstract understanding of cognitive or linguistic systems. At this time the factors affecting syntactic development have not been so well explored as those affecting general cognitive development (Piaget, 1970), but the findings from the many studies of spontaneous and induced formation of logical concepts provide a tentative model against which to test factors influencing syntactic development.

This analysis of possible similarities in some of the concepts which develop during the early school years is offered as a spur to further theoretical and empirical analyses, and an invitation to others to join in the difficult and intriguing task of exploring some of the relationships of language and thinking.

ACKNOWLEDGEMENTS

My thanks to Stephen Straight for his insightful comments, to Susan Schulman for her bibliographic assistance, and to Jacquelyn Goldman, who first encouraged my interest in language and thought.

REFERENCES

- Beilin, H. and Spontak, G. Active-passive transformations and operational reversibility. Paper presented at the Biennial Meeting, Society for Research in Child Development, March 26-29, Santa Monica, California, 1969.
- Bloom, Lois. Language development: Form and function in emerging grammars. Cambridge, Massachusetts: M.I.T. Press, 1970.
- Brown, Roger. A first language: The early stages. Cambridge, Massachusetts: Harvard University Press, 1973.
- Bulium, Nissan; Rynders, John; Wolf, Judith; and Turnure, James. Early maternal linguistic environment of normal and abnormal language learning children. Paper presented at American Psychological Association, Montreal, August, 1973.
- Carter, Anne L. Development of the pre-syntactic system: A case study. Paper presented at Society for Research in Child Development, Philadelphia, March, 1973.
- Cazden, Courtney B. On individual differences in language competence and performance. Journal of Special Education, 1: 135-150, 1967.
- Chomsky, Carol. The acquisition of syntax in children from 5 to 10. Research Monograph No. 57, Cambridge, Massachusetts: M.I.T. Press, 1969.
- Chomsky, Noam. Review of Skinner's (1957) Verbal Behavior. Language, 35, 26-58, 1959.
- Chomsky, Noam. Formal discussion. In U. Bellugi and R. Brown (Eds.), The acquisition of language, Monographs of the Society for Research in Child Development, 29, 92: 35-39, 1964.
- Chomsky, Noam. Aspects of a theory of syntax. Cambridge, Massachusetts: M.I.T. Press, 1965.
- DeVilliers, Peter A., and DeVilliers, Jill G. Early judgments of semantic and syntactic acceptability by children. Journal of Psycholinguistic Research, 1, 299-310, 1972.
- Fillmore, Charles J. The case for case. In E. Bach and R. T. Harms (Eds.), Universals in linguistic theory. New York: Holt, Rinehart and Winston, 1968, pp. 1-87.
- Flavell, J. C. and Hill, J. P. Developmental psychology. In P. Mussen and M. Rosenzweig (Eds.), Annual Review of Psychology. Palo Alto California: Annual Reviews, Inc., 20, 1-56, 1969.

- Flavell, J. H. and Wohlwill, J. F. Formal and functional aspects of cognitive development. In D. Elkind and J. H. Flavell (Eds.), Studies in cognitive development: Essays in honor of Jean Piaget. New York: Oxford University Press, pp. 67-120, 1969.
- Fodor, J. A. Could meaning be a r_m ? Journal of Verbal Learning and Verbal Behavior, 4, 73-81, 1965.
- Fodor, Jerry and Garrett, Merrill. Some reflections on competence and performance. In John Lyons and Roger J. Wales (Eds.), Psycholinguistic papers. Edinburgh: University of Edinburgh Press, pp. 115-120, 1966.
- Fraser, C., Bellugi, U., and Brown, R. Control of grammar in imitation, comprehension, and production. Journal of Verbal Learning and Verbal Behavior, 1963, 2: 121-135.
- Gagne, R. M. The conditions of learning. New York: Holt, Rinehart and Winston, 1965.
- Golinkoff, R. M. Semantic development in infants: The concepts of agent and recipient. Paper presented at Society for Research in Child Development meetings, March, 1973.
- Goodnow, J. J. and Levine, R. A. The grammar of action: Sequence and syntax in children's copying. Cognitive Psychology, 4, 82-98, 1973.
- Gowie, Cheryl J. Effects of children's expectation on mastery of the minimum distance principle. Paper presented at the American Educational Research Association, New Orleans, Louisiana, 1973.
- Greenfield, P. M., Nelson, K., and Saltzman, E. The development of rulebound strategies for manipulating seriated cups: A parallel between action and grammar. Cognitive Psychology, 3, 291-310, 1972.
- Hutson, Barbara A. How abstract is a young child's knowledge of syntax? Paper presented at American Psychological Association, Montreal, August, 1973.
- Hutson, Barbara A. and Powers, James E. Reversing irreversible sentences: Semantic and syntactic factors. Journal of Reading Behavior, in press.
- Hutson, B. A. and Shub, J. Developmental factors involved in choice of conjunctions. Paper presented at American Educational Research Association, Chicago, 1974.
- Huttenlocher, Janellen. Children's language: Word phrase relationships. In A. Bar-Adon and W. F. Leopold (Eds.), Child language: A book of readings, Englewood Cliffs, New Jersey: Prentice-Hall, Inc. 1971.

- Jenkins, James J. Language and thought. In J. F. Voss (Ed.), Approaches to thought. Columbus, Ohio: Charles E. Merrill, pp. 211-237, 1969.
- Jenkins, James J. Remember that old theory about memory? - Forget it. Paper presented at American Psychological Association, Montreal, August, 1973.
- Katz, J. J. and Postal, P. M. An integrated theory of linguistic descriptions. Cambridge, Massachusetts: M.I.T. Press, 1964.
- Lakoff, G. and Ross, J. R. Is deep structure necessary? Cambridge, Massachusetts: Author, 1967.
- Macnamara, John. Cognitive basis of learning in infants. Psychological Review, 79: 1-14, 1972.
- McCawley, J. D. The role of semantics in grammar. In S. Bach and R. T. Harms (Eds.), Universals in linguistic theory, New York: Holt, Rinehart and Winston, 124-169, 1968.
- McNeill, David. Developmental psycholinguistics. In E. Smith and G. Miller (Eds.), The genesis of language: A psycholinguistic approach, Cambridge, Massachusetts: M.I.T. Press, 15-84, 1966.
- Menyuk, Paula. Syntactic rules used by children from preschool through first grade. Child Development, 35: 533-546, 1964.
- Nelson, Katherine. Structure and strategy in learning to talk. Mono-graphs of the Society for Research in Child Development, 1973, 38 (1-2).
- Olson, D. R. Language acquisition and cognitive development. In H. C. Haywood (Ed.), Social cultural aspects of mental retardation - proceedings of the Peabody-NIMH conference. New York: Appleton-Century-Crofts, 113-202, 1970a.
- Olson, David R. Language and thought: Aspects of a cognitive theory of semantics. Psychological Review, 77: 257-273, July, 1970b.
- Palermo, David S. and Molfese, Dennis L. Language acquisition from age five onward. Psychological Bulletin, 78, 1: 409-428.
- Perfetti, C. A. Psychosemantics: Some cognitive aspects of structural meaning. Psychological Bulletin, Vol. 78, 240-259, No. 4, October, 1972.
- Piaget, Jean. Intellectual evolution from adolescence to adulthood. Human Development, 15: 1-12, 1972.

- Piaget, Jean. Piaget's Theory. In P. H. Mussen (Ed.), Carmichael's manual of child psychology, New York: John Wiley and Sons, Inc., 1970, 1, 703-732.
- Piaget, Jean. The origins of intelligence in children (1st edition, 1936). New York: International University Press, 1952.
- Piaget, Jean. The psychology of intelligence. Translated by M. Piercy and DiBertyne, London: Routledge, 1950.
- Pufall, P. B. Acquisition and generalization of spatial order conservation in young children. Unpublished Ph. D. thesis, Catholic University, Washington, D. C., 1966.
- Ricciuti, Henry N. Object grouping and selective ordering behavior in infants 12 to 24 months old. Merrill-Palmer Quarterly, 11, 129-147, 1965.
- Rommetveit, R. Words, meanings and messages: Theory and experiments in psycholinguistics. New York: Academic Press, 1968.
- Schlesinger, I. M. Production of utterances and language acquisition. In D. I. Slobin (Ed.), The ontogenesis of grammar, New York: Academic Press, 63-101, 197.
- Shatz, M. and Gelman, R. The development of communication skills: modifications in the speech of young children as a function of listener. Monographs of the Society for Research in Child Development, 1973, 38 (5).
- Sinclair-de-Zwart, H. Developmental psycholinguistics. In D. Elkind and J. H. Flavell (Eds.), Studies in cognitive development: Essays in honor of Jean Piaget, New York: Oxford University Press, 315-336, 1969.
- Slobin, Dan I. Comments on "developmental psycholinguistics." In F. Smith and G. A. Miller (Eds.), The genesis of language, Cambridge, Massachusetts: M.I.T. Press, 85-91, 1966.
- Slobin, Dan I. Universals of grammatical development in children. In G. B. Flores d'Arcais and W. J. M. Levelt (Eds.), Advances in psycholinguistics. Amsterdam: North Holland, 1971.
- Slobin, Dan I. Suggested universals in the ontogenesis of grammar. Report #WP-32, Language and Behavior Research Laboratory. Berkeley: University of California, April, 1970.
- Smedslund, J. The acquisition of conservation of substance and weight in children: III. Extinction of conservation of weight acquired "normally" and by means of empirical controls on a balance. Scandinavian Journal of Psychology, 1961, 2: 85-87.
- Zimmerman, B. J. and Rosenthal, T. I. Observational learning of rule-governed behavior by children. Psychological Bulletin, 1974, 81: 29-42.

A SYSTEMS AND SET THEORY MODEL FOR THE EMPIRICAL INVESTIGATION OF PERCEPTION AND STRUCTURED KNOWLEDGE

M. Sue Bogner
Catholic University of America

Various terms such as perception, cognition, structure, knowledge, behavior have been used in the psychological literature in reference to the person's functioning in the environment. These terms have not been rigorously interrelated and empirically defined, and because of that have hindered the scientific study of the functioning person as an integrated whole. It is to that problem that this paper is addressed.

As logic represents man's ordering of knowledge, and set theory provides the algebra and geometry for logic, it follows that set theory provides a valid means of expressing the processes by which a person orders knowledge, the cognitive processes. Assuming that a person is not endowed with innate ideas, but a capacity for understanding his environmental context, then the process of ordering knowledge includes the process of internalizing aspects of the environment by the process of perception. These processes which have long been debated by philosophers have also been considered by psychologists. Because the processes of cognition and perception are internal to the person and not objectively quantifiable per se, many psychologists refer to them as hypothetical constructs; still others, the psychologists of the behaviorist school, refer to such processes as "explanatory fictions". Yet, those processes may indeed be systematically assessed and quantified. This is accomplished by considering the person as a "black box" system.

By defining the person as a system in terms of set theory, a viable model for the characterization of knowledge structures which underlie behavior and their relation to perception is herein explicated. This model provides the means for the empirical investigation and understanding of perception and structured knowledge.

The person as a system

A system is described by its internal and external quantities. These quantities are considered analogous respectively to the subjective "I" and the objective "me" of the duplex self (James, 1890). The person as a system is denoted P . The "I", the phenomenological, intra-personal, subjective, active self as knower which constitutes the internal quantities of the system is denoted P' . The "me", the extrapersonal self as known, the object of other people's perceptions which constitute the external quantities of the system is denoted P'' .

The system is activated by the input of the external, extrasystemic stimulus quantities of environmental events which occur independent of the system. The input interacts with the internal quantities, or the "I", P' , to cause intrasystemic events or processes. These processes which are dependent upon the nature of the system, P , produce the output or response quantities which define external quantities or the "me" of the system, P'' . The relation between external quantities as activities of the system defines behaviors of the system. The behaviors are indicative of properties of the system's, the person's organizational nature, its structure.

Internal quantities

Basic to the consideration of the person as a system is a set of assumed characteristics of the nature of that system. Each assumption (such as those previously

stated herein) is denoted p_i , $i = 1, 2, \dots, n$; and contained in the set of all assumptions P :

$$P = \{p_1, p_2, \dots, p_n\}.$$

The assumptions are incorporated in the specification of the internal quantities in space and time. Let $p'(t)$ represent the assumed internal quantities which specify P in space at time t contained in the set of all such times T ; the set of all $p'(t)$ is denoted $P'(t)$ and is expressed:

$$P'(t) = \{p'(t) | \forall t \subseteq T\} \quad p'(t) \leftrightarrow p \in P$$

An example of such a set of assumed and specified internal quantities, $P'(t)$, is the presumed process of internalizing the environment at time t .

The quantities of the set $P'(t)$ are further defined by the values of a specific discipline denoted by the subscript i . Those values are of psychology, more specifically those of theoretical psychology. The set of all such psychological values of internal quantities denoted $P'_i(t)$, represent the psychological "I" of the person. This set contains the values of the assumed internal quantities of the system as specified in space and time, $p'_i(t)$, such that for every t contained in T there exists a psychological value of internal quantities which correspond to specified assumed quantities of the system for all quantities $i = 1, 2, \dots, n$:

$$P'_i(t) = \{p'_i(t) | \forall t \subseteq T\} \quad p'_i(t) \leftrightarrow p'(t) \subseteq P'(t), \quad i = 1, 2, \dots, n.$$

The psychological values applied herein are those of the theory of personal constructs (Kelly, 1955). Such values are: a person's processes are psychologically channelized by the anticipation of events, events are anticipated by construing their replications, the constructs are related to form a system, the construct system may vary as the person successively construes the replications of events.

Because the person is considered a "black box" system, the internal quantities, indeed the psychological values of those quantities are only conjecture until evidenced in the output of the system, those independently observable external quantities representative of the "me" of the system.

External quantities

The output is those quantities produced by the system which exert on the environment. Various aspects of the output of the system, denoted x_ℓ , $\ell = 1, 2, \dots, m$; are contained in the set of all output quantities denoted X :

$$X = \{x_1, x_2, \dots, x_m\}.$$

Each member of the set is essentially a subset, i.e., x_1, x_2, \dots, x_m , each representing a class of output such as verbal statements, physical gestures, movements. The output of the system represents the external quantities of the system when specified in space and time in terms of a resolution level. The resolution level determines the accuracy and frequency of observation of the external quantities. The accuracy of observation is the consideration of a quantity with respect to an ideal value of that quantity. For example, for a gesture, the accuracy would be the extent to which that specific

gesture x_1 is representative of the set of all gestures X_2 produced by the system, the person. The frequency of observation is the time for sampling the considered quantities. The set of all such resolution levels, L , contains the output quantities of the system observed according to standards of accuracy and frequency. The resolution level is formally expressed:

$$L = \{X_1, X_2, \dots, X_m, T\}$$

and serves to define the external quantities, the "me" of the system at time t :

$$P'' = \{x_1, x_2, \dots, x_m, t, X_1, X_2, \dots, X_m, T\}.$$

By substitution of the sets X and L , the general expression of the external quantities, P'' , is simplified to the terms of output, X , observed at the resolution level L at time t :

$$P'' = \{X, t, L\}.$$

The aspects of the external quantities defined by X , t , L , are denoted p''_ℓ with $\ell = 1, 2, \dots, m$, and are contained in the set of all such aspects, P''_ℓ which is expressed:

$$P''_\ell = \{p''_\ell \mid \forall p''_\ell \exists p''_\ell \in X \text{ for all } \ell = 1, 2, \dots, m\}.$$

An example of P''_ℓ is further defined by values of a discipline which are denoted by the subscript j . In this case, those values are of psychology, more specifically the values of psychological methodology. The set of all such values of psychological methodology for the external quantities of the system is denoted P''_j and contains the values of the external quantities p''_j , such that for every p''_j there exists an external quantity p''_ℓ contained in the output X of the system:

$$P''_j = \{p''_j \mid \forall p''_j \exists p''_\ell \in X\}.$$

The set P''_j operationally defines the external quantities of output. An example of P''_j is the comment of P'' represented by the numerical characteristics of an adjective in a given adjective check list, i.e., 2 = descriptive, 1 = not descriptive, 0 = not applicable. It is to be noted that the values of the external quantities may or may not represent or point to consistency in the external quantities of the system. To correct for this, the values of several quantities are sampled in time as activities of the system.

The activity of a system represents the "me" of the system by the ensemble of variations in time of some of the values of the external quantities, p''_j , under consideration. The value of the activity of the system representing the psychological "me" as resolved at time t , $p''_j(t)$, contained in the set of all values of the resolved external quantities, denoted $P''_j(t)$, such that for time t contained in T , there exists a value of external quantities p''_j , for all $j = 1, 2, \dots, m$, corresponding to an external quantity p''_ℓ contained in the set of all output X is expressed:

$$P''_j(t) = \{p''_1(t), p''_2(t), \dots, p''_m(t) \mid t \in T, p''_j \in P''_j \text{ for all } j = 1, 2, \dots, m\}.$$

For example, the psychological values are 2 = descriptive, 1 = not descriptive, 0 = not

applicable, and the set P''_ℓ contains $p''_1 = \text{happy}$, $p''_2 = \text{sad}$, $p''_3 = \text{warm}$, $p''_4 = \text{cool}$, with the values of $p''_j(t)$ respectively, as 2, 0, 2, 1 (see Table 1).

Table 1
Example of activity matrix for $P''_j(t)$
with values 0, 1, 2 at time t for $\ell = 1, 2, 3, 4$

$P''_\ell =$		t
	p''_1 happy	2
	p''_2 sad	0
	p''_3 warm	2
	p''_4 cold	1

Activity is further defined in terms of a considered time interval, $t(0 \leq t \leq t_{\max})$ as a time-invariant relation. This is expressed as the set:

$$P''_j(0 \leq t \leq t_{\max}) = \{p''_j(t) \mid \forall t(0 \leq t \leq t_{\max}) \in T \exists p''_j \leftrightarrow p''_\ell \in X, j = 1, 2, \dots, m\}.$$

The time-invariant relations are between the values of members of P''_ℓ which are satisfied within a considered time interval $t(0 \leq t \leq t_{\max})$. For the case (Table 1) in which the considered time interval is t , the time-invariant relation is the relation between the values of the members of $P''_j(t)$ or the relation between $p''_1(t)$, $p''_2(t)$, $p''_3(t)$, and $p''_4(t)$.

A time-invariant relation between the values of the subset members of $P''_j(t)$ is a set contained in the Cartesian product of the subsets. Considering the subsets $p''_1(t)$ and $p''_2(t)$, their time-invariant relation r is a subset of their product:

$$r \subseteq P''_1(t) \times P''_2(t).$$

More specifically, the considered relation, $r(P''_1(t), P''_2(t))$ represents the mutual investment of each of the given values of adjectives in a set of ordered pairs such that for every $p''_1(t)$ contained in $P''_1(t)$ and every $p''_2(t)$ contained in $P''_2(t)$, there exists an ordered pair $(p''_1(t), p''_2(t))$ contained in or equal to the set which comprises the relation:


$$r(P''_1(t), P''_2(t)) = \{(p''_1(t), p''_2(t)) \mid \forall (p''_1(t), p''_2(t)) \exists p''_1(t) \in P''_1(t) \text{ and } p''_2(t) \in P''_2(t)\}.$$

The relation for the two sets of the example as illustrated by Venn diagrams in Diagram 1 appears as the area of commonality or similarity defined by the intersection of those sets.

The time-invariant relations between the external quantities with respect to the time of the consideration of the system, and the ensemble of variation, the chosen pattern of samples of the activity, is represented by the set of all mutually different samples of a certain pattern that belong to a considered time interval. The chosen pattern of activity is the set P''_ℓ , the ensemble of variation is the sampling of that pattern within the time interval $t(0 \leq t \leq t_{\max})$.

Diagram 1

Venn diagram of the mutual investment of the sets $P''_1(t)$, $P''_2(t)$ defined by their intersection contained in their Cartesian product



It is to be noted that a given time t defines a specific environmental context. As time is sampled, ($0 \leq t \leq t_{max}$), a variation may occur in the environment. Because input to the system is the exertion of the environment on the system, variation in the environment suggests variation in input. The converse of this also holds, variations in the input to the system implies variations in the environment, hence variation in the considered times.

The chosen pattern of activity of a system is expressed as a matrix with the columnar headings as the considered times, t , (implying input) and the row headings of the external quantities, p'_l . The values $p''_j(t)$ are entered in the cells. For example, the activity of a system is sampled at time $t_1 = 1$ minute, $t_2 = 2$ minutes, $t_{max} = 3$ minutes with respect to comments as adjectives in the set P''_l : $p''_1 =$ happy, $p''_2 =$ sad, $p''_3 =$ warm, $p''_4 =$ cold. The values of those adjectives, p''_j , are 2 = descriptive, 1 = not descriptive, 0 = not applicable which when entered in the cells of the matrix become $p''_j(t)$. The matrix for the example appears as Table 2.

Table 2
Example of activity matrix for $P''_j(t)$ with values 0, 1, 2, when ($0 \leq t \leq t_{max}$) and $l = 1, 2, 3, 4$

		t_1	t_2	t_{max}	
		1 min.	2 mins.	3 mins.	
$P''_l =$	p''_1	happy	2	2	2
	p''_2	sad	0	0	0
	p''_3	warm	2	1	2
	p''_4	cold	1	1	1

The time-invariant relation in the example in Table 2 is the pattern of relations among the members of P''_l sampled at each t . Such a sample may be taken for past values, present values, and future values of P''_l . The time of such sampling, denoted ($t + \alpha$), is contained in the set of all time T and refers to the samples of instantaneous values when $\alpha = 0$, past values when $\alpha < 0$ and future values when $\alpha > 0$. The

time-invariant relations representing the "me" of the system, $P''_j(t + \alpha)$, is the set of resolved values of external quantities sampled for $(t + \alpha)$ such that for every $(t + \alpha)$ contained in T , with $\alpha <, =, > 0$, there exists a corresponding activity at a given time:

$$P''_j(t + \alpha) = \{p''_j(t + \alpha) \mid \forall (t + \alpha) \in T\} \quad p''_j(t + \alpha) \leftrightarrow p''_j(t) \in P''_j(t),$$

$$j = 1, 2, m; \quad \alpha <, =, > 0\}.$$

To meet the conditions of $(t + \alpha)$, the activity of the system might be sampled for given input quantities presented at the t_1, t_2, t_{\max} under varying instructions. Instantaneous values of P''_ℓ , for the condition of $(t + \alpha)$, $\alpha = 0$, are found by instructing the system, the person, to produce activity as describing the input of objects, events and/or persons given at t_1, t_2, t_{\max} as they are "now". The future values of P''_ℓ , the condition for $(t + \alpha)$, $\alpha > 0$, may be found by instructing the system to describe the input at t_1, t_2, t_{\max} as they would be a year from now. The past values of P''_ℓ , the conditions for $(t + \alpha)$, $\alpha < 0$, may be found by the instructions to describe t_1, t_2, t_{\max} as they were a year ago. The values of P''_ℓ for the conditions of $(t + \alpha)$, $\alpha <, =, > 0$ may also be found by actually sampling the activity of the system elicited by the input t_1, t_2, t_{\max} at three different times. These times are defined as the past, $\alpha < 0$, and the future, $\alpha > 0$, with respect to the intermediate sampling designated as the present, $\alpha = 0$.

For example, the activity at t described as Table 2 becomes a sample of activity of the present values of P''_ℓ at $(t + \alpha)$, $\alpha = 0$. The time-invariant relation is described by three sets of relations among the values of P''_ℓ : those among $p''_j(t_1)$, $p''_j(t_2)$, and $p''_j(t_{\max})$ represented by $P''_j(t_1)$, $P''_j(t_2)$, $P''_j(t_{\max})$ respectively. The relations defined by the intersection of those three sets within each sample for $(t + \alpha)$, $\alpha <, =, > 0$, serve to define three classes of time-invariant relations. In terms of future values of P''_ℓ represented as the matrix of Table 3, that intersection of $P''_j(t_1)$, with $P''_j(t_2)$ and $P''_j(t_{\max})$ for $(t + \alpha)$, $\alpha > 0$ is the value of $P''_j(t + \alpha)$ or happy as descriptive.

Table 3
Example of activity matrix for $P''_j(t + \alpha)$ with values 0, 1, 2 at time $\alpha > 0$ for $\ell = 1, 2, 3, 4$

		t_1	t_2	t_{\max}
		1 min.	2 mins.	3 mins.
$P''_\ell =$	p''_1 happy	2	2	2
	p''_2 sad	0	0	1
	p''_3 warm	2	0	2
	p''_4 cold	1	2	0

The three time-invariant relations distinguished by the activity of the values of the external quantities are those: 1) within a short time interval of a particular activity, the local relation, containing values of P''_ℓ at $(t + \alpha)$ when $\alpha <$ or $=$ or > 0 ; 2) anywhere within a particular activity, the relative relation, containing values of P''_ℓ at $(t + \alpha)$ for $\alpha <$ and $= 0$, or $\alpha <$ and > 0 , or $\alpha =$ and > 0 ; 3) over the entire time interval of every possible particular activity containing the quantities, the

absolute relation containing values of P''_{ℓ} at $(t + \alpha)$ for $\alpha <$ and $=$ and > 0 . An example of the relative relation is the relation between the samples for $t + \alpha$ at $\alpha = 0$ in Table 2 and at $\alpha > 0$, Table 3. That relative relation contains the value 2 (descriptive) for p''_j happy. An example of the activity matrix for past values of P''_{ℓ} appears as Table 4.

Table 4
Example of activity matrix for $P''_j(t + \alpha)$ with values 0, 1, 2 at time $\alpha < 0$ for $\ell = 1, 2, 3, 4$

		t_1	t_2	t_{\max}
		1 min.	2 mins.	3 mins.
$P''_{\ell} =$	p''_1 happy	2	2	2
	p''_2 sad	0	0	0
	p''_3 warm	2	1	0
	p''_4 cold	1	0	0

The time-invariant relations are incorporated into the behavior of the system which serves to define the link between the psychological "me" and psychological "I" of the system, the person. The behavior of the system is the correspondence of a member of the set of psychological values of the external quantities in terms of time-invariant relations, $p''_j(t + \alpha)$ with a member of the set of psychological values of the internal quantities of the system, $p'_i(t)$. This one-to-one correspondence is denoted $(j, \alpha) \leftrightarrow i$ and allows the expression of the intrinsic relation of external to internal quantities:

$$p''_j(t + \alpha) = p'_i(t) \text{ when } (j, \alpha) \leftrightarrow i,$$

where $p''_j(t + \alpha)$ is the psychological value of the quantity p'_i at time t , $(t + \alpha) \in T$;

$$p''_j(t + \alpha) \in P''_j, p'_i(t) \in P''_j, i = 1, 2, \dots, n \text{ and } j = 1, 2, \dots, m.$$

The behavior of the system denoted B , in terms of the "I" of the person, the psychological value of the internal quantities, is now expressed in terms of corresponding value of external quantities:

$$B = \{b \mid \forall b \exists p'_i(t) \leftrightarrow p''_j(t + \alpha), t(t + \alpha) \in T \text{ when } (j, \alpha) \leftrightarrow i, p''_j(t + \alpha) \in P''_j \text{ and } p'_i(t) \in P''_j, i = 1, 2, \dots, n, j = 1, 2, \dots, m, \text{ and } \alpha <, =, > 0\}.$$

The value $P''_j(t)$ as behavior when partitioned by time represents the traits of the system as a past value of P'_i at time t (for $(j, \alpha) \leftrightarrow i, \alpha < 0$), the instantaneous value of P'_i at time t (for $(j, \alpha) \leftrightarrow i, \alpha = 0$), or a future value of P'_i at time t (for $(j, \alpha) \leftrightarrow i, \alpha > 0$). Because $p'_i(t)$ is contained in P''_j and P'_i represents the set of all possible values of the quantity p'_i , then $P'_i = P''_j$ assuming $(j, \alpha) \leftrightarrow i$ for any α .

Behavior is defined by a relation of values of the external quantities. Such a relation is a subset of the Cartesian product of $P''_1 \times P''_2 \times \dots \times P''_m$ and is formally expressed:

$$B = \bigcup_{1 \leq i \leq n} r(P^i_j) \subseteq \bigcup_{j=1}^m P^j$$

It is determined from the temporal context of the traits of the system, whether the behavior is permanent, relatively permanent, or temporary.

Three basic types of behavior, B_1 , B_2 , B_3 , derived from the three types of time-invariant relations of the activity of the person are distinguished:

1. Permanent behavior, denoted B_1 is indicated by the absolute relation comprised of the set of all local relations of every possible activity containing sets of values of instantaneous and past and future values of external quantities expressed by psychological methodology which directly correspond to or operationalize the sets of values of internal quantities of the person:

$$B_1 = \bigcup_{1 < i < n} r(P^i_j) \subseteq \bigcup_{j=1}^m P^j(t + \alpha) \text{ when } (j, \alpha) \leftrightarrow i \text{ with } \alpha < 0 \text{ and } \alpha = 0 \text{ and } \alpha > 0.$$

An example of B_1 is offered when examples of activity at time t (Table 2) which implies $t + \alpha$ when $\alpha = 0$, and at time $(t + \alpha)$ with $\alpha > 0$ (Table 3) and at time $(t + \alpha)$ with $\alpha < 0$ (Table 4) are considered as an absolute relation. This relation is the area of commonality between the three activity matrices, or in terms of the example, the use of p^1_1 , happy as value 2 or descriptive across all considered times for each $t + \alpha$, $\alpha <, =, > 0$.

2. Relatively permanent behavior denoted B is indicated by the relative relation comprised of the set of all local relations of a particular activity which operationally defines the sets of psychological values of the internal quantities:

$$B = \bigcup_{1 \leq i \leq n} r(F^i_j) \subseteq \bigcup_{j=1}^m P^j(t + \alpha) \text{ when } (j, \alpha) \leftrightarrow t \text{ with } \alpha < 0 \text{ and } \alpha = 0, \text{ or } \alpha < 0$$

and $\alpha > 0$, or $\alpha = 0$ and $\alpha > 0$.

An example of B_2 is provided by the relative relation between $(t + \alpha)$ for $\alpha < 0$ and $\alpha = 0$, for the samples of activity matrices in Tables 2 and 4. It is evident that B_2 contains $p^j_1(t + \alpha)$ of the value of 2 (descriptive) for p^1_1 or happy and the value 0 (not applicable) for p^2_2 or sad for $\alpha < 0$ and $\alpha = 0$.

3. Temporary behavior denoted B_3 is indicated by the local relation comprised of an aspect of a particular activity which operationalizes the sets of internal quantities of the person:

$$B_3 = \bigcup_{1 \leq i \leq n} r(P^i_j) \subseteq \bigcup_{j=1}^m P^j(t + \alpha) \text{ when } (j, \alpha) \leftrightarrow i \text{ with } \alpha < 0, \text{ or } \alpha = 0, \text{ or } \alpha > 0.$$

Table 2 provides an example of B_3 in the local relation of the value of 2 (descriptive) for p^1_1 or happy, the value 0 (not applicable) for p^2_2 or sad, and the value of 1 (not descriptive) for p^4_4 or cold for $\alpha = 0$.

The behaviors, B_1, B_2, B_3 , define three classes of properties of the system. The properties of a system comprise its organization, the nature of the person. The properties of the system are the properties of the internal quantities, the "I" or "P", as manifest in and inferred from the "me", "P" of the system. As previously defined, the behaviors of the system may be consistent over time as well as transitory. To allow inference from the behavior, through the activities, $P''_j(t + \alpha)$, to the properties, $P'_j(t)$, the properties of the organization of the system are differentiated into constant and variable components which, respectively, are the structure and the program of the system.

The program is inferred from behaviors, B_3 , sampled at time $(t + \alpha)$ when $\alpha = 0$, temporary behaviors. The structure is further differentiated into components designated as the real structure which is inferred from B_1 and the hypothetic structure which is inferred from B_2 . In terms of the dimension of time, the real structure is indicated by behavior across all time for the system, $(t + \alpha)$, $\alpha <, =, > 0$; Whereas the hypothetic structure is manifest in behavior in the past and present, $(t + \alpha)$, $\alpha <, = 0$; past and future, $\alpha <, > 0$; present and future $\alpha =, > 0$.

An example of the program is not only the consistency of the relation of P''_1 to P''_2, P''_3 and P''_4 for a given t , but also P''_j across t_1, t_2 , and t_{max} for $j = 1, 2, 3, 4$. In Table 2 this relation indicates that for temporary behavior, the adjective happy is considered descriptive while the adjective sad is considered not descriptive and the adjective cold is not applicable. This suggests the use of a dimension defined by happy and cold, and the inapplicability of sad in the temporary behavior of the system in its transitory processing of input. Even though the program is transitory, it provides an index for the psychological value that a person's processes are psychologically channelized by the anticipation of events. This process is considered as the initial process by which the system of the person internalizes the environment, the process of perception.

The hypothetic structure is conceptualized as an hypothesis, based on behavior in the past which influences present behavior by the anticipation of that past behavior in the present.

An example of the hypothetic structure is the relative relation between the activities in Tables 2 and 4. The descriptive value of happy and the not applicable value of sad defines that relation. Inferring to internal quantities, it appears that the construct of happy is used as a unipolar construct and not as a bipolar dimension of happy-sad as in the psychological literature. Inference from the hypothetic structure is consistent with, hence supports, the value of the internal quantities that events are construed as replication of past events.

An example of the real structure is the relations within and among the activities in Tables 2, 3, and 4. These relations produce the descriptive value of happy. This construct apparently is critical to this sample of the person's processes. Thus the real structure indicates the internal quantities used over time, those which are central to the person functioning.

The real and hypothetic structures are considered as indices to the structured knowledge of the person. This model illustrates two aspects of knowledge, that of

the construction of the environment which serves to construe the environment. By knowledge serving to construe the environment, it also influences the program of the system, perception. Knowledge, because it is acquired through perception, is influenced by perception. Thus by using the model explicated herein, in empirical investigations, an understanding of perception and structured knowledge which underlie a person's behavior may be reached.

References

- James, W. Principles of psychology. New York: Holt, Rinehart and Winston, Inc., 1890. 2 Vols.
- Kelly, G. The psychology of personal constructs. New York: Norton, 1955. Vols. I & II.
- Klir, G. An approach to general systems theory. New York: Van Nostrand Reinhold Company, 1969.

AN INTERDISCIPLINARY THEORY OF ADAPTING BEHAVIOR

Rosalind Schulman and Doreen R. Steg
Drexel University

ABSTRACT

Steg has defined adaptive behavior as behavior where a system adapts itself to the requirements of the environment, and adapting behavior, where the system changes the environment to suit itself using a variety of tools. A wide variety of possible mixtures of both also becomes evident.

Schulman has stated that the result of economic affluence is such "adapting behavior" in the demand for goods and services, and Steg has stated that dynamic sensory feedback provides an intricate means of regulating motion in relation to the environment; while both have stressed the necessity of programs that communicate -- Steg in learning theory and Schulman in economics. However, these phenomena of human behavior, today evincing themselves in economics and learning cannot be solely restricted to one or two fields. It is postulated that they apply to all facets of human existence (economic, educational, ethical, psychological, cultural, social and political) in all material-reward-oriented societies, when each society has attained a modest but adequate standard of living for the majority of its population.

Introduction

Since 1953 we have been observing an economic phenomenon for which there was no apparent explanation. That phenomenon was declining price elasticity of demand as measured by income elasticity of consumption. In other words, an ossification of purchasing pattern was spreading through every income class of America except the poor. Businesses were going bankrupt without apparent cause in the midst of unprecedented prosperity.

By 1963 when accurate data became available, ⁽¹⁾ it became evident that there had been, over the period, a change in elasticity (the ratio between a difference in expenditure for a particular good and an income change). Elasticity was declining due to increase in income and education (no other variable proved significant); of the two, education showed the most significant pattern.

It has always been postulated by economic theorists that mean elasticity of consumption declines for goods ⁽²⁾ as income rises. But, this time the data showed an even greater decline with an increase in education. This was contrary to all predictions and could not be explained on the basis of any previous postulate, (Keynesian, Friedman's permanent income hypothesis, or the adaptation of the permanent wealth theory; or Duesenberry's previous standard of living hypothesis). When Steg's papers (1966 et al.) were finally digested it became evident that they were an explicative formula for this economic behavioral phenomenon. The consumer was "learning" and exhibiting an "adapting behavior" cybernetic mechanism.

(1) U.S. Department of Labor, 1963-70, BLS report 237-8.

(2) It increases for services.

In a series of papers Steg has treated deviation-counteracting feedback in human behavior, i.e. negative feedback, and suggested that at least two distinctive human behaviors become operative in society with possible mixtures of both. These are adaptive behavior, or behavior where a system adapts itself to the requirements of the environment, and adapting behavior where the system changes the environment to suit itself, using a variety of tools: social, economic, psychological, physical, and even political. (Steg, 1966, 1970, 1971, 1973, 1974).

In the paper on "Communication and Feedback in the Technology of Consumption," Schulman (1973) has shown that in a society where discretionary purchases and consumption are available to the majority of the population, the consumer becomes a "least cost" buyer for necessities, and refuses to follow previously established patterns of authority in purchasing discretionary goods. In a free or semi-free market this plays havoc with fashion's dictates; causes individualistic consumer reaction in the marketplace - to the point of purchase refusal if desires are not met; and expresses itself in overt criticism, lobbying and myriad other group and non-group activities to influence the market and the manufacturers, economically, politically, and socially.

In other words, the consumer is exhibiting an "adapting behavior" pattern, which will increase in intensity as the society becomes more affluent.

Behavior: Adaptive versus Adapting

Adaptive

Automatic activity of man, animal or machine is an adaptive control system, by its very nature. It is safe to assume that, as with the laws of physics, the laws governing control systems apply equally to animal, man or machine. In the language of the system engineer, this is a closed-loop control system. The control system pattern consists of (1) an input signal that triggers some action, (2) a feedback signal of the result of this action to compare with the input signal, (3) a closing of the loop and a summation of the two signals and (4) effective action to counteract this summing signal. A persistent residuary signal can be made to affect memory which results in "learning". In a control system, work is triggered as a result of an actual error input.¹ The error is essential to the activity of any control system. These mechanical patterns apply equally to automatic machinery, animal behavior, and man's everyday automatic activity.

¹ The term "error input" is an engineering term commonly accepted to mean a disturbance.

Adapting

An important deviation from the automatic pattern occurs when the automaticity of a system is eliminated. Non-automatic activity will not necessarily be subject to the adaptive nature of the control system and trigger its energy to cancel the disturbance.

With the automaticity eliminated the response to a disturbance is chosen after the disturbance has been analyzed as to its source, the energy involved in the disturbance, the possible response and resulting consequences, including analysis and assessment of energy sources and energy balances. In other words, understanding is replacing automatic response.

To recapitulate, an adaptive control system is subject to the effect of the environment on its sensing elements and has no freedom to control the effect of the environment on its sensing elements. It can only adapt the system by using its own energy to satisfy the requirement from the environment conveyed through the sensors.

Opposed to this automaticity is the human ability of adapting an environment by means that extend human reach in a specific fashion, including in the process the use of tools, machines, and psychological, socio-political, economic, educational and other instruments. Specifically, the human mechanism directs the signal-triggered action with a view to the adaptation of the environment to eliminate the differential between the fed-back signal resulting from the modified environment and the original input signal. The mechanism involved in the latter system or disturbance is subject to the filter of intelligence, thus creating an art image of the environment to serve as a blueprint for the adapting process. The system involved in specifically human activity is operable only when an action is triggered to adapt the existing, "given", "objective" environment to an art or dream image.

Adapting Behavior of the Consumer in Economic Life

At a point of economic affluence in any society (where a large majority of the population is living at a level considered by its culture to be "modest but adequate", and has major discretionary purchasing power in terms of whatever costs are being considered within the parameters of the culture) consumers of goods and services begin to exhibit patterns of "adapting behavior". These patterns are quantifiable and measurable in terms of price and elasticity of demand (\sum_q^p) and income elasticity of consumption (\sum_c^y), both over time and at specific single periods of time, for specific characteristics of both goods and prices. (Schulman, 1973)

The data show that consumer demand becomes inelastic as income and education increase, even for discretionary purchases; and that the substitution effect thus becomes a more important part of the change in purchasing patterns than the income effect. The substitution effect causes the buyer to follow a "least-cost" pattern of purchase, whatever the important component of cost may be for the individual - money, time or convenience.

Declining elasticity causes a more rigid purchasing pattern on the one hand (I want what I want, when I want it), and on the other a more flexible willingness to switch from one good to another (when the two goods have almost identical characteristics) on the basis of price in terms of money, time or convenience. It also shows that the consumer has learned to say "No".

Consumers purchase bundles of "characteristics" (Lancaster, 1965, 1972) and not individual goods and services. Demand loses elasticity for many characteristics with rising income and education (no other variables proving significant). This loss of elasticity and consequent increase in the importance of the least-cost (substitution) effect extends to almost all characteristics and costs with the exception of those falling within the "responsive" and "adapting" behavior patterns where the human system changes the environment to suit itself.

As economic affluence increases, the effects of inelasticity of demand becomes evident in all material-reward-oriented cultures, no matter what their political, economic or social systems. When a dictatorial society decides not to make desired consumer goods available - the most profitable industry in such a society will be underground, outside-the-law or smuggling oriented.

Adapting Behavior: Thinking and Education

In an adapting control system, the response to an input signal is not necessarily proportional to the input. This is because of differential individual perception and valuation. Perception and understanding are shown by empirical data to be altered by education.

In 1960-61, for education of head of household of eight years or less, \sum_c^y for "gifts and contributions" reaches an inflection point* at income of about \$10,000; for head of household having graduated college, the inflection point is reached at income from \$5,000 to \$6,000. (The relationship of "gifts and contributions" to income is inelastic until the inflection point is reached and elastic thereafter.) Similarly, the inflection point for \sum_c^y for "medical care" is reached at \$5,000 in 1960-61 income for households having non-grade-school graduate heads; but at \$2,000 for households with college graduate heads.²

Although other consumption patterns are differentiated by education, the two cited above are important because the first ("gifts and contributions") is an example of a feeling of social responsibility existing at lower income levels for highly educated families - to be specific at roughly half the income level of low-education families; and the second ("medical care") is an example of appreciation of the necessity for preventive medical advice occurring again at a lower income level for the more highly educated family. When the income elasticity of consumption (\sum_c^y) is greater than 1, or elastic, it specifically means that as family income rises, an increased proportion of such income is spent on the good or service. Thus, a college-graduate-headed family, with an increase in income of \$5,500 to a mean of \$6,750, will increase its charitable contributions by more than double its percentage income rise, and the same is true of medical care expenditures.

* point of slope change

² Schulman, 1972, pp.403-8.

The input signal is identical for both groups of families - an increase in income; but the response by each group is not a function of the signal above, it is a function of the signal and the educational level of the family. Incidentally, for both goods and services studied, the response at each income level is not a function of race.

We have thus a model of thinking which contains quality as an essential element and operates pragmatically as a closed self-organizing loop. It accounts for teleological processes like problem-solving, "planning", and mechanistic behavior. It allows for an infinite variety of awareness-cognition-response feedback systems.

As defined by Dewey, art is "to select what is significant and to reject by the very same impulse what is irrelevant and thereby compressing and intensifying the significant."³ We should add to the statement that both the "significant" and the "irrelevant" are dynamic concepts that continually change position.. Because machines have only automatic, adaptive responses, and thus have built-in "significant aspects", "creativity" is impossible.

Education (formal and/or informal) is the phenomenon which initiates a control activity, triggered by the element of relation, association or construction that appears, for example, when an artist produces an image unlike the one achieved by a camera. It also appears in all scientific discovery, as a change from the accepted previous concept. In other words, education centers on the "art" created image and its involvement in control system activity.

Adapting behavior depends on education and not training alone. Training involves learning some specified pattern of behavior, be it prestidigitation or tightrope walking, while education is new concept formation. The result of education is creativity, while the result of training is performance involving skill.

If the adapting control process "filters" disturbances or input signals in the closed-loop servo-system which controls human action, education is then taking place.

The servo-mechanism of the human control system continuously develops and grows as thinking develops and grows. Inquiry and correlation of experience are tools used in this process of education; they are elements which trigger the controls. As for experience itself, we can no more know what a particular "experience" will do to education than what a "pencil" will write. Experience, of course, is a pre-requisite, just as one needs a pencil or something to write with.

Any realization of something being wrong is a discovery. It contradicts the previously assumed satisfactory order. Anything that has been logical up to this point becomes illogical, becomes wrong, becomes an error, and will make room for the elimination of error - for a new logic - for the "ought" instead of the "is". This realization that something is wrong (which initiates the process) is a prerequisite required for a new concept formation.

³ Dewey, 1934, p. 208.

There is a difference between man and animal or man and a machine which is made to simulate man's behavior. The computer essentially accomplished its function by operating on a multitude of types of problems with techniques for solving them. Thus, a problem fed into the computer in a sense triggers the answer that was originally built into it. But, to reiterate, human problem solving is a matter of education and growth. It creates or formulates problems and at times their solutions.

Adapting Behavior and Learning

Learning in education is the possibility of going outside of a frame of activity. The difference between man and animal or machine is specifically that a machine that has "automatic" activity has, of course, been programmed to so act. It can automatically perform activities which it was designed to perform. An animal or man can also be programmed, i.e., the responses are limited to the programming or designing, just as in behavioral terms, persons automatically respond as experience, reinforcement or "programming" has determined that they shall. The responses are the result of training. Brainwashed man is as programmed as a machine. The learning in this case is programmed, hence automatic. But it is questionable whether one can train all men. The possibility of training may be inversely related to the distance the individual has progressed from the animal state.

Similarly, what is happening in the marketplace is the development of a responsive environment for consumers with increasingly adapting behavior patterns. Those markets that are not responsive (in its cybernetic sense) environments are going bankrupt.

A system is an organized whole of parts. Hence: $\sum_{i=1}^N E_i$ is a system. However, is the system the same if E_2 comes first and E_1 comes second? The answer is no.

Just as in a responsive environment, learning occurs when the individual controls and influences his environment (Steg, 1973). Similarly, the environment, not only a learning environment but the total environment, must be made responsive to the individual's actions. It may be that we have now reached the stage, wherein, when an individual gets the time to learn in an environment responsive to his desires* he may now begin to think he can influence his environment, and attempt to influence society to give him what he wants. Patterns of consumption acquired in consumer behavior seem to indicate this. Note the meat boycotts, rise of indigenous buying groups, refusal to accept authority, changes in purchases of clothing, furniture, cars and the rise of the consumer veto.

Once choice is present, once options are available, the environment must become responsive to the individual's desires (note, needs are necessities and inelastically demanded, wants and desires are elastically demanded) for the environment to survive. In this case, however, the environment is the usable environment, in the sense of goods and services, and in political, moral and/or social behavior, or education. Furthermore, an individual's gain need not entail another's loss. We are faced here with non-zero sum games. Adapting behavior on the part of the individual implies maturity and non-degradation of the environment.

* Princeton experiment with negative income tax

To create a properly functioning social entity that is active in animal husbandry, agricultural pursuits and other activities under the adapting control system (adapting the environment instead of being adapted to it) a system of communication is required in which understanding, as an element of an adapting control system, plays a major role. The communication system suitable for an adapting process requires means of communicating elements leading to understanding. A characteristic specific to an adapting system is a type of communication, the nature of which goes beyond transmittal of information. It involves an element of possible reaction to the signal on the part of the receiver that permits understanding of this reaction by the originator of the communication. This particular phenomenon implies a closed loop communication.

In addition, it is an adapting communication system in its own right. Thus the early establishment of the U.S. Agricultural Experiment Station was sited within a radius of a single day's travel there and back for the average farmer in the vicinity. As modes of transportation became capable of encompassing longer distances, the Stations could move further apart; but the necessity for communication was the main determinant of the time-distance.

In the same fashion, industrial and technical research has always tended to cluster about centers permitting constant interaction and communication between individuals. "Science centers" have always existed since the temples of Mesopotamia and Egypt. Knowledge and invention cannot be pursued in a vacuum.

Adapting Behavior and Ethics

There is a further phenomenon that arises when the individual subjugates the environment rather than developing understanding and control. The meaning of the word "rule" as used herein is to describe a hypothetical relationship of cause and effect as applied to behavior in a control system. It implies a mathematical formula relationship as an end product. Thus, the law or rule of gravity translated into a mathematical equation by Newton used the word gravity in describing a phenomenon, the nature of which was, and still is, a mystery.

When we use the word "ethic", we are describing a causal relationship. Anthropomorphically, a rule was a causal relationship derived from the "Lord". But this is not our concept of ethics.

Communication in an adapting system of man-man stands for a relationship between man and man as described by the term "Ethics", and comprises a special condition that implies mutual understanding, awareness, consciousness and reasoning, in contrast with a relationship lacking these ingredients.

We have distinguished between an adaptive system, which is a self-organizing system, and an adapting system, where there is organized control, requiring consciousness. Consciousness is the acquired characteristic of an adapting control. Thus, training is not sufficient nor satisfactory for moral behavior.

As previously noted, communication in an adapting control system relates to a relationship between man and man and comprises a special condition that implies mutual understanding, awareness, consciousness and reasoning. This characteristic is acquired by each new generation from the previous one by means of the educational process. Thus, the educational process has a prerequisite, mutual understanding.

The social form of government under an oligarchy, ruling according to rules of slavery or domination of any sort, is the expression of undeveloped understanding by the few in their effort at using the adapting controls to tame the many. Slavery is the expression of the system of adapting controls characteristic of man, but paradoxically enough, so is freedom from slavery.

To put the above in technical control language would sound something like the following: Moral rules contain the desired quiescent state of a system. This implies that control action does not take place when the moral rules conditions are satisfied. The reaction of the system when the moral rules are not satisfied can thus be considered of two kinds:

1. To satisfy the requirement of a moral rule.
2. To eliminate the requirement for a moral rule.

When control is vested in an oligarchy, the moral rules to satisfy that oligarchy are to be found in category 1. The reluctance of human nature to follow such rules falls in category 2. Control of the adapting type is required in order to eliminate this conflict between the oligarchy and the reluctance of people to follow.

Free social forms are only possible with the overwhelming majority understanding the nature of the adapting control system of man.* This makes equal opportunity, voluntary cooperation and competition for all a satisfactory social environment. Communication in such a system establishes a relationship between man and man as described by the term ethics and consists of decisions concerning continuous choices. Education for ethical behavior is education for choices to be continuously made.

Control Systems versus Reinforcement Control

It has been observed experimentally that providing knowledge of results, rather than reducing or withholding knowledge, may lead to more effective learning. Immediate knowledge is more effective than delayed knowledge, but it will not automatically enhance efficiency of performance and learning. Yet, it is generally assumed that learning can be enhanced if it is followed by reinforcement.

Dynamic sensory feedback provides an intrinsic means of regulating motion in relation to the environment, while knowledge of results, given after a response, is a static after-effect which may give information about accuracy, but does not give dynamic regulating stimuli. Dynamic feedback indication of "error" would thus be expected to be more effective in performance and learning than static knowledge of results.

* Today there are no free social forms in our society. Just because the majority has the power does not mean that they are using it with understanding.

Furthermore, the efficacy of reinforcement assumes an active need or drive state, while feedback theory assumes that the organism is built as an action system and thus energizes itself. Hence, body needs and wants are satisfied by behavior that is structured primarily according to perceptual organizational mechanisms, and require programs that communicate. We can now judge why reinforcement of child turning his head to the right being reinforced by a sucrose solution sucked from a bottle takes hundreds of tries, and Bruner's baby with the \$20,000 pacifier takes only a few tries, about five seconds, before he learns to focus a picture of his mother, and he isn't even hungry.⁴ The bottle experiment is a stimulus-response model, while the pacifier experiment is a true cybernetic feedback model.

Pacifier Experiment

Bottle Experiment

Cybernetic Control Model

Stimulus-Response Model

- | | |
|--|---|
| <ul style="list-style-type: none"> . closed loop . internal control . voluntary control . intrinsic means of regulating motion . means and ends not bonded . systematic relation to the learned behavior . learning requires <u>no</u> reinforcement . behavior is the control of input . dynamic continuous feedback . few trials for learning . self-determined learning . primitive adapting system | <ul style="list-style-type: none"> . open loop . external control . stimulus control . extrinsic reinforcement schedule . means and ends are bonded . no systematic relation to the learned behavior . if learning occurs, it is transient requiring reinforcement . behavior is the control of output . static after-effect of knowledge . many trials for learning . doubtful feasibility of conditioning . system adaptive only if and when successfully engaged |
|--|---|

To summarize: Use of linear programs (including branching) in all teaching deliberately limits the media of communication, the experiences of the student and thus the depth of understanding that he achieves. Instead the student should be provided with a broad context of experience by resorting to all of the activities and to all of the communicative media at our disposal. This includes verbal and non-verbal material. Thus, the student learns by responding to the perceptual organization of his environment.

Beyond deviation-counteracting feedback or negative feedback, there is also operative a deviation-amplifying parameter, or positive feedback. (Steg, 1972).

The world of advertising media in our present "free society" has been geared to the development of reinforcement stimulus-response models and not cybernetic control, because the media has usually assumed adaptive behavior on the part of the consumer. By assuming that the consumer is an adaptive personality and therefore learns what is being taught without wanting to use the learning as a means of further expression, the media in advertising have assumed that constant repetition would cause the consumer to learn, without thinking either of the repetitive method or of the application.

⁴ Kalnins, I.V. and Bruner, J.S., "Infant Sucking Used to Change the Clarity of a Visual Display," unpublished, 1973.

Reinforcement control without obtaining consumer reaction (other than in the most general fashion of like or dislike, I or not I, percentage listening versus percentage not listening, percentage tuned in versus percentage not tuned in, and the entire world or Nielsen Ratings) is solely for an adaptive behavior society.

Cybernetic control on the other hand, assumes that the response of the media to the needs of the consumer dictates the type of approach to the consumer and that this approach is changeable as the consumer responses are obtained. The chart showing the differentiation between reinforcement and cybernetic control has been given above.

If consumer response to product differentiation can be looked at as a form of adaptive behavior (inelastically demanded product for which according to empirical evidence and theory the substitution effect of least cost is greater than the income effect) no manner of reinforcement control can influence the buying of that good which is cheapest in time, money or convenience. As some of our manufacturers have found to their sorrow (the latest bankruptcy in men's clothing being Botany Industries) no amount of reinforcement control could possibly influence an adapting society, which is exactly what happened.

For the first time in the early nineteen-sixties, President John F. Kennedy sent a message to Congress on Consumer Rights:

1. The Right to Safety - To be protected against the marketing of goods which are hazardous to health and life.
2. The Right to be Informed - To be protected against fraudulent, deceitful or grossly misleading information, advertising, labeling or other practices, and to be given the facts needed to make an informed choice.
3. The Right to Choose - To be assured, wherever possible, access to a variety of products and services at competitive prices; and, in those industries where government regulations are substituted, an assurance of satisfactory quality and service at fair prices.
4. The Right to be Heard - To be assured that consumer interests will receive full and sympathetic consideration in the formulation of government policy and fair expeditious treatment in its administrative tribunals.⁵

It should be noted that of these four Consumer Rights first mentioned in John Kennedy's message, three are solely for an adapting behavior society: the right to be informed, the right to choose, and the right to be heard, which are of course the basis of cybernetic control or feedback.

⁵ Magnuson, "Consumerism and the Emerging Goals of a New Society," in Ralph M. Gaedeke and Warren W. Etcheson, eds., *Consumerism: Viewpoints from Business, Government and the Public Interest*, (San Francisco, Canfield Press, 1972), pp. 3-4.

Behavior and Societal Development

While there have been numerous studies that have attempted to integrate the negation of material or cognitive reward by the substitution of conditioning, only a material-reward approach can be successfully projected. If deviation-amplification sets in, cognitive development occurs. What appears to follow is social and affective development. (Steg, Oxford, 1972)

At this time it must be maintained that the individual in a materially oriented society that has not attained the "modest but adequate" pattern of living of that society in an economic sense is adaptive. As the individual becomes less financially restricted (more and more able to obtain the "normal" accepted level of material life), behavior becomes adapting. In other words, he can use ever greater economic, political or social leeway to change his environment.

Thus, Florence prior to the Renaissance was the richest country (in terms of time and access) in Italy because it was the one area with sufficient food (which in the Middle Ages was the equivalent of a "modest but adequate" level of material living). Hence, they had discretionary purchase power and from this a positive time correlation with the start of the Renaissance, which was an integrated scientific-cultural-social development within the limitation of the technology of the Middle Ages.

The only two other non-slave based pre-technological societies which attained discretionary purchasing power in terms of their respective cultures were the Inca society in central America and the land of Israel at the time of Solomon. Although the former society is presently undocumented except from tales of the Spanish conquistadores and ruins, the latter has been aptly described in the Old Testament as the 40 years of fulfillment of the promise of the Lord to Solomon at Gibeon (Chronicles II, 12) "I will give thee riches and wealth, and honour, such as none of the kings have had that have been before thee, neither shall there any after thee have the like." And for 20 years Israel, in a land having perhaps less than 3 million inhabitants, built the Temple, with an equivalent man-year labor expenditure of twenty times 153,600 men.

Our present technology is a system of three types of production of which the second has only become technologically feasible in the most industries producing discretionary goods within the past 50 years.

1. Mass Production - infinite runs of identical goods.
2. Differentiated Production - combination of mass produced parts forming an infinite variety of permutations and combinations - (note that an individual hardly ever, if ever, sees two cars alike) possible combinations are myriad. Similarly, clothing parts are mass produced, but assembly is individual. Therefore, a particular production run can be as short as one wants to make it.
3. Hand Production - Here, there are, of course, infinite combinations possible.

The second and third types of production act to allow for discretionary income. The first type answers the living base. Naturally, combinations of the three types of production run the same gamut as combinations of goods and their characteristics.

As demand becomes inelastic, as items become necessities or individual reaction to cost changes become inflexible, the substitution effect takes over. The black box makes no distinction between equivalent characteristics except for cost to the individual, be it in dollars, time, convenience, longevity and so on.

As the United States approaches in the 1970's and 1980's the same relative level of discretionary income as Florence or Solomon's time, we see burgeoning the means to grow from this country an indigenous cultural-technological-sociological expression which may well be the beginning, within our own technological ability, of a new integration.

It has only become possible with the work of Kelvin Lancaster to distinguish elasticity of demand for characteristics of goods by components of price in such a way that integrated comparisons of different mores become scientifically feasible. Consequently, for the first time it may be possible to generalize from single country experience on the probability that comparable changes for characteristics (even if goods are non-comparable) occurs in materially oriented countries, no matter what their system of ownership, political power, or social structure. However, we cannot compare patterns in a material and non-materially* oriented culture.

In primitive societies, and in societies where non-material rewards operate, the mass of the people are adapted to the requirements of the few that are leading. Material rewards must be of this life at this time and cannot be a credit for the next life or transformation.

As mentioned above, there are material-reward societies where even though scientific development has occurred we still have a majority adapted to the requirements of a minority, but with consequences of inefficiency, bottlenecks, breakdown and resource waste.**

In the U.S., the condition of greater interdependence arose with the age of rail. Today, almost instantaneous communication creates interdependence. Therefore, with a more educated population one gets group dynamic adapting behavior. No matter how small a group (2%, 20% or 90% of the population) each can create as much hell as the other. Interdependence carries over into social action. An individual transacts with the closest individual with whom he is already involved in some capacity, thus leading rather easily into cooperative venture (i.e. students getting together to picket a laundry that ruined John's shirt).

This kind of social interdependence occurs over the entire range of income groups except where the individual is absolutely indigent. It extends from the very top to the very bottom. It changes and is amorphous, since one is dealing with a cybernetic flow a situation having social feedback, dynamic give and take occurring between people. As understanding replaces automatic response social transactions develop. The number of possible permutations is infinite!

* i.e. Zen Buddhism

** The largest industry in the Soviet Union is smuggling.

One can never predict people's actions with certainty as people have become accustomed to behaving adaptively, but the laws of probability permit approximation.

As noted above, art is something no one can teach. One cannot choose what is significant for another. One cannot make a person enhance or distort something in a way that one does not himself know how to distort or enhance. Yet, such oblique or surrealist views and disorderly processes are a necessity for adapting behavior to occur. Furthermore, the selection of the "significant" depends on choice being present. Otherwise no selection can occur. All scientific discovery depends on such processes, as a change from the accepted previous concept.

We now postulate that countries such as the USSR and China will not grow unless there is enhancement of adapting behavior. And the U.S. as we know it, will probably continue to change. The present state of the United States is that of an adapting society, rich today and tomorrow, given the technological possibility of unlimited non-polluting power and water within a few generations.

In the U.S. the governmental system has worked more on intra-party accommodation than on inter-party rivalry. If there is an infinitely possible variety of permutations of feedback effects, then there is going to be inter-party accommodation in the future. The effect of division into parties becomes much less important. What emerges is a form of consensus called the pragmatic position. Presently in many questionnaires, checks are asked for the following categories in voting: Democrat, Republican, Leftist, Rightist, Independent, and Pragmatic (votes on issues). The Republican Turks band together with the Democratic Turks on particular actions and do not gang up on one another, thus getting a shift into permanent tailspin, yielding one party on any particular issue. What is occurring is that the two party system is disintegrating into a kaleidoscope of issue-dominated, shifting time, accommodative groupings. 'A' may agree with 'B' on issue 1 and with 'C' on issue 2, but this does not prevent 'A' from disagreeing with both 'B' and 'C' on issue 3. As stated above, the system also changes according to order as well as quantity in a non-zero sum game.

Since there is no life without order or rule or ethic, what we are developing is a continuous-shift pragmatic grouping society.

Feedback, in the cybernetic sense (as opposed to that feedback which means knowledge of results) occurs not only in economics but in politics, in the family, in life. (Note the emergence of communication as a descriptive and explicative framework in psychiatry.) It is affecting the behavior of an entire population. We now have evidence and measurable data in economics that indicate the presence of this phenomenon in consumer behavior. The presence of this adapting behavior does not mean that it is limited to economics only. Why should it be isolated to this field? It is a phenomenon that is pervasive, be it in education, the psychology of learning, social development, ethical behavior, or political development.

Adapting behavior is much more than just adapting economic behavior.

In 1910, pre-World War I United States, there was a small middle class and a large lower class. By 1960, the Newport estates, the yachts and the 27 servants had just about disappeared. The shift occurred in that pre-World War I the majority was adapted to the requirements of the minority. It led to where the minority became adapted to the majority. The shift is still going on and consolidations continue to take place.

At present there are societies that are technologically and scientifically advanced and are still slave societies (Slavery df = control of the individual's ability to work, think, move, and establish familial relationships.) The question arises as to whether in such a society technology and science can be used for or by the slaves. If the slaves are the master of technology and begin to have discretionary income, time and purchasing power, then agitation will set in. The move towards adapting behavior will occur.

There can also be a technologically advanced society that may not provide a sufficiency of goods for a majority of people. This can also be a master-slave society. But once a group of people, be it a slave class or not, has discretionary purchasing power in terms of income or time they become adapting creatures as exhibited in consumption patterns, where slavery cannot co-exist.

In a technological society where the masters use the technology for the slaves and the slaves "receive" the "good life", there are several problems that become associated:

1. The masters will tire out sooner or later if they do all the work.
2. The masters become the working class. Eventually no one will want to be in it.
3. Such society is similar to a bee hive, or an ant heap where the masters become the workers and the queen and the slaves become the drones. Inbreeding would eventually kill off the masters, if the slaves would not be permitted to mix.

There is evidence that learning in the sense of new concept formation develops with conscious individual growth and assertion (Steg et al., 1973). If such learning is not enhanced, education is not taking place and only training is allowed, such society will show no discoveries. At best it will have innovations, but growth will not occur and the society will be static and deteriorate.

For about 7,000 years of recorded history, since Shub Ad of Ur, the vast majority of people have worried about one thing...food. The entire human energy output has been expended on trying to eat. Today this is still true for the vast majority of the world, for most of the third world, the greater portion of the Middle East, South America and Asia. This is also still the case in some parts of Europe. Only two centuries, 1/14th of the time span, has given us progress to the point where some people stopped starving in some countries. Basic progress has occurred since the time of mass production, during the second to the eighth decades of the 20th century.

In the decades since the 1950's the greatest strides ever have been made. We thus have major developments in less than 1/10th of the time of recorded history. In the U.S. the discretionary class became a majority after 1964 - less than ten years ago. (We now have evidence that the majority of the blacks have entered the middle class (Wattenberg and Scammon, 1973).⁶ The black leadership that led the agitation of the blacks were from the lower upper and upper middle class. It is thus understandable why this is not the case yet with the Chicanos and the Amerinds (American Indians) and not quite the case with the Puerto-Ricans. But that is beginning.

Conclusion

It is, therefore, postulated that all societies which have attained freedom from abject want will eventually approach the point at which affluence in time, goods and/or money will make it necessary for the societies to respond to adapting behavior.

⁶ Wattenberg, Ben J. and Scammon, Richard M., "Black Progress and Liberal Rhetoric," Commentary, April 1973, pp. 35-44.

Bibliography

- Bergstrom, T.C. and Goodman, R.P.: "Private Demands for Public Goods," AER Vol. 63, No. 3, June 1973.
- Bruner, Jerome S.: Studies in Cognitive Growth: Infancy (Volume III Heinz Werner Lecture Series). Worcester, Mass.: Clark University Press with Barre Publishers, 1968.
- Bruner, Jerome S.: Up From Helplessness, Readings in Educational Psychology. DeCecco, John P. (ed.) Del Mar, California: CMR Books, 1970, pp. 71-75.
- Bruner, J.S. and Kalnins, I.V.: "Infant Sucking Used to Change the Clarity of a Visual Display," unpublished, 1973.
- Dewey, John: Art as Experience. New York: Minton Balch, 1934.
- Lancaster, K.J.: "Change and Innovation in the Technology of Consumption," AER, LVI:2, May 1966.
- Lancaster, K.J.: "A New Approach to Consumer theory," J.P. Econ., LXXXIV, 1966.
- Lancaster, K.J.: Consumer Demand. New York: Columbia University Press, 1971.
- Magnuson, Senator W.G.: "Consumerism and the Emerging Goals of a New Society," in Gaedeke and Etcheson: Consumerism: Viewpoints from Business, Government, and the Public Interest. San Francisco: Canfield Press, 1972.
- Schulman, R.: The Economics of Consumption for a Changing Society. Philadelphia: Drexel University Press, 1972.
- Schulman, R.: "Bibliography for Consumer Economics," Philadelphia: Drexel University Press, 1969 (mimeo.).

- Schulman, R.: "Communication and Feedback in the Technology of Consumption," VIIth International Congress on Cybernetics, Namur, Belgium, September 1973. (to be published in 1974).
- Scriven, Michael: "Teaching Ourselves by Learning Machines," Journal of Philosophy, Volume LXVI, Number 21, November 5, 1970, Sixty-Seventh Annual Meeting of the American Philosophical Association, Eastern Division, pp 898-908.
- Steg, D.R. and Schulman, R.: "A General Theory of Adapting Behavior," VIIth International Congress on Cybernetics, Namur, Belgium, September 1973. (to be published in 1974).
- Steg, D.R.: "Programmed Teaching and Learning," Proceedings of Philosophy of Education Society, Chicago, Illinois. (Southern Illinois University, Edwardsville, Illinois, May 1963).
- Steg, D.R.: "Some Aspects of Teaching and Learning," Proceedings of Philosophy of Education Society, Chicago, Illinois. (Southern Illinois University, Edwardsville, Illinois, 1964).
- Steg, D.R.: "Systems, Rules and Ethics," Proceedings of Philosophy of Education Society, Chicago, Illinois. (Southern Illinois University, Edwardsville, Illinois, pp. 20-25, 1966).
- Steg, D.R.: "Effects of Individual Programmed Instruction on Initial Reading Skills and Language Behavior in Early Childhood," International Reading Association, 6 Tyre Avenue, Newark, Delaware, April 1969 (co-authored).
- Steg, D.R.: "Some Theoretical and Experimental Considerations of Responsive Environments, Learning and Social Development," J. Rose (ed.) Progress of Cybernetics, Volume 3 (Proceedings of the First International Congress of Cybernetics, London, England, 1969), New York: Gordon and Breach Science Publishers, 1970, with A. D'Annunzio.
- Steg, D.R.: "The Limitations of Learning Machines and Some Aspects of Learning," Focus on Learning, Volume 1, No. 1, Davis Hall, Indiana University, University of Pennsylvania, Indiana, Pennsylvania, Spring, 1971.
- Steg, D.R.: "A Philosophical and Cybernetic Model of Thinking, or a Feedback Analog to Thinking and Some Consequences," Society for Structural Learning, University of Pennsylvania, April 1973 (to be published in 1974).
- Steg, D.R.: "Deviation-Amplifying Processes and Individual Human Growth and Behavior," World Organization of General Systems and Cybernetics, Blackburn College of Technology and Design, Feilden Street, Blackburn, BB2, 1LH Lancashire, England, August 1972 (with A. D'Annunzio and C. Fox, to be published in 1973).
- U.S. Department of Labor, Bureau of Labor Statistics: Survey of Consumer Expenditures: Consumer Expenditures and Income, Urban U.S., 1960, 1961 and 1960-61. B.L.S. Report 237-38, Washington, D.C.: Government Printing Office, 1963-70 .
- Wattenberg, Ben J. and Scammon, Richard M.: "Black Progress and Liberal Rhetoric," Commentary, April 1973, pp 35-44.

OPEN-INFORMAL EDUCATION AND LEARNING CENTER CHOICES BASED BEHAVIOR MODIFICATION TECHNIQUES WITH YOUNG CHILDREN

Thomas D. Yawkey
University of Wisconsin

Dawn M. Jones
Montgomery County Schools, Md.

By definition, analytic behavior analysis is the process of applying sometimes tentative principles of behavior to the improvement of specific behaviors and simultaneously evaluating whether or not any changes noted are indeed attributable to the process of application, and if so to what parts of that process. Providing a frame of reference and systematic procedures for modifying specific behaviors, behavior analysis has been applied to a wide variety of socially relevant problems. More specifically, it has been employed in institutions, clinics, and schools to modify and maintain basic behavior patterns.

Central to the application of behavioral analysis has been the reinforcement contingencies or consequences used to facilitate the development of behavioral repertoires. Experimentation reinforcement contingencies have also been varied. For example, token reinforcements, as trinkets and food; or symbols as stars and A's, have been successful reinforcements. Social reinforcements, such as teacher attention, peer pressure, praise and rate of academic behaviors have also been systematically manipulated to modify specific behaviors. Of considerable importance to education is the potential use of systematic manipulation of teacher attention to analyze and modify the behaviors of normal children in public school classrooms. Teacher attention, however, as a means of identifying, controlling, and shaping young children's behavior in a classroom settings has been examined. Hasazi and Hasazi (1972) were successful in modifying digit reversal behavior of a primary grade child through teacher attention. In a regular self-contained primary grade classroom, Yawkey (1971) described and operationally used rules, ignore, praise to significantly increase independent work behaviors in reading with seven year old children. The results of these studies suggest the efficacy of teacher reinforcement as facilitating the development of academic behavioral repertoires in young children. Given the results of these studies, it would seem advisable to analyze other classroom settings with young children that involve teacher attention.

One type of classroom setting in which the reinforcing event of teacher attention is most likely operative is the open education classroom, or open space settings, suggested by recent developments in some English Infant Schools. These open-informal settings have been characterized by:

1. A lack of interior partitions in which the visible and acoustical separation between learning centers are limited or eliminated.
2. The assumption that the understandings are abstracted from a child's environmental encounters through learning experiences based upon manipulation.
3. An orientation toward the use of many avenues or ways of mastering a skill, i.e. individual learning styles based upon a student's own learning rates.
4. Utilization of learning centers or stations which are areas of the school devoted to academic and nonacademic skills based upon the children's specific manipulation of the stimuli at the centers.

5. Locations where teachers and children both contribute to the decision making process, i.e. teacher direction to and control of learning is noticeably absent.
6. The concern for interest and expression as sources of learning have legitimate places in the classroom.
7. Evaluation of performances in terms of the total intellectual, social, emotional and psychomotor growth of children.

The present study was concerned with the effects of teacher contingencies on academic behaviors of normal children at cognitively oriented centers in an open education classroom. The main question under investigation was, "What effects, if any, does teacher attention have on the number of academic learning centers chosen and completed by kindergarten children in an open education setting?"

Method

Twenty-six kindergarten children, ages 5 1/2 through 6, were used as subjects in this experimentation. These children as a group had a mean I.Q. of 115.9 based upon the Peabody Picture Vocabulary Test. The experimenter randomly divided the children into two independent groups using a random number table. One group was labeled BAB and the other ABA with BAB and ABA referring to the types of experimental treatments. The BAB and ABA groups contained B, A, and B and A, B, and A phases respectively. Across groups, BAB and ABA experimental designs have been referred to as "replication treatments". Across groups, subjects in the B phase of BAB and ABA groups were independently given teacher attention contingent upon completion of academic or subject oriented learning centers. Across groups, subjects in the A phase of BAB and ABA groups were not given teacher reinforcement. That is, the teacher ignored all the children's behaviors and performances originating at both the academic and non-academic centers in phase A of the BAB and ABA groups. At the end of ten consecutive days or two school weeks one treatment phase was immediately terminated and the other began across both BAB and ABA groups. The entire experiment lasted thirty days. Two examples of teacher reinforcement elicited upon completion of academic centers in phase B were:

1. "Sally has done an excellent job at the ___ academic center."
2. "Sam, I'd like to put the paper you did at the ___ academic center up for the other children to see."

In phase A, the teacher ignored all behavior at the academic and nonacademic center unless physical attacks were observed. When questioned by an ABA and BAB child in phase A, the teacher did not respond or replied quickly and briefly.

Results

Analysis of I.Q. scores indicated no significant differences in intelligence between the BAB and ABA subjects. There were a maximum of 78 possible nonacademic and 78 academic centers. The total number of academic centers successfully begun per child per phase in BAB and ABA were employed as response measures and summed for all children within phases. In the BAB group t-tests were run between phases B and A, A and B, and B and B, in ABA between A and B, B and A, and A and A. For BAB, statistically significant differences were observed between phases B and A, and A and B. For ABA, statistically significant differences appeared between A and B, and B and A. No statistically significant differences between phase B and B (BAB) and A and A (ABA) were noted. The results of the research suggested that young

children chose academic centers to a significantly greater degree than they chose nonacademic centers contingent upon teacher attention for successful completion of academic center choices.

Discussion

Central to this kindergarten and other open education classrooms are academic centers where the children are essentially free to manipulate and discover their environments, i.e. build curriculum from materials provided. With this approach to instruction, teachers generally make use of many concrete resources and numerous materials in the classroom. The children interacting with other children and learning from one another are regarded as further aides to learning in open environment classrooms. In a sense, both children and teachers become "primary senders", "receivers", and "discoverers" of knowledge. This research seemed to suggest that young children chose academic centers to a significantly greater degree than they chose nonacademic centers, contingent upon teacher attention for successful completion of academic center responses. It must be noted that nonacademic centers were viable classroom options available to both ABA and BAB children. Criticism often voiced at open education as "mere play" may be unjust when work behavior of children reinforced by operant techniques is observed to increase over time.

Poor or not, evaluation of young children's outcomes in learning are further criticism of open education. Academic centers are usually built upon certain skills or interests specified by the teacher and/or children. If all possible skills at or interests in academic centers can be described and thought of as outcomes, evaluation in open schools and centers could be observed and explained. A corollary to this might be that all children become interested in some skills taken from a set of all possible skills at the particular center. The important developmental principles of comparisons of an individual with the same individual over time could be possible in open education systems. At the applied level, teachers would not feel that certain numbers of academically-oriented centers had to be required per day if academic choices were reinforced through behavioral analysis techniques and recorded.

Then too, the free atmosphere, social interchange, and other basic characteristics of open education can be facilitators of one another to enhance learning of young children. A major responsibility of teachers, regardless of philosophical orientation is to play for positive and continued changes in students' behavior. This research illustrated that productive classroom work can be increased with contingent teacher attention.

Behavioral analysis can be an effective teaching technique with normal children using academic and nonacademic centers in open education classrooms. The children received praise for work at academic centers and responded by seeking academic centers and similar work forms of behavior producing reinforcement from the teacher. The children not receiving reinforcement eventually responded by not responding, i.e. ignoring the teacher.

References

- Hasazi, J.E. and Hasazi, S.E. "Effects of teacher attention on digit reversal behavior in an elementary school child." Journal of Applied Behavioral Analysis, 1972, 5, 157-162.
- Yawkey, T.D. Alternative response measures for mathematical behaviors on numerical tasks in preschool subjects. Journal of Experimental Education, 1971, 40, 88-96.

THE ORGANIZATION AND ACQUISITION OF FORMALLY DEFINED CONCEPTS

Jed Lewis

University of Minnesota

One goal of instruction in high school and college classrooms and educational textbooks is the teaching of subject-matter concepts and their associated meanings. A large proportion of these concepts are taught verbally from formal definitions. Each formal definition embodies an intensive set of properties which characterize the corresponding concept and which convey a set of meanings. Such concepts are usually taught in units consisting of sequences of interrelated concepts, with the meaning of each concept in the sequence being dependent upon an understanding of both natural language and subject-matter terminology.

In order to investigate the acquisition of formally defined concepts, four questions must be addressed and answered:

1. What is learned when formally defined concepts are acquired? Included within a proposed response to this question must be the detailing of: a) how such a concept can be interpreted; b) how meaning can be interpreted; c) how concepts and meanings are related, if at all; and, d) the relation between the logical and psychological natures of the formally defined concept.
2. What independent variables affect what is learned, and in what ways?
3. What psychological processes are involved in this kind of learning, in what order, and relating to which independent variables, i.e. how does learning occur?
4. How is what is learned organized in long term memory?

Existing accounts of the psychology of concept acquisition (e.g., Bruner, Goodnow, & Austin (1956), Suppes (1970), those reported in Bourne (1966) have almost exclusively been restricted to analyses of the inductive learning of non-verbal task materials. These materials have consisted of attributes which are primarily perceptual and which are very simply related. They are not novel attributes, but are taken from the realm of a subject's everyday experience. Moreover, only samples of such concepts have been used in these kinds of experiments, rather than a network of related concepts. Considered in this way, these analyses are suitable primarily as accounts of conceptual development and concept formation outside of the classroom.

Such accounts of concept acquisition processes will not do, therefore, for the class of formally defined concepts. These concepts, nonperceptual in nature, novel, and complexly interrelated, are acquired, for the most part, through meaningful reception learning (Ausubel & Robinson, 1969). In this form of learning, all that must be learned is presented in its final form to the learner in such a way that he can sensibly relate it to what has already been learned. Furthermore, the acquisition of these concepts within this educational setting is primarily deductive, rather than inductive (Carroll, 1964). While a subclass of the concepts may be primitive, in the sense that they cannot be defined in terms of other concepts which have already been introduced and may only be learned inductively through the repeated presentation of instances and non-instances of the concept, the remainder of the class are precisely defined, each usually being presented via its definition and several examples.

In order to adequately respond to the four questions posed above, therefore, an investigation must be made regarding what kind of psychological orientation is needed, and what methodological approaches are most appropriate for confirming and extending such an orientation. This is the aim of the paper.

The research questions have been explored by many investigators, each of whom has viewed them from a certain perspective. On the basis of similar theoretical assumptions, interpretations of empirical phenomena, and methodologies, studies of concept acquisition can be grouped into several different classes of approaches.

Behavioral approaches to concept acquisition share the common view of considering psychological phenomena as behaviors or behavioral events which are shaped and controlled by the environment. These events are observed as overt performances and are represented, or described, using a language system of stimulus-response associations (Kendler, 1964; Tulving, 1968). Learning is considered to be a "process by which the behavior of an organism is modified by experience" (Jung, 1968), or a "change in human disposition or capability" (Gagne, 1970).

Answers to the question of what is learned in the acquisition of a formally defined concept would include one or more of the following: a set of criterial attributes and rules for relating them (Bourne, 1966; Bruner *et al.*, 1956), a set of positive concept instances, an associated concept name (Archer, 1966), a linguistic experiences (Carroll, 1964), a behavioral capability (Gagne, 1966, 1970; Scandura, 1973), or a set of associations between dissimilar stimuli and a common response (Kendler, 1964).

Within this class of approaches, words are considered to be convenient labels of concepts. In formal definitions, therefore, content words function as representational signs which, through contiguous pairing, can come to stand for particular concepts.

The meanings of words can be understood in several ways. In the first sense, the meaning of a word is its referent (Jung, 1968) which could either be a structural description (Brown, 1958) or a societally-standardized concept (Carroll, 1964). In a second sense, the meaning of a word is expressed by the associated word which is elicited by it (Bousfield, 1961) or the total distribution of associated responses (Deese, 1968). Alternatively, it can be considered as a logical relation between a stimulus term and a response term (Noble, 1952) or as a part of the response produced by the concept signified by the word (Mowrer, 1954).

Given that one or more of the above are learned, the question that must be addressed next is how a learner proceeds from a formal definition to a criterial state of learning. In one behavioral approach, a sentence is interpreted as a string of words which is considered to be a subject-predicate unit (Mowrer, 1954; Osgood, 1956). The chief function of the sentence is to serve as a conditioning device to produce new learning in the form of new meaning reaction associations between the subject and the predicate. The linking of meaning reactions defines the act of sentence comprehension. As a consequence, constituent classes within sentences, i.e. subject and predicate units, would be stored in memory in terms of their associated meaning reactions.

The comprehension of a formal definition is only part of the learning process, for the component elements of the formally defined concept which are represented by words would first have to be learned. The processes involved in the acquisition of these elements have been considered in several ways: associationistic accounts (Bourne, 1966; Gagne, 1966, 1970); transfer accounts (Underwood, 1966); mathematical learning theory description models (Bourne & Restle, 1959; Suppes, 1970); structural task analyses (Suppes, 1967); and hypothesis testing and strategy models (Bruner *et al.*, 1956).

The independent variables which affect the acquisition process have been identified and studied using the reception and selection procedures (Bourne, 1966). These variables can be divided into the categories of stimulus variables (properties of concepts, hence affecting what is learned), task variables (environmental variables affecting the degree and progression of acquisition), and learner variables. Stimulus variables include difficulty of the concept rule that must be learned, the number of irrelevant or relevant stimulus dimensions, degree of concept abstractness, and word dominance. Task variables include the extensiveness and specificity of instructions, the sequence of stimulus presentation, and feedback conditions. Learner variables include age, IQ, learner style, and previous concept knowledge.

During acquisition, the concepts that are learned are organized in memory in various ways, with reorganization also occurring during the learning process

(Mandler, 1972; Tulving, 1968). According to Gagne (1962), all knowledge is organized as hierarchies of subordinate knowledges. These hierarchies describe the structure of sets of interrelated concepts and act as guides for the expected sequences of learning and memorial organization. When the attributes and rules of concepts are associated with their respective verbal labels, these labels can be viewed as verbal units which are encoded in memory (Voss, 1972). Hence, formally defined concepts would be organized with respect to feature dimensions of their associated verbal labels.

As opposed to the behavioral orientation, cognitive orientations are interested in accounting for the establishment of internal representations of external events. These events may take a behavioral form as overt actions, but they also encompass sensory and verbal data.

Cognitive approaches (Collins & Quillian, 1972; Kintsch, 1972; Rumelhart, Lindsay, & Norman, 1972; Simmons, 1973) have the same concern with characterizing concepts in terms of the relations that hold among other concepts. What is learned in concept and concept meaning acquisition is a network of interrelationships among concepts and words. Meanings are interpreted denotatively, and meanings and concepts are defined coextensively. When formally defined concepts are acquired, semantic relations among other concepts would be learned, where these latter concepts can be depicted as criterial properties of the concept being defined, and the semantic relations as rules, in order to suggest that the cognitive interpretations of a concept are, in essence, also structural descriptions.

In keeping with the overall orientation of cognitive approaches, concept acquisition involves the interiorization of originally external events into internal representations, or cognitive structures (Michon, 1972). The process of acquiring formally defined concepts includes the processes of sentence comprehension (Chomsky, 1965; Kintsch, 1972; Quillian, 1968; Simmons, 1973), encoding (Anderson & Bower, 1973), and cognitive structure construction and/or modification (Anderson & Bower, 1973; Ausubel & Robinson, 1969; Collins & Quillian, 1972; Wickelgren, 1972).

Stimulus variables directly affect the processes of comprehension and cognitive structure construction and/or modification. They include syntactic sentence structure complexity, lexical density of a sentence (Perfetti, 1969), meaningfulness of the words in the formal definition (Ausubel, 1966), and the internal and external connectedness of a structure (Greeno, 1972). Task variables include the sequence of concept presentation, instructional materials, and practice variables. Learner variables are the most important class of variables affecting learning. The most significant learner variable is the individual's existing cognitive structure. Other learner variables include developmental readiness and intellectual ability.

Approaches to the organization of concepts in memory would follow from the ways in which concepts, meanings, comprehension, encoding, and learning are interpreted. For concepts and meanings are defined with respect to a hypothesized memorial organization, and learning involves the progressive organization and reorganization of concepts and meanings in memory. Furthermore, if comprehension implies the simultaneous encoding of sentential semantic information, then that encoding implies that the verbal content of cognitive structures consists of semantically represented lexical items and semantic relationships between those items. Cognitive structures, then, can be variously interpreted as competence (Chomsky, 1965; Flavell & Wohlwill, 1969), semantic structure (Simmons, 1973), lexical memory (Kintsch, 1972), and semantic memory (Quillian, 1968). Logical structure is a systematic reorganization of psychological structure and is the final form the structure can take (Ausubel, 1966).

Methodologies that have been used to explore the nature of cognitive organization and to determine how it changes over time or during learning include the scaling of relatedness coefficient matrices which are computed from word association test data and the calculation of a distance metric between these matrices (Shavelson, 1972), the scaling of sorting data (Anglin, 1970; Miller, 1969), and the gathering of response latency measures from verification tasks (Collins & Quillian, 1972; Rips, Shoben, & Smith, 1973).

In the last major class of approaches, the interactionist orientations, researchers are interested in both behavioral and cognitive psychological events and how they interact to ultimately effect and affect external, overt performance.

According to Piaget (1966), a concept is a plan of action or of operation. It is not an internal reflection of anything in the external world. Behaviorally, this notion could be interpreted as an experience, but it is an active, mental construction instead of an associated behavioral capability. Cognitively, because actions are coordinated and organized in units which are related, concepts could also be interpreted as relations among other concepts. As psychological constructs, concepts are represented by symbols which are idiosyncratic (Piaget, 1967). When words represent concepts, they function as linguistic signs that are arbitrary and consensual. Both signs and symbols are signifiers which are independent of the concepts. Hence, a full description of a word concept would include both its cognitive and verbal form (Furth, 1969).

To Piaget, the internal representation of a concept is a function of the structure of actions in which the concept is assimilated. That representation can also change over time. As a result, a logical description of a concept is taken to be the formalization of the final representation.

Newell and Simon (1972) do not specify any particular conception of a concept, but they do discuss its psychological internal representation. In their view, external events are encoded in symbol structures which function as linguistic deep structures. This implies that if symbols are isomorphic with words, a sentence could be represented as a symbol structure with its deep structure serving as the input to a semantic interpretive system.

According to Piaget (1971), the meaning of a concept is defined by the cognitive structure into which it is assimilated. For Newell and Simon (1972), a concept's meaning is defined by the symbol structure representing the concept.

The conclusion that can be extracted from these conceptions is that what is learned during concept acquisition is either a plan of action or operation, or a reflected, internal symbol structure. Words are not concepts and only designate them, while conveying their denotative meaning in the form of an implied set of interrelationships with other concepts which have already been acquired.

For Piaget (1966), all thinking consists of the linking of meanings. Since a sentence consists of the linking of signifiers, this suggests that the surface linking leads to the extraction of the meaning linkings which are then assimilated into the appropriate cognitive structure. The process which enables a learner to acquire word meanings is an organismic capability called the symbolic function.

In information processing approaches, the comprehender is an information processing system and sentence comprehension can occur either through the syntactic and semantic analyses of a sentence (Winograd, 1972) or through the production of semantic structures (Hunt, 1973). Either of these two conceptions of comprehension would be incorporated within an overall information processing account of formally defined concept learning which would include encoding and representation modification processes. Although Newell and Simon (1972) do not provide such a complete systematization, they do provide enough clues in their theory in order to construct a reasonable extrapolation.

Learning, in their theory, involves the heuristic use, conversion, and extension of previous experiences. The performances or behaviors which must be learned are defined by a goal which is located in an external, objective task environment and which is interpreted by the learner as a problem. The environment provides a description of behavioral constraints and structures the context in which the goal is located. The task instructions (including the formally defined concept) which define a given problem (the learning of that concept) also provide a specific external representation of that task (concept) which is abstracted from the task environment and internalized by the learner as a symbol structure located within what is called the learner's problem space. An initial problem space representation would include representations of the task environment and the goal concept itself.

The collection of symbol structures stored in the problem space is called a knowledge state. The state is modified by information processing operators which manipulate the structures in order to create new state representations. After the creation of each new knowledge state, the currently created state (representation) is compared with the goal state (representation) and the results are evaluated in order to decide which operator to apply next. This is a feedback system in serial operation. The new operators are selected through the application of heuristics which are constrained by the nature of the internal representations. The functioning of this system is described using a state and process language. States and processes are really two different perspectives of the same series of events (Newell, 1972).

Three independent variables that stand out in this class of orientations are the goal problem, the task environment, and the learner's current cognitive structure.

Within the interactionist approaches, once concepts are acquired, they are organized within some form of relational structure. For Piaget, as actions, concepts would be organized according to coordinated schemes; as operations, according to qualitative groupings. Simon (1972) defines concept organization in terms of directed associations within the symbol structures.

Methodological tools that have been used within this class include individual observational protocols, the clinical method, eye-movement data, and simulation.

From this review of the psychological literature, it can be seen that many alternative answers have been proposed to the four questions that were raised at the beginning of the paper. The obvious question that must be addressed next is which approach, or what components of each of several possible approaches, should be accepted.

Beginning with the "what is learned" problem, the most serious difficulty with the various descriptions of formally defined concepts is that attention is restricted, within most of these conceptions, to only one type of psychological data. The facts that psychological data must ultimately depend upon performances, and that behaviors are performances, imply that behaviors are valid psychological data and that any psychological problem can be (at least partially) describable within this framework.

On the other hand, a basic assumption made in scientific research is that an investigator can make independent observations of events which are external to him and which can be recorded without prejudice. In viewing events this way, the observer divorces himself from the stream of events and dichotomizes the world into observer and observed. Yet, this is a misleading assumption, for an observer cannot truly be distinguished from that which is observed. Because the observer is a sentient being in, and not of, the world, he creates and defines it just as it creates and defines him. His perceptions and conceptions are constructed and serve to define his perceptions and conceptions. Therefore, rather than speaking of this dichotomous situation, we must speak about a process (or pattern) in which every organism, including an observer, mutually interacts with its environment as an organism/environment field. As a result of this conclusion, and the requirement that descriptions correctly classify and order all the experimental observations which occur within a performance domain of interest, the notion of a concept cannot be reduced solely to environmental stimuli and/or behaviors or behavioral capabilities.

The notion of a concept which is adopted has direct implications as to which interpretations of concept meaning should be accepted. If an interpretation is given which is to be consistent with the notion of a concept stated above, then that interpretation should include both behavioral and cognitive components. Concept meaning, identified as performance capabilities and the set of defining interrelationships between concepts, on one level, and, between the signifier verbal labels, on another level, would be such a consistent interpretation.

Besides the requirement that accounts of concept acquisition attend to the patterns of the organism/environment unity, another reason for needing to consider both behavioral and cognitive perspectives is that if a necessary (e.g., causal,

probabilistic, functional, or historical) and/or sufficient (e.g., explanatory adequate (Chomsky, 1965), constructive (e.g., Simon, 1969) explanation of the learning process is desired, then a behavioral approach alone will not suffice. This point has been made many times in connection with associational approaches (e.g. Deese, 1968; Tulving, 1968): associations are only descriptions of the fact that when a certain stimulus is presented to an organism, a certain response results. Explaining concept acquisition by means of associations amounts, in the end, to stating a tautology.

More specifically, behavioral accounts violate the following criteria for accepting proposed descriptions and explanations:

(1) failure to provide a thorough analysis of the mechanisms which underlie the processes of abstraction, induction, association chaining, hypothesis selection, and strategy acquisition.

(2) incompleteness of the proposed descriptions. Several mechanisms are omitted from the different behavioral accounts. For example, associationistic accounts do not provide an explicit description or analysis of the mechanism whereby verbal responses are produced when physical stimuli are presented to the learner, except for the unexplained notion of coded response (Mandler, 1968). Also, quantitative mathematical models only characterize the learning process to the extent of calculating the probability that the concept will be learned, given a specific trial, and ignore the roles of the independent variables and memory. Similar comments can be made about the structural description and hypothesis testing and strategy models.

(3) limitations and implications of the proposed descriptions which are exclusive of those stated in the introduction.

(i) Do the experimental findings extend to situations where other independent variables and task materials are manipulated?

(ii) The learning of abstract verbal concepts cannot be explained if meanings and concepts are ultimately reduced to physical experiences. Also, if it assumed that an organism only has a finite set of specific behaviors or behavioral capabilities, this cannot account for the spontaneous performance of other behaviors unless the unexplained mechanism of generalization is posited.

(iii) Accounts of comprehension imply that there is no easy way to estimate how the concepts are organized in memory, if they proliferate in one-one correspondence with organismic reactions, unless a gross overlapping of reactions and associated signs is assumed. Also, what is the meaning reaction of an abstract noun? Does context affect comprehension? How does a comprehender know how to decompose a novel sentence into its component units? How can paraphrase be handled? Are there different meaning reactions for sentences in which the surface structures are different, but where the deep structures are the same? If so, the number of associated meanings proliferates interminably (no meaning equivalence mechanisms are specified). In addition, an implication of the fact that comprehension is identified with a set of behaviors is that it would be identical with performance. Yet, as Bem (1970) and Sigel (1969) indicate, such a correspondence does not always hold.

(iv) Gagne's (1970) conception of organization implies a one-one correspondence between the logical organization of learning set capabilities, an organization which could depend upon the arbitrary decisions of a curriculum designer, and the nonarbitrarily organized behavioral capabilities of a learner. Theoretically, however, psychological organization could vary as a function of the task and the learner. Also, such hierarchies, organized according to the chaining of appropriate concepts, have difficulty accounting for the learning of paraphrased definitions.

(v) An environmental context of learning implies that relations between units rather than isolated units must affect the learning process and/or be part of what is learned.

(vi) Except for Scandura (1973), the notions of a concept are either logical or behavioral. The bridge between these two interpretations is not gapped, as it must be, if the relations between them are to be understood.

(4) the vagueness and lack of testability of terms such as associations, and mediated stimuli and responses.

Many of the objections raised against the behavioral approaches can be met by the cognitive ones. The hypothesized cognitive structures are explanatorily adequate (Chomsky, 1965). The hypothesized processes of comprehension and encoding can successfully account for part of the learning process. Meaning and referent are separate, abstract words and thoughts can be comprehended, paraphrase and segmentation are accounted for, and the mechanisms in these processes are better defined and more explicit than in the behavioral approaches. Either the syntactic or semantic component can be taken as primary, but the works of Sachs (1967), Perfetti (1969), and Bransford and Franks (1971), in free recall, suggest that the semantic component is more important. Cognitive approaches can also account for experimental results such as free recall data, memory confusion data, recognition failure data, and generalization data, in ways that are more acceptable than the behavioral approaches (Hunt, 1973).

Cognitive approaches still leave some problems open, however. The acquisition mechanisms postulated by Collins and Quillian (1972), Kintsch (1972), and Ausubel and Robinson (1969) have not been completely specified and analyzed, although simulations written by Anderson and Bower (1973), Quillian (1968), and Simmons (1973) make up for this deficiency. The cognitive approaches are incomplete, in that they ignore the behavioral aspects of concept acquisition, so that the relationships between cognitive structure and behavior are not examined. Many of the independent variables which have been seen to influence performance, thus probably influencing the acquisition of the contents of cognitive structures, are neglected. Except for Ausubel and Robinson (1969) and Shavelson (1972), these accounts do not consider the fact that the cognitive organization of verbal concepts appears to change during learning, and they also assume that the logical and psychological structures of a formally defined concept are in one-one correspondence, thus raising the same objections as in Gagne's approach plus the question of how the same concept can be represented differently in different semantic memories (as would be shown by sorting measures and similarity judgments). The comprehension explanations cannot account for the understanding of grammatically inadequate sentences and provisions are not included for the use of information outside of the dictionary (lexicon) and the semantic component. Also, it is unclear how a comprehender enters the stream of definitions, or where he stops, in Quillian's (1968) system, a system which neither accounts for the fact that comprehenders have a faster search time when more hints are given nor the problems of syntactic ambiguity, and paraphrase. Lastly, several psychological terms, such as internal and external connectedness, progressive differentiation, and potentially meaningful learning, are vague and lack empirical testability, the empirical verification of the structures posited by Rumelhart *et al.* (1972), Kintsch (1972), Chomsky (1965), and Simmons (1973) is unclear, and the experimental verifications provided by latency data, hierarchical clustering procedures, and computer simulations can be questioned unless it is specified that the assumptions involved in the use of these techniques have been carefully followed.

In this class of approaches, an underlying competence (cognitive structure) is postulated which is implied to be independent of performance limitations. Yet, how can an experimenter get at this structure and the independent variables which affect it, without at the same time confounding the experimental investigation with performance factors (Fillenbaum, 1970; Greene, 1972)? A possible escape from this dilemma is to assume that there is either an isomorphism between psychological

competence and performance and that idealized competence rules are only useful fictions from which research can originate, or that competence is a complete fiction which is irrelevant and can be dispensed with (since psychologists are really only interested in performance anyway), so that competence and performance are defined along the same psychological continuum. In the former case, the validity of psychological experimentation is preserved; the only alternative would be to say that psychology cannot be experimental, but only empirical (Greene, 1972). In that event, only intuitions would be acceptable psychological data, a result clearly contrary to any scientific paradigm. In the latter case, competence would be valid as a structural description of performance which is represented by a set of clear cases and specified independent variables. In this case, performance data would provide feedback for a competence model and for future predictions.

Turning to the interactionist orientations, they along, out of the three classes, attend to both the organismic and environmental aspects of acquisition, as well as the relationships between them. Concepts are interpreted as having both attribute and relational components. Functional and historical explanations are offered, as are sufficient explanations (simulations). The information processing explanations are explicit and testable, and Winograd's theory accounts for the mechanisms used in comprehension, as well as the paraphrase and segmentation problems. The questions of reorganization during learning and individual differences in internal representations are also successfully handled.

Yet, no theory is perfect. Problems within this class include the observation that Piaget's explanations are vague and do not explicitly mention the processes which connect mental activity and performance, the lack of empirical replicability of some of Piaget's results, the lack of incorporation of independent variables, the assumption of a one-one correspondence between logical and psychological structure in Winograd's theory, and the excessive reliance on protocol and simulation data.

In summary, then, no one approach or orientation best satisfies all of the objections raised above. Therefore, any approach which attempts to answer the questions posed at the beginning of the paper must involve some sort of synthesis of approaches from the three classes. Furthermore, the expositions and criticisms imply that such a synthetic approach should meet the following requirements:

- (a) the theory is an interactionist one;
- (b) both internal representations and performance behaviors, and their interactions, are accounted for;
- (c) the notion of a concept includes relations;
- (d) learning is viewed from both cognitive and behavioral perspectives;
- (e) all significant independent variables uncovered by research within the three classes are incorporated, with the relationships to learning being explicitly stated. In particular, some form of feedback mechanism is included;
- (f) contextual constraints are seen to affect the organizational structure of memory, i.e. what is learned is seen to affect performance;
- (g) organization (therefore, semantic structure) is seen to change during acquisition, and semantic relations should have primacy over semantic ones;
- (h) the approach is experimentally testable and verifiable.

A theory which incorporates these components shall now be proposed. First, this account assumes axiomatically that any partition of the organism/environment unity into a dichotomy is an artificial discrimination, but recognizes that some form of partitioning is necessary for the purpose of psychological analysis. However, because emphasis upon a particular aspect of this unity is ultimately subjective and up to the individual investigator's digression, this account assumes a neutral position, asserting that both components are equally important.

Therefore, psychological events of interest will be interpreted as whole events which can be viewed from both cognitive and behavioral perspectives.

As cognitive events, concepts can be described as abstract, gestalt representations of intuitions, percepts, or physical actions which are, themselves, undefinable psychological events. Special classes of concepts are called objects and relations. The referents of objects are single, primitive percepts, collections of such percepts (which are identified as material entities), other objects, or collections of other objects. As gestalts, objects are whole representations which are in a one-one correspondence with their referents. Relations are defined with respect to nonprimitive objects. They characterize the system of relationships among the elements of an object concept and exist only when these concepts can be known. Relations may also be either single and primitive, or more complex. Both objects and relations can refer to either a class of representations or to an individual representation.

This conception of a concept is ultimately reductionistic, for it assumes that all concepts have their roots in organism/environment experiences and that representations are just shorthand ways of referring to these experiences.

Concepts are, therefore, psychological constructs which arise through the interaction between an organism and the environment. As objects, they are either primitive (such as a visual perception) or they are composed of other relationships. Objects plus their conjoined properties are themselves objects. Thus, object concepts consist of a set of component concepts that are related in a set of ways. Relations are also composed of component objects which have certain properties conjoined to them and relationships between them. Some component relationships may also be primitive and undefined. For nonprimitive objects, then, relationships are a necessary part of a concept and must be included in a characterization of a concept. Even in the case of primitive objects, relationships must ultimately be attended to, since a primitive object can only be defined in relation to some other object. In that event, the primitive concept is the field, and the other one, the ground, and the primitive concept is characterized by not consciously attending to the ground concept.

This interpretation of a concept implies that word concepts can be considered to be particular forms of object and relation concepts, i.e., ones that symbolically signify their referents. By extension, then, formally defined concepts, as verbal representations of psychological events which are expressed in the form of sentences, are specifiable by the component objects, represented by nouns and adjectives, and by the relations, represented by nouns, adjectives, and verbs, which are included within the sentence. Sentence conjunctions are used to aid the incorporation of more component concepts into the definition.

The notions of word meaning and concept meaning adopted within this system follow directly from the above remarks. If a sentence is used to define a concept, then the meaning of the defined concept is given by the component concepts and by the interrelationships of these components, both of which are defined within the sentence. In both cases, the implied meaning is denotational, and the meaning of a concept is synonymous with the characterization of the concept itself. There may be affective, connotational meaning also attached to a concept, but as this would be the result of extra-sentential experiences, it is not of concern here.

If a sentence is assumed to be a particular characterization of a concept, and if this characterization is synonymous with the concept's meaning, then it follows that the meaning of a sentence is defined within the sentence's semantic structure, and that this meaning characterizes the concept. This characterization is the one which is posited to be internalized as a cognitive event.

Turning to the behavioral aspect of concept acquisition, it is assumed that a correct concept characterization allows, as a performance criterion, for the

correct selection of an extensional instantiation of that concept. This means that the instantiation embodies the intensional characteristics which have been internalized and that the learner can correctly produce or recognize positive and negative concept instances, as well as discriminate between them.

Viewed from these two perspectives, concept learning is posited to involve the acquisition of an internal representation and the development of the ability to correctly demonstrate some required performances. New concepts are defined with respect to existing concepts. What is learned during this acquisition process would, therefore, not be a set of performance behaviors, for these are assumed to have already been acquired, but, in the case of the meaningful reception learning of relations and nonprimitive objects, the concept's component objects, component object properties, and component object relationships.

The processes involved in the acquisition of formally defined concepts are assumed to be a function of the interaction of four elements: the context created by the task environment, the existing cognitive structure, the task performances exhibited by the learner during the course of acquisition, and an evaluation of these performances followed by some form of information feedback. In this system, learning is represented as a control system which is initiated by the goal concept that the learner must acquire. The goal concept exists within a particular task environment consisting of a set of concepts which are logically related to it. This logical organization is referred to as the content structure of a particular subject matter domain, and it is the context within which the goal concept is input into the learner's cognitive structure. It directly affects what components will be retrieved from long term memory and utilized during the acquisition process. The cognitive structure, in its turn, affects performance on the behavioral task given to measure the degree of acquisition. This performance is then evaluated, either by the learner himself or by another person, and information feedback is given in order to modify what is being learned. Considered together, this system can be interpreted cybernetically as a feedback loop which proceeds from the task environment, to the learner's internal cognitive representation, to his overt performance, to the evaluation of that performance, and, to the provision of appropriate feedback which feeds back, again, into the internal cognitive representation that is being acquired. In this system, cognitive structure acts as a transfer function within the black box (called the mind) that intervenes between the content structure and task performance. This process would terminate, with learning being said to have occurred, when the specific performance criterion given to the learner is reached.

More specifically, because the input to the acquisition process is a formally defined concept, the cognitive components of the learning process consist of comprehension, memory retrieval, encoding, matching, and correction processes which are interpreted as cognitive algorithmic information processing procedures. This sequence of mental operations does not operate serially. Instead, the processes mutually initiate and control each other at different points within the learning process.

The comprehension process is initiated upon the presentation of a formal definition to the learner. This definition is a logical one, defined within some content structure, and consists of component objects, component object properties, and component object relationships. This process, at the present time, consists of a syntactic parsing of the input sentence followed by a semantic sort of that parsing which defines a possible meaning of the sentence, i.e., a candidate, cognitive, learner-interpretation of the goal concept. The parsing process is based upon a modification of the systemic grammar used by Winograd (1972), and consists of a set of phrase structure rules which decompose a sentence into a finite number of sentence units. Because formally defined concepts are characterized as

declarative sentences which are complete units of discourse, only a limited set of phrase structure rules is needed. These rules are identified with the psychological operations that are presumed to occur during parsing. They are, therefore performance rules, and the grammar formally stating them is assumed to be a set of competence rules which represent the performance rules under conditions that are not limited by perceptual, memorial, attentional, or other similar kinds of constraints. The grammar is, as a result, a theory of performance, and no assumptions need to be made regarding the existence of a deep structure and transformational rules. As performance rules, the phrase structure rules can be sequenced and represented as a procedure which outlines, in detail, how the parsing process occurs.

The output of the parsing routine is assumed to be stored in an area of memory called, for want of a better name, a parsing store. This is not part of long term memory, corresponding more closely to the working memory of Greeno (1973) or the episodic memory of Tulving (1972). The contents of the parsing store serve as input to a routine which semantically sorts the parsed sentence units in order to extract the formally defined concept (i.e., the sentence meaning). This routine is context-specific and operates on the assumption that the sentence units are equivalent to the cases of a case grammar. The semantic sort routine knows how to assign the component object properties and relations; this is a function of how the parsed sentence units were stored in the parsing store.

After the words in the parsing store are semantically interpreted, they are placed in an area of memory called temporary concept store (TCS). This area can also be considered part of a working or episodic memory. It consists of the categories of component concepts, component concept properties, and component concept relationships. The contents of the TCS serve as the first approximation to the goal concept. They are idiosyncratic to an individual learner. These contents are modified in the course of an induction process which comprises the remainder of the learning process. During this process, the contents of the TCS change, so that when learning ends, the contents resemble the contents of the logical structure of the concept being acquired.

In almost every case of concept acquisition reviewed above, concept examples were used to aid in the induction of the concept. In meaningful concept learning, however, they serve to illustrate the concept, and help to provide feedback regarding whether or not it has been learned. Subjects are not told whether the instance is a positive or negative one; rather, they must make some behavioral response, such as classifying it. This behavioral aspect of concept acquisition is directly affected by the output of the comprehension process, i.e., the contents of the TCS.

Because the components of the TCS are concepts, they are also composed of other component concepts, component concept properties, and component concept relationships which are stored in long term memory and which must be retrieved in order to complete the meaning of the concept being acquired. Rather than be open to the charge of requiring that all the component concept's components be retrieved regressively until the primitive concepts are reached, it is assumed that component concepts, properties, and relationships two, or at most, three times removed, are retrieved. In this process of recall, they are transferred to the TCS, entering in the same format as the output of the semantic sort routine.

A matching process occurs following the transfer of information from LTM to the TCS and the presentation of a concept exemplar. It consists of testing to see whether or not all of the contents within the TCS are satisfied by the exemplar. If they are, then the exemplar is declared to be positive, while if they are not, it is stated to be a negative instance. This process is, essentially, a comparison routine, defining how the hypothetical internal characterization of the concept is linked to the behavioral performance which indicates whether or not the concept has been learned.

Once this performance has been made, it is evaluated in the external environment with respect to the required performance criterion, and if an incorrect response was given, self-regulated or external feedback is assumed to modify the contents of the TCS, thus altering the learner's concept characterization. Once this correction is made, the new representation is assumed to be used as input to the matching process when a new exemplar is presented.

The part of long term memory where the new concept is stored is interpreted in this approach as a section of a semantic memory which is segmented into separate areas, each one consisting of a particular domain of interest. Concepts within each area are linked together in a relational network. The links between the concepts are the properties and relationships which define them. Thus, in semantic memory, a concept would be stored as the cluster of its defining elements. Since these elements also define the concept's meaning, it follows that only the semantic meaning is stored in semantic memory, and that the intercomponent relationships are semantic relationships. The organization of semantic memory is sequential, but a hierarchical organization for objects is not precluded.

The contents and structure of this section of semantic memory can be formalized using graph theory (Harary, Norman, & Cartwright, 1965). Each concept in memory is identified as a node and the defining relationships between concepts are identified as lines. The concept nodes can be both objects and relations, and the node links can be unidirectional or bidirectional. The semantic memory formulation lies in direct correspondence with the formulation of a concept defined by the contents of the TCS.

The description of semantic memory using this notational system has the advantage that an individual's cognitive structure can be described using parameterized variables. The complexity of a concept can be defined by the number of links that occur with other concepts in memory. Two concepts can be considered reachable if there is a series of connecting links between the nodes which identify them. Furthermore, if every two nodes are mutually reachable, then the structure is said to be strongly connected; while it is weakly connected if at least two nodes are reachable. The distance between any two nodes is defined as the minimum number of links which are needed in order to go from one of the nodes to the other.

This proposed approach also allows for the direct inclusion of the stimulus, task, and learner variables. Stimulus variables enter into a consideration of how the task environment affects the encoded representation and the subsequent organization of semantic memory. Task variables affect the learner's performance during the testing of an exemplar, as well as the evaluation of that response and the determination of the required feedback. Learner variables affect the encoding and performance process and organization.

A special set of stimulus variables, called subject-matter variables, logically characterize the formally defined concept and directly affect what is learned, comprehension and retrieval, while indirectly affecting the behavioral process. Concept depth is a semantic variable referring to the number of component concepts, component concept properties, and component concept relationships which characterize a concept. The degree of concept connectedness refers to the extent to which the logical structure of the concepts presented within a unit can be interconnected. The degree of concept completeness refers to the extent to which the concept components exist and are stored in memory. Concept depth affects the comprehension, initial encoding, matching, and feedback processes and concept connectedness affects the ease of encoding and storing all the concepts in the unit. Lastly, if a goal concept is incomplete, the relevant subcomponents would first have to be learned or relearned.

As can be seen, the proposed theory takes various elements from several of the approaches in all three orientations and fuses them, with some modifications.

The notion of a concept and cognitive structure is akin to the ones offered by Collins and Quillian (1972), Kintsch (1972), Simmons (1973), Newell and Simon (1972), and Hunt (1973), and is not far from that of Piaget (1966), the notion of a task environment is close to Newell and Simon's (1972), the comprehension and encoding processes are modifications of Winograd's (1972) ideas, the independent variables are taken from all three orientations, and the subject-matter variable of connectedness was suggested by Greeno (1972), while the over-all model is formulated within the interactionist information processing paradigm.

This approach has many advantages. It answers the initial questions of the paper, attempting to meet the requirements of a synthetic approach which were given above. It is an explanatory theory which provides a necessary (functional) explanation and sufficient (explanatorily adequate and constructive) explanations. Concept acquisition is a goal-directed process. The internal representations can change during learning. The representations are idiosyncratic, thus allowing for individual differences. There is provision for both logical and psychological forms of the goal concept. Competence is considered to be an idealization of performance. Several previously unclear independent variables are made precise. Independent variables can be manipulated in order to see their effects upon the system's operation. The mechanisms of learning and the organization of memory are explicitly characterized. The comprehension and encoding process accounts for segmentation. Clustering in recall can be understood. And, most importantly, the failure of learning can be explained as a result of one or more of the following factors: incorrect comprehension, incorrect encoding, an incomplete semantic memory, failure in the testing process, or failure to receive and/or utilize feedback information correctly.

The model, however, also has a few drawbacks. It is a somewhat specialized model which is limited to non-primitive formally defined concepts. It is a bit ad hoc in parts. Not all of the mechanisms are fully explained, such as how the learner locates the relevant exemplar components in order to compare them to the concept constituents stored in the TCS. Lastly, it cannot handle the paraphrase problem. This is not viewed as irreconcilable, however, since some rules, similar to those used by Simmons (1973), could be written into the parsing procedure which would cause the same set of syntactic elements to be input into the semantic sort.

Given this description and theory of acquisition, the question of empirical testability still remains to be considered. A *reductio ad absurdum* argument is proposed which asks: If concept components, component properties, and component interrelationships are not learned, then what else could be learned? The suggested alternatives might be that (1) the formal definition is rote learned verbatim; (2) the syntactical structure of the definition is learned; (3) only a form of PA pairings between exemplars and the concept label occurs, with no accompanying comprehension. These alternatives can be assessed by subjecting four groups of subjects (one for each of the four alternatives) to training sets. One group would be trained to learn the definition verbatim, another to perform PA pairings, another to perform a syntactical analysis, and the last to extract relevant concept constituents. Each group is first given a pretest in order to insure that previous knowledge is equated, followed by the same set of concepts. When a common criterion is satisfied by all the groups, including a control group (also given a pretest) which until now has been resting, are given a new unit of concepts and the rate and ease of acquisition is measured for each group. All of the stimulus, task, and learner independent variables are equated across groups. The dependent variables that could be used include the number of trials to learn each concept in the unit and the number of concepts learned for a fixed number of trials. A trial would consist of the successive presentation of each concept, where each presentation constitutes the formal definition of the concept followed by exemplars of that concept. The number of

concepts to be learned would be an independent task variable, as would be the presentation sequence and the form of verbal definition used.

This methodology is developed under the assumption that if concept constituents are learned, fewer trials are required for acquiring all the concepts in the unit, and more concepts are learned for a fixed number of trials, than if one of the alternative approaches to what is learned is accepted. The control group is used in order to see which alternative group performance it is closest to, thus suggesting that untrained Ss use similar processes during learning.

The methodology also involves the assumption that this conception of what is learned is most effective in terms of promoting retention, transfer to new concept acquisition, and problem solving. Retention (recognition) can be tested by requiring that Ss perform the same learning behaviors when new concept instances are presented. Confusion data would be gathered, as well as the number of correct recognitions. Recall tests would involve the production of new instances. Transfer could be tested by the presentation of additional concept units to Ss in each group on successive days. It would be revealed by the increased ease of learning, as measured by the number of trials to criterion, and would be a function of the subject-matter variables. Problem solving tasks would be knowledge application tasks. They could include translating a verbal geometric concept into spatial or symbolic terms, extrapolating a conclusion of a theorem from given antecedent conditions, and proving a theorem, given the relevant definitions and theorems. Independent variables affecting these tasks would be manipulated and held constant across groups.

Turning to the organization of semantic memory, sorting tasks could be used in order to see how the currently acquired unit is encoded, how that encoding affects the organization of previously acquired concepts, if at all, and whether or not concepts are organized with respect to their constituent components. For both group and individual data, sorting data would be converted into proximity data which is transformed into an adjacency matrix that characterizes the digraph formulation. The adjacency matrix is then converted into a digraph-distance matrix (Harary *et al.*, 1965), and scaled multidimensionally or clustered hierarchically. Euclidean distances can also be computed between the digraph-distance matrices and a digraph-distance matrix constructed from the logical content structure. Other methods that can also be used to investigate semantic memory organization are the use of partition metrics (cf. Arabie & Boorman, 1973) and the scaling of similarity judgment data (Johnson, 1970).

The experimental formulation of this organization could also help to provide empirical support for the existence and specification of the hypothetical operators which are part of the information processing processes. If different concept representations can be identified at various points of the learning process, with each such representation corresponding to a different learning state, a sequence of these states would implicitly define a transformation process. The reason for this is that if structure and process are just different abstract descriptions of the behavior of a system of events, then they are relative to, and mutually define one another. Process would then be the change in states over time. If the change is limited to two successive states, then an operator can be categorized as the transformational relation between these states.

Two other methodological tools are simulations and confirmation tasks. In the latter case, reaction time data could be used in order to lend support to the claim that concept constituents are tested during learning to see if they are instantiated in the test exemplars, as well as to test the depth of concepts in memory.

The successful development and execution of this methodology would satisfy the criterion that proposed explanations be testable and empirically verifiable. They would be the final steps in the initial development of a theory of the acquisition of formally defined concepts, providing feedback for the further development and modification of that theory.

References

- Anderson, J. R., & Bower, G. H. Human associative memory. Washington, D.C.: V. H. Winston, 1973.
- Anglin, J. M. The growth of word meaning. Cambridge, Mass.: M.I.T. Press, 1970.
- Arabie, P., & Boorman, S. A. Multidimensional scaling of measures of distance between partitions. Journal of Mathematical Psychology, 1973, 10, 148-203.
- Archer, E. J. The psychological nature of concepts. In H. J. Klausmeier & C. W. Harris (Eds.), Analyses of concept learning. New York: Academic Press, 1966.
- Ausubel, D. P. Meaningful reception learning and the acquisition of concepts. In H. J. Klausmeier & C. W. Harris (Eds.), Analyses of concept learning. New York: Academic Press, 1966.
- Ausubel, D. P., & Robinson, F. G. School learning. New York: Holt, Rinehart & Winston, 1969.
- Bem, S. L. The role of comprehension in children's problem solving. Developmental Psychology, 1970, 2, 351-358.
- Bourne, L. E., Jr. Human conceptual behavior. Boston: Allyn & Bacon, 1966.
- Bourne, L. E., Jr., & Restle, F. Mathematical theory of concept identification. Psychological Review, 1959, 66, 278-296.
- Bousfield, W. A. The problem of meaning in verbal learning. In C. N. Cofer (Ed.), Verbal learning and verbal behavior. New York: McGraw-Hill, 1961.
- Bransford, J. D., & Franks, J. J. Abstraction of linguistic ideas. Cognitive Psychology, 1971, 2, 331-350.
- Brown, R. Words and things. Glencoe, Ill.: Free Press, 1958.
- Bruner, J. S., Goodnow, J. J., & Austin, G. A. A study of thinking. New York: Wiley, 1956.
- Carroll, J. B. Words, meanings and concepts. Harvard Educational Review, 1964, 34, 178-202.
- Chomsky, N. Aspects of the theory of syntax. Cambridge, Mass.: M.I.T. Press, 1965.
- Collins, A. M., & Quillian, M. R. How to make a language user. In E. Tulving & W. Donaldson (Eds.), Organization of memory. New York: Academic Press, 1972.
- Deese, J. Association and memory. In T. R. Dixon & D. L. Horton (Eds.), Verbal behavior and general behavior theory. Englewood Cliffs, N.J.: Prentice-Hall, 1968.
- Fillenbaum, S. On the use of memorial techniques to assess syntactic structures. Psychological Bulletin, 1970, 73, 231-237.
- Flavell, J. H., & Wohlwill, J. F. Formal and functional aspects of cognitive development. In Elkind & J. H. Flavell (Eds.), Studies in cognitive development. New York: Oxford University Press, 1969.
- Furth, H. G. Piaget and knowledge. Englewood Cliffs, N.J.: Prentice-Hall, 1969.
- Gagne, R. M. The acquisition of knowledge. Psychological Review, 1962, 69, 355-365.
- Gagne, R. M. The learning of principles. In H. J. Klausmeier & C. W. Harris (Eds.), Analyses of concept learning. New York: Academic Press, 1966.
- Gange, R. M. The conditions of learning. New York: Holt, Rinehart & Winston, 1970.
- Greene, J. Psycholinguistics. Middlesex, England: Penguin Books, 1972.
- Greeno, J. G. On the acquisition of a simple cognitive structure. In E. Tulving & W. Donaldson (Eds.), Organization of memory. New York: Academic Press, 1972.
- Greeno, J. G. The structure of memory and the process of solving problems. In R. L. Solso (Ed.), Contemporary issues in cognitive psychology. Washington, D.C.: V. H. Winston, 1973.
- Harary, R., Norman, R. Z., & Cartwright, D. Structural models: An introduction to the theory of directed graphs. New York: Wiley, 1965.
- Hunt, E. The memory we must have. In R. C. Shank & K. M. Colby (Eds.), Computer models of thought and language. San Francisco: Freeman, 1973.
- Johnson, P. E., Cox, D. L., & Curran, T. E. Psychological reality of physical concepts. Psychonomic Science, 1970, 19, 245-247.

308 Organization and Acquisition of Formally Defined Concepts

- Jung, J. Verbal learning. New York: Holt, Rinehart & Winston, 1968.
- Kendler, H. H. The concept of the concept. In A. W. Melton (Ed.), Categories of human learning. New York: Academic Press, 1964.
- Kintsch, W. Notes on the structure of semantic memory. In E. Tulving & W. Donaldson (Eds.), Organization of memory. New York: Academic Press, 1972.
- Mandler, G. Organization and recognition. In E. Tulving & W. Donaldson (Eds.), Organization of memory. New York: Academic Press, 1972.
- Michon, J. A. Multidimensional and hierarchical analysis of progress in learning. In L. W. Gregg (Ed.), Cognition in learning and memory. New York: Wiley, 1972.
- Miller, G. A. A psychological method to investigate verbal concepts. Journal of Mathematical Psychology, 1969, 6, 169-191.
- Mowrer, O. H. The psychologist looks at language. American Psychologist, 1954, 9, 660-694.
- Newell, A. A note on process-structure distinctions in developmental psychology. In S. Farnham-Diggory (Ed.), Information processing in children. New York: Academic Press, 1972.
- Newell, A., & Simon, H. A. Human problem solving. Englewood Cliffs, N.J.: Prentice-Hall, 1972.
- Noble, C. E. Analysis of meaning. Psychological Review, 1952, 59, 421-430.
- Osgood, C. E. Behavior theory and the social sciences. Behavioral Science, 1956, 1, 167-185.
- Perfetti, C. A. Lexical density and phrase structure depth as variables in sentence retention. Journal of Verbal Learning and Verbal Behavior, 1969, 8, 719-724.
- Piaget, J. The psychology of intelligence. Totowa, N.J.: Littlefield, Adams & Co., 1966.
- Piaget, J. Language and thought from the genetic point of view. In D. Elkind (Ed.), Six psychological studies. New York: Random House, 1967.
- Piaget, J. Biology and knowledge. Chicago: University of Chicago Press, 1971.
- Quillian, M. R. Semantic memory. In M. Minsky (Eds.), Semantic information processing. Cambridge, Mass.: M.I.T. Press, 1968.
- Rips, L. J., Shoben, E. J., & Smith, E. E. Semantic distance and the verification of semantic relations. Journal of Verbal Learning and Verbal Behavior, 1973, 12, 1-20.
- Rumelhart, D. E., Lindsay, P. H., & Norman, D. A. A process model for long-term memory. In E. Tulving & W. Donaldson (Eds.), Organization of memory. New York: Academic Press, 1972.
- Sachs, J. D. S. Recognition memory for syntactic and semantic aspects of connected discourse. Perception and Psychophysics, 1967, 2, 437-442.
- Scandura, J. M. Structural learning I: Theory and research. New York: Gordon & Breach, 1973.
- Shavelson, R. J. Some aspects of the correspondence between content structure and cognitive structure in physics instruction. Journal of Educational Psychology, 1972, 63, 225-234.
- Sigel, I. E. The Piagetian system and the world of education. In D. Elkind & J. H. Flavell (Eds.) Studies in cognitive development. New York: Oxford University Press, 1969.
- Simmons, R. F. Semantic networks: Their computation and use for understanding English sentences. In R. C. Shank & K. M. Colby (Eds.), Computer model of thought and language. San Francisco: Freeman, 1973.
- Simon, H. A. The sciences of the artificial. Cambridge, Mass.: M.I.T. Press, 1969.
- Simon, H. A. The heuristic compiler. In H. A. Simon & L. Siklossy (Eds.), Representation and meaning. Englewood Cliffs, N.J.: Prentice-Hall, 1972.
- Suppes, P. Some theoretical models for mathematics learning. Journal of Research and Development in Education, 1967, 1, 5-22.
- Suppes, P. On the behavioral foundations of mathematical concepts. In Cognitive development in children. Chicago: University of Chicago Press, 1970.
- Tulving, E. Theoretical issues in free recall. In T. R. Dixon & D. L. Horton (Eds.) Verbal behavior and general behavior theory. Englewood Cliffs, N.J.: Prentice-Hall, 1968.
- Tulving, E. Episodic and semantic memory. In E. Tulving & W. Donaldson (Eds.), Organization of memory. New York: Academic Press, 1972.

- Underwood, B. J. Some relationships between concept learning and verbal learning. In H. J. Klausmeier & C. W. Harris (Eds.), Analyses of concept learning. New York: Wiley, 1966.
- Voss, J. F. On the relationship of associative and organization processes. In E. Tulving & W. Donaldson (Eds.), Organization of memory. New York: Academic Press, 1972.
- Wickelgren, W. A. Coding, retrieval, and dynamics of multitrace associative memory. In L. W. Gregg (Ed.), Cognition in learning and memory. New York: Press, 1972.
- Winograd, T. A program for understanding natural language. Cognitive Psychology, 1972, 3, 1-191.

A TEST-THEORETICAL APPROACH TO STRUCTURAL LEARNING

Wilhelm F. Kempf

*Institut für die Pädagogik der Naturwissenschaften
an der Universität Kiel, West Germany*

In the first chapter of his book on "Structural Learning" SCANDURA (1973) has expressed his belief that deterministic theorizing about complex human learning may actually be easier than stochastic theorizing. There is also a number of papers (i.e. SCANDURA 1971, SCANDURA & DURIN 1971) in which SCANDURA and his co-workers have argued that deterministic theorizing may be even more fruitful.

Nevertheless, there are some reasons, why stochastic theorizing might be preferred, and there arises the question whether and how the typical problems connected with stochastic theories of human learning can be solved.

The present paper is concerned with some of the strongest of SCANDURA's arguments in favour of the deterministic approach. It focuses on the question how interindividual differences of competence and knowledge can be buildt into a theory of learning in such a way, that the individuals' competence can be measured in a satisfactory manner and that the structure of a task can be inferred independently from the sample of individuals being tested.

As will be shown, most of SCANDURA's methodological postulates can not only be fulfilled by deterministic models but also within the framework of stochastic and dynamic test theory. On the basis of test-models suggested by RASCH (1960, 1961), FISCHER (1972) and KEMPF (1974a, b) a stochastic theory of structural learning will be proposed. This theory, however, is not merely another methodological

approach, but it diverges from SCANDURA's theory in some basicly psychological assumptions too. In contrast to SCANDURA's theory

- it will be assumed that an individual is *more or less* capable of performing a rule;
- it will be assumed that an individual does learn something from the performance of a task, that this learning is dependent from the structure of the task and that it is generalized to other tasks as well;
- it will not be assumed that an individual will solve any item which is composed of rules that are known by the individual, regardless of the number and difficulty of the rules to be applied.

As compared with traditional statistical methods, deterministic theorizing about human learning involves three advantages, at least. In deterministic theories of structural learning

- an individual's behavior potential can be inferred easily from the individual's response pattern;
- the structure of a task can be inferred independently from the sample of individuals on which the empirical analysis is being based, so that every sample will lead to essentially the same results; and
- on the basis of this the behavior of individual subjects in specific situations can be predicted.

The domain of deterministic theorizing, however, is restricted to rather simple and - sometimes - even artificial situations. In the analysis of complex tasks one would prefer to include *failures by chance* which are due to fluctuations of attention, memory, and so on. Nevertheless, SCANDURA & DUPNIN (1971) have adopted the view that there is *no need* to resort to probability. Although perfect prediction is impossible in complex situations, the authors argue that the uncertainty, rather than being an explicit part of the theory or prediction

model, has its source elsewhere.

"The reason that perfect prediction is impossible, in principle, is that in addition to the role played by memory, there is always some residual uncertainty about the capabilities and motivations that an S brings to any given task. And, this is exactly where this uncertainty should be put" (SCANDURA & DURNIN 1971, p.4).

Although this is one possible approach to finish with the uncertainty-problem, I would not agree with it. If we add some source of uncertainty or failure to a deterministic theory, then it cannot be decided whether a deviating observation contradicts the theory or whether it is due to the uncertainty. Consequently, the testability of the theory gets lost, and there are but two possibilities to restore it. We either might try to build the psychological sources of uncertainty explicitly into our theory and by this to return to perfect predictions and to a new and more elaborated deterministic theory, or, if we do not succeed in this, we might leave the errors in prediction to inadequacies in observation and measurement. In this case, however, we would have to look for some methods to estimate the failure component in our observations, and this is exactly what classical test theory has tried to do. Classical test theory has tried to estimate the uncertainty *without* making it an explicit part of the theory itself. And classical test theory has failed in this. All of the well known indices of classical test theory are highly dependend on the arbitrary choosen sample of individuals and hence cannot be applied to the single observation or single subject (cf. FISCHER 1968).

Another of SCANDURA & DURNIN's (1971) arguments in favour of the deterministic approach arises directly from the shortcoming of classical test theory. As classical test theory has neglected to ask for a meta-theory of measurement, the traditional approach to the measurement of individual differences has remained highly unelaborated. Predictions about an

individual's behavior use to be made on the basis of the individual's performance on a random sampling of test items and usually result in such statements as "On the average, he should get eight out of ten items correct" (cf. SCANDURA & DURNIN 1971, p.4).

It is obvious, that statements like this are worthless for a theory of human learning or any other kind of complex human behavior. Most of the traditional stochastic learning models, therefore, have neglected the interindividual differences at all, and this has caused another of SCANDURA's arguments:

"For example, in stochastic models of paired-associate learning it is usually assumed that each subject has the same probability of learning on each trial. Even the most superficial analysis of relevant data, however, indicates clearly that the probability of success for different subjects may vary greatly. And one cannot attribute this to the fact that the probability of learning is a random variable. This would still not explain the fundamental fact that the probability of success of many subjects tends to be either uniformly high, or low, over different trials" (SCANDURA 1971, p.23).

As far as I know, the first psychological application of probability in which the interindividual differences have been treated in a satisfactory manner was stochastic test theory which describes the probability of an individual's (v) success on a task or item (i) as a function of the individual's competence (ξ_v).

$$(1.) \quad p\{+/v, i\} = f_i(\xi_v)$$

The functions f_i are called *trace lines* (LAZARFIELD 1950) or *item characteristic curves* (LORD & NOVICK 1968).

In contrast to stochastic learning theory, however, the item characteristic curves are assumed to be *locally stochastic independent* so that all item intercorrelations vanish if ξ_v is held constant. The individual responses' functional dependency on the latent variable ξ is assumed

to be the only source of the item intercorrelations.
With the notation

$$a_{vi} = \begin{cases} 1 & \text{if } S \text{ v solves item } i \\ 0 & \text{if } S \text{ v does not solve the item} \end{cases}$$

the formal definition of local stochastic independence is

$$(2.) \quad p\{a_{v1}, a_{v2}, \dots, a_{vg}\} = \prod_{i=1}^g p\{a_{vi}\} \\ \text{for all } g=1, \dots, k,$$

in which k is the total number of items in the test. The probability of an individual's success on an item is assumed to be independent from whether or not the individual has solved any of the preceding items.

From this it becomes clear that stochastic test theory gives the methodological framework for a theory of competence only, but not for a theory of learning. Similarly as in SCANDURA's (1971, 1973) deterministic partial theory of structured knowledge, the individuals' competence is measured under the idealized condition that no learning will occur during testing. Differently from SCANDURA's theory, however, the individuals' competence is assumed to be a quantitative variable, and the behavior of individual subjects in specific situations cannot be predicted deterministically but only with a well known probability.

This probability is assumed to be a function of the individual's competence as well as a function of the situation or task and if the item characteristic curves can be described to be functions of the atomic rules into which the items are to be decomposed, then stochastic test theory can also be used in inferring the *structure* of a task.

With the introduction of individual parameters ξ_v , however, there arises a considerable statistical problem. Each S that

is added to the sample causes the introduction of another parameter and according to a well known theorem by NEYMAN & SCOTT (1948) the traditional methods of parameter estimation fail if the number of parameters to be estimated does not tend towards a fixed numeral while the number of observations grows to infinity. The only models in which consistent estimators exist in such a situation are those in which the individual parameters can be separated from the structural parameters of the items by use of *conditional inference methods* (cf. ANDERSEN 1973a) and as ANDERSEN (1973b) has shown, within the framework of stochastic test theory, these are the models suggested by RASCH (1960, 1961) only.

In case of binary data the *Rasch*-model has the form

$$(3.) \quad p\{+/v, i\} = \frac{\xi_v}{\xi_v + \sigma_i},$$

in which σ_i stands for the difficulty of the task, which, because of the separability of the parameters, can be estimated independently from the underlying sample of individuals by maximizing the *conditional likelihood*

$$(4.) \quad L = p\{A/(a_{v0})\}$$

in which A stands for the response-matrix of n individuals to k items. (a_{v0}) stands for the vector $(a_{10}, a_{20}, \dots, a_{n0})$ in which

$$a_{v0} = \sum_{i=1}^k a_{vi}$$

is the number of items solved by the individual v .

If the model (3.) holds, then (a_{v0}) is a minimal sufficient statistic of the parameter-vector $(\xi_1, \xi_2, \dots, \xi_n)$ and the

conditional likelihood (4.), therefore, does not depend on the individual parameters.

Similarly, the individual parameters ξ_v can be estimated from a conditional likelihood function which does not depend on the item parameters.

By use of *conditional likelihood ratio tests* (ANDERSEN 1973a), finally, it can be tested whether the model fits the data and whether the assumed structure of the tasks conforms with the empirical results (cf. FISCHER 1973). The *Rasch*-model, therefore, is as good a tool for research on structured knowledge or competence as any deterministic model can be. It is a *testable* model that refers to *inter-individual differences* of competence and in which the structure and difficulty of the items can be inferred independently from the arbitrary chosen sample of individuals. The behavior of individual subjects in specific situations can be predicted in probability.

The first one to apply the *Rasch*-model in research on structured knowledge was SCHEIBLECHNER (1972) who made use of a linear extension⁺⁾ of the model (3.).

His answer to the question, how reasoning processes can be broken down into atomic rules or operations was to factorize the item difficulties σ_i into subfactors η_j which represent the difficulties of the single operations $j=1, m$ of which the tasks are composed.

According to this model, a set of tasks can be broken down into atomic rules *iff* the task difficulties can be broken down into corresponding factors, so that

$$(5.) \quad \sigma_i = \exp\left(\sum_{j=1}^m f_{ij} \eta_j + c\right) ,$$

+) Most of the mathematical and statistical foundations of the *linear logistic test model* have been worked out by FISCHER (1972) and FISCHER & FORMANN (1972).

in which f_{ij} denotes the weight of rule j in the solution process of item i .

In the most simple case f_{ij} will be the number of times that the rule has to be applied in order to solve the task (cf. SPADA 1973).

Another interpretation of the model (cf. FISCHER 1973) defines

$$f_{ij} = \begin{cases} 1 & \text{if rule } j \text{ has to be applied on item } i \\ & \text{in order to solve the item} \\ 0 & \text{if rule } j \text{ is not needed for the solution} \\ & \text{of the item} \end{cases}$$

and thus assumes, that the difficulty of a task is primarily determined by the combination of different rules to be applied and is not increased significantly when the same operation occurs repeatedly within the task.

Both interpretations of the model have successfully been applied to empirical data. In the analysis of reasoning processes underlying elementary differential calculus as taught in secondary school mathematics, for instance, the second interpretation of the model was to be preferred (cf. FISCHER 1973), while the first one worked better in the analysis of reasoning processes underlying elementary mechanics (cf. SPADA, FISCHER & HEYNER 1973).

Some more empirical applications of the model have been carried out by SCHEIBLECHNER (1972) and HEINRICH (1973) who analyzed syllogistic reasoning.

In some of these studies, slightly modified versions of the model (5.) were used, which also describe some training effects that arise from the individuals' prior trials on the atomic rules. These modifications of the model make the assumption that the training effect which arises from an individual's trial on a task can also be broken down into atomic fractions which correspond to the rules that

are to be applied on the task.

As an example of this type of models we may discuss the model of a rule specific training transfer (SPADA 1973)

$$(6.) \quad \sigma_i = \exp\left(\sum_{j=1}^m f_{ij}(\omega_j - H_{ij}\beta_j) + c\right),$$

in which β_j is the training transfer corresponding to the rule j and H_{ij} is some function of the training frequency

$$h_{ij} = \sum_{g < i} f_{gj}$$

of the rule j in the preceding tasks $g=1, i-1$. Evidently, the item difficulties are assumed to be dependend on *variable operation difficulties* $\eta_j = \omega_j - H_{ij}\beta_j$ now, which are a function of

- the initial difficulties ω_j of the rules,
- the training frequencies h_{ij} of the rules, and of
- the training effect parameters β_j .

As compared with the simple decomposition model (5.) and with SCANDURA's deterministic theory, therefore, the training effect model has the important advantage to bridge the gap between learning and testing. The testing of an individual is no longer assumed to leave the individual unchanged.

All of these training effect models, however, have to make one assumption which is not very plausible from a psychological point of view: the training effect arising from an individual's trial on a task must be assumed to be independent from the individual's success. It must be assumed to be the same when an individual gives a correct solution of the task as when the individual cannot solve it. Any other assumption would violate the condition of local stochastic independency which is a *fundamental* presupposition of the *Rasch*-models. If a stochastic model of structural *learning* is to be constructed, then the framework of stochastic test theory must

be extended to what I have called *dynamic test theory* (KEMPF 1974c).

The basic conception of dynamic test theory is to replace the assumption of local stochastic independency by what may be called *local serial dependency*

$$(7.) \quad p\{(a_{vi})\} = \prod_{i=1}^k p\{a_{vi}/s_{vi}\},$$

in which $(a_{vi}) = (a_{v1}, a_{v2}, \dots, a_{vk})$ stands for the individual's response vector on the total of k items and s_{vi} stands for the partial response vector $(a_{v1}, a_{v2}, \dots, a_{vi-1})$. The item characteristic curves $f_i(\xi_v)$ are replaced by *conditional item characteristic curves*

$$(8.) \quad f_{i.s_{vi}}(\xi_v) = p\{a_{vi}=1/(a_{v1}, a_{v2}, \dots, a_{vi-1})=s_{vi}\}$$

which describe the individual's probability of success on item i under the side condition of the individual's responses to the preceding items.

On the basis of this formalism dynamic test models that have essentially the same properties as the *Rasch*-model can be constructed. Two of these dynamic models have been elaborated to some statistical detail by KEMPF (1974a,b). Some of the most important results are summarized in KEMPF (1974c).

Initially, I had suggested dynamic test theory as a tool for research in social psychology (KEMPF 1974a,c) and the model that was developed in these studies was based on the assumption that the *transfer* arising from the individual's responses to the preceding items is dependent on the *number of positive responses* only. Although the model can be applied to some simple learning situations also, it will not be discussed here.

The other model was suggested by KEMPF (1974b) and treats the conditional item characteristic curves

$$(9.) \quad f_{i.s_{vi}}(\xi_v) = \begin{cases} \frac{\xi_v + \psi_{s_{vi}}}{\xi_v + \sigma_i} & \text{for } \psi_{s_{vi}} < \sigma_i \\ 1 & \text{for } \psi_{s_{vi}} \geq \sigma_i \end{cases}$$

under the assumption that the learning effect $\psi_{s_{vi}}$ is a function of the cumulative weights

$$g_{vij} = \sum_{g < i} a_{vg} f_{gj}$$

that the atomic rules have in those of the preceding items to which the individual has responded positively.

Similarly as in the linear logistic model, an exponential function

$$(10.) \quad \psi_{s_{vi}} = \exp\left(\sum_{j=1}^m \frac{g_{vij} \cdot \alpha_j}{1 + g_{vij}} + c\right) + c'$$

is used. α_j is a rule-specific learning parameter and as

$$g_{vij}/(1 + g_{vij})$$

tends towards 1 when g_{vij} grows to infinity, α_j is to be interpreted as the upper bound of the rule-specific learning effect. The item difficulties σ_i are decomposed as in the training transfer model (6.). Conditional maximum likelihood estimators for the parameters have been derived by KEMPF (1974b).

From inserting the estimates into formula (9.) the probability of an individual's success on the various items can be predicted and on the basis of these predictions an individualistic selection of items can be carried out in such a way that the

individual will have a maximum learning effect after a minimum number of trials.

An individual will be predicted to solve an item with probability 1, if the item difficulty is estimated to be smaller than the learning effect. By this, the model bridges the gap between stochastic and deterministic theorizing.

REFERENCES

- ANDERSEN, E.B.: Conditional inference and models for measuring. Copenhagen: Mentalhygiejnisk Forlag, 1973a.
- ANDERSEN, E.B.: Conditional inference for multiple-choice questionnaires. *British Journal of mathematical and statistical Psychology* 1973b, 26, 31-44.
- FISCHER, G.H.: *Psychologische Testtheorie*. Bern: Huber, 1968.
- FISCHER, G.H.: Conditional maximum-likelihood estimation of item parameters for a linear logistic test model. *Research Bulletin 9, Psychologisches Institut der Universität Wien*, 1972.
- FISCHER, G.H.: The linear logistic test model as an instrument in educational research. *Acta Psychologica* 1973, 37, 359-374.
- FISCHER, G.H. & FORMANN, A.K.: An algorithm and a *FORTRAN* program for estimating the item parameters of the linear logistic test model. *Research Bulletin 11, Psychologisches Institut der Universität Wien*, 1972.
- HEINRICH, I.: Veränderung des logischen Schlußprozesses bei semantischer Variation seiner Elemente. Unveröff. Diss., Psychologisches Institut der Universität Wien, 1973.
- KEMPF, W.F.: Dynamische Modelle zur Messung sozialer Verhaltensdispositionen. In: KEMPF, W.F. (Hrsg.): *Probabilistische Modelle in der Sozialpsychologie*. Bern: Huber, 1974a (in print).
- KEMPF, W.F.: Einige dynamische Meßmodelle und ihre Anwendung in der Denk- und Lernpsychologie. In: KEMPF, W.F. (Hrsg.): *Drei Arbeiten über verallgemeinerte Rasch-Modelle. Three papers on generalized Rasch-models*. IPN-Arbeitsbericht 10, Institut für die Pädagogik der Naturwissenschaften an der Universität Kiel, 1974b.
- KEMPF, W.F.: A dynamic model for measuring individual and situative influences on social behavior. *Proceedings of the COMPSTAT symposium on computational statistics, Vienna (Austria) Sept. 2-6, 1974c* (in print).
- LAZARUSFELD, P.F.: Logical and mathematical foundations of latent structure analysis. In: STAUFFER, S.A. et al.: *Measurement and Prediction*. Princeton: Princeton University Press, 1950.

- LORD, F.M. & NOVICK, M.R.: Statistical theories of mental test scores. Reading, Mass.: Addison-Wesley, 1968.
- RASCH, G.: Probabilistic models for some intelligence and attainment tests. Copenhagen: Nielsen & Lydiche, 1960.
- RASCH, G.: On general laws and the meaning of measurement in psychology. Proceedings of the fourth Berkeley symposium on mathematical statistics and probability, Vol.IV, Berkeley: University of California Press, 1961.
- SCANDURA, J.M.: Deterministic theorizing in structural learning: three levels of empiricism. *Journal of Structural Learning* 1971, 3, 21-53.
- SCANDURA, J.M.: Structural learning I. Theory and research. New York: Gordon & Breach, 1973.
- SCANDURA, J.M. & DURVIN, J.H.: Assessing behavior potential: adequacy of basic theoretical assumptions. Report 62, University of Pennsylvania: Structural Learning Series, 1971.
- SCHEIBLECHNER, H.: Das Lernen und Lösen komplexer Denkaufgaben. *Zeitschrift für experimentelle und angewandte Psychologie* 1972, 19, 476-506.
- SPADA, H.: Denk- und Lernmodelle der *Rasch*-Meßtheorie unter psychologischen, formalen und didaktischen Aspekten. In: SPADA, H., HAUSSLER, P. & HEYNER, W. (Hrsg.): Denkoperationen und Lernprozesse als Grundlage für lernerorientierten Unterricht. IPN-Arbeitsbericht 5, Institut für die Pädagogik der Naturwissenschaften an der Universität Kiel, 1973.
- SPADA, H., FISCHER, C.H. & HEYNER, W.: Die Analyse von Denkoperationen und Lernprozessen bei der Lösung von Problemstellungen aus der Mechanik mittels des linearen logistischen Modells. In: SPADA, H., HAUSSLER, P. & HEYNER, W. (Hrsg.): Denkoperationen und Lernprozesse als Grundlage für lernerorientierten Unterricht. IPN-Arbeitsbericht 5, Institut für die Pädagogik der Naturwissenschaften an der Universität Kiel, 1973.

ROUNDTABLE DISCUSSION (How Does Learning Take Place?), Sunday, April 21, 1974, 4:30-5:30

Joseph M. Scandura: The topic of the roundtable this afternoon is "How does learning take place." After listening intently for two days, I suspect that many of you would like an opportunity to get into the act. Let's start by allowing each of the panelists up to three minutes to make a preliminary statement expressing his general view. To allow an opportunity for as much discussion as possible, let's limit all subsequent comments from both participants and panelists to a maximum of one minute. Dave would like to start.

David Merrill: I'd like to make a couple of remarks that differ from some of the things I said this morning. Since the topic of this panel is how learning takes place, let me suggest that, in spite of the excellent work on models that is being done, we are unable at this point in time to make very adequate predictions, especially in a practical learning kind of situation, about how learning takes place. Therefore, I would propose that some of the work that Dr. Pask described and some of the work that we've done with computer assisted instruction capitalizes on the notion that the learner can be taught to manipulate his learning environment or the learning instructional system in such a way as to optimize or maximize learning efficiency or effectiveness for him. Perhaps the optimal approach to the question of learning is to teach learners learner control.

Many of the assumptions we've made in C.A.I. systems and even in developing models is that we ought to eventually be able to prescribe learning programs for individual learners. I would suggest that that's an inappropriate assumption and maybe what we ought to be doing is teaching learners to adapt the environment to their own needs rather than talking about adapting the environment to learners.

Coltan P. Dienes: I think there are essentially two kinds of learning which ought to be distinguished. One is a rather simple kind of learning which relies largely upon memory and a small amount of organizing. That type of learning is what is most investigated, as evidenced by the vast amount of learning literature that exists today. The moment you go beyond ordinary, simple things, and you have to learn how to operate a system, then you get very much more complex situations. You need to split up the learning into a number of different stages. I don't know if you can in practice follow the actual learning dynamics under laboratory conditions, simply because it's too expensive to have somebody there for so long to take down everything. It would take months possibly -- certainly weeks -- before someone could effectively observe learning in a real situation. And so then, you go to a sort of Piaget type of observation or teacher observation or psychologist observation. In fact, this is how I developed my rather vague model of six stages, which I will just mention very briefly. The first stage is free interaction with the learning environment. Second, the discovery and manipulation of the rules apparent in the environment. Third, a comparison of one rule structure with another rule structure, one embodiment to another embodiment, etc. Fourth, the mapping of rule structures onto visually distinguishable maps, that is, representations. Fifth, the development of a symbolic language for description of these mapped rule structures. Sixth, comes the formalization stage when you put some order into the symbolized way of describing rules. So, pre-interaction, rule learning or discovery, comparison of rules, representation, symbolization, and formalization, I think, are essential parts of learning. Obviously, this is not exhaustive. I am suggesting this as a [general descriptive] model.

Anan Banerji: I have a problem. I don't think it is a problem about learning; I think it is a problem about guessing. I have been showing a guy a bunch of two-digit numbers written in ordinary decimal notation and I have been calling the numbers good or bad. For example, I show someone 72, 54, and 86, and I have called them good. Then, I have shown him 51 and 73; I have called them bad. At this point, he may wonder whether any time he sees 2, 4, or 6 [units digits], it is good. But he really doesn't have any confidence in saying that because, after all, he has seen a 2 only once, a 4 only once, and a 6 only once. So, he really can't guess at all. I do the same experiment with the smart aleck who has learned about odd and even numbers, and he sees three even numbers

to be good and two odd numbers to be bad. So, he's more sure that I'm talking about even and odd numbers as being good and bad. Now, both of these guys have really seen the same experiment and just have a slightly different basis of knowledge. Does that really make the second guy's guess any better? And, if so, what do we mean by better? That's my problem.

Marvin Minsky: That set of six stages, why not put them in the opposite order? When you get a problem, surely you approach it with a high level language and a formalism, and you try it for a while, and if it doesn't work, then you change it and back off. It seems to me ontogeny doesn't recapitulate phylogeny every time you solve a problem or else you would be as dumb as a baby. It seems to me [that] the important things about learning to solve a new problem has to do with [what] adults [do], or what Piaget calls the formal stage. It requires that there be a great deal of organized knowledge about how to select the right stereotype for a situation and then how to change it. I think it is naive to think that one observes a problem in terms of naive observations and then does this great systems programming job, and invents Fortran and Algol and then LISP, and so forth. One works the opposite way. One starts with English, perhaps, and when that doesn't work, one uses a great deal of knowledge that smart people have and dumb people don't about how to modify it to make another system that will squeak through.

Harry Beilen: I start with the assumption that I don't really know anything about learning. I cannot very well say what learning is. I'm not even sure that there is such a discrete entity called learning. I guess the difficulty that I have is that I'm a developmental psychologist who deals with change over short and long periods of time. The characterization of [behavioral] change cannot be easily identified as either due to learning or due to some other kind of entity that you could give some name to. The problem, of course, has become pretty acute in recent years in regard to explaining the acquisition of language. In fact, talking about acquisition, and not language learning per se gives some insight into the fact that there is difficulty over assigning the source of the change. It may simply be that revealing the different levels of explanation is a function of the facts that reveal the change over short periods of time as against dealing with it over long periods of time. And, it is not altogether clear exactly what the relationship is between one kind of change and the other.

Kurt Fischer: I think I'm not really going to say what learning is because I agree with what Dr. Beilen just said. When we say what learning is, we are talking about a basic process, assuming that it is the same across many situations, across many different organisms, and across many different ages. It seems to me that we have to recognize, first of all, that there are probably a number of different things that are going on in different kinds of organisms. Somehow, we are going to have to take that into account if we are going to come up with a general model for what we put under the umbrella of learning. I also agree with Dr. Dienes that a lot of the things that we study deal with memory and transfer sorts of phenomena, and are really very different from that which has [the most] significant effect on us when we learn. [I refer here to] the sort of things that can happen very quickly and not the sort of things that happen, for instance, in cognitive development. [One more point.] When we are dealing with these broad problems, there are some things which we tend to leave out most of the time, myself included: [for example, like having to] recognize that there is a problem before you can solve it. Most of us don't even know how to think about that.

Gordon Pask: Well, of course, I think I know what learning means, and I have a fancy formalism in which I choose to express it. I'd like to agree with Dr. Merrill about learner control, saying only that I prefer to regard it as learning to learn, and I'd like to give some interpretation to this. Marvin, you were talking about general skills applicable in a number of situations. Undoubtedly, these exist and they may be exhibited in cases where problems are defined. Perhaps we're more concerned with cases where problems are not specifically defined. I have, incidentally, a slight disagreement with you, Dr. Dienes, in the distinction between sorts of learning. Essentially, I believe that there is an enormous span of stuff covered by essentially the same kind of processes

If we think about learning to learn, it has another aspect distinct from the bundle of specific heuristics, or specific algorithms, or whatever, which have been brought to bear on a wide variety of different interpretations of a problem. That is, there is an aspect, which we can detect and reliably exhibit by means of learning pathology, which consists in learning by analogy, learning by looking at morphisms and their interpretation or, in other words, by using analogical argument to get from one notion to another. The pathology exhibited is well known; I call it globetrotting. It consists in a topological word chain, that a beehive is like a city, is like an anthill, is like a ... Well, [how do we find out whether the learner knows] what is a beehive. [We can] require explanations. The system I chatted about is obsessive in this respect, requiring an explanation at every node, which is clearly unnecessary if it happens that the topics are isomorphically related.

The other type of learning which becomes manifest, and the pathology to which it is subject, is a thing I would like to call operation learning. It consists in making a series of models and then renaming, for example, learning various concepts in physics and never recognizing that they are analogous, as a matter of fact, isomorphic, even though symbolic identity with respect to these equations is recognized. The ability to perform these operations without the pathology of, on the one hand, wasteful relearning and, on the other hand, globetrotting, appears to be fundamental to efficient learning. The ability to apply them perhaps has something to do with the clever tricks which are applicable in various domains. What is also essential to learning to learn, I submit, is an ability to build descriptions. There is some evidence that people who see how they learn do transfer an ability to build descriptions to hitherto otherwise unstructured subject matter once they have experienced the feedback which shows them how they learn themselves, whether or not they have matched these conditions themselves.

Tom Landauer: I pass this round.

Scandura: This question of how learning takes place is probably the most difficult in all structural theories. In this regard, I'd like to propose the hypothesis that "The more precise the theory, the more difficult the question is to answer."

Hienes (re: Minsky): I'd like to reply to this idea of starting at the last stage. Of course, you're not really starting at the last stage because what you're playing with is systems, and you have that weaponry at your disposal in the same way that in a previous stage, you had some little objects and things on which you were acting merely perceptually at your disposal. So, depending on how much you know and how far you've developed...

Minsky: Who is the you. It's as though there is a you operating these learning things. I don't quite see where that you is and what it does.

Hienes: Well, there is a person who is learning.

Minsky: No, there is not. There is a learning machine; there is no other person who is running it.

Hienes: There is a person who is learning. Let me define learning as being a personal adaptation to the environment at the end of which you are able to cope with something more effectively from the point of view of your own goals than you were before. According to that definition, one person is learning vis a vis the environment. You may use the inputs that are perceptual inputs and you may use tools that are handy and you still start at the first stage when you are fooling around with things. You are trying to put them together as an initial adaptation stage even if the objects that you use are mental objects. Even if they are sophisticated systems, like Einstein's theory of relativity, you are still at the first stage.

Minsky: You wouldn't hold that to be true for all kinds of encounters with the world, would you?

Hienes: Well, I wouldn't like to regard it as [gospel] for everything. I would imagine that the first thing you do with anything that you don't know something about,

and you feel that you need to learn, is to somehow play around and fool around with it. After all, when you get a puzzle in your hand, what do you do? You don't measure the curvature; you fiddle.

Pask: I think this is true.

Minsky: I'd like to make the same point I think two people mentioned, that one of the problems in problem solving is recognizing the problem. I take the same position on that. Namely, when I come into a situation, the first thing I do is to set up a stereotype which says, "Well, this is one of those, and the problem is this." Now, I may be wrong; I may have formulated a bad problem, but I think one starts out by assuming a problem, and if the problem turns out to be uninteresting, or too interesting, or too hard, you give it up. But, I think, an adult at least, enters situations with a pretty standard format of saying, "If it's going to be Chinese puzzles, I assume it's going to be that Chinese puzzle. I know about five of those Chinese ones; I assume it's the one with three sticks." And, after a while, I find it isn't. It isn't that I collect evidence for a long time. I start out by assuming that this is puzzle type 37, period. And the problem is to find the pairs of corresponding sticks. Then, when I can't, I start falling apart. In other words, I start as an integrated ego knowing just what I'm doing, and then I crumble. It's not that I started as a pile of sand and make myself out of nothing.

Question: What happens when you crumble?

Minsky: That's the interesting part. If I'm smart, I call one of 50,000 procedures for dealing with such and such a situation by characterizing one that doesn't work. If none of them works, then I really crumble.

Dienes: So, you teach your child how to swim by giving him a textbook on hydrodynamics.

Minsky: I teach a graduate student how to build a ship by giving him a textbook on hydrodynamics. It's not absurd to give humans textbooks on hydrodynamics. That's why they're written. But, not for swimming.

Scandura: Do any of the panelists have a position on what aspects of knowledge are teachable and what aspects of an organism are innate, or at least common, say, to ten-year-olds and above?

Landauer: Why is that an interesting question?

Scandura: Basically, I view learning as resulting from interactions among [the environment and] rules that are active, or immediately available, in the processor. [Showed slide.] About seven of these things [rules, goals, internal stimuli], or eight or six, depending on whether you believe Voorhies' data, may be active in memory at one time. These things constitute, I think, the source of all cognition, all learning, retrieval, etc. All of them take place on the basis of interactions among the [available] discrete entities according to some fixed mechanism which is assumed to be built into the processor. What I'm saying, in effect, then is that I believe the [individual rules] to be teachable. [On the other hand, for example,] the interaction mechanism is not teachable. My question is what kinds of things fall into each category: what is teachable and what is not?

Landauer: I've never been able to get it straight why some things are easier to accomplish by experience than by evolution or vice versa. And if it is teachable then it could have been learned by the experience of the species over a thousand years just as easily as by the individual over ten or vice versa.

Scandura: Suppose you're talking about something like number conservation, which takes place over a long period of time. Why is it you can't teach it in a way that would make Piaget happy? He will concede [that you can] on narrowly defined tasks [but not when number conservation is broadly defined]. The reason is fairly simple, I think. In typical American studies, the tasks as they have been taught are narrow and very highly prescribed. It is easy to devise simple procedures [rules] for solving such tasks and for teaching how to solve them. It is much more difficult, if not impossible, to devise

procedures which parallel what Piaget would have a conserver do. [Tape changed.]

[Another point: If a person doesn't know a procedure for solving a given task, then he must learn one. If he must do this on his own, then the ease with which he can learn such a procedure will depend on what he already knows.] For example, a problem in automata theory that Marvin might find quite easy is likely for me to be very difficult. The reason is that we enter [the problem] with a quite different degrees of familiarity with the subject matter.

Minsky: Even for conservation, the fact that you don't have a complete theory of education means you have to be [inaudible] because it takes 3 or 4 years for conservation to appear in the aspects that Piaget recognizes. You can't be sure it will appear in three months if you [simply] tell the child. The evidence is that people don't learn by instruction. There is an annoying chapter in one of Bruner's books on verbal instruction. All it shows is that bad advice makes people worse at something.

Dienes: Why is it in New Guinea that the children get on to the conservation stage, and even the concrete operational stage, at least two or three years later than they do in other cultures? What's more, those that do go to school move backwards about two years on the developmental scale, and take about two or three [years] to catch up to where they were when they [entered] school. In other words, trying to get them to verbalize puts them back.

Minsky: Well, presumably, people are telling them bad things. What else would you expect? Surely, the school is doing them more harm than good by our standards. I don't see any conflict.

Dienes: But they still don't "make it" in the same way that their brothers and sisters do on the neighboring islands. Why don't they?

Minsky: Because their education may damage them.

Dienes: They don't go to school.

Minsky: Well, they're being educated. I didn't say [in] school. Maybe they have a superstition against that. My image of so called primitive people is extremely primitive, and I imagine that there are these doctrines around, about what causes things to happen, that are very bad for the normal Genevan-Piagetian sequence. One might want to send developmental psychologists there, instead of anthropologists -- or, is there a good theory on that [already]?

Beilen: There is Kohlberg's theory of the dream concept which discusses this problem of regression.

Landauer: There is a good deal of evidence that what you say is true. If you ask your questions the right way, or if you teach the child the right thing, you can get children to show a lot of the Piagetian things years earlier than you would if you used the prescribed Piagetian questions.

Fischer: But there is a real question about all the data that I know. Giving a "yes" or "no" answer is easy; all you have to do is reinforce a child for correctly reciting a sentence, but that's not the same as understanding the task, and the controversy and misunderstanding revolves around that issue.

Neilogg Wilson: My opinion is that, outside of the emotional system, all learning is propositional. Dr. Landauer seems to be the only one to disagree that learning is propositional.

Landauer: It is propositional. The propositions are little.

Melix Kopstein: Traditionally, in psychology [the question of] "what does learning amount to" has been answered by "it's a change in a performance capability." What causes that? It isn't due to growth, due to drugs, due to maturation, etc. What is it that is learned? It's new performance capability. The next question I have is "what is that?" You can solve problems you couldn't solve before. What do you mean by that?

Minsky: When you ask a question about propositional learning, I don't think that that's an adequate idea. Surely, people learn procedures for using propositional information, and the acquisition of this capability, if you look at a person as a kind of computer, involves acquiring both fact and procedures. In my view, a lot of the important things that make smart people smart are learning how to debug procedures. There is a special kind of knowledge about how to take a procedure that is not quite working and fix it, which is really vital.

Wilson: Are not these procedures, propositions?

Minsky: Mathematically, there is a sense in which a machine can be described as a set of propositions. But, not in the sense of mathematical logic and, in fact, one thing logic systems don't have is a rule saying what to do next, and that makes a tremendous difference. If a lot of our knowledge is about which proposition to use next, and if you leave that in propositional form, you lose, because you have to try everything to find which [proposition] to use next, and then it's too late. You'd have to get the proposition out to say which to use next before you use the one you really should use next. It has got to be in a procedure, so that it will work without being asked.

Kopstein: Accepting the notion that these are propositional, I think you're also saying that you then have something that seems like a hierarchy. You have propositions about propositions; you have further propositions about the second level propositions and so on. The next question is: Where is the apex of this damn thing? What determines what the apex is? One is tempted to think that ultimately this should lead to a proposition that is universal, which obviously doesn't make sense, and obviously doesn't exist. Therefore, what do you mean by that?

Minsky: In Landauer's theory, there is a sort of universal thing which, whatever state you're in, looks for something in memory which tells what to do if it is what matches best. Most of the modern theories, like the complicated Newell-Simon production theory, end up with one basic operation, which says if there is something in long term memory which is pretty similar to what is in short term memory now, pull out what is associated with it, and it will tell you what to do. There is no real circle there. If you get bad data in you, you will go around in a circle.

Banerji: That's the whole point. If we talk of learning as learning to perform better, what Marv says makes a lot of sense in that you start at the highest level of knowledge that you've got. The only reason that works, and you don't have to start building from the bottom again, is because it is true that some very abstract thing which has worked in the past will work again. It so happens that at a certain level of abstraction, this becomes true. The point is that the very fact that we're alive seems to indicate that something like that [universal thing] is out there which is independent of what is learned and is unchanged, and is effective, and is used. Once you have simplified the planetary motions through the idea of the "inverse-gravity" law, you stick to the "inverse-gravity" law, and you don't start going back to epicycles again.

Scandura: I'd like to propose a resolution of the problem of how much structure there is in memory. Does the answer depend on whether the theory is to account for individual behavior or for average behavior? If you are dealing with average performance, I suspect that a truly random model, such as Tom has proposed, might work. If, on the other hand, you are concerned with the behavior of a particular individual in a particular situation, I think the ballgame is quite different, whether it be memory or whatever.

Landauer: I don't see why.

Banerji: I don't either.

Dienes: Well, the chances of Newton discovering all the theories he discovered, or Einstein discovering the theory of relativity, by randomly going around in his brain are practically nil.

Pask: You are assuming we were talking about learning; you're talking about storage.

Minsky: Tom, why don't you take the next step which is, instead of this thing being random, there is this area over here which is good at chess, and this area here which is good at this and that? You keep that model, but you [include] context by being able to steer the memory pointer into whatever.

Landauer: I think you can allow it to be steered a little bit, but it's dangerous because [with] much steering, it loses its properties. I don't know how much yet, but if it doesn't act like it's pretty nearly random, then, for example, it doesn't generate realistic learning curves.

Minsky: Do experts have realistic learning curves?

Landauer: For simple learning they have realistic learning curves.

Pask [re: Minsky]: It seems to me that this wouldn't give realistic learning curves for rote learning, and you would get very realistic curves for actual individual learning.

Banerji: Marv said if you have all the chess facts in one place and you want to play chess, you are supposed to push your pointer that way. But if you never put the pointer anywhere for chess [during storage], the chess facts will be all over the place anyway. So what's the difference?

Minsky: With a good chess player, you make a move, and he says "oh, so and so did that in 1956," and there isn't any learning curve at all. If it stays in the same region that it uses for playing chess, then you only need one trial.

Banerji: No, no. That guy has said that to himself millions of times.

Landauer: And, of course, he has rehearsed it.

Banerji: It's all over his brain.

Question: For one shot learning, you might build in some sort of replicating device which scatters the memories through space.

Landauer: The only way you can do that is by rehearsing them. If you try to remember, for example, somebody's name that you met at a party by saying it over and over to yourself right after you meet him, you are wasting your time. Once is as good as doing that. The only way you can scatter it over your memory is to say it to yourself once, then wait until you've just about forgotten it, then say it again. That really works.

Pask: The right way to do it in fact is to build multiple paths of reconstruction. You remember the name in slightly different contexts.

Landauer: That's another way of doing it.

Pask: It's a rather effective procedure for teaching people to recall, and they do it, as a matter of fact, without a good deal of interference. If you give them specific interference, it does get rubbed out.

Question: You store the meaning and the context?

Pask: That would be one way of doing it.

Question: Individual differences may be modeled by different radii of Dr. Landauer's circle. An individual who can recall a fact all the time may have a very large radius. The register length may also collect a lot of other associated things. So, I think there is a lot of flexibility in that theory which still can be expounded

Scandura: I don't know of any data that pertains to the behavior of individuals to which this theory has addressed itself.

Landauer: I gave you several examples of individual behavior. I'll give you some qualitative examples. For example, you remember things at one moment and not at another; that's a random process that I think is true at least of me.

Scandura: You think it's random!

Landauer: At least it is fluctuating. That's something that comes from the theory that's simply a counterexample to your assertion that there are no individual facts.

Scandura: But, is that fluctuation a function of the way memory works or [of the] random nature of the environment [particular reference to standard memory experiments]?

Dienes: Could you explain how extinction takes place? Because, according to your theory, the only way extinction takes place is by superimposition of something else.

Landauer: That is what the data in [the] experimental psychology of learning suggest. Extinction is, in fact, the learning of contrary things rather than the erasure of old things.

Dienes: If you didn't learn anything at all, if you kept yourself in a glass cage, you mean you wouldn't forget anything at all?

Landauer: If you teach a pigeon to peck and put it away in its cage for -- the longest experiment I know of is six years, and put it back in the place where it learned to peck, it starts pecking immediately.

Wilson: I see a very wide gap between Landauer and Minsky. And I find myself disagreeing with both of them for entirely different reasons. The only way I think that Landauer's model would work ...[inaudible]..., for example, to induce such a memory to produce a paraphrase of a sentence, [would be] to have an executive that had an awful lot of knowledge stored in it that in some way affects kernels of memory. But in Minsky's position, it would appear to me that our memories are procedures or programs. This almost seems to me to be overly structured in the sense that I believe there should be some capacity within the intelligent system to analyze its programs -- as a program that is more or less external to the program. And, if we do have procedures, I think the procedures have to have a lot of semantic characterization, so that they can be appropriately assembled, decomposed, or analyzed, rather than just run through.

[Editor's note: The above distinction between the contents of memory and control mechanisms which operate on memory came up on numerous occasions in the discussion. A specific executive (control mechanism) is proposed and applied to a wide variety of phenomena in Scandura's paper (this volume). The distinction also plays an important role in Pask's theory, and, as Minsky and Banerji mentioned, is implicit in Newell's production systems as well.]

Merrill: I have a point of clarification, Tom. I didn't perceive you to be saying, as some of the comments [indicate that you had], that what is stored is an association. You're not saying that. Can't you store an instruction or even a program?

Landauer: I think you can store a lot of different things.

Question: In what space? In a linked bin?

Landauer: In the brain.

Minsky: Computers have a fetch memory. A lot of computers have a couple of memories. I'm on the lookout for the kind of thing that a chemist looks for from physics. Here is a very beautiful, clean model that exhibits a lot of interesting phenomena, I don't think Tom himself would want all of those assumptions to be preserved intact, because if this sphere is part of the brain, it's going to be denser here or there, and maybe there are things to steer it. But, if you have a computer, it is possible to have a single associative memory and run big programs. It raises the question of where the developmental structures are located that cause the right things to be learned. I'd rather, I think, have two memories, one of which has painfully developed interpreters and the other of which has a lot of data.

Landauer: I would prefer that there be one large associative memory. I may be forced to believe that there is an executive, but I don't know that anything has to be stored in the executive. It all may be stored in the big memory and its developmental processes how it gets the stuff into it and learns how to do things, I would like to believe are

historical. [I would like to believe] that it is the individual who has had a lucky history who turns out to be smart, as well as maybe one who has got a lot of storage registers.

Scandura: If you allow big chunks in [memory, then] I think your theory actually corresponds to] one part of mine. In particular, if that expanding "red bomb" can vary [in size] according to the individual, say, that in one individual [it] might expand further. For example...

Landauer: [I think it's] really little when you're drunk.

Scandura: If you allow the individual elements [in memory] to be anything you want, that is, big chunks [rules, networks, etc.] as well as little chunks, and, if you ignore the time lapse in your search [from Landauer's arrowhead], then your moving circle corresponds [in my theory] to that portion of memory that is active at a given point in time. Our respective control structures differ greatly, of course. In my theory [the control structure operates so that] the active elements make it possible to recreate elements outside the active sphere; in yours, it essentially involves expansion from your pointer.

I know you want expansion [time lapse in your theory] because you want to account for latencies. But, the problem with that, I think, is that the latencies involved [in many memory experiments] may in part be an artifact of averaging over tasks and/or individuals. Tom, of course, is thoroughly familiar with this [literature] and may want to disagree.

Landauer: One of the experiments I showed you had six subjects. What do you mean by averaging? I average over several different words in a given frequency for one individual. You have to do that because of the statistical properties of the system. Each reaction time is going to be different from almost every other one. But, if I'm answering your question correctly, the individual subjects show those reaction times [averaged over different words].

Scandura: It's not just a question of "Did you average [over items] within an individual or over a group?" But, I'd rather not say more at this time.

Ask: I think there is a difference between how much structure there is in learning and how much structure there is in memory. What you were talking about is how much structure there is in memory. And there is a whole lot of apparatus for inputting, accessing, putting in little pointers, and so on. And to some extent, I have tried to reconcile it with a diagram I have drawn on the board.

Here we have what I guess Dr. Scandura and I would call a representation of concepts, the sort of thing I was talking about as an entailment structure. You can see, of course, that these are external entities; they are entities possibly revealed to a student and possibly used by the student in some kind of map to get around. But, presumably, they also refer internally to some region which can be demarcated in the brain, in storage.

In this situation there are people who behave as follows: One sort of person behaves by defining an area in that map, and then operates very much like your model operates, with rehearsal in a task structure rather than recall tasks. A program in this case is a rehearsal program. [Such people] choose one area, and [in this case] your model, I think, would apply perfectly. Other people behave rather like this: They choose an area, then they choose other areas, and then they interlink, presumably by links which are stored, from one area to another. In other words, they're giving substance internally to this external structure by that method.

It seems to me these two types are manifested in learning. None of this contradicts the model you have put forward, but I have taken the liberty of shifting the domain of interpretation from that of simple storage to that of learning. I have tried to discriminate two types of learning, and I think there's no real conflict as far as I can see. If you address tasks where it is possible to exhibit structured learning [i.e., structured attention against word association], I submit that you would find the same result, same model, but with that amendment.

Drill [to Landauer]: How do you face the problem if you extend it in an indefinite way; you can think of this space as filled with a whole bunch of little spaces -- you get

a big ball filled with a bunch of marbles. Now, what kind of a pointer do I have to get from marble to marble?

Landauer: The connection between different topics in physics or different physical facts would presumably not be that they are all in the same part of the brain. They are scattered all over, but they are connected by shared links so that you can find B by using A as the tag and then using B as the tag, find C, and so forth. This will allow you to walk around in your memory, say, of Paris, but it does have the nice feature that every once in a while, by sort of a random process, you'll pop out of it. You'll find when you are thinking about Paris that every once in a while, you'll have a thought about New York or your dog. I assert that that's realistic -- that, in fact, when one sits and cogitates about some field, he doesn't stay neatly packed into that area. At least if he's me, he doesn't.

Minsky: However, we in artificial intelligence haven't had to lean over backwards to get the system to make mistakes. This nice feature of your theory is also a feature of all suitably complicated theories.

Question: Isn't there a fundamental difficulty in being able to delineate what is truly memory from all other kinds of psychological tasks? There isn't a domain of psychological processes that doesn't have implicit in it some memory processing. In the tasks that you decide to use, you try to strip away as much as possible the other domains, or the other complicating facts, to expose the memory feature. But, in effect, you really cannot; in every kind of memory experiment, you are really involved with some other kind of phenomenon which may be rather difficult to identify.

Landauer: Agreed. The theoretical strategy that I have been following, or would like to be able to follow were I intellectually strong enough to do so, is to assume that God had made one good invention, and had replicated it a lot, and by so doing, produced the miracles. If a few hundred thousand years ago, human evolution had invented a really keen little storage device that could store a big hunk of something in a tiny place and had a great communication system like an executive, where the executive can just say "I want information about dogs" and, if this little device has anything about dogs stored in it, it will return it without [having to] connect wires. Had that been invented, it might exhibit all kinds of wonderful properties. I want to see how much of the complexities and wonderment of our performance can be accounted for by such kinds of simple notions.

Scandura: ...[Lost in changing tape]... and so it basically becomes a question of representation.

Minsky: I should think both are true. Namely, it might take a thousand genes to make an apparently useful memory, but then there's this other million genes to debug it, and make various things to prevent loops, and adjust the way it matches, and stuff like that. I suspect that a model like that might work out very well, and maybe one only needs 20 or 30 such theories of memory, but then one needs 1,000 other structures to keep them from clobbering each other.

John Fox: There are only so many genes.

Minsky: There are so many nuclei in the brain, and lots of them seem to do things.

Kopstein: Why do they have to be genes. You see, I think good old Mother Nature is usually much cleverer and more efficient than to use up that many. Why use all of these just for purposes of memory. Why not use them for rules which in turn build up, in the maturing individual, a growing capability to construct the software which handles some of that.

Landauer: How many new genes are there for humans relative to our less capable ancestors?

Minsky: Humans don't have an unusually large number of genes for an animal.

Landauer: How many genes have changed [through evolution] that are responsible for the structure of memory?

Minsky: You can't tell because there are too many redundant genes.

Dienes [to Landauer]: One thing that's been bothering me. If I have understood your theory (or if I haven't, tell me where I've gone wrong), you have these "things" knocking around, say you have about five of these, and you have a certain probability of being able to recall these. Now, if these things are linked, then these links presumably are stored somewhere else.

Landauer: No. Look. You store A-B here and B-C over here. You look for A and it returns B; now look for B and this one returns C. It's a sequential process. The links are not represented in the structure of memory.

Dienes: So the link itself is not in any particular place?

Landauer: No.

Dienes: Well, in that case, it doesn't just have things knocking around that this red balloon is looking for. You've got some mythical things; A triggers off B.

Landauer: Imagine that you did this with little cards -- you throw them into a wastebasket. Maybe that'll make it easier.

Dienes: What worries me is this: According to [your theory], linked things would be more difficult to recall and remember than unlinked things. Things with lots of links are in fact easier to remember. If you have a lot of nonsense syllables to learn, it's obviously more difficult to remember than a poem. So, according to your theory, it seems to me that you would get the opposite prediction.

Landauer: If what you are asking subjects to do is to get them all back, then they have to have some way of following a path, which means that the more links you give them, the easier it is going to be to make something that will allow them to go from here to here and to there. If you make that hard, then once they have gotten to [an item] they come back with something that won't allow them to get to the next one. That's not a problem.

Question: There seems to be an equivalence here between talking about the structure of memory, with the structures represented explicitly in the design of the memory, and, on the other hand, talking about an unstructured memory, where the structure is induced by the items that are stored in the memory. They seem to me to be equivalent notions in a strong sense.

Banerji: If they are equivalent, they have to be worked out with mathematical care.

Landauer: I think there may be equivalences, but the demands that it makes on the person who is building the system are very different. I have a file system that works according to my theory. Man, is it easy to operate. Every letter I get I throw into the drawer. It doesn't take a secretary. It doesn't take anything. I don't even have any tab cards. And it works. I can prove to you that, under some assumptions of distributions of frequencies of letters and reprint requests and so on, my system works better than yours.

Question: But you put the whole burden of discovering relations in memory upon a processor, and essentially you are going to do the same thing in many instances. You are going to be tracing the same paths that might in somebody else's memory be represented equivalently. You haven't escaped the problem; you have just ignored it temporarily.

Landauer: That's not correct. Let me give you one argument. If you have to represent all the relations that you are going to need in knowledge by storing them in a way that represents those relations, then you have to do it for everything conceivable that you may need in the future. If you just put raw data in and store it, and compute the relations later, then you only have to compute those relations that you in fact ever need. That saves you a hell of a lot of effort.

Scandura: Tom, how do you explain the kind of thing Don Voorhies was talking about. Suppose you use a rule, a procedure, say, for adding numbers. First, two digits come in, and you add -- and continue as new digits come into the processor. At first, you're

dealing with just a simple problem but [if] you keep adding numbers you [eventually] fail. Now, people might differ in how [many digits] they can add, but sooner or later everybody fails. If you have a fixed capacity processor, then after too much gets into it, it bumps something out, and that thing is in some sense lost. But it's not very far away because it was just there. For example, in adding, I may retain the given digit along with the partial sums. After a while, they're going to overload my capacity. Maybe I can add two more digits, but when I get to, say, the next pair, I'll [have to] dump something out. The thing that is bumped out seems not to be retrievable, and yet according to your model, you ought to be still very close [to it]. This is, of course a quite different kind of experiment.

Landauer: Could I pass on that?

Question: Isn't that what is happening after a long sequence of numbers; you forget?

Scandura: Exactly. But what I'm saying is, how would Tom [account for] this? It's not a question of random movement. As a matter of fact, there are certain things that [seem] to be eliminated first, simply because of the nature of the particular rule being used by the person processing the information. [In our theory] there is a very explicit structure that determines what is forgotten and what is retained at each point in time [e.g. after having been added, given digits are erased before partial sums]. A structure-less kind of memory wouldn't seem to account for this.

It's getting late. Why don't we retire to the lounge for some wine and cheese.

Postsession: Theoretical Issues in Structural Learning -- Monday, April 22, 1974

[At the beginning of the session, the chairman polled the participants to determine topics of most interest.]

Landauer: The first issue to be discussed is: "How are rules induced from examples?" The second question will be: "How do know rules interact with the environment to produce useful knowledge?"

Wilson: I'd like to give an example of what started me thinking first about this problem. It was a paper given at this conference three years ago which, narrowly considered, was a quite adequate paper, but it left such enormous questions begging that I felt very unsatisfied. The theme of the paper was an extension of Newell and Simon's GPS [General Problem Solver] on the learning of syntactic rules. The author of the paper had a child applying GPS to strings of nouns, adjectives, adverbs, plurals, etc. This is all fine and good, but how does the child learn all these grammatical categories prior to the acquisition of the rules themselves? How will the child know what a word is? How will the child even know what a phoneme is? What are the elementary units that are involved in building up the rules? I don't have any terribly smashing insights to impart to you, [but] one slight insight is that, in looking at syntactic rules, Chomsky came to the conclusion, I think quite correctly, that there are similar types of rules operating down to the phonological level.

Minsky: Why do you think he is right?

Wilson: What I'm asserting, on the basis of terribly flimsy intuition, is that we have to learn not only the units, but the rules involving the generation of strings of units, and that [the units and rules] are somehow learned at the same time.

Minsky: To learn a unit and the rule at the same time seems bizarre. It's too difficult.

Wilson: I don't think a unit would really mean anything aside from the rule in which it enters. To say you learn the units before you learn the rules means that you have to have some kind of pre-knowledge about what the rules ought to be. That is to say, if I want to learn rules; then I [must] first learn some units appropriate to the rules. That strikes me as bizarre.

Minsky: In evolution, you develop a structure for a staple purpose and then get more complicated things to do with it.

Dienes: Could you tell us what you mean by having learned a rule? You could mean being able to operate a rule correctly, or you could mean being able to explain one correctly.

Wilson: [I mean operating one correctly. But,] to learn a syntactic rule, you need to have some equipment for identifying what the rules might be, and what they might apply to. Essentially, what I'm asking is "what is the minimal equipment?" I'd like to mention briefly an experiment: Sequences of computer generated sounds were presented to subjects. They were supposed to be words in a language that was not familiar. Subjects were asked to identify where the word boundaries were. After a relatively small number of presentations, subjects quite often were able to segment sequences that co-occurred together. This implies there is some kind of inductive capability. Interestingly, these were sound sequences that no one could pronounce. We have had very little experimentation on this kind of induction.

Minsky: The question of how complicated it is to learn rules of grammar ought to be raised before one regards it as a monstrous problem. In the current Newell-Simon theory (which I think may be too simple, but which certainly makes programs that do things), all rules are in the form of productions, and long term memory is organized in such a way that if the first part of a production is in short term memory, then it goes and matches [what is] in long term memory, and out comes the second part. So, if you believe Newell and Simon, and I think it's an important first step toward a theory, there isn't any serious problem about how rules of syntax are learned. You don't need much in the way of

mechanisms; you need induction. The procedure for long term memory in general is the same as the one for building grammatical rules.

The second point is that Chomsky's transformational rules are rather complicated. I would say that none of the people in artificial intelligence right now believe that [his] was a good formulation of how grammar is understood or produced. So, one shouldn't worry seriously about Chomsky's opinions about how hard it is to learn that stuff; [it] doesn't appear to be a practical theory. The practical theory right now in computational linguistics is something like case grammar, which says that when a word is mentioned, a number of questions are pulled out of memory associated with that word. Induction involves learning what questions to ask when you hear a word. Then those questions are usually answered as you hear the rest of the sentence or utter it. For example, if I say "He kicks," then [you] look for a kickable object coming later in the sentence. So modern theories of grammar use simpler ingredients and are more powerful than most of the Chomsky things. Modern experiments on the learning of language in children show that there's a long period in which they don't have grammar; they have two-word sentences. The words in the sentences call on the semantic case, but not grammar. Then there is a period of a year or two in which [children] learn words, learn their cases, and then learn the rules of grammar. The important thing is that the kid has a chance to learn the semantic quality of case grammar, what the words mean, what questions are associated with what, [and] how they are answered before he has to learn any syntactic rules at all.

Question: I was wondering, Dr. Minsky, if induction is a program.

Minsky: There isn't any one such [program]. It's easy to make a machine that will learn all kinds of smart things. One of the questions we're not discussing today is which kinds of things should people learn. That's a serious question.

Question: Let me ask about rule induction, then.

Minsky: There are lots of kinds of rule induction. As soon as you formulate the question, you can see that there are trivial ways to do it. And if you think of a hard enough way, then nobody's done it yet. There isn't any question of principle about making machines do learning of any kind; prototypes of all sorts have been done.

Wilson: I wasn't making a strong case for Chomsky's transformational grammar. I was just saying that the idea that grammatical rules can extend down to the phonological level is, I think, an interesting one. I think the real question would be, "can you get learning? I would agree with you about the importance of case grammar. [But,] in any known system now, could you get identification of cases, the identification of words even, from strings of arbitrary noises, which may be lawfully generated, and events in the environment that are in some way associated with them?"

Minsky: I'd be terribly neutral about that. There is an argument in the Artificial Intelligence Laboratory because there's a lot of money in doing research on speech recognition. Quite a few of the laboratories are doing that, but we're not. My reason for not encouraging our people to do [speech recognition] is that when you understand how phonology is handled in the human ear, [I think that] it's the least valuable knowledge you could get, because it's probably a special purpose machine. Laryngeal control and the development of special purpose gadgets in the ear probably evolved together. I'm not saying we know it's a special purpose machine, but it seems like a very dangerous thing to try to generalize from. I would much rather work on something like vision. It's very plausible that in an input device like the ear, there are special purpose things for firing the atoms of phonemes in them. People can learn to recognize simple combinations of [phonemes], but I don't see that that has any connection with language.

Scandura: Your observation is probably one of the reasons that psychologists are so interested in speech recognition. In general, they seem to be less interested in understanding [procedural] competence than in discovering behavioral correlates of man's physiological nature.

I would like to turn the discussion toward the earlier comments on the Newell-Simon system. Those of you who are familiar with their recent book will recall that at the

end of almost 1,000 pages, they have just a few pages devoted to the topic of learning. Their basic conclusion is that they are talking about performance, and that the theory as it stands has little to say about learning. Any comments [on this]?

Fox: The basic idea of a production system is that if a thing in the world is true, or an input to a machine [satisfies a condition], then a particular action will be carried out. That's very much an S-R theory, but it has one rather nice advantage, and that is that every production in the production system is completely independent. If a condition becomes true, this action will occur. It will change the machine, and whatever other productions are playing at the time won't be affected by it. The most worrying thing about all production systems, so far, is that in order to implement this workably, the actual condition side, a statement of what causes a production to be applied, is never modeled. If you have any complexity, any depth, it's a very, very complicated thing. The control structure is very deep in the guts of the language or the machine. An enormous amount of knowledge and an enormous amount of structure is in there, but never appears in the discussion of the program. What is enormously important is where that control structure for any kind of developing animal comes from, how it gets from being a naive animal up to being able to evaluate complex conditions. Whether or not production systems are a good model for human behavior -- for how the animal learns to do things. [or] whether they are given or universal, or whether they are things which are built up from elementary operations is, it seems to me, the fundamental learning problem. It, at least, has the same status as the problem of how, once you've got these elementary functions and operations which work together, you then learn more of them.

Scandura: The problem as you have stated it is precisely the problem; that is, how is new knowledge acquired as a result of the basic units that happen to be available at some given point in time, whether [this be the time of] birth [or later on]. My own feeling is that [learning may be governed by] a surprisingly simple mechanism. [Such a mechanism] has for at least four years now provided us with a basis for dealing with a wide variety of different [kinds of behavior]. Indeed, [it is possible that this mechanism] may be innate, and that it governs the way in which available elements [rules] interact to produce both behavior and new knowledge, including making selections when two or more elements might be used, breaking problems down into subgoals, and so on. Very simply put, though there may be some slight enrichments that might be necessary, the mechanism operates as follows: When you present a subject with a problem, he tests available rules, in particular, the rules that are available in short term memory [to see if they apply]. If none of them apply, control automatically goes to a higher level goal of generating a rule which will apply, again, on the basis of the rules that are [in short term memory]. Very specific tests are performed. [The mechanism] can be mechanized and has been in a limited way. If one succeeds in deriving such a rule [which applies], control goes back down. Newly derived rules become part of the knowledge base, and [may be] used to solve the [original] problem. If [the subject does not succeed], control continues to go up [until success is achieved or memory is overloaded]. As I said, there are minor variations on this theme that allow us to deal with [such things as] rule selection and subgoal formation.

Dienes: Can I bring things down to brass tacks again. I'd like to propose a task, a rule learning situation, and I'd like to ask you a few questions about it. You have a red card and a yellow card in the window of a machine. The subject plays either a red card or a yellow card. If he plays a yellow card, the window of the machine shuts and the same color card appears again. Now, if the subject plays a red card, the window shuts, and the other card appears in the window. This is all. The thing that I cannot explain is that some subjects never learn to predict; that is, before the card appears, they have to predict which card is going to appear in the window after they've played their card on the table. Some university mathematics students after 200 trials don't succeed in learning the rule. That's an interesting fact which I'd like some comments on. Now, suppose at the end of the test, they have succeeded, success being ten consecutive correct predictions. Then, you ask a subject, "What is the rule?" Some subjects say.

"If I play a red, it changes the color; if I play a yellow, the same color comes up." Other subjects say, "If I play a card the same color as the window, yellow comes up; if the card I play is different, then red comes up." Both these explanations are, of course, correct. Other subjects enumerate all four possible cases, "Red against red gives yellow; red against yellow gives red," and so on. Other subjects say they don't know the rule, even though they've given ten correct consecutive responses. So one of my problems is, "why is such a simple thing not learned by people that have I.Q.'s maybe of about 130?" The other problem is that people that give the higher order explanation, that is, the first explanation I gave, do the task with fewer trials and errors than those that give the second explanation. Those that give the third explanation make even more errors and take even longer. There are [also] some in-between explanations. In fact, recently, we simply counted the number of words in the explanations, correlated this with the number of errors they made, and the correlation was over .9 consistently. This is also something I'd like explained, because it's really a very high correlation. How did it happen? ...Also, why is one subject more effective than another?

Pask: Could I ask some questions about protocols of the people who fail to learn altogether, those who hadn't the remotest idea what the rule was? For example, was there any confusion as to whether red and yellow were the only things?

Dienes: No.

Pask: This is no more surprising than the protocols of people in guessing experiments where there isn't a rule. In other words, the subject is often reminded that there are only two cards, but he doesn't really believe it. It's rather like work on using and getting information. People are told there's a limited universe, and they will verbalize and repeat and say, "Yes, this is so." But, as a matter of fact, when they begin to formulate hypotheses about the matter, they don't really believe it. They think there's a trick in the thing.

Dienes: We have had no evidence of this. Two cards are two cards.

Minsky: Let me say the same thing Gordon did from a much lower level. Although there are only two cards, the universe of rules which a subject conceives of as possible is uncertain. And it would be nice to ask those failing subjects what they think is in the environment. How many of you know the thesis of Julian Feldman? It's in Computers and Thought, and he attacked exactly this problem. The experiment is different. In Feldman's it's a choice experiment. The subject is given a series of 0's and 1's and these are in fact random, but the subject is not told they are random. He is told to see if he can predict what the next one will be. Feldman makes a computer model of the situation, in which there are six hypotheses or molecules. One is "change." [One is "is this thing a repeating sequence."] What happens is that Feldman's model, when it sees a run of, say, 1's says, "Well, it's in a run of 1's." If it sees a certain number of them, it may make the hypothesis, "Well, he's going to change soon." It's all the common sense things; there are only five or six of them. But the idea is that Feldman's computer model looks back over the recent history -- it doesn't remember much about what happened a long time ago -- but it looks back, and as soon as one of these predictors fails, it remembers what little it had of the sequence and says, "Oh, he might be alternating now; he just did that one to fool me." The result is that in one sequence of 200 of these trials... Feldman's model correlated with what the subject was doing [20 times]. What Feldman's model does is, it keeps correlating its own predictions with the subject's until it has converged on a model of that subject.

Dienes: That doesn't answer my question.

Minsky: The point is that in this situation, Feldman throws in a certain collection of rules, of micro rules, from which the strategy is made. Now your student is starting in with some collection [of] maybe half a dozen of them. As the thing comes up and he fails he recombines them. Apparently, in this situation he never throws away this bunch of atoms and puts in a new one, probably because he is getting some sort of partial reward and thinks he's on the right track. He thinks he's got the right rule... So he tries it

some more. The only way to find out is to ask him [what he is doing]. If you had taken protocols of the same guy after each try, [you might] get him to say a little bit about what he did. It seems to me that it's worthless in the present state of the world to run experiments like this, in cognitive psychology, without finding out what the people think they're doing.

Scandura: There are a couple of comments that occurred to me while all of you were talking. One relates to something that Marv Levine told me about 12 years ago. It has to do with the following very simple observation: You have two cards, one with A on it and [the other with] B. You hold up one at a time and you ask subjects to predict the next card. The rule is, [they] alternate; first an A then B, A, B, and so on. If [however] you confuse the subject by using a random basis for selecting the first 3 to 5 cards, the subjects may go on for literally hundreds of trials, and never catch on the A and then B [rule], simply because in the beginning, he began to look for something rather esoteric. He's already ignored the simple possibilities....

Dienes: Everything was completely regular, though.

Scandura: I understand. But, if, on the other hand, the subject is unsure, and in the beginning he remembers one card [incorrectly], it would have precisely the same effect on his behavior. Insofar as he was concerned, he remembered something that was different from the rule you had in mind.

Dienes: But he has been told [that] this experiment will work like a machine, and it never makes a mistake.

Scandura: That's precisely what Levine did.

Dienes: The subject would say, "I have forgotten," not, "the machine went wrong."

Scandura: But, the point I'm making is that this is precisely why the Levine experiment worked; the subject was not expecting to be fouled up; the memory [was] faulty [but the subject may not know it]. The other point I would want to make is: If there are differences in processing capacity, [with] one person, say, being able to retain six things at one time, and another, eight at one time, is there a possibility that individuals who can retain more different things are more apt to learn, for example, the three case [rules] rather than the two rules -- simply because their processing capacity is such as to make that a more feasible learning strategy?

Pask: It is very difficult, because you choose a miniscule situation and because you, for whatever reasons, [didn't] interrogate the subject or engage [him] in discourse during the experiment to say exactly what is happening. Although the experiment is simple enough, it would be very difficult to "winkle" out the effects of fixity and the effects of style. With the people who have learned the three different rules, you are getting a manifestation of the effects of style in learning....

Banerji: I think an interpretation is possible. I haven't worked it all out. The point is this. If you look at the basic statement, "This changes when this happens," and write the student's verbalization of the rule as an "and" or "or" statement. Now, if [there is] an "or" [in it, then] my student is saying that his experimental observation can either be in this cell or in that cell. Then he needs a certain number of objects in each cell before he is sure that he is right. If the statement is done in such a language that there is only one cell, then every experiment is tested to see whether it's in the same cell or not. It gives him a much faster way of getting a significant sample. It's much easier to believe a statement which always happens and is a short statement than to believe in a statement which often happens and is a very long statement.

Minsky: It's a matter of style, of whether you conceive of putting it that way.

Banerji: I agree, but Dienes seems to indicate that irrespective of what the person's style is, this seems to happen. This may be false.

Minsky: No. He's saying there are several populations.

Dienes: What we've found is that people who have effective explanations also have a very effective strategy, in playing the same card several times to see what was happening, whereas the others played a lot of different cards and, of course, got immediate interference with their learning.

Minsky: If you put these three ideas together, then these people who do very well might be somewhat bad on statistical problems where you do have to collect a lot of evidence. There are the other people, like in a "Bruner" style experiment, who pick on one hypothesis at a time and throw it away if it fails. They might be worse at one where you do have to fill four cells.

Banerji: Marv, what a hypothesis is depends on what your basic predicates are, and that's why it's not clear.

Minsky: Yes. And on the question of whether you entertain several hypotheses at once. I don't know if any psychologists believe anyone does that anymore.

Pask: Yes, I do.

Kopstein: I haven't heard anyone ponder [the following] issue. With respect to the kind of problems Dr. Dienes has described, there are a couple of objects, the card the student has, the card that is displayed to him, and each is capable of being in one of two states. One could ask very simply what relates the states of these two objects. One would ask further, what sort of algorithm is appropriate, or what set of algorithms is appropriate to the determination of what relates these two states? I could then ask, in what sense does the student traverse these algorithms and at what point does each student, or each group of students, if there are patterns in which they behave, diverge from it, and in what sense? This may be a clue as to what's going on. For instance, you might find loop in the algorithm. You might find a student that automatically always goes around and around in the loop and, therefore, he'll never reach a sensible solution. This is really what you're trying to determine; that is, [you are trying to] diagnose, if you will, what is going on. This doesn't answer the question, necessarily, in a satisfying way as to why it's going on, but it's at least a step in that direction.

Pask: In an effort to get at similar sorts of things, we recently devised a test, called the Cartoons Test, a device for seeing how people come to grips with a problem. The task is an old one in psychological history, the spy network problem, where information is presented in the form of lists of ordered pairs, for example A-B, B-C, C-D, A-C, A-D, which purport to represent the connections between spies in a spy network. A, B, and so on are spies. The arrows indicate A can communicate with B or B with C, etc. Now, a finding in the past literature is that people who can successfully do a test of this sort, in fact, are able to draw directed graphs and are able to memorize some sort of cognitive representation. The test we used consisted of a series of five graphs which were related in a cartoon, that is, an ordered set of graphs with a relation imposed on them. It represents the spy network in the years of 1880, 1885, 1890, 1895, and 1900. There's a predicate on the graph called "countries" which divides the spies into belonging to different countries. (I have a series of 5 graphs. As a matter of fact, the cartoon is periodic; eventually going one further stage, comes back to the original.) In all cases, the information is presented by saying, "Here is the typical sequences of messages among spies," presented mechanically, or else...in the form of lists. [These, then] have to be fabricated into nets [and] the nets...combined into a cartoon. The relations concern [such questions as] at the end: What about the political history of these countries? What major disruptions occurred to the spy network? And after that, of course, [the subjects must] recall the actual network.

Now, you get two rather grand manifestations of behavior with all sorts of messy cases in between. One sort of person says, all right, I've got a global view of a regular periodic function, I can predict that the spy network went through a disruption point, and will go through another later. From this, they begin to infer what the countries were and what the predicates were. From those, they begin to infer a graph, and from the graph, they reconstruct the list; they can read the list off the graph. On the other hand, you get people who attempt to do something like listing pairs. They will

try to recall the list and build the whole thing out, and make an inference from that. These two types are distinct enough. And the chap who does the building up in this manner, rather than remembering local graphs, of course, nearly always fails, because he gets in a mess.

These types are manifest in an environment rich enough to see what is going on. The point is the messy case. If you did it in a conversational way, in sort of a dialogue way, so that actually you do it with the subject, then you don't get messy results at all; you get very clearcut results. People do fall into definite categories if this is done in experiments or you engage in dialogue. You can ask what the nature of the things that were learned are and what the hypotheses were, and so on.

The point I'm trying to make is twofold, therefore. First of all, you get the same manifestation of the one [pertaining to Dienes' Experiment?] you observed in miniscule, where it is very difficult to see what is going on. Secondly, unless you do a continual protocol, it's extraordinarily difficult to get a clearcut result because people get into all sorts of aberrations which are irrelevant to the main target of the whole exercise. [We want] to ascertain how they do store these pieces, and how they recall them, and what sort of strategy they use, whether they use a global function -- recalling strategy, or whether they recall mini-functions and then put them together.

Landauer: I think I want to make a methodological remark that has to do with believing subjects' protocol. It is conceivable that the whole thing works the other way around. For example, in this yellow and red card problem, you have some subjects who are lucky in getting the right sequence of events; they put the right card down with the right one and made the right prediction; and by being lucky, they have simplified what they have to remember about the history. That simple history makes it easier for them to infer the rule. So the rule isn't telling you how they perform the task; it's telling you something about their inference based on their own behavior. The point is that there are some chance events involved in what the subject gets in the way of the history. I am suggesting that [with] some of those histories [it] may be easier to discern the pattern than with others. Your subjects who do very well, and give you what you call the best rule, may be simply those who are lucky to have a nice, patterned run in which it was visible. They may not be special people; they might be just lucky.

Minsky: It may be that some people are consistently lucky.

Landauer: I am making a methodological remark that you have to do something to demonstrate that to Dr. Dienes. The methodological remark that I make to you has to do with a lot of procedures that are used in simulation theories which are, to a psychologist's point of view, extraordinarily naive, in believing that the subject knows what is going on in his head.

Minsky: But that's exactly the reason why protocols are good, not bad. Because you can ask him on the second guess what he is doing. He might say "I'm going to guess a bunch of yellows for a while." That kills the randomness and exposes the structure.

Landauer: But you see, you're willing to believe that when the subject says "This is my strategy," that is in fact his strategy. I think that's naive.

Minsky: I don't believe in absolutes, but that's better than nothing.

Landauer: Do you want me to believe that you're just lucky that the words you're saying make sentences, or did you have a plan? I'm not quite sure.

Minsky: You don't mistrust everything that people say. Why do you mistrust them in these cases?

Landauer: I distrust everything that people say. I particularly distrust someone when they're asked to tell me what the machinery in their head is doing. At the very least, the monitor that is allowing them to do that ought to be part of the theory. It's a remarkable performance for a machine to be able to know what it's doing at all times with absolute accuracy. When, in fact, I've done experiments of this kind, guessing experiments, and asked people to tell me their strategy after each move, I've discovered quite

often that they tell me a strategy that simply couldn't be true. They say, "I have been alternating; every time there's a red, I say black." I look back at their record and on the last trial they didn't do that. I believe from this that what they're telling me, when I ask them what their strategy is, is their description of what they think they're doing, not the rule that they are in fact using.

Pask: Whereas I totally agree with your comment, I think it's only fair to remark that you can believe a decent sort of protocol only if it's confirmed, that is, if there's some correspondence between what the subject does and what he says he's doing. That, after all, is the whole point of the dialogue method, otherwise, you wouldn't need the dialogue. You would simply have a tape recorder and ask him when the light comes on what he thinks he's doing. The whole point about the method of paired experiments is that you have a chance to confirm the stated hypothesis, whether there is evidence for execution of the stated predicate or whatever. Secondly, I think it is necessary to [use a rather complex situation because often it is extraordinarily difficult in a card test experiment to formulate what you are doing. I don't know about card guessing; I do know about code learning, and it is extraordinarily difficult to give a coherent account. Whereas, in a thing like the cartoon, I can say I'm thinking something about a structure which I can verbalize about very easily. A spy network -- it's got some guts to it. Furthermore, [in the experiment, the subject knows (but I don't know if he believes)] that he is required to ultimately learn both kinds of rules, both, if you [re: Scandura] like the higher order rule and the lower order rule, and that he's going to be questioned on both of them. This, I think he knows in the sense that he's told at the outset. But, [in a classroom situation, I'm not very convinced that he believes it, because the [current] bias in education is very much toward asking him to output the lower order rules. He can use the higher order rules if he likes, but he better keep them concealed because they don't have any examination credit. This in fact is a very realistic comment; it's not by any means a stupid one. At any rate, Great Britain has a built-in bias to this effect. So it's necessary also, I think, to have the possibility exhibited during dialogue that either sort of cognitive mechanism [higher or lower order rule] or both of them can be manifest and tested and be subject to interrogation. Under these circumstances, I believe protocols can be reliable.

[Change tape]

Scandura: I wonder if it might be appropriate to turn the discussion to a more general consideration of the problem. As Tom and several others were talking, I had the impression that it might be useful to [be more precise] about what is assumed to be random. It's not necessarily the case that thought processes are random, and certainly in [Dienes type of task, I would agree with Marvin and others that it probably is not random. On the other hand, Tom may be right, from the standpoint of the type of experiment run, that what is random is the order in which the instances are presented. Suppose you [present subject with the number 1], and you tell him that he is supposed to generate the appropriate output. If you wish, you can tell him even more, for example, that [the output] involves the operation of addition, subtraction, multiplication, or division. Now, you have constrained the problem quite a bit for him already. On the first problem before he sees the output, he can come up with a number of possible answers, depending on whether he thinks [the correct operation is] addition, subtraction, etc. After he is given the correct response, 3, however, his effective goal becomes quite different. It's not just a question of generating a response [to the new input, but rather of finding an operation which is consistent with the first instance, 1→3, as well.] [In general] the order of instances presented will affect how the subject's effective goal changes over trials.

My own view is that the process of rule discovery can be interpreted in terms of higher order rules. Suppose, for example, that the subject must see 1→3, 2→6, and 3→9 before he can finally predict the correct response [12] to the input 4. In this case, a higher order rule that works looks at the pair 2→6, say, divides the 2 into 6, giving 3 and compares this [3] with the output of 1→3. Since they match, the desired rule is "multiply by 3." In this case, the three instances correspond to [degenerate] lower order rules. The control mechanism works as before. I suspect this may not be inconsistent with what several of you are saying about thought processes not necessarily

being random.

[Given the time], let's go now to the next topic on educational implications, and under that "memorizing versus organizing." Who among you [wants to] put yourself on record on some aspect of these issues?

Dienes: Let's take the famous multiplication tables. [I don't know any psychology] that would give me any kind of a suggestion as to how to predict, over the average run of children in the world, if it would be more effective to teach the tables only up to say 3 or 5, or up to 8, or up to 10, or up to 15, or even up to 20, in order to carry out compound multiplication as the final task. [What is] involved, say, in memorizing things when you have to absolutely know for certain, certain facts, and work out the rest - what is the proportion of what to what?

Scandura: My immediate reaction is that that's going to have a large number of different answers depending on exactly what you're trying to accomplish.

Minsky: Suppose you were a Martian and you came and saw people asking about this problem. Then, I think a very clear question would come to mind; namely, let's assume that we know that 9×5 is 5×9 . If you take that [commutively], then the 10 by 10 multiplication table has 55 entries. It appears that human children take the order of 100 hours to learn these pairs [so they can respond] fairly quickly. And the question the Martian would ask is "What on earth is going on here?" There's something wrong with this whole system because these children learn more than 55 things a day, [having] all sorts of characteristics. In fact, the first thing he would think of is "Is there a plot to make this difficult?" Quite seriously, there's something wrong. My theory is, [it takes] too much effort to make kids understand the meaning of 6×3 is 18. It's a very shallow fact, and so kids have to invent fantasies for each of these, which very, very badly interfere with their performance. But I'd like to hear the other theories of why it takes 100 hours to learn 55 facts, a result unprecedented in experimental psychology. There is no experiment known in which learning is that difficult.

Landauer: One [important] point is that those facts are being learned at a very high criterion compared to other facts. You're supposed to be able to do them real fast; you're supposed to be able to say 45 to the problem 9×5 faster than you can say your own name.

Minsky: It's only a few percent faster, if at all, and there's no need for it to be that fast. You're saying that it takes a long time to overlearn it, but you're wrong. Lots of people don't know that 7×6 is 42 at all, without figuring it out at the end of 50 hours, [even] at the end of 60 hours. It's not that you are getting slow responses. In other words, the overlearning is a problem but people don't seem to know the things for a long time. One of my graduate students thought 7×6 was 43, and I had to prove that it couldn't possibly be an odd number.

Landauer: You're familiar with a classical explanation of why that's so hard? It's the most interfering set of things that you could hope to try to learn. Generally, if you have things in which the stimuli and responses are all very much alike, it makes [them] very hard to learn. [This is the type of case that is] the hardest thing there is to learn, in addition to which you have to learn the addition tables and subtraction tables using the same pairs, which creates even more confusion.

Minsky: Can I tell [a funny] anecdote, even though it's out of turn. A friend of mine had a son who was having a lot of trouble with the multiplication tables. I interviewed him (This was introspection so it doesn't count), and said, "Michael, what's the trouble?" He said, "Well, I hate it." So I went on for a while and he said, "What's the use, after I learn this one, there will just be another table, and another table." Nobody had ever told him that there were only two tables and his picture was that mathematics was hopeless because each year there would surely be another one.

Fox: Could I just raise a question? It seems to me that there are lots of reasons why the acquisition of information is hard. The problem is how you use that information. The problem is also giving them structure.

Minsky: Well, why can't they acquire this then. You're missing the point.

Fox: I'm suggesting they do acquire it.

Minsky: I'm suggesting they don't. They don't know the products. Of course, there's a problem in doing n-place multiplications. The [basic] phenomenon is that they don't know 7×6 is 42. The phenomenon is not that they can't use it when they've got it.

Fox: But I'm still concerned that the learner has enormous problems in understanding. And this seems to me related to the question of how you traverse this knowledge structure

Minsky: Well, I'd like to ask Dr. Dienes that, if students know the tables, do they have major problems in doing other multiplications as well? Out of the 1,000 hours spent on arithmetic, how much is spent learning to do the two or three place multiplication algorithm after they know the tables?

Dienes: I think it's intuitively reasonably obvious, although correct me if I'm wrong. To learn multiple digit multiplication, and the rules about which way to do it, where to put the zeros, and shift the digits to the right or to the left, etc., there are many things [that may] be confused which are different but very similar. So if you don't know what you're doing, if you don't know that when you multiply 10 by 10, that the next digit has to go in the hundreds place, then you might easily put it somewhere else. We find in our project at Sherbrooke, for example, that [our experimental] children will score well under the norm in the first two or three years, but then they catch up by about fourth grade. By about sixth grade, we'll certainly beat everybody in the province in problem solving and actual techniques of multiple digit multiplication. We haven't done any kind of controlled experiments on this but we do have the actual results in provincial tests. For example, in the Quebec province, we almost invariably get the top provincial grades in our school, yet we only spend half an hour or an hour a week on actual arithmetic. In the other schools, they spend most of the time on it, and they don't do as well. So there must be something in the organizing that we give them that enables them to do these things. I think the answer to your question is that if you do know what it is that you're doing in the sense that you have a reason in your mind, then you'll make less mistakes.

Scandura: I wanted to comment on the question you [re: Minsky] raised. Could I perhaps very briefly outline what I think may be going on here. In talking about learning the tables, I'll be talking about addition, but it could just as easily be multiplication. Take something like $3 + 2$, [where] the output is 5. You can teach the child a general algorithm for doing this. Indeed, when we talk about teaching them the "meaning," we are in fact teaching them an algorithm. In particular, if you think of addends in terms of sets that you put together, there are basically three things [sub-operations]: giving a set interpretation to a numeral; once you've done that, forming the union of those sets; and, finally, attaching a numeral to the set union. You can formalize this as an algorithm, of course, and it will work with any pair of numbers, whether big ones or small ones, but it's not feasible except for very small numbers. So we teach something like that [algorithm for numbers] up to maybe $9 + 9$.

[In adding larger numbers], it's not just a question of learning that [set union] algorithm. If you analyze, for example, the load on memory in using this [set union algorithm] with a problem like, $457 + 398$ [in column addition form at the blackboard], it's quite high. It's higher than one would want to put up with in actually adding. So, what one does is present the basic facts [e.g., $3 + 2$] over and over; the idea is that sooner or later, hopefully sooner, the child will learn to drop out some of the irrelevant steps in the [set union] algorithm, and ultimately end up with a direct connection, or degenerate rule, or association between the addends and the sum. Once he has got those [facts], we can go about teaching him how to solve the [column addition] tasks.

A final remark that I would make on this is that we have run training experiments where we have ignored the meaning of the algorithm. We have had subjects who came in knowing the facts, and taught them to do things like add or subtract with arbitrary numbers without their even knowing what they were doing. This was literally true of first graders. They didn't know what they were doing, [but] if they came in knowing th

facts; [and we could teach them to] do column addition [with arbitrary numbers] in about 10-12 hours of time. Subtraction took about half as long. It was a game; they didn't know what they were doing or why they were doing it. [But, they could do it.] The moral, I suppose, is that given any task, it is sufficient to teach an algorithm that works. That algorithm, of course, may not cover all that you're interested in.

Edward Haupt: I'd like to go back a little to what Tom Landauer said. In the case of learning the multiplication tables, even much more [so than] the addition tables, it's a consummate case of cognitive interference.

Minsky: No. Addition is more interfering because the numbers are the same.

Haupt: Well, the algorithm for learning multiplication involves addition, and the substitute method which kids usually use for constructing [multiplication tables] often involves addition. The one thing that teachers do that I think makes a large amount of difference in the difficulty of learning the addition table or multiplication table is that they will take something which they call the "four facts" or the "five facts" or the "ones," and they will give the kid this whole raft of things. If you work out an oral drill, so that a child just has to learn one new fact, practice it a couple times, and alternate it with things that he has already learned, then you have a system that will get somewhat faster learning. I have some data on kids who are relatively poor at memorizing, who get better retention [by this method]. I know an even better way to get good retention, and that is to use the operant procedure called fading, but that's not typical of what people do in a classroom. [In effect, if you simplify the] interference problem by giving them one new fact at a time, they can perform algorithms rapidly. If you [use] a slower kind of algorithm, set union for addition, repeated counting for multiplication...then, when a child can figure out all the [sums], you really don't need to teach him anything more. It might well be that the child has learned what multiplication means in general. Once he has satisfied that criterion you could go to training him on the [fact] drill procedure so that he will have the speed necessary for [doing the usual] algorithms. [Drilling is needed] to cut down on the number of steps that are involved.

Scandura: That's precisely the reason. If you ask him to do addition by going through the set union addition procedure, plus the [usual addition] algorithm, the load on memory becomes astronomical and the number of mistakes will increase.

Landauer: I'd like to state a general and very extreme view on the matter we're discussing. I'll rephrase it just slightly as a distinction between memory and computational logic, rather than between memory and organization. I think there is very good reason to believe that what the human is good at is the storage of vast amounts of elementary pieces of data. As John [Fox] noted a minute ago, a person can look at thousands of pictures in a small space of time and then recognize each one of them afterwards. The average human knows a quarter of a million different words, etc. On the other hand, when you ask someone to do a simple computation, like multiplying two 2-digit numbers, it's beyond his capability usually. A human can only keep one or two things in mind at the same time. The human is a crappy calculator. He's got an unbelievably fantastic core storage system and an unbelievably bad CPU. Now, it's my view that if we go about devoting a lot of education to trying to make good logical calculators out of people instead of spending the time loading up their billions of memory store registers with good things that they could use for the rest [of their lives], that we are in fact fighting nature, and doing children a grave disservice. What we should do is to teach them a lot of little tiny facts and save our creative energies with regard to computational logic for inventing machines to do that for people.

Scandura: Just because people have a small processing capacity and a large core storage, it doesn't follow that we should [just] load up that store with little pieces of data. The most important bits of knowledge may consist of higher order rules or strategies. This is what we need to create good machines. Are these the "good things"?

Minsky: Most of what people know is what the facts are for, and for each of the 250,000 words (I don't really think there are that many); you probably learn 100 or more bits of procedure for when to use them and how. I think you're just plain wrong. Well, you're right that there's a lot of memory but it's not in the form of facts. Every fact has to be accompanied by a lot more information about how to use it.

Dienes: Just to contradict, may I show you how the Sherbrooke children do their arithmetic. To save someone from having to retain all these little facts, which are interfering with each other all the time, I do [some] thing like this. They have a matrix, the rows of which go up by doubling, successive doubling. (By the way, that's something else I'd like to know; why is it that children find doubling so easy?) They can go as far as they want to on this matrix; then in the columns you have trebling, which is a lot harder to learn. You don't learn each doubling or trebling separately, but how to do them. So you have doubling going on [in rows] and trebling going on [in columns]. They don't necessarily write this down. They have some kind of a mental picture of this. No along [the row], they not only have multiplication by two but they also have it by four, by eight, etc., but structurally joined together. Now, they know, sort of in their bones, the associative principle. That is, you multiply something by two and then multiply that by two, you have in fact multiplied it by four. They know these kinds of things. The same with the 3 and the same with the 6. In other words, to get 6's they can double and then treble.

My problem is the following. Obviously you're taking a certain amount of time up while you're working these things out. Kids in other schools have learned them all by heart but they still don't get as good results as we do. We learn them this way and the children work them out, but they work them out extremely rapidly. Now, what is the balance between this type of organizing [and memorizing]?

Minsky: The fact that you use the word "treble," which is not in my vocabulary, reduces the interference by a good factor. It might be that if you had ten other words for the numbers, you'd sidestep most of that problem.

Dienes: Well, possibly. But they do realize it's three times.

Minsky: But it doesn't interfere with memory so quickly.

Dienes: Yes, possibly. You may have something there.

Minsky: Numbers don't have synonyms, whereas many other words do.

Dienes: Now, children seem to be able to have hold of this and have pictured this kind of numerical table, or image, which they hold in their mind's eye, much better than ten rather disconnected and structurally unrelated multiplication tables.

Pask: I quite agree. In response to Dr. Landauer, if you had said that the human being processor was extremely good, so that it retains a lot of little procedures, which it is capable of executing, and these procedures are instincts and a perceptual act is predication [i.e., looking at these pictures and finding what in the world the appropriate features are and likewise making a linguistic statement], with propositions encoded in the form of procedures, (Even in your mode), I assume that they get learned, instead of simply stored because, after all, they're all lists somewhere which are going to be interpreted, by expanding a balloon or something) -- if you had said that, well, then, I'd agree with you. That, of course, is compatible with pattern recognition, that is, essentially a pattern which stands for a procedure. And it doesn't stand for one procedure; it's called a pattern because it stands for a whole lot of little procedures. Now, if you had said that, I'd agree with you. As it is, however, I don't agree with you at all. I mean, the storage of large amounts of data I don't think is a crucial point actually. I think the crucial [points are] the predication problem, and the fact that you are storing those procedures.

Fox: Can I try and join them together? There are, of course, a lot of issues but it seems certainly true that people can store an enormous amount of information. And it's also true that people can operate on that information. And there's very good reason to

believe that the way they operate is to look at two of those bits and try to construct a relation. Now, the processes of short term memory, or whatever, seem to be that every piece of information has exactly the kind of structure that would make that kind of [construction] possible. If all that were true, the educational implications are (this is the MIT line) that technology should be trying to phrase knowledge structures in terms that are clearly utilizable given the cognitive facilities of the average person, maybe in some sort of relational triple. Until you have gotten there, you are stuck with remembering. You can only give people facts until you understand the problems. When you understand the problems and can cast them in ways people can act on, they can attend to small parts of them and not overload. In other words, you give them one thing at a time because they take one new thing and attach it to an old thing. You've [then] got a single symbol that you can expand later if you need to. You have extended your knowledge.

candura: I suspect that many of you have a lot more you would like to say, but I am afraid that time is about up for now. See you after lunch.

Navy

- 4 Dr. Marshall Farr Dir.
Personnel and Training
Research Programs
Office of Naval Research
Arlington, VA 22217
- 1 ONR Branch Office
495 Summer St.
Boston, MA 02210
ATTN: Research Psych.
- 1 ONR Branch Office
1030 East Green Street
Pasadena, CA 91101
ATTN: E.E. Gloye
- 1 ONR Branch Office
536 South Clark Street
Chicago, IL 60605
ATTN: M.A. Bertin
- 1 Office of Naval Research
Area Office
207 West 24th Street
New York, NY 10011
- 6 Director
Naval Research Laboratory
Code 2627
Washington, DC 20390
- 1 Defense Documentation Center
Cameron Station, Building 5
5010 Duke Street
Alexandria, VA 22314
- 1 Special Assistant for Manpower
OASN (M&RA)
Pentagon, Room 4E794
Washington, DC 20350
- 1 LCDR Charles J. Theisen, Jr.
4024
Naval Air Development Center
Warminster, PA 18974
- 1 Chief of Naval Reserve
Code 3055
New Orleans, LA 70146
- 1 Dr. Harold R. Booher
Naval Air Systems Command
NAVAIR 04A4
Washington, DC 20361
- 1 Commander
Naval Air Systems Command
Department of the Navy
AIR-413C
Washington, DC 20361
- 1 Dr. Lee Miller
Naval Air Systems Command
AIR-413E
Washington, DC 20361
- 1 CAPT John F. Riley, USN
Commanding Officer
U.S. Naval Amphibious School
Coronado, CA 92155
- 1 CAPT Ouida C. Upchurch, USN
Program Coordinator
Bureau of Medicine & Surgery
Washington, DC 20372
- 1 Chairman
Behavioral Science Department
Naval Command & Management Div.
U.S. Naval Academy
Luce Hall
Annapolis, MD 21402
- 1 Chief of Naval Educ.
and Training
Naval Air Station
Pensacola, FL 32508
ATTN: Capt. Bruce Stone
- 1 Mr. Arnold Kubinstein
Naval Material Command
Room 820, Crystal Plaza #6
Washington, DC 20360
- 1 Commanding Officer
Naval Medical Neuropsychiatric
Research Unit
San Diego, CA 92152
- 1 Director, Navy Occupational Task
Analysis Program (NOTAP)
Navy Personnel Program Support
Activity
Building 1304, Bolling AFB
Washington, DC 20336
- 1 Dr. Richard J. Niehaus
Office of Civilian
Manpower Management
Code 06A
Washington, DC 20390
- 1 Dept. of the Navy
Office of Civilian
Manpower Management
Code 263
Washington, DC 20390
- 1 Superintendent
Naval Postgraduate School
Monterey, CA 93940
ATTN: Library (Code 2)
- 1 Commander, Navy Recruiting
4015 Wilson Boulevard
Arlington, VA 22201
ATTN: Code 015
- 1 Mr. George N. Graine
Naval Ship Systems Command
SHIPS 047C12
Washington, DC 20360
- 1 Chief of Naval Tech.
Naval Air Station
Millington, TN 38055
Attn: Dr. Norman Ke...
- 1 Commanding Officer
Service School Command
U.S. Naval Training
San Diego, CA 92115
ATTN: Code 3030
- 1 Dr. William L. Maloy
Principal Civilian Analyst
for Education & Training
Naval Training Command
Pensacola, FL 32508
- 1 Dr. Alfred F. Smoder
Training Analysis &
Evaluation Group
Naval Train. Equip.
Code N-00T
Orlando, FL 32813
- 1 Dr. Hanns H. Wolff
Tech. Dir. (Code N-00T)
Naval Training Equip.
Orlando, FL 32813

Chief of Naval Training Sup.
Code N-21, Bldg. 45
Naval Air Station
Pensacola, FL 32508

1 Dr. Ralph Dusek
U.S. Army Research Inst.
1300 Wilson Blvd.
Arlington, VA 22209

1 AFOSR/NL
1400 Wilson Blvd.
Arlington, VA 22209

Dr. Robert French
Naval Undersea Center
San Diego, CA 92132

1 Mr. Edmund F. Fuchs
U. S. Army Research Inst.
1300 Wilson Blvd.
Arlington, VA 22209

1 Commandant
USAF School of Aerosp. Med
Aeromedical Library
Brooks AFB, TX 78235

CDR Richard L. Martin
Fighter Squadron 124
NAS Miramar, CA 92145

1 Dr. Leon H. Nawrocki
U. S. Army Research Inst.
1300 Wilson Blvd.
Arlington, VA 22209

1 Dr. Sylvia Mayer
HQ Electronic Systems Div.
LG Hanscom Field
Bedford, MA 01730

Navy Personnel R&D Center
San Diego, CA 92152

1 Dr. J. E. Uhlaner
U.S. Army Research Inst.
1300 Wilson Blvd.
Arlington, VA 22209

1 Capt. Jack Thorpe
Flying Training Div.(HRL)
Williams AFB, AZ 85224

Marine Corps

Navy Personnel R&D Center
San Diego, CA 92152
ATTN: Code 10

1 Dr. Joseph Ward
U.S. Army Research Inst.
1300 Wilson Blvd.
Arlington, VA 22209

1 Mr. E.A. Dover
Manpower Measurement Unit
Arlington Annex, Rm 2413
Arlington, VA 20380

U.S. Army Admin. Center
Personnel Admin. Combat
Devel. Activity
ATCP-HRO
Ft. Ben. Harrison
IN 46249

1 HQ USAREUR & 7th Army
ODCSOPS
USAREUR Director of GED
APO New York 09403

1 Commandant, Marine Corps
Headquarters
Code MPI-20
Washington, D.C. 20380

1 Dir. Office of Manpower
Utilization
HQ, Marine Corps
MCB (Bldg. 2009)
Quantico, VA 22134

Armed Forces Staff Col.
Norfolk, VA 23511
ATTN: Library

Air Force

1 Research Branch (AF/DPXYR)
Pentagon, Room 5C428
Washington, D.C. 20330

1 Dr. A.L. Slafkosky
Scientific Advisor RD-1
HQ, Marine Corps
Washington, DC 20380

Director of Research
U.S. Army Armor Human
Research Unit
Bldg. 2422 Morade St.
Fort Knox, KY 40121
ATTN: Library

1 Dr. G. A. Eckstrand (AFHRL/AS)
Wright-Patterson AFB
Ohio 45433

1 Dr. Ross L. Morgan (AFHRL/AST)
Wright Patterson AFB
Ohio 45433

Coast Guard

1 Mr. Joseph J. Cowan
Psychological Research
Branch (G-P-1/62)
Coast Guard Headquarters
Washington, DC 20590

Commandant
U.S. Army Infantry Sch.
ATTN: ATSH-DET
Fort Benning, GA 31905

1 AFHRL/DOJN
Stop # 63
Lackland AFB, TX 78236

Other DOD

Deputy Commander
U.S. Army Inst. of Admin.
Fort BEN Harrison
IN 46216
ATTN: EA

1 Dr. Martin Rockway (AFHRL/TT)
Lowry AFB
Colorado 80230

1 Lt. Col. Henry L. Taylor
Mil. Asst. for Human Res.
OAD (E&LS) ODDR&E
Pentagon, Room 3D129
Washington, DC 20301

Dr. Frank J. Harris
U.S. Army Research Institute
1300 Wilson Boulevard
Arlington, VA 22209

1 Major P.J. DeLeo
Instructional Technology
AF Human Resources Lab
Lowry AFB, CO 80230

1 Mr. William J. Stormer DOD Computer Inst. Washington Navy Yard, Bldg. 175 Washington, DC 20374	1 Dr. David Bowers University of Michigan Inst. for Social Research Ann Arbor, MI 48106	1 Dr. M. D. Havron Human Sciences Research 7710 Old Spring House Rd West Gate Industrial Pk. McLean, VA 22101
1 Col. Austin Kibler Advanced Research Projects Ag. Human Resources Res. Office 1400 Wilson Blvd. Arlington, VA 22209	1 Mr. Kenneth M. Bromberg Manager- Wash. Operations Information Concepts, Inc. 1701 North Fort Myer Dr. Arlington, VA 22209	1 HumRRO Division 3 P.O. Box 5787 Presidio of Monterey, CA
1 Mr. Thomas O'Sullivan Advanced Res. Projects Agency Human Resources Res. Office 1400 Wilson Blvd. Arlington, VA 22209	1 Dr. Ronald P. Carver Am. Institutes for Research Foxhall Square 3301 New Mexico Ave., N.W. Washington, D.C. 20016	1 HumRRO Division 4 Infantry P.O. Box 2086 Fort Benning, GA 31905
1 Helga L. Reich Advanced Res. Projects Agency Manpower Management Office 1400 Wilson Blvd. Arlington, VA 22209	1 Centry Research Corporation 4113 Lee Highway Arlington, VA 22207	1 HumRRO Division 5 Air Defense P.O. Box 6057 Fort Bliss, TX
<u>Other Government</u>		
1 Dr. Eric McWilliams Technology & Systems, TIE National Science Foundation Washington, DC 20550	1 Dr. Allan M. Collins Bolt, Beranek, and Newman 50 Moulton St. Cambridge, MA 02138	1 Dr. Lawrence Johnson 200 S. Street, N.W. Suite 502 Washington, DC 20009
1 Dr. Andrew R. Molnar Tech. Innovations in Educ. National Science Foundation Washington, DC 20550	1 Dr. H. Peter Dachler University of Maryland Dept. of Psychology College Park, MD 20742	1 Dr. Milton S. Katz Mitre Corporation Westgate Research Center McLean, VA 22101
<u>Miscellaneous</u>		
1 Dr. Scarvia B. Anderson Educational Testing Service 17 Executive Park Drive, N.E. Atlanta, GA 30329	1 ERIC Processing and Reference Fac. 4833 Rugby Ave. Bethesda, MD 20014	1 Dr. Steven W. Keele University of Oregon Dept. of Psychology Eugene, OR 97403
1 Dr. John Annett The Open University Milton Keynes Buckinghamshire ENGLAND	1 Dr. Victor Fields Montgomery College Dept. of Psychology Rockville, MD 20850	1 Dr. David Klahr Carnegie-Mellon University Dept. of Psychology Pittsburgh, Pa. 15213
1 Dr. Richard C. Winson Stanford University Department of Psychology Stanford, CA 94305	1 Dr. Edwin A. Fleishman Am. Institutes for Research Foxhall Square 3301 New Mexico Ave., N.W. Washington, D.C. 20016	1 Dr. Alma E. Lantz University of Denver Denver Research Institute Industrial Economics Division Denver, CO 80210
1 Dr. Gerald V. Barrett University of Akron Department of Psychology Akron, OH 44325	1 Dr. Robert Glaser Learning R&D Center University of Pittsburgh Pittsburgh, PA 15213	1 Dr. Robert Mackie Human Factors Research 6780 Cortona Drive Santa Barbara Research Goleta, CA 93017
1 Dr. Bernard M. Bass University of Rochester Management Research Center Rochester, NY 14627	1 Dr. Henry J. Hamburger University of California School of Social Sciences Irvine, CA 92664	1 Mr. A.J. Pesch Eclectech Associates, P.O. Box 178 North Stonington, CT

- 1 Mr. Luigi Petrullo
2431 North Edgewood Street
Arlington, VA 22207
- 1 Dr. Joseph W. Rigney
University of Southern California
Behavioral Technology Laboratories
3717 South Grand
Los Angeles, CA 90007
- 1 Dr. Leonard L. Rosenbaum, Chairman
Montgomery College
Department of Psychology
Rockville, MD 20850
- 1 Dr. George E. Rowland
Rowland and Company, Inc.
P.O. Box 61
Haddonfield, NJ 08033
- 1 Dr. Arthur I. Siegel
Applied Psychological Services
404 East Lancaster Avenue
Wayne, PA 19087
- 1 Dr. C. Harold Stone
1428 Virginia Avenue
Glendale, CA 91202
- 1 Mr. Dennis J. Sullivan
725 Benson Way
Thousand Oaks, CA 91360
- 1 Dr. Barton J. Underwood
Northwestern University
Department of Psychology
Evanston, IL 60201
- 1 Dr. David J. Weiss
University of Minnesota
Department of Psychology
Minneapolis, MN 55455
- 1 Dr. Anita West
Denver Research Institute
University of Denver
Denver, CO 80210
- 1 Dr. Kenneth N. Wexler
University of California
School of Social Sciences
Irvine, CA 92664
- 1 Dr. Donald A. Norman
University of California, S.D.
Center for Human Information Processing
LaJolla, CA 92037
- 1 Dr. Dexter Fletcher
University of Illinois
at Chicago Circle
Dept. of Psychology
Chicago, IL 60680

MERGE Research Institute

Reports

(Place of publication is shown in parentheses.)

BEST COPY AVAILABLE

1. Scandura, J. M. An exceptionally rare athlete. 1959. (Mentor Magazine, 1959, 30-31.)
2. Scandura, J. M. Abstract card tasks for use in problem solving research. 1964. (Journal of Experimental Education, 1964, 33, 145-148.)
3. Scandura, J. M. An analysis of exposition and discovery modes of problem solving instruction. 1964. (Journal of Experimental Education, 1964, 33, 149-154; also in P. C. Burns & A. M. Johnson (Ed.) Research on Elementary School Curriculum and Organization. Boston: Allyn and Bacon, 1969; and in Dissertation Abstracts, 1963, 23, No. 8.)
4. Scandura, J. M. The other role of schools of education. 1964. (Western New York Study Council Journal, 1964, 15, 11-13.)
5. Scandura, J. M. Fractions--names and numbers. 1964. (The Arithmetic Teacher, 1964, 11, 468.)
6. Scandura, J. M. Probability theory--basic concepts. 1965. (New York State Mathematics Teachers Journal, 1965, 112-118.)
7. Scandura, J. M. & Nelson, J. L. The emerging research role of the subject matter educator. 1965. (Journal of Research in Science Teaching, 1965, 3, 51-53.)
8. Scandura, J. M. Educational research and the mathematics educator. 1965. (The Mathematics Teacher, 1965, 58, 131-138; also in The Mathematics Education, 1967, 1, 1-14.)
9. Scandura, J. M., Johnson, D. A., & Thomason, G. M. Federal funds for the improvement of mathematics education. 1965. (The Mathematics Teacher, 1965, 58, 551-554; also in The Arithmetic Teacher, 1965, 12, 496-499.)
10. Scandura, J. M. Familiarization and cue selection. 1965. (Psychological Reports, 1965, 17, 27-38.)
11. Scandura, J. M. Familiarization and one-trial learning: A positive interaction. 1965. (Psychological Reports, 1965, 17, 27-38.)
12. Scandura, J. M. Problem solving and prior learning. 1965. (Journal of Experimental Education, 1966, 34, 7-11.)
13. Scandura, J. M. Algorithm learning and problem solving. 1966. (Journal of Experimental Education, 1966, 34, 1-6.)
14. Scandura, J. M. Prior learning, presentation order and prerequisite practice in problem solving. 1966. (Journal of Experimental Education, 1966, 34, 12-18.)
15. Scandura, J. M. & Behr, M. Prerequisite practice and criterion form in mathematics learning. 1966. (Journal of Experimental Education, 1966, 35, 54-55.)
16. Scandura, J. M. Problem solving and prior learning. 1966. (The Reporter, Central New York School Study Council, School of Education, Syracuse University, 1966, 17(3), 1-5.)
17. Scandura, J. M. Teaching--technology or theory. 1966. (American Educational Research Journal, 1966, 3, 139-146; also in N. I. E. Journal, 1, 11-17; and in Scientia Paedagogica Experimentalis, 1967, 4, 231-243.)
18. Scandura, J. M. Precision in research on mathematics learning: The emerging field of psycho-mathematics. 1966. (Journal of Research in Science Teaching, 1966, 4, 253-274.)
19. Greeno, J. G. & Scandura, J. M. All or none transfer based on verbally mediated concepts. 1966. (Journal of Mathematical Psychology, 1966, 3, 388-411.)
20. Scandura, J. M. & Roughead, W. G. Conceptual organizers in short-term memory. 1966. (Journal of Verbal Learning and Verbal Behavior, 1967, 6, 679-682.)
21. Scandura, J. M. A new doctoral program in mathematics education research. 1966. (The Mathematics Teacher, 1966, 8, 729; also in American Mathematical Monthly, 1967, 74, 72-74.)
22. Scandura, J. M. Concrete examples of commutative non-associative systems. 1966. (The Mathematics Teacher, 1966, 59, 735-736.)
23. Scandura, J. M. Training leaders in mathematics education: A contrast with science education. 1966. (Journal of Research in Science Teaching, 1966, 4, 227-229.)
24. Scandura, J. M. The nature of research in mathematics education. 1967. (The High School Journal, 1967, 50, 227-237.)
25. Scandura, J. M. An alternative approach to research in mathematics learning--a reaction to Suppes's paper. October, 1967. (Journal of Research and Development in Education, 1967, 1, 33-43.)
26. Scandura, J. M. Critique of "Research in the improvement of verbal problem solving." December, 1967.
27. Scandura, J. M. Training in mathematics education: A proposal for action. 1967. (The Mathematics Teacher, 1967, 60, 714 & 719.)
28. Scandura, J. M. & Wells, J. N. Advance organizers in learning abstract mathematics. 1967. (American Educational Research Journal, 1967, 4, 295-301.)
29. Scandura, J. M., Woodward, E., & Lee, F. Rule generality and consistency in mathematics learning. 1967. (American Educational Research Journal, 1967, 4, 303-319.)
30. Scandura, J. M. The basic unit in meaningful learning--association or principle? 1967. (The School Review, 1967, 75, 329-341; also in Proceedings, 75th Annual Convention APA, 275-276, 1967.)
31. Scandura, J. M. Learning verbal and symbolic statements of mathematical rules. 1967. (Journal of Educational Psychology, 1967, 58, 356-364.)
32. Scandura, J. M. Learning principles in paired-associate lists. 1967. (Psychological Reports, 1967.)
33. Scandura, J. M. Concept dominance in short-term memory. 1967. (Journal of Verbal Learning and Verbal Behavior, 1967, 6, 461-469.)
34. Scandura, J. M. & Sattov, J. An analysis of existing curricular materials in mathematics. 1968.
35. Scandura, J. M. (Ed.) Report of a conference on elementary mathematics education and the Pennsylvania State Department of Public Instruction. 1968.
36. Scandura, J. M. Con Ebel's "Staw man". 1968. (Phi Delta Kappan, 1968, 49, 286.)
37. Scandura, J. M. & Anderson, J. Research in mathematics and science education. 1968. (School Science and Mathematics, 1968, 58, 185-192; also in Educational Research and the Science Educator, 1968, 52, 353-358.)
38. Scandura, J. M. Research in psycho-mathematics. October, 1968. (The Mathematics Teacher, 1968, 61, 581-591; also in The Mathematics Education, 1968, 3; and in R. B. Ashlock & W. L. Herman (Ed.) Current Research in Elementary School Mathematics. New York: Macmillan, 1970.)
39. Roughead, W. G. and Scandura, J. M. "What is learned" in mathematical discovery. 1968. (Journal of Educational Psychology, 1968, 59, 283-289.)
40. Scandura, J. M. Expository teaching of hierarchical subject matters. 1968. (Journal of Structural Learning, 1969, 2(1), 17-25; also in Psychology in the Schools, 1969, 6, 307-310.)
41. Scandura, J. M. The formulation of research on subject matter learning. July, 1968. (Psychology in the Schools, 1968, 5, 330-341.)
42. Scandura, J. M. and Durnin, J. H. Extra-scope transfer in learning mathematical strategies. 1968. (Journal of Educational Psychology, 1968, 59, 350-354.)
- Scandura, J. M. New directions for theory and research on rule learning: 1. A set-function language. 1968. (Acta Psychologica, 1968, 28, 301-321.)

44. Scandura, J. M. New directions for theory and research on rule learning: II. Empirical research. 1968. (Acta Psychologica, 1969, 29, 101-133.)
45. Scandura, J. M. New directions for theory and research on rule learning: III. Analysis and theoretical direction. 1968. (Acta Psychologica, 1969, 29, 205-227.)
46. Scandura, J. M. Federal control and basic research in education. 1968. (School & Society, 1969, April, 227-228.)
47. Scandura, J. M., Barksdale, J., Durnin, J., & McGee, R. An unexpected relationship between failure and subsequent mathematics learning. 1968. (Educational Studies in Mathematics, 1969, 1, 247-251; also in Psychology in the Schools, 1969, 4, 379-381.)
48. Scandura, J. M. Rule-governed behavior in mathematical structures. February 2, 1969.
49. Scandura, J. M. A research basis for mathematics education. November 10, 1969. (High School Journal, 1970, 53, 264-280; Educational Studies in Mathematics, 1971, 3, 229-243; and Journal of Structural Learning, 1970, 2(3), 1-18.)
50. Scandura, J. M. The role of rules in behavior: Toward an operational definition of what (rule) is learned. November 25, 1969. (Psychological Review, 1970, 77, 516-533.)
51. Scandura, J. M. & McGee, R. An exploratory investigation of basic mathematical abilities of kindergarten children. December 1, 1969. (Educational Studies in Mathematics, 1972, 4, 331-345.)
52. Scandura, J. M. A theory of mathematical knowledge: Can rules account for creative behavior? January 28, 1970. (Journal for Research in Mathematics Education, 1971, 2, 183-196.)
53. Ehrenpreis, W. & McGee, R. Research implications for the elementary school: From the laboratory to the classroom. March 1, 1970.
54. Scandura, J. M. & Voorhies, D. The role of irrelevant attributes and irrelevant operations in rule learning. May 20, 1970. (Journal of Educational Psychology, 1971, 62, 352-356.)
55. Scandura, J. M. Deterministic theorizing in structural learning: Three levels of empiricism. May 20, 1970. (Journal of Structural Learning, 1971, 3(1), 21-53.)
56. Scandura, J. M. What is a rule? May 21, 1969. (Journal of Educational Psychology, 1972, 63, 179-185.)
57. Scandura, J. M. Problem solving ability in elementary mathematics -- its nature and measurement. February 8, 1968.
58. Scandura, J. M. Theoretical note: S-R theory or automata? (A reaction to Suppes' reaction). In Structural Learning II.
59. Scandura, J. M. A plan for the development of a conceptually based mathematics curriculum for disadvantaged children. October 15, 1970. (Instructional Science, 1972, 247-263.)
60. Scandura, J. M. An algorithmic base for curriculum construction in mathematics (A theory of mathematical knowledge). October 29, 1970. (In Lamon, W. E. (Ed.) Learning and the Nature of Mathematics. Chicago: Science Research Associates, Inc., 1971.)
61. Scandura, J. M. On the value of status studies in education. (Journal for Research in Mathematics Education, 1971, 2, 115-117.)
62. Scandura, J. M., and Durnin, J. H. Assessing behavior potential: Adequacy of basic theoretical assumptions. November 30, 1971. (Journal of Structural Learning, in press.)
63. Durnin, J. H., and Scandura, J. M. An algorithmic approach to assessing behavior potential: Comparison with item forms and hierarchical technologies. January 1, 1972. (Journal of Educational Psychology, 1973, 65, 262-272.)
64. Ehrenpreis, W., and Scandura, J. M. An algorithmic approach to curriculum construction in mathematics: A field test. (Journal of Educational Psychology, 1974, in press.)
65. Loverre, G., and Scandura, J. M. Conceptually based development of individualized materials for critical thinking based on logical inference. May 1, 1972. (Reading Research Quarterly, IX, 186-205.)
66. Scandura, J. M. Mathematical problem solving. June 1, 1972. (American Mathematical Monthly, 81, 273-280.)
67. Scandura, J. M. The role of higher order rules in problem solving. (Journal of Experimental Psychology, 1974, 6, 984-991.)
68. Scandura, J. M. Structured knowledge: Static or dynamic? (Contemporary Psychology, 1973, 18, 23-24.)
69. Scandura, J. M. Competence, knowledge, and their relation to human behavior. (Talk given at Psychonomic Society Meeting in St. Louis, Mo., Nov. 2, 1972. Catalog of Selected Documents in Psychology, 1974, in press.)
70. Scandura, J. M., Durnin, J. H., Wulfek, W. H. II. Higher order rule characterization of heuristics for compass and straightedge constructions in geometry. (Artificial Intelligence, 1974, in press.)
71. Scandura, J. M. Structural learning and the optimization of open education. (Pending.)
72. Scandura, J. M. Structural learning and the design of educational materials. (Educational Technology, August, 1973, 7-13.)
73. Scandura, J. M. The structure of memory: Fixed or flexible? (Catalog of Selected Documents in Psychology, Spring, 1974, Abstract on pp. 37-38.)
74. Scandura, J. M. On higher order rules. (Educational Psychologist, 1973, 10, 159-160.)
75. Scandura, J. M. On conceptual foundations for mathematical education. (Pending.)
76. Scandura, J. M. A book review of O. Milton, Alternatives to the traditional. (Educational Studies, in press.)
77. Scandura, J. M. The resolution of instructional problems: Rhetoric, research, or theory? (Invited address, Division 15, APA, 1974. Pending.)
78. Scandura, J. M. A short review of Arthur R. Jensen's Educability and Group Differences. (New York, Harper & Row, 1973) (In Choice, February, 1974, p. 1919.)
79. Scandura, J. M. Problem definition, rule selection, and retrieval in problem solving. (Pending.)
80. Scandura, J. M., Durnin, J. H., Wulfek, W. H. II, & Ehrenpreis, W. Diagnosis and instruction of higher order rules for solving geometry construction problems. (In 1974 Conference Proceedings; see B-5 below.)
81. Scandura, J. M. On hiring and promotion. (University of Pennsylvania: The Almanac, 28, No. 21, February 5, 1974.)
82. Scandura, J. M. Empirical evaluation of rule based theories.
83. Scandura, J. M. "Reaction to Symposium. The Learning Research Laboratory and the Classroom." (Educational Psychologist, 1974, in press.)
84. Voorhies, D. & Scandura, J. M. Determination of memory load in information processing. (In 1974 Conference Proceedings; see B-5 below.)
85. Scandura, J. M. Human problem solving: A synthesis of content, cognition, and individual differences. (In 1974 Conference Proceedings; see B-5 below.)
86. Scandura, J. M. How does mathematics learning take place? (Invited talk given at General Session, NCTM meeting, Bozeman, Montana, June 14, 1974.)

Books

- B-1. Scandura, J. M. Mathematics: Concrete Behavioral Foundations. New York: Harper & Row, 1971.
- B-2. Scandura, J. M., Durnin, J. H., Ehrenpreis, W., & Luger, G. An Algorithmic Approach to Mathematics: Concrete Behavioral Foundations. Harper & Row, 1971.
- B-3. Scandura, J. M. Structural Learning: The Theory and Empirical Research. New York: Gordon & Breach Science Publishers (One Park Avenue, New York, New York 10016), 1973.
- B-4. Scandura, J. M. (Ed.) with contributions by Arbib, M., Corcoran, J., Domotor, Z., Greeno, J., Lovell, K., Newall, A., Rosenbloom, P., Scandura, J., Shaw, R., Simon, H., Suppes, P., Wittrock, M., & Witz, K. Structural Learning: Views and Approaches. Pending.
- B-5. Scandura, J. M., Durnin, J. H. & Wulfek, W. H. II (Eds.) 1974 Proceedings: Fifth Annual Interdisciplinary Conference on Structural Learning. Office of Naval Research (Available from National Technical Information Services, Springfield, Va.), 1974, in press.

Curriculum Materials

- C-1. Scandura, J. M., Loverre, G. P., & Scandura, A. M. Critical Reading Series (A self-diagnostic and instructional series of workbooks for teaching critical reading based on logical inference). Ann Arbor Publishers, Inc. (P.O. Box 388, Worthington, Ohio 43085). (Workbooks A, B, C, D and Teachers' Manual.)