DOCUMENT RESUME

ED 361 176 SE 053 562

TITLE To Strengthen American Cognitive Science for the Twenty-First Century. Report of a Planning Workshop for the Cognitive Science Initiative at the National Science Foundation (Washington, D.C., April 20-21, 1991).

INSTITUTION National Science Foundation, Washington, D.C.

REPORT NO NSF-92-4

PUB DATE [92]

NOTE 19p.

PUB TYPE Reports - Research/Technical (143) -- Collected Works - Conference Proceedings (021)

EDRS PRICE MF01/PC01 Plus Postage.

DESCRIPTORS *Artificial Intelligence; *Cognitive Development; *Cognitive Psychology; Concept Formation; *Educational Research; Higher Education; *Interdisciplinary Approach; Neuropsychology; Prediction; Schemata (Cognition); Science and Society; Science Education; Workshops

IDENTIFIERS Cognitive Research; *Cognitive Sciences; National Science Foundation; *Neural Networks

ABSTRACT

Cognitive science, the study of both biological and artificial intelligent systems, is an inherently interdisciplinary activity that embraces aspects of psychology, linguistics, artificial intelligence, neuroscience, engineering, and other behavioral and social sciences. This document reports the results of a workshop designed to provide advice to the National Science Foundation regarding plans for achieving the potential of cognitive science in the 21st century. The report is separated into three sections. The introduction presents three directions in which cognitive science can be developed: (1) the sciences of agents interacting in social and physical environments; (2) the science of individual agents; and (3) the sciences of natural processes. The second section presents brief discussions of a sample of research topics that are active foci of work in cognitive science. Research topics include language, conceptual knowledge, abstract neural networks, and analyses of cognitive activity in social and physical environments. The third section discusses recommendations for enhancing research in the cognitive sciences. Recommendations include: (1) extending the working corps of interdisciplinary cognitive scientists; (2) utilizing technology resource for collaboration with new methodologies and across disciplines; (3) enhancing institutional structures in support of interdisciplinary research; and (4) funding levels. A list of workshop participants is provided. (MDH)

REPORT OF A PLANNING WORKSHOP

To Strengthen American Cognitive Science for the Twenty-First Century

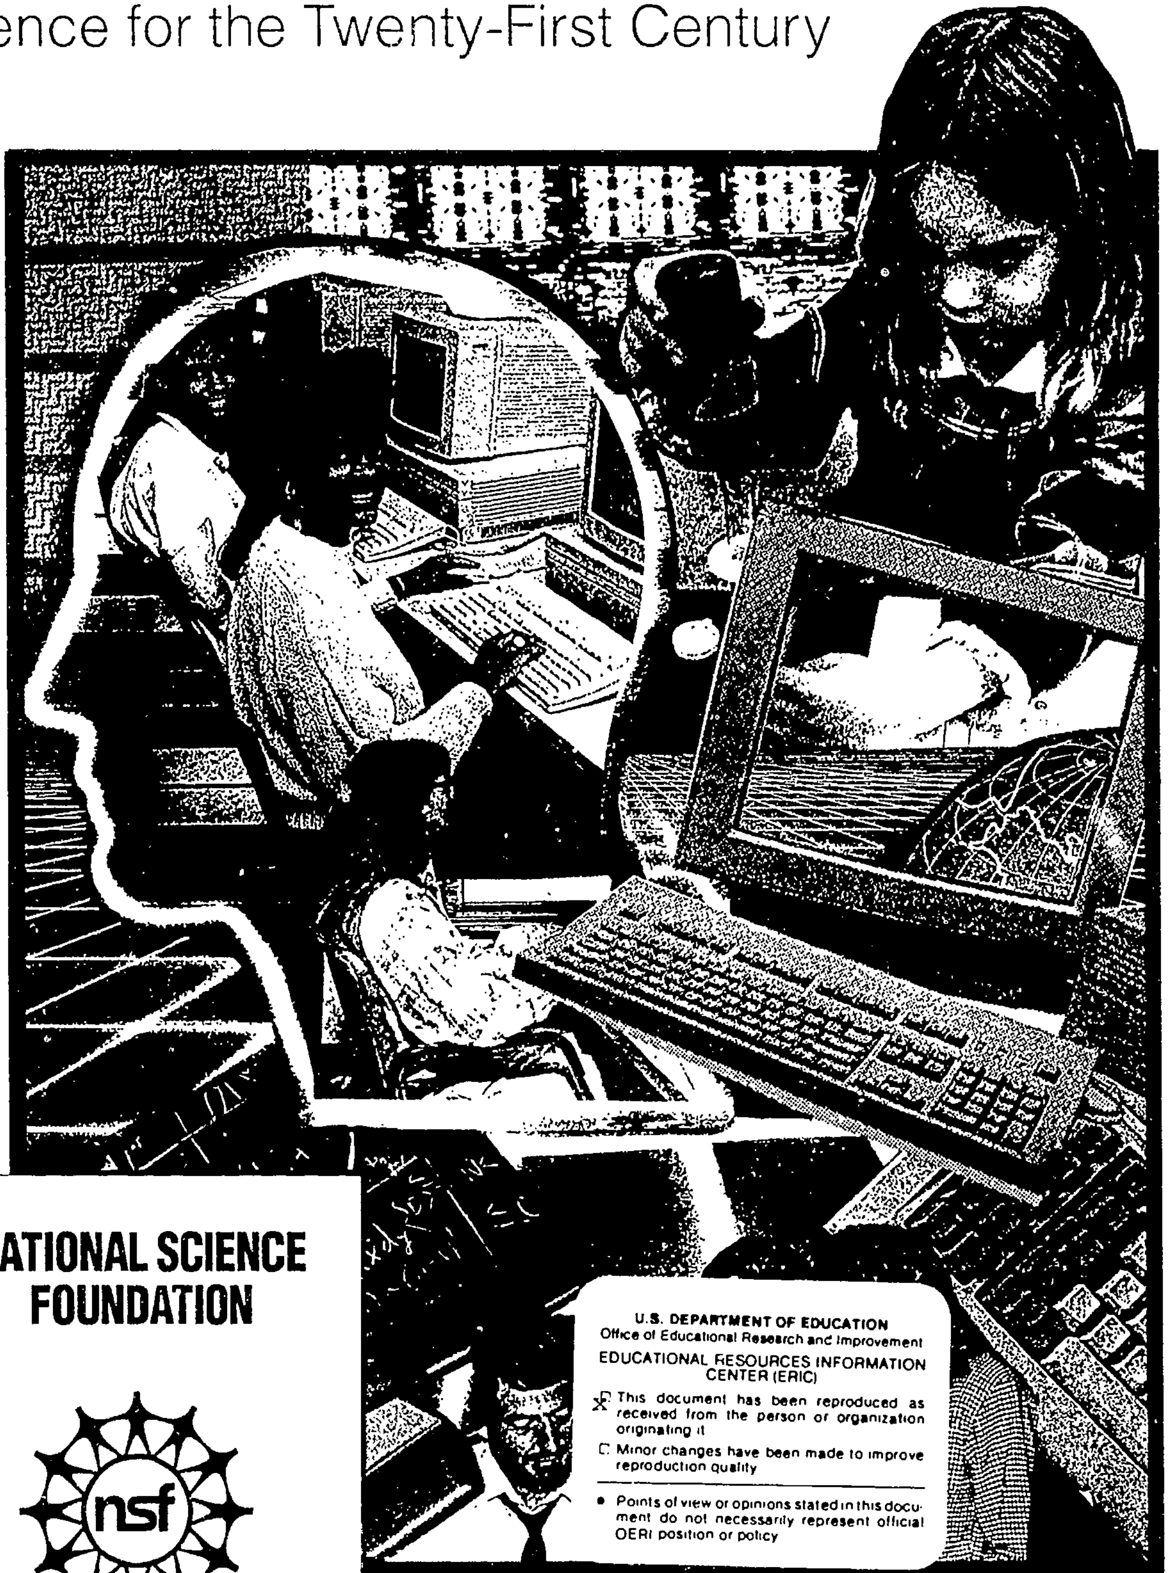

NATIONAL SCIENCE FOUNDATION

nsf

Washington, D.C. 20550

SPONSORED BY:

Directorate for Biological, Behavioral & Social Sciences
Directorate for Computer & Information Science & Engineering
Directorate for Education and Human Resources
Directorate for Engineering

Cognitive Science

COGNITIVE SCIENCE, THE STUDY OF INTELLIGENT systems, both biological and artificial, can provide fundamental new knowledge supporting America's efforts to regain world leadership in the education of its citizens and in industrial competitiveness. In the educational sphere, understanding of human intellectual processes provided by cognitive science research can help us develop methods, strategies, and technologies to help each student learn as much as he or she can. In the industrial sector, cognitive science approaches can help develop intelligent manufacturing systems and can guide computer design. For all these reasons, cognitive science research can have a significant economic impact.

Cognitive science is an inherently interdisciplinary activity, embracing aspects of psychology, linguistics, artificial intelligence, neuroscience, engineering, and other behavioral and social sciences. Some representative foci of cognitive science research include how people make decisions, how memories are stored, how computers can best be constructed and programmed, and how humans can interact optimally in an increasingly computer-oriented environment. Cognitive science extends from basic knowledge about the mental development of children to the study of teaching and learning in the classroom. Knowledge models and network theory have developed as a result of cognitive science research and are now being used in robotics and other industrial applications. Such applications are aiding the economic expansion of both Western and Eastern countries.

Cognitive science is, however, being prevented from reaching its potential for two reasons: lack of resources; and lack of any unified, focussed plan for bringing maturity and coherence to the enterprise. This workshop report recommends first steps toward such a unified plan for achieving the potential of cognitive science.

The findings and views disclosed in this workshop report are those of the workshop participants and are offered as guidance for future discussions of the goals and directions of the cognitive sciences. They do not reflect the views or intentions of the National Science foundation.

Executive Summary

A WORKSHOP TO ADVISE THE NATIONAL SCIENCE Foundation (NSF) regarding plans for a Cognitive Science Initiative was held 20-21 April, 1991, in Washington, DC. The 17 invited participants were charged to identify directions which NSF should take in Fiscal Year 1993 and beyond, by answering the questions:

- What can cognitive science do for the country in meeting the needs of society?
- What should the NSF ideally do to facilitate the development of cognitive science?

GENERAL AREAS OF ACCOMPLISHMENTS AND PROSPECTS

Cognitive science's major accomplishments have been analyses of cognition by individual agents, including language, problem solving, representations of knowledge in memory, perception, and decision making. Prospects for continued progress on these important problems are excellent. At the same time, the field is on the threshold of building new integrative understanding in two important directions. One involves understanding individual cognitive agents as participants in larger social and physical systems. This involves integrating the study of individual cognition with research about social practices and about systems in which natural resources and technology contribute to cognitive functioning. The other integrative direction involves connecting the study of cognitive processes with the growing body of research results in neuroscience.

The participants decided to prepare brief descriptions of research accomplishments and prospects in four topical areas, rather than trying to survey all important areas of research in cognitive science. The report, therefore, presents overviews of current and prospective research about four general topics: (1) language, (2) conceptual knowledge, (3) abstract neural networks, and (4) analyses of cognitive activity in social and physical environments.

RECOMMENDATIONS

Prospects for continued growth of cognitive science require vigorous interdisciplinary research efforts. Continued progress in the study of individual cognition involves psychology, linguistics, computer science, mathematics, philosophy, and education, and progress in broader integrative directions requires interdisciplinary studies involving these fields with the social sciences and neuroscience.

- The committee assigned highest priority for new funding to programs that would extend the working corps of interdisciplinary cognitive scientists, to expand opportunities and incentives to enter the field, and for individuals in the field to acquire capabilities to contribute to interdisciplinary research progress.
- Needs were also identified for development of enabling technologies, including augmentation of shared computational resources and computational technologies in support of new research methodologies.
- Structures in support of interdisciplinary research should include program level activities and interdisciplinary institutes.

The level of funding that NSF provides for support of cognitive science could be increased productively from its present level of approximately $10 million to at least $50 million by the beginning of the Twenty-First Century. This would include about $15 million to extend the working corps of interdisciplinary researchers and to develop and support research resources for interdisciplinary and interlaboratory collaboration, $10-12 million for new program-level activities, and $15-20 million to enable current programs to grow at a rate that is modestly higher than normal.

Workshop Participants

James Greeno (Chair)
School of Education, Stanford University, and the Institute for Research on Learning

James Anderson
Department of Cognitive and Linguistic Sciences, Brown University

Kathryn Bock
Department of Psychology, Michigan State University

Gail Carpenter
Center for Adaptive Systems, Boston University

John Carroll
Sloan School of Management, Massachusetts Institute of Technology

Allan Collins
Bolt Beranek and Newman Inc., and Institute for the Learning Sciences, Northwestern University

Gary Dell
Department of Psychology, University of Illinois, Champaign

Alan Gevins
EEG Systems LaboratoryRichard Granger, Department of Information and Computer Science, University of California, Irvine

Max Henrion
Rockwell Research Laboratories

Edwin Hutchins
Department of Cognitive Science, University of California, San Diego

William Hutchinson
BehavHeuristics

Aravind Joshi
Department of Computer and Information Science, University of Pennsylvania

Frank Keil
Department of Psychology, Cornell University

Jill Larkin
Department of Psychology, Carnegie-Mellon University

Alan Lesgold
Learning Research and Development Center, University of Pittsburgh

Paul Smolensky
Department of Computer Science, University of Colorado

1. Introduction

WE HAVE ENTERED AN ERA WHERE PEOPLE THINK and learn in increasingly computational environments, and the development of technology is far outpacing our understanding of how humans best function in such environments. Nevertheless, we now have in the cognitive sciences the theoretical and technological tools to build a coherent understanding of human cognition and the design of computational environments that facilitate productive human cognition.

During the second half of the Twentieth Century, cognitive science has developed as an interdisciplinary science, bringing resources of artificial intelligence, linguistics, psychology, philosophy, and educational research to bear on the scientific study of information, intelligence, and learning. The development of cognitive science has been reciprocally interdependent with the development of electronic information technologies in these decades. Advanced information technologies and computational formalisms have enabled theoretical investigations in cognitive science that explore hypotheses of realistic degrees of complexity about the organization of information and processing of language, memory, and problem solving. Important theoretical and empirical progress has been accomplished by viewing the human mind as an information-processing system. And the design of complex information systems in business and education has been informed by the general principles of information-processing developed in studies of human cognition and artificial intelligence, as well as by specific studies of properties of intelligent human-computer systems.

It is not likely that the problems of understanding human cognition or the design of computational environments can be solved within the perspective of any one of the traditional disciplines. Instead, we need analyses that map between and integrate scientific understanding at the different levels that address how social interaction of multiple agents is related to and constrained by the cognitive functioning of individuals, and how that in turn is related to and constrained by the interaction of neural elements.

This report focuses on prospects for accelerated growth of American cognitive science during the remaining years of the Twentieth Century, to maintain U.S. scientific leadership in this field and to ensure that cognitive science will have strong capabilities for continuing to inform and collaborate in development of information technology that leads the world.

Figure 1 depicts three directions in which cognitive science can be developed strongly in the next several years.

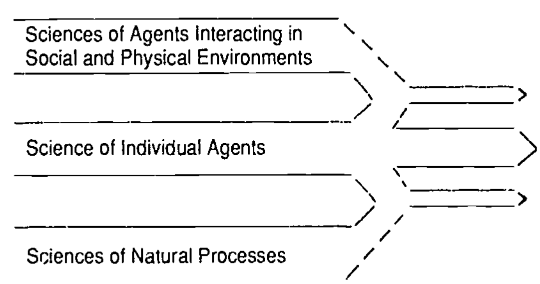

Figure 1. Three vectors of potential growth.

Cognitive science, as it has been developed, combines resources of cognitive psychology, linguistics, and artificial intelligence in analyses of cognitive processes of individual agents. Three prospects that we present will extend the scientific capabilities and contributions of the field by (1) understanding individual cognition as a component of larger systems, (2) understanding relations between processes of cognition and the neural substrate, and (3) extending the utilization of several interdisciplinary vectors within individual cognitive science that have proven validity and potential but have only begun to be exploited.

COGNITION IN SOCIAL AND PHYSICAL ENVIRONMENTS

Individual agents function as parts of larger systems. Human individuals participate in social groups. The situations in which people act include physical and technological resources and impediments that affect cognitive activities in critical ways.

While understanding of cognition in contexts has always been a goal of cognitive science, recent developments have made significant progress toward that goal, and accelerated development of these directions would greatly strengthen the scientific and practical contributions of the field. One promising development involves the study of cognitive activities of groups and individuals in everyday settings of work and other social interactions, using methods of cognitive anthropology and videotaped records from which group processes of communication, reasoning, and decision making can be analyzed in detail. Another promising direction is the attention being given to distributed intelligent systems, focussing on the differing roles that can be played by individuals and groups of persons and by computational systems that are designed to serve as resources for human agents.

RELATIONS TO NEURAL PROCESSES

Remarkable advances in neuroscience have been made in recent years, providing challenging opportunities for developing significant coordination of scientific knowledge between psychological analyses of processes such as pattern recognition and memory and analyses of the neural structures that constitute the biological medium of cognition. Important theoretical progress on this task is being made using abstract neural networks (ANNs), and further important progress would result from funding that would encourage the needed interdisciplinary research.

INTERDISCIPLINARY ANALYSES OF INDIVIDUAL COGNITIVE PROCESSES

Important progress has been made in developing methods for analyzing individual cognition using artificial-intelligence tools for formulating hypotheses, conducting empirical and theoretical analyses of human-computer systems, and designing artificial information systems that augment human cognitive capabilities. Major advances in the understanding of human and distributed cognition would be achieved through full utilization of these methods, which would be facilitated by augmented funding of the interdisciplinary research that is needed.

2. Illustrative Topics

IN THIS SECTION WE PRESENT BRIEF DISCUSSIONS of a sample of research topics that are active foci of work in cognitive science. The first two examples illustrate research accomplishments and prospects in the sciences of individual agents. The third example involves the relation of cognition and neuroscience, and the final example presents some prospects involving cognition in social and physical environments.

LANGUAGE

One hallmark of an intelligent system is the ability of use language. The past twenty years have seen the development of a theory of the nature of language and how it is learned, comprehended, and produced. The results of these investigations have led to advances in language processing technology that have the potential to revolutionize the way that people communicate with machines and with each other. For example, cognitive science research has enabled the development of systems that can recognize and produce speech and, even more significantly, can understand and produce meaningful sentences. In addition to these technological advances, this research has led to powerful theories of the structure of natural language. These theories illuminate the way that humans learn language and, more generally, offer insight into the nature of the human mind.

Some accomplishments

Question-answering systems. It is now possible to ask spoken questions of a machine and receive spoken answers. Such a machine must recognize spoken words, determine the linguistic structure of the question (parsing), and the semantic interpretation of that structure, search its data base for the answer, put that answer in the form of a sentence, and generate speech. Progress on each of these steps has been made possible by groups of researchers who have combined linguistic theory and computational principles with a growing understanding of how people produce and understand language and speech.

Machine translation. One profound barrier to human communication and commerce is the diversity of languages used by the peoples of the world. However, after many painful beginnings, language translation by machine is enjoying some success with the development of systems that enable rudimentary translations from one language to another. Such systems are currently used by various governments and their agencies to translate thousands of pages of text per day, with usable results. These systems have developed from collaborations between computer scientists and linguists. They are made possible, in part, by new understandings of languages as integrated representational systems and of the role of the communicative context in language understanding. New developments in this area can be informed by ongoing work in how the human users of languages understand them, speak them, and use them in conversation and collaboration.

Language acquisition. The study of language acquisition was one of the first targets of the interdisciplinary collaboration among cognitive scientists, including linguists and cognitive psychologists. Before about 1960, language was widely viewed as a catalog of elements and structures that children could simply memorize. With the emergence of the view that major components of language are novel, productive, and not exhaustively available to children, the problem of language acquisition acquired unprecedented interest. That new interest has led to discoveries about the kinds of things that children are prepared to learn by virtue of their biological endowments, the kinds of things they learn without explicit teaching, and, just as important, the kinds of things that they do not learn. As a result of these discoveries, we now have a very different view of language itself, a view that has radically changed our ideas about the languages of the deaf, about the biological substrates of language, and about the environments in which language is learned. We are in a position to go beyond these discoveries with the emergence of powerful computational models of the acquisition process, the development of new methodologies for the study of infants, and the invention of technologies that reveal the neurobiological structures and processes that underlie language acquisition and use.

Some prospects

Although cognitive science has enabled us to achieve the success mentioned above, it is clear that systems for language processing developed so far are limited in their coverage of language phenomena. The accuracy of performance is not satisfactory. The systems are brittle because they are unable to handle variations in the input. They are limited in their ability to deal with pragmatic aspects of interaction, both in single user and multi-user interactions. They do not exploit the role of prosody (i.e., speech, rhythm, and intonation) in disambiguation and in conveying pragmatic information. The systems are essentially handcrafted in the sense that syntactic, semantic, and pragmatic knowledge are built in. Finally, the systems have no learning capabilities.

In order to overcome these deficiencies, major thrusts in cognitive science are required. Theoretical contributions in these areas will lead to computational models that will be the basis for the next generation of natural language/speech technology. This will promote truly flexible and helpful human-machine interfaces, to be sure, in conjunction with interfaces in other modalities, such as graphics and animation. Some of the major thrusts are as follows:

- We must develop an understanding of how prosody interfaces with syntax, semantics, and pragmatics, and how prosodic information can be extracted from speech input and added to speech output.
- Handcrafting of grammar in the end does not work; the resulting systems are too brittle. We need to find out how to automatically acquire grammar from very large corpora, on the order of 10^8 or even 10^{10} words. Availability of such corpora for researchers is an important issue in the resources for the field.
- It is becoming increasingly clear that statistical information (e.g., the distribution of verbs and their arguments, the distribution of construction, etc.) is critical both in making systems robust and in supporting the automatic acquisition of grammar.
- At the theoretical level, it is crucial to develop theories that will integrate structural and statistical information in meaningful ways. Developments in neural net theory are starting to achieve this integration. At the same time, cognitive modeling of language provides a domain in which neural net theorists can begin the effort to incorporate structural information into their framework. Thus, we have a wonderful opportunity to integrate symbolic and nonsymbolic systems in a very specific and well understood domain.
- We must develop a deeper understanding of discourse structure, both in the single- and multi-agent environment (e.g., the role of planning in discourse).

We have described the above thrusts in the monolingual situation. However, they are applicable in the bilingual or multilingual environment, with the added component of how syntactic, semantic, and pragmatic information is transferred from one language system to another. These theoretical developments will be crucial to advancing the machine translation technology, a technology that is attracting attention again and needs some theoretical underpinning.

CONCEPTUAL KNOWLEDGE

How are concepts acquired and used? How can we make it possible for humans to use ever larger amounts of information? How can we better teach the formal conceptual systems of mathematics and science to human learners? These questions are not only scientifically fascinating but also constitute pressing challenges to our society.

Traditionally, concepts were seen as isolated mental structures consisting of lists of singly necessary and jointly sufficient features. Thus the "dog concept" consisted of a list of features like "four legs" and "furry." But, of course, people readily recognize a dog lacking a hind leg as a dog. Early experimentation turned to artificial concepts (e.g., "small blue triangle") which could be characterized by lists of features.

A new phase of research on concepts (beginning about 1975) first reflected the then-new discipline of cognitive science. Experimental work indicated that the number of features matched by a concept exemplar corresponded to how "good" or typical the exemplar is. Robins and canaries (small, flies, sings, feathered, beaked) are better examples of the concept "bird" than is a chicken or a penguin. Simultaneously, computational methods were developed for representing objects defined by probabilistic or varied clusters of features, and these techniques were quickly used to build computational models relating and predicting data on both concept application and concept learning. These models have accounted for some nonintuitive observations.

The development of technology in the form of expert systems has required and contributed to the scientific understanding of concepts. Knowledge representation systems in artificial intelligence have provided a setting for designing organizations of conceptual structures that are computationally tractable and useful in domains of practical work, such as medical diagnosis and engineering design.

As is common in science, these advances in understanding of conceptual knowledge also clarified what was still to be understood. For example, these models failed to capture causal knowledge and beliefs that were not repeatedly observed in human conceptual knowledge. For example, although people associate "curvedness" comparably to boomerangs and bananas, the two cases evoke strongly different beliefs about the causal functionality of curvedness. A straight banana is still a fine banana. A straight boomerang violates peoples' beliefs about requirements for a functional boomerang.

This theme of causal connections in conceptual knowledge and the embedding of concepts in large belief structures receive complementary support from experimental psychology, cross-cultural studies, and characterization of scientific thought.

Constraints on the Models of Cognitive Science

As illustrated by the preceding description of our deepening scientific knowledge of concepts in cognition, theories in cognitive science are constrained by the following factors:

- *Computability*. Theories expressed verbally may have gaps in logic that are quickly revealed when the postulated mechanisms are reified in computer models that must simulate the behavior of humans.
- *Learnability*. There can be many ways to describe an intact, fully learned knowledge structure. There are far fewer ways of describing and producing computational models that start with plausible initial states and construct the knowledge in a manner consistent with that observed in humans.
- *Unification*. Increasingly, cognitive scientists reject micro-models that describe one phenomenon or part of a phenomenon. The field pushes toward models that can account for not only internal mental structures for concepts, but also use of language to describe this connection of symbolic mental structures with measurements of neurological function.
- *Social construction and use of knowledge*. Knowledge is communicated and shared. As we have come to understand much about individual knowledge, there is increasing interest in unifying and extending this understanding to the process by which social groups come to construct consistent knowledge structures about which they can communicate with language.

Formal Concepts

In addition to natural concepts, there are artificial concepts occurring in formal domains such as mathematics, logic, and science. Understanding the mental mechanisms by which these concepts are learned and used is central both to understanding the nature of scientific reasoning and the processes of learning and teaching science.

Traditionally, both physical and social scientists have believed that people are innately logical, and that formally defined concepts such as "acceleration" or "multiplication" should be learned easily through careful attention to clear explanation, together with adequate practice. However, many well-replicated studies have shown that students in wide varieties of high schools and universities score at a level of about 40% correct on typical tests of mechanics concepts both before and after taking a course in mechanics. In other studies, misunderstandings of probabilistic reasoning robustly survive statistics courses.

Starting in the early 1970s, cognitive scientists began building computational models of reasoning with formal concepts. Early models were based on general reasoning methods applied to domain-specific knowledge bases, and emphasized the importance of decomposing problems into subproblems. Continued observation of human reasoners, however, revealed behavior that could not be modeled with general methods, and initiated still-continuing research elucidating the psychological mechanisms of reasoning in specific domains including, in particular, physics and medicine.

Models began to clarify why formal reasoning is difficult. For example, many physics students believe that when the velocity of an object is zero the acceleration must also be zero. This is a plausible inference based on everyday belief systems, in which typically zero acceleration (e.g., for a car at rest) results in zero velocity. Avoiding such misconceptions and appropriately using the physics concept of acceleration requires separately considering the velocity at one time and the velocity shortly later, and then inferring that if the velocity changes, then acceleration (which is the rate of change of velocity) is nonzero. This task is not easy—even well-trained physicists easily make mistakes in some problem situations.

Most recent work on formal reasoning has focused on how a rich body of redundant knowledge allows a reasoner to cross-check the accuracy of formal procedures. Such knowledge consists of processes for reasoning qualitatively, or with mental models, so as to place constraints on the properties that a formal result must have. Also, the diagrams and highly formatted equations used in science have properties that facilitate accurate reasoning. For example, equations are typically written with the order of their subparts based on either a convention (e.g., alphabetical) or tied to meaningful elements of the equation (e.g., dimensionless ratios).

These models of formal conceptual knowledge have had profound and exciting applications in science and mathematics education. Computer-implemented tutors have regularly improved performance substantially, sometimes by a full letter grade or as much as three standard deviations and/or produced somewhat smaller improvements in learning with savings in learning time as large as a factor of two. In tests of ability to use the concepts of force and acceleration, students performing at the 40% level improved to the 80% level after instruction based on the kinds of research discussed above. A computer tutor in geometry produced not only this kind of improved learning, but also immense changes in motivation. High school students who always came late to class now arrived early every class period, rushing to work with the geometry tutor.

Other Benefits of Growing Understanding of Conceptual Knowledge

The previous discussion of formal concepts suggests one set of current and potential benefits of cognitive science research on conceptual knowledge. The following is a brief list of other practical benefits on which applied research is underway.

- Representation of concepts in data bases, in computer interfaces, and in information management systems.

- Language understanding and production by computers.
- Complex feature detection; for example, computer processing of images of assembly-line items in order to detect flaws.
- Robotics, including studies of navigation, planning, and processing visual and verbal input.
- Expert systems, including systems in which specific knowledge of human experts is encoded in computer systems that can then perform some tasks (e.g., configuring computer parts) better than humans.

ABSTRACT NEURAL NETWORKS

In this section we describe what is emerging as a central methodology in cognitive science: the mathematical and computer systems we shall refer to as abstract neural networks (ANNs). ANNs in cognitive science lie at the crossroads of a multidisciplinary scientific revolution that is integrating the sciences of brain and behavior and creating new technologies.

Neural networks provide the mathematical language and techniques that make possible this multidisciplinary synthesis. With this methodology, principled integrative theories lead to predictions that generate critical experimental tests. Remarkably, ANN models derived from analysis of individual behavior serve both as a model of the underlying neural machinery and as a blueprint for technology transfer. In the neural domain, the model illuminates the function of individual neural elements. Model components are refined by comparison with existing neural data; predictions emerge from analysis of the functional role of individual components.

In the direction of technology transfer, an abstract neural network, specified as a system of nonlinear differential equations, becomes a design specification that is translated by an engineer into working hardware and software. Network derivation as a cognitive machine allows the engineer to appreciate its function and to map that function onto specific application domains. Thousands of engineers are now engaged in development projects that use ANNs in a computational role that complements traditional technology. One such application area is robotic design. A robot needs to integrate input from one or more sensors. Neural models of binocular vision, adaptive eye movement, pattern recognition, and attention are now used in working robots. Similarly, mechanisms of robotic action employ neural network designs derived from analysis of planning, conditioning, prediction, and motor control.

ANN is rapidly building an impressive body of technical results

The past few years of ANN research have yielded a wealth of mathematical results built on foundations laid by work carried out over the past several decades. Many of these results can be discussed under the three rubrics: theory of computation, theory of dynamical systems, and theory of probability and statistics.

Formal theories concerning what can be computed by ANNs, under various formal definitions of ANN and the computational complexity of ANN algorithms, are now taking shape thanks to a rapidly growing body of theorems. The new field of computational learning theory, intimately connected with the theory of learning in ANNs, has formally defined notions of learnability and decided the learnability of many important classes of functions. These recent theoretical results are already driving the development of more powerful learning algorithms.

Regarding an ANN as a dynamical system, like those traditionally studied in mathematical physics, has produced a complementary body of formal results. These typically concern the convergence properties of the dynamical behavior of ANNs, and analysis of other global system properties such as the structure of attractors and chaotic behavior. These results have been instrumental in the design of ANNs and in their application.

Finally, an increasingly important view of ANNs is as complex probabilistic inference systems. Many results documenting the power and limits of ANN learning algorithms to develop accurate statistical models of their training data have been proved, and used both for the analysis of current ANN algorithms and the development of more statistically optimal procedures.

In each of these areas, the relation between ANN and the other disciplines is bidirectional. While mathematical techniques continue to be effectively imported into ANN, the field has contributed much in return; new models of computation, new classes of dynamical systems, and new statistical techniques have all resulted from the interaction. The vigor with which ANN problems have been attacked by computer scientists, mathematical physicists, and statisticians attests to the significance of the novel ideas that ANNs have contributed to these allied fields.

ANN has enormous consequences both for fundamental scientific research and for technology.

Scientific research. The mathematical techniques being developed in ANN can be applied in many ways not only within cognitive science but also in other scientific disciplines. For example, since many ANN techniques can be regarded as powerful statistical methods, all the problems for which statistics is a valuable tool are potential applications for ANN techniques.

Within cognitive science, ANN techniques have been successfully applied within and between all three levels of analysis. Within the neural level proper, ANNs form the basis of biological neural models, in which detailed properties of actual neural circuits are modeled mathematically. Within the individual agent level, ANNs are used to model networks of mental processes, showing how intelligent behavior in an agent can emerge from the interaction of a large number of simple numerical processes. Within the sociocultural level, ANNs, like simpler statistical techniques, are used to model the interaction of economic variables (e.g., predicting future values of economic indicators from present values), and model the interaction of multiple agents in a social group.

Even more critically for cognitive science, ANNs serve to model the relations between the levels. The raison d'etre of ANNs has historically been the explanation of how interacting neural processes can give rise to mental phenomena at the individual agent level. Ideas developed through the study of ANNs about how simple neuron-like processing elements can organize to achieve complex high-level cognitive functions provide the current best hope for understanding neural function and organization. And ANN insights into how cognitive processes can be realized in networks of abstract neural elements has generated many new mathematical theories of cognitive function.

Similarly, the collective behavior of groups of individual agents, each modeled as an ANN, has been used to model important phenomena that crucially involve interactions between the individual-agent and socio-cultural levels; e.g., the Baldwin effect, by which learning in individual organisms dramatically increases the rate of "learning," i.e. evolution in the species. The potential for many equally fruitful applications of this type is great.

Technological implications of ANNs. A striking fact about the mathematics of ANN is that, in addition to its important uses in pure research, it has wide applicability to technological problems. ANN techniques lend themselves well to unusually rapid technology transfer: transformation of pure research results into developed tools and products for well-defined tasks, in applications ranging from sensors and robotics to organizational decision making.

Cognitive science research at the intersection of neuroscience, computation, mathematics, and psychology has led to fundamental findings about the nature of learning, and ANNs embodying these findings comprise adaptable systems that change their behavior in response to the situational environment. This allows the creation of applications systems that improve their behavior with use, and can therefore be trained on sample problems to perform desired input-output behavior that often turns out to be superior to the behavior of competing systems preprogrammed to perform the same tasks. Examples include visual imaging processing systems that can be trained to recognize specific categories of objects, e.g., particular vehicles, of specified shapes, traveling at particular speeds, or particular flaws in manufactured objects on an assembly line. Visual image processing systems have been taught, via ANN learning rules, to recognize handwritten characters and thereby enable automated sorting processes for use in, e.g., the postal service. Additional specific applications include the discovery of novel classifications, database organization, manufacturing, scientific computing, biological analysis, adaptive robotics and control, and novel statistical packages for analysis of complex multidimensional data; adaptive data analysis of this latter sort has been applied successfully in automated recognition of sonar data and speech signals.

Most of the ANN learning rules to date have emerged from mathematical models of rules underlying synaptic plasticity in brain circuitry; the methods by which cell-cell connections in the brain are changed during learning. Recent discoveries in neuroscience have uncovered novel physiological processes involved in the induction of forms of synaptic plasticity which appear to underlie rapid learning of the sort typified by human recognition memory, in which brief exposure to a new input can be retained, generalized, and subsequently recognized in a long-lasting way. Computational modeling of these identified forms of neurobiological learning rules is resulting in new ANN learning algorithms with the potential capability for scaling up to very large networks.

A less immediate but equally important application area of ANNs is to computer design. Almost everyone now believes that the next advance in computation will come from massive parallelism. The Japanese have now launched the "6th generation" project: The development of massively parallel computers. Some of the most prominent Japanese neural network researchers are advisers to the 6th generation project. The Japanese are well aware of the great potential of cognitive science for organizing parallel computer software and architecture and suggesting problems for them.

There is also great interest in ANNs within the engineering community. A number of different laboratories, including such large corporations as Intel, Siemens, AT&T Bell Laboratories, and Bellcore have fabricated and tested ANNs in VLSI. McDonnell-Douglas engineers have chosen ANN learning algorithms for controlling the "space plane" they are designing. The complex problems involved have quantifiable objectives, such as keeping the aircraft stable, which require many interrelated actions to achieve. This very practical technology is based directly on ANN learning algorithms that derived from cognitive science models of conditioning in humans and laboratory animals.

Frontiers and future

There are a number of major intellectual challenges ahead for the field of abstract neural networks. Perhaps the most important involves the kinds of computation that they do. At present neural networks and the more classical symbolic models of cognition, characteristic also of traditional artificial intelligence, are very different. There is no doubt that symbol-processing accounts of cognition do capture some of the important aspects of mental computation. Symbol processing, for certain kinds of problems, is fast, flexible, and avoids certain serious problems found in ANNs. Several groups are now trying to combine the flexibility and power of symbol processing with the parallelism and biological plausibility of ANNs. A major problem in the next decades will be to integrate some symbol processing with networks to make hybrid systems that effectively combine the virtue of symbolic structure processing with those of ANN statistical computation. Current exploratory research along these lines seeks to understand, for example, the emergence of knowledge in learning ANNs that can be expressed in symbolic rules and explanations.

Development of ANN-based theories will strengthen the scientific relations between cognition and the growing knowledge of brain function. Technological advances during the last decade have resulted in the first detailed images of the metabolic and electrical processes of thinking. This new generation of PET and EEG scanners, which embody these technologies, were produced by interdisciplinary teams of physical, psychological, and biomedical scientists and engineers working together to integrate advances from many disciplines. Further development and application of these brain scanning technologies during the coming decade promises to contribute to understanding of fundamental questions about the neural mechanisms underlying memory, learning, and other cognitive processes.

ANALYSES OF COGNITIVE ACTIVITY IN SOCIAL AND PHYSICAL ENVIRONMENTS

Cognitive science seeks to understand the nature of human intelligence, how knowledge is communicated, learned, represented, and used by humans, and how aspects of intelligent activity can be produced by computational systems.

Examples of notable achievements in understanding human intelligence have included the study of human decision making in reasoning under uncertainty, the organization of information in memory, problem solving and planning, and individual learning. This work in human cognition has proceeded in tandem with corresponding developments of formal theories of intelligent behavior, such as decision theory, search theory and automated problem solving, formal theories of knowledge representation, and the theory of complexity and learnability. Much of the success of cognitive science has resulted from the interaction of descriptive models based on observation of human behavior and the development of formal models of rational agents.

This scientific work has already led to significant successes in practical applications. Examples include decision analysis as a successful practice for guiding high-stakes business and government decisions, expert and knowledge-based systems, advanced representations and interfaces for information retrieval, improved technologies for human-computer interaction, including windows, icons, and other graphics-based interface techniques, and intelligent tutoring systems.

In recent years, there has been a growing effort to understand individual cognitive agents as participants in larger social and physical systems. Regarding social interactions, studies of work practices are leading to characterizations of decision-making, reasoning, and memory that occur in group activity. Regarding physical and computational systems, it is becoming clear that computer-based systems are most effective when they play a complementary role supporting and augmenting the human decision maker or planner. Many of the most important applications of this technology, such as software engineering, computer-aided design, and most human decision making, involve groups of people, not just individuals. Consequently, recent research is starting to look at computer support for cooperative work. Electronic mail, shared workplaces, vision support technology, and modular structures, are all computational techniques that can sup-port group processes in these endeavors. However, unlike individual decision making, there is not a well-developed theory of group process available to support the engineering work. We see this as a major theoretical challenge with important practical implications facing cognitive scientists concerned with the social level of analysis.

The fact that human cognition always takes place in material and social contexts has long been recognized. In recent years, however, a view is being developed that the very nature of the cognition of individual agents depends on an interaction with a complex environment. This realization is reflected in changes that are now occurring in research on various kinds of cognitive processes.

Decision making

The initial attempts to apply computer technology to the problem of decision making led to the creation of expert systems. These are computer systems that attempt to simulate the kind of reasoning that experts do, and they are intended to replace the human expert. More recently, however, attention has shifted to the development of decision support systems. This reflects a realization that decisions are quite often made collaboratively rather than individually. The design of such systems requires that integration of knowledge from computer science, organizational science and decision theory.

Planning and problem solving

The emphasis in planning research is shifting from a model of an agent acting in a static problem space to a conception of a continuous interaction between planner or problem solver and a dynamic world. This research, including development of systems that engage in reactive planning, draws together knowledge from the psychology of action as well as from computer science.

Memory

Increasing attention is being given to the interaction between internal and external memory systems, including the cognitive impact of fundamental symbolic technologies such as literacy. This work is mapping new territory in the world of cognitive phenomena looking at the way that even simple technologies transform human cognitive capabilities. This knowledge will be absolutely essential to an understanding of the cognitive consequences of more complex "thought technologies" that are a goal of some computer-science research.

Learning

Increasing attention is being given to ways that learning happens in everyday settings, including activities that are not normally considered as learning settings, including both work and play activities. This research focuses on understanding of ways that learners participate in communities of practice. This work is informed by, and has direct application in the domain of on-the-job learning. To really understand learning will require the integration of findings and methods from many current disciplines, from neural models of learning that are needed to account for phenomena like implicit learning in individuals to the social science of communities in which learning activity is situated.

3. Recommendations for Enhancing Research

EXTENDING THE WORKING CORPS OF INTERDISCIPLINARY COGNITIVE SCIENTISTS

The highest priority is to extend the working corps of cognitive scientists with interdisciplinary capabilities. The field of cognitive science has developed methods and problems for which many additional research personnel could be deployed profitably, and significant progress in the field now depends on recruiting and training sufficient researchers to utilize the resources for scientific activity that have been developed. Programs should be established that will expand opportunities for individuals to enter the field and for individuals in the field to acquire capabilities to contribute to interdisciplinary research progress.

At the present time, doing cognitive science can be likened to speaking pidgin. A pidgin is a simplified communication system that springs up when speakers of many different languages are forced to communicate. It mixes words from the languages and has few rules. Consequently, each speaker of the pidgin uses structures from her or his own native language when forming pidgin sentences. Cognitive science research and theories may include concepts from more than one discipline, but they strongly reveal the home discipline of the originators.

A linguistic pidgin either dies out or it evolves suddenly into a creole, a new and genuine language with all the expressive power and systematicity of any human language. The process of creolization has a remarkable mechanism: If a single generation of children grows up hearing the pidgin as their primary linguistic input, the children, in a sense, create the creole.

We need, then, to create the conditions in which cognitive science can become a creole, a new coherent discipline. Many indications are that the needed creolization is happening, and, just as the analogy would suggest, it is through the students, those who have been exposed to the interdisciplinary pidgin.

The present situation in cognitive science is similar in some ways to the situation in computer science in the early 1980's, when there were only three institutions that had the necessary infrastructure and equipment to continue to do first-rate research in computer science, especially in experimental computer science. There were researchers in many other institutions who were doing excellent work but it was recognized that unless 20-25 departments were brought into the top level as far as the environment for research was concerned, computer science in the U.S. would fall behind very rapidly. Experimental computer science has kept pace with the marketplace as a result of NSF initiative.

Predoctoral and postdoctoral training

The traditional mechanism of training grants for predoctoral and postdoctoral fellows seems ideally suited to the present needs of cognitive science. The special needs of interdisciplinary training often require that students spend more-than-average time in gaining research experience in multiple laboratories, which is difficult to support with research assistantships committed to specific projects.

Extending interdisciplinary capabilities among current researchers

Researchers already working in one of the disciplines of cognitive science should be encouraged to acquire cross-disciplinary experience and training. Support for visits to laboratories and enabling interdisciplinary working groups at institutes such as the Center for Advanced Study in the Behavioral Sciences, as well as at universities, would be important resources. It would also be appropriate to support remote interdisciplinary collaborations in which "virtual working groups" could be established.

Undergraduate preparation

Because cognitive science is just reaching critical mass, its infrastructure lacks several key components. Among these is a communally shared sense about the undergraduate education that is appropriate for students planning graduate work in cognitive science. Some of the disciplines that contribute to cognitive science require virtually no preparation in mathematics and the formalisms of computer science; some require minimal preparation in the biological and neural sciences; some require no exposure to the social science methodologies that have allowed cognitive science's formalisms to connect with data. The result is that many students desiring to begin graduate work in cognitive science face a choice between spending a couple of years doing more undergraduate work or finding some of the most important parts of their field inaccessible.

NSF should support a study group to develop recommendations for programs of undergraduate study that would constitute preparation for graduate work in cognitive science, and encourage submission of proposals for curriculum development in cognitive science at the undergraduate level.

ENABLING TECHNOLOGIES: RESOURCES FOR COLLABORATION WITH NEW METHODOLOGIES AND ACROSS DISCIPLINES

Multidisciplinary research in cognitive science requires a wide range of tools, especially computational and information-processing tools. Often, one person's re search project becomes the tool for important work in a different area. For example, language generation and knowledge base tools from linguistics and computer science can be very useful in facilitating richer and more clearly recorded cognitive interactions among humans and between humans and machines. computer-controlled and monitored experimentation provides a level of replicability potential that is very important. In research on natural language and speech it is now recognized that one must have large bodies of actual texts and transcripts of speech in actual interactions, annotated in a variety of ways. Large-scale corpora are needed in a form that can be shared by the entire research community, including industrial laboratories. NSF and DARPA are taking initiatives in supporting resources in this area.

Shared computational resources in cognitive science

Computational systems such as SOAR, ACT, BOXER, and others can provide a basis for collaborative and coordinated theoretical and developmental work at multiple sites. Yet another area where tools are needed is for multi-person interactive problem solving studies over local area networks (where each person sits in front of some sort of work station and communicates with others partly through a network).

In general, a small number of core cognitive science research centers have prototypes of such tools, but those prototypes are not very portable. More often than not, they are accessible only to their developers, those peoples' students, and close acquaintances who can keep asking for help in using undocumented, brittle and buggy software.

To develop such resources requires effort beyond those needed for specific research projects. Development of the systems has to extend to making them usable in sites distant from their development and for maintaining the systems, consulting with users on trouble, and advising local groups in their efforts to utilize or expand functions of the systems. Funds to cover the cost of good software engineering to produce portable, user-friendly versions of basic tools are extremely important.

Video Technology

Another critical developing tool area is video technology and especially the ability to annotate, sort, and catalog video segments. Computational technologies are emerging for analyzing video records of cognitive and social activities, and for managing the large and diverse collections of records, commentaries, and data that are essential to achieve an adequately rich and accurate interpretation of social and cognitive activity. The existing computational systems are useful enough to provide valuable resources in the laboratories that have developed them, and resources for making them sharable and maintaining their use in networks for researchers are needed.

In addition, more powerful systems are needed, which require solution of significant problems in computer science involving coordination and management of data represented in multiple forms and to provide multiple indexing for data records, some of which are in real time (video or audio) and others of which are discrete annotations of the real-time records.

ENHANCING INSTITUTIONAL STRUCTURES IN SUPPORT OF INTERDISCIPLINARY RESEARCH

Funding activities at the level of programs are needed to provide review of proposals that is not only competent but also committed to the advancement of interdisciplinary research. At least three program-level activities should be initiated and charged with review of proposals that require interdisciplinary foci. One of these would focus on problems that relate cognition with neural processes; one would focus on problems that relate individual cognition to its social and material contexts; and one would focus on individual cognition using methods drawn from computer science, linguistics, and psychology. The scientists recruited to constitute review panels should be individuals who are particularly committed to the advancement of interdisciplinary dimensions of research.

Especially for the new areas of research that require interdisciplinary capabilities, there is a need for resources to develop perspectives that have been brought into the problem area. This requires significant commitments of time by researchers dedicated to forming the required syntheses, in which they need to interact across disciplinary boundaries for significant periods, working on specific scientific problems. Institutes need to be established for this kind of work.

LEVELS OF FUNDING

NSF funding for research in cognitive science could be increased productively from its present level of approximately $10 million to at least $50 million by the beginning of the 21st Century. A program of predoctoral and postdoctoral fellowships and support for extending the interdisciplinary capabilities of researchers in cognitive science should be built to a level of about $10 million per year over the next three to four years. At least $1 million per year should be allocated to the development of undergraduate curricula and teaching in cognitive science. $3 to $4 million per year should be allocated to the development and support of research resources that are needed for interdisciplinary collaboration and interlaboratory collaboration. The three proposed new program-level activities should quickly reach a level of funding of about $3 to $4 million each, per year. And existing programs that support cognitive science should grow at an above-normal rate; if these programs were to double in real dollars — a fairly modest projection— their allocations would require additional funds of $15 to $20 million.

ELECTRONIC DISSEMINATION

You can get information fast through STIS (Science and Technology Information System), NSF's online publishing system, describe in NSF 91-10, the "STIS flyer."

To get a paper copy of the flyer, call the NSF Publications Sectionat 202-357-7861. For an electronic copy, send an E-mail message to stisfly@nsf.gov (Internet) or STISFLY@NSF (Bitnet).

ORDERING BY ELECTRONIC MAIL

If you are a user of electronic mail and have access to either BITNET or Internet, you may prefer to order publications electronically. BITNET users should address requests to pubs@nsf. Internet users should send requests to pubs@nsf.gov. In your request, include the NSF publication number and title, number of copies, your name, and a complete mailing address. Publications will be mailed within 2 days of receipt of request.

TDD #: 202-357-7492

NATIONAL SCIENCE FOUNDATION
WASHINGTON, D.C. 20550

OFFICIAL BUSINESS
PENALTY FOR PRIVATE USE $300

RETURN THIS COVER SHEET TO ROOM 233, IF YOU DO NOT WISH TO RECEIVE THIS MATERIAL ☐ , OR IF CHANGE OF ADDRESS IS NEEDED ☐ , INDICATE CHANGE, INCLUDING ZIP CODE ON THE LABEL. (DO NOT REMOVE LABEL).

BULK RATE
POSTAGE & FEES PAID
National Science Foundation
Permit No. G-69