ABSTRACT

        Trends in emphases in the study of clinical reasoning
are examined, with attention to three major branches of research:
problem-solving, knowledge engineering, and propositional analysis.
There has been a general progression from a focus on the generic form
of clinical reasoning to an emphasis on medical content that supports
the reasoning process. This has led to investigation of the nature of
physicians' understanding of medical problems and the role of the
basic biomedical sciences in clinical thinking. Studies of expert and
novice diagnosis in pediatric cardiology have shown the importance
of: medical knowledge in clinical thinking; and hypothesizing in case
solution. Other research has assessed problem representation in a
medical field and how it influences radiological diagnosis. Bill
Clancey, a researcher in artificial intelligence at Stanford, has
studied expert diagnostic performance rules and traced back
underlying cognitive structures that give rules their conceptual
meaning. Clancey used MYCIN heuristic programming in studying
clinical diagnosis. Patel and colleagues at McGill University have
focused directly on clinical understanding (based on theories of
comprehension and the field of discourse processing) and have
examined inferences by experts aimed at supporting clinical actions.
Other lines of research are also noted. Four pages of references are
included. (SW)

ERIC

# THE PURSUIT OF UNDERSTANDING IN CLINICAL REASONING[*]

Paul J. Feltovich
Southern Illinois University School of Medicine


Vimla L. Patel
Faculty of Medicine, McGill University

# THE PURSUIT OF UNDERSTANDING IN CLINICAL REASONING

Paul J. Feltovich
Southern Illinois University School of Medicine

Vimla L. Patel
Faculty of Medicine, McGill University

The study of clinical reasoning ha      m a general progression from an
emphasis on the global form of clinical      soning to an emphasis on medical
content which supports the reasoning prc_ess.  This has led to investigation
of the nature of physicians' understanding of medical problems and the role of
the basic biomedical sciences in clinical thinking.  This paper will trace
some of this development through selected lines of research.  It is not inten-
ded to be comprehensive, but will focus on three major branches of research
contributing to this trend, with presentation of selected studies represen-
ting:  problem-solving, knowledge engineering, and propositional analysis.

Early, classic studies of clinical reasoning established important charac-
teristics of the clinical reasoning process but also introduced some paradoxes
(Elstein, Shulman, & Sprafka, 1978; Barrows, Feightner, Neufeld, & Norman,
1978).  A level of regularity in the form of clinical reasoning was estab-
lished.  Clinicians use early case cues to generate sets of tentative alterna-
tive explanations, hypotheses, for the patient's condition.  These hypotheses
then structure and guide further interrogation of the case.  Hypotheses pro-
vide expectations of further clinical manifestations that should be present if
a hypotheses is true of the case, and actual patient findings are compared to
expectations to select among the alternatives.  Hypotheses and hypothesis sets
can change and be restructured as diagnosis progresses.

The unexpected and somewhat paradoxical result of these studies was that
most of the measures of this process which were studied did not differ among
clinicians, either those with greater or less experience, or those of differ-
ent externally judged expertise.  These were largely measures of timing (e.g.
percentage of cues to first hypotheses) and number (e.g. number of hypotheses
generated and maintained in active consideration) which did not take account
of what was being considered, rather the form of its consideration.  The only
measures that discriminated levels of expertise were measures that took
account of the content of reasoning, i.e., what was reasoned about (Barrows,
et al., 1978, "hypothesis aggregate scores").  Experts and more experienced
subjects differed in the "quality" (by reference to the hypotheses of other
experts) of the hypotheses they considered throughout the case and in the
quality of the diagnostic outcome.  Experts in a sense chose better "moves" in
their diagnostic explanations for a case.

These findings in medicine were consistent with studies of problem solving
emerging from other fields, particularly the classic studies of chess
(deGroot, 1965; Chase & Simon, 1973; Simon & Chase, 1973).  Chess masters and
novices did not differ in the number of alternative moves they considered from
a board configuration or the depth with which they considered the implications

of a possible move (interestingly, intermediate level players did the most search -- following up good and bad moves alike). Experts simply chose higher quality moves from a given position.

In both chess and medicine these findings were attributed to the knowledge base of the practitioner. In chess the expert was proposed to have a large differentiated store of memory models which enabled recognition of familiar board configurations. Once recognized in a presented board, these models had associated with them moves found to be effective through past experience. This store of recognizable configurations was assumed to be immense; the expert seldom has to deal with novelty, having brought much of his work-world into the realm of the familiar (see also, Charness, 1979; Egan & Schwartz, 1979). Although the work in chess was originally done using visual displays, Chase and colleagues have found similar results using verbal presentations of game situations (Chase & Chi, 1980).

## Pediatric Cardiology Studies

Studies in pediatric cardiology at the University of Minnesota in the late seventies set out to find some clues to expert diagnostic quality in the knowledge base of expert clinicians, since prior studies in medicine, although suggesting its importance, had not concentrated directly on the knowledge base.

Several theoretical ideas contributed to the studies: (1) "Frames" (e.g., Minsky, 1975) as organized packets of knowledge representing things and events in the world. Frames specified "slots," sets of characteristics or features to be expected in the represented events. (2) Generalization hierarchies (e.g., Winograd, 1975) of frames in which frames are organized from general abstract descriptions down to highly specific instances, and (3) interactive (top-down/bottom-up) (Bobrow and Norman, 1975) activation of frames by which frames are activated either by events in the world or other frames, and once activated, actively seek to find and fill their own slots.

A study by Feltovich (1981/84) was a study of the knowledge base of expert and novice physicians. Interest was not so much on the general process of diagnosis. A basic process of generating and testing hypotheses was assumed and, indeed, the study made use of this assumption in order to achieve its goals.

Feltovich speculated on a knowledge structure for the organization among disease models (frames/schemas) in memory (Figure 1). These speculations were supported by some findings in medicine (e.g., Wortman, 1972) and some psychological findings regarding the nature of knowledge base change and development (e.g., Wortman & Greenberg, 1972). Disease frames were hypothesized to be organized at three levels. At the most global level, individual disease frames were organized within frames representing general disease categories (e.g., "obstruction lesions"). At an intermediate level were "classical" (cf. Rosch & Mervis, 1975, "prototype") disease instances of the category, diseases that were most characteristic of the category. Branching off each disease was a set of frames specifying variations and subtypes of the disease, where the disease variants could be defined on many different dimensions, e.g., severity, time-course development, slight anatomical difference.

Level 1:
    Disease categories
    e.g., "Obstruction
    lesions"

Level 2:
    Classical
    instance
    (prototype)

Level 3:
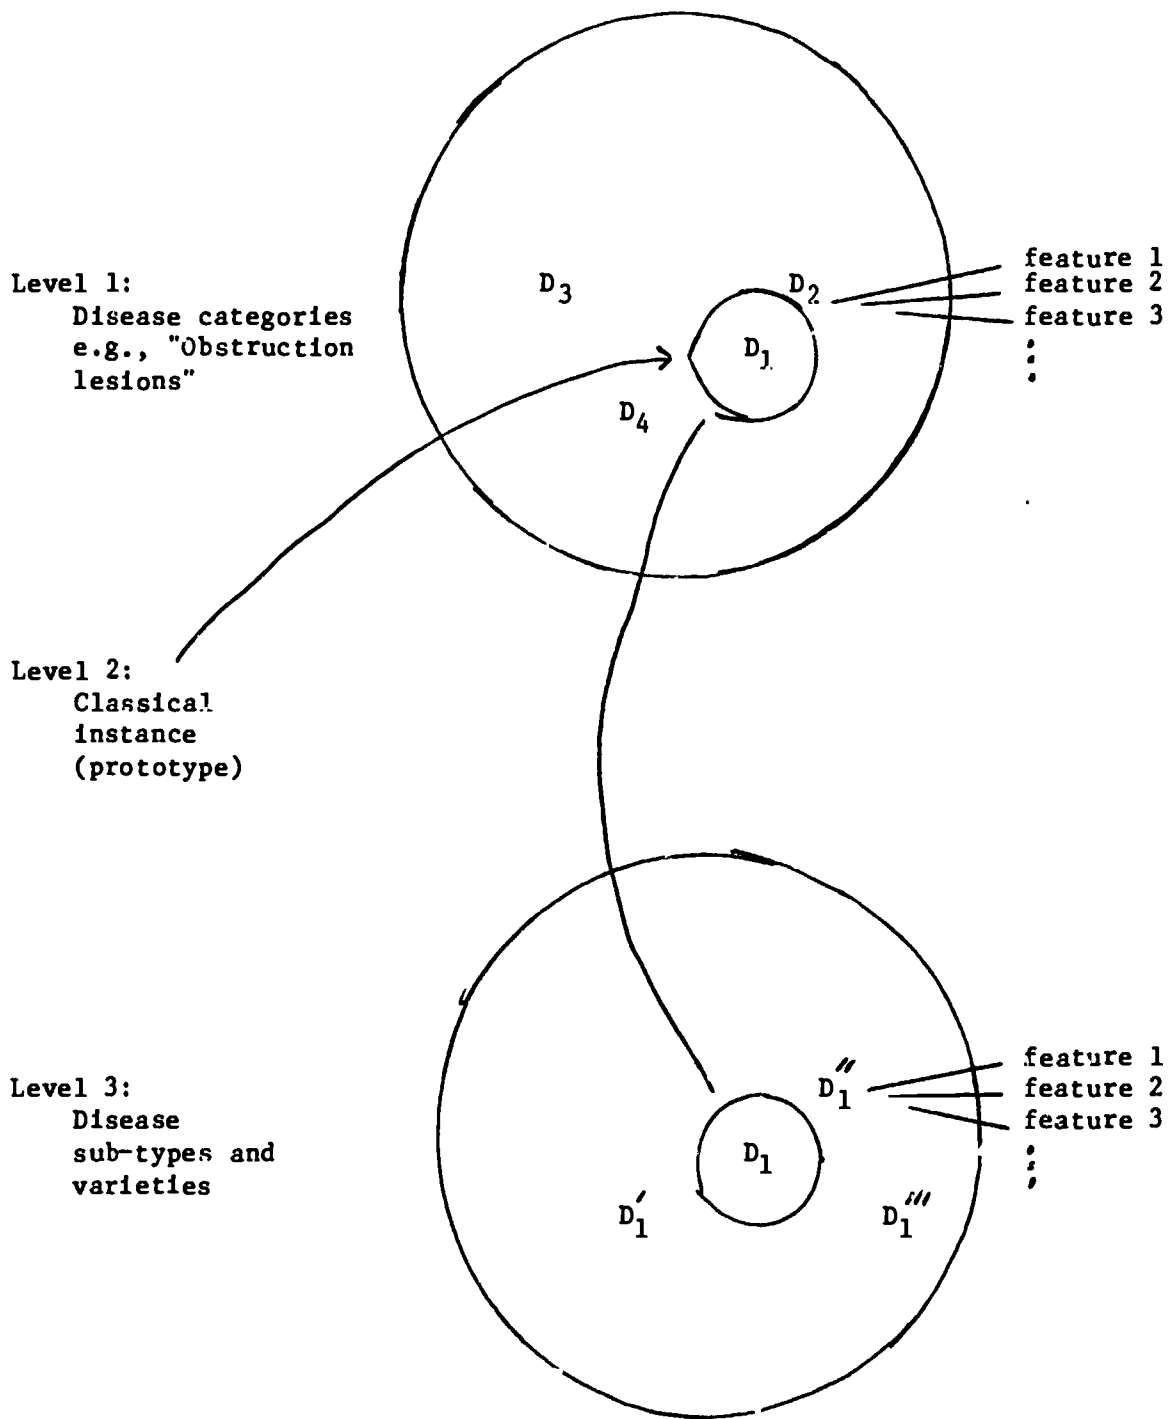    Disease
    sub-types and
    varieties

Fig. 1.  Organization of disease frames in memory.

Several general hypotheses about the nature of knowledge-base differences between experts and novices were made: (1) The knowledge base of experts was hypothesized to be dense. Diseases in a category were tightly packed and interactive; when one was activated others would be. (2) Precise. The values allowed for filling features in frame slots would appropriately reflect the naturally occurring variability in patients with that disease. (3) Penumbral. Each disease would be differentiated into many distinct variations. Novice knowledge bases, in contrast, were speculated to be sparse (not much cross-linkage and interaction among diseases), imprecise (expected features would at times be overly restrictive, overly lenient, or wrong), and classical (knowledge of the prototype within a category would develop first).

The study utilized the structure of cases, in terms of where in the case critical patient information arose, and assuming a general frame activation and instantiation process, to test these ideas. Early case data would "push" subjects (through frame activation) into level 1 (category level) or level 3 (disease variants level) of the disease hierarchy. Furthermore, the early data would cause generation of the prototype at each level -- but this classical hypothesis would be wrong as a diagnosis for the case. This control of entry into the knowledge base would allow assessment of density, by showing which related hypotheses within the cluster were generated concurrently; classicality, by showing which subjects "got stuck" on the erroneous classical hypotheses; and precision, since certain critical data items would have to be interpreted in precise ways if the subject were to extract himself from the initial "foil" interpretation.

Since the study was interested in the bases of "quality," we defined a set of "good" frame clusters for each case - sets of diseases that were reasonable interpretations for the case and easily were confused among each other. This was done through expert consultation and reviews of medical literature. These predefined "logical competitor sets" (LCS) for a case were disease sets that shared an underlying "deep" pathophysiological mechanism, e.g., "right to left shunts at the atrial level" of the heart.

Subjects (students, residents and specialists in pediatric cardiology) diagnosed fixed order cases in small information segments while thinking aloud. Resulting "thinking aloud" protocols were analyzed for the generation of LCS hypotheses together in clusters and for evaluations of data items critical for discriminating LCS diseases.

Results were generally consistent with predictions. Density. When experts generated one of the LCS diseases at either the disease or variant level they usually generated them all. Classicality. Novices tended to focus on the prototype at either level, and often fixated on these to the exclusion of the correct interpretation. Precision. Experts had more precise feature expectations within disease frames that enabled them to discriminate better among diseases.

We came out of these studies thinking something had been learned about the keys to quality in clinical diagnosis. These included: (1) Tightness, compactedness of competitor set diseases in memory - interframe activation; (2) Organization of clusters by deep pathophysiological principles; (3) Many

variations on each disease theme ·- each specified as a separate frame in experts; (4) Well "tuned" clinical expectations within frames at any level. Yet the studies raised questions and left others unanswered:

(1) <u>Need for relational features</u>. The conceptualization of each frame within the study was as a memory structure that specified expected clinical features. As new clinical data were encountered, these were compared to frame-specified features. The study did not, a <u>priori</u>, attend to the importance of <u>relations</u> among features. Yet subjects, especially experts, showed instances of using such relations to discriminate among diseases. Relations among expected features of a disease such as causal, conditional, or temporal order relations, provide a whole order of increased discriminability among hypotheses. For example, two diseases with exactly the same feature <u>set</u> may yet differ in the relational structure among them. The relations subjects used seemed based in the underlying pathophysiology that provided the means for connection.

(2) <u>Prespecified frames</u> . Disease frames were construed as prestructured entities even though ranges and variables were allowed in feature expecta- tions, and frames could vary in their levels of abstraction, say from general categorical structures down to highly refined diseases. There was no mechan- ism by which frames could be constructed, as situationally required, from more basic principles -- as might be needed in novel situations, unusual combina- tions of abnormalities, etc. The study's approach to strangeness, especially in experts, was in proposing lots of frames -- including many disease <u>varia- tion</u> frames (level 3). While this approach is consistent with many concep- tions of expertise - the expert is one who has brought most of his task-world into the realm of the familiar - a generative capacity for frames is clearly needed for a more complete model of clinical reasoning.* It should be said that even though the study under discussion did not utilize a generative capa- city in frames, the more general theory of clinical reasoning which provided a backdrop for the study did provide for generativity (e.g., Swanson, Feltovich, Johnson, 1977). Swanson, in his computer modeling of clinical reasoning in pediatric cardiology (Swanson, 1978), incorporated generative processes which could take any combination of structural/ functional basic cardiovascular dis- orders (e.g., stenotic valves, shunts) and create an associated frame of cli- nical expectations by applying principles of cardiovascular physiology.

(3) <u>Diagnostic Strategy</u>. The stu`y contributed little to a conception of strategical issues in clinical reasoning, in terms of strategies for collec- ting and using clinical information and managing the space of hypotheses. Strategy (planning) and procedures in general can be characterized at many levels of abstraction -- from the highly general and abstract to highly detailed descriptions of particular steps (e.g., Sacerdoti, 1974; 1977). Prior studies of clinical reasoning had proposed strategic conceptualizations at the most global levels of generality -- that is, as a process of generating and testing hypotheses. This level of abstraction did not discriminate

* However, anybody proposing to study the expert in non-adapted situations (e.g., strange, novel, or "atypical" cases) must exercise great care. Just because a case is atypical to a novice (or the experimenter), this does not imply it is necessarily atypical or strange to an expert.

-6-

7

strategy among individuals, does not provide much guidance for training (telling a student he/she should create hypotheses does not help much if we cannot also tell him/her which to generate when and _how_, e.g., using what kinds of information in what ways), and is not very satisfying as explanation for clinical reasoning. If hypothesis-testing was too general, the present study was perhaps too specific. In terms of strategy, it amounted to saying that for a particular type of case one must use particular sets of data to create particular clusters of hypotheses, and specified a discrimination plan within each competitor set. While this case-specific and case-_class_ specific orientation to strategy is probably not far wrong as a conceptualization of skilled clinical reasoning, it would be useful to have a conception of strategy cast at intermediate levels of abstraction -- between global form and detailed steps. This is necessary for any kind of useful notion of strategical generality and for teaching any kind of cross-case approach to working with clinical problems.

## Radiology Studies

The pediatric cardiology studies had shown the importance of the medical knowledge base in clinical thinking, and like other work in medical reasoning, it had shown the importance of the hypothesis set, indicating that the way a clinician construes and mentally represents the set of basic medical issues in a case has major implications for case solution. Other work in medicine has shown that the earliness with which a good hypothesis set is created is predictive of quality of diagnostic outcome (Barrows, Norman, Neufeld, & Feightner, 1982). Good clinicians have good problem models early. The cardiology work had indicated that the basis for a good hypothesis set (in its domain of study) was in the "deep structure" of a problem, i.e., the underlying key pathophysiological principle applicable in a case. This was consistent with other work in psychology showing the importance of early problem representation for problem solution. Work in physics, especially, has indicated that experts conduct a "qualitative analysis" for a problem before attempting solution (de Kleer, 1977; McDermott & Larkin, 1978). In this qualitative analysis, a model of the problem is built, an internal mental "problem representation" that represents problem elements and their interrelationships, along with embellishments and inferences drawn from the knowledge base of the individual. Greeno (1977) has argued that characteristics of this problem representation define degrees of "understanding" of a problem; a problem must be understood well before it can be solved well, efficiently and with elegance, etc. Furthermore, Greeno has argued that this representation is built up progressively over a span of time, from an initial representation to one that is modified and grows as problem solution progresses. Work in physics (Chi, Feltovich, Glaser, 1981) has shown that experts represent problems more abstractly than novices, organizing problems by deep structure, i.e., major underlying physics laws and principles, rather than more apparent features of a problem, given in the problem statement.

The work to be discussed in radiology by Lesgold and colleagues (Lesgold, Feltovich, Glaser, & Wang, 1981; Lesgold 1984; Lesgold, Rubinson, Feltovich, Glaser, & Klopfer, forthcoming) involved the study of problem representation in a medical field and how it influences and interacts with radiological diagnosis.

In the early studies from this project, subjects from first year radiology residents to experts in radiology interpreted x-rays while thinking aloud the whole time they viewed the film. As a way of getting a better notion of the subjects' early representation of the film, the procedure included a two second "flash-viewing" of the film before more deliberate analysis. Subjects viewed a film for two seconds, reported on what they saw during that period, then proceeded to full analysis while thinking aloud. Later follow-up studies tested initial hunches about the process of radiological diagnosis, extracted from initial protocol studies, in more focused experimental procedures in which subjects were required to trace on x-ray films various anatomic and diagnostic features they saw in the films.

Discussion here will focus on the primary findings and conclusions from these studies.

Initial representations: Two-Second Viewing.

(1) Experts encoded more of the important features from a film and in a way that captured their interrelationships.

(2) Experts captured the main deep underlying pathophysiological dynamics of a film that enabled a global characterization of the film (e.g., "hyper-expansive chest," a chest characterized by obstructive lung dynamics and the main consequences of such obstruction).

(3) Global characterizations set <u>constraints</u> for further interpretation. It was not as though experts were <u>seeing</u> ultimate diagnoses in two seconds. Rather, their initial gross characterizations set restrictions and limits for further more detailed interrogations. They were capturing the overall pathophysiological/anatomical "gist" of the film.

The process of film analysis.

(1) Analysis/diagnosis of the film was conducted within the constraints set in initial encoding. Initial encodings to some degree set boundaries for more detailed interpretations.

(2) The diagnostic progress was interactive. Frames (schemas) for diseases and other abnormal processes partly directed what was seen and what was attended to. On the other hand, more accurate "seeing" of components of anatomy and abnormal film features both contributed to the generation of interpretations and tempered/ modulated expectations provided by frames. While novices saw an x-ray film and its features, experts built a model of a defective human body, using the x-ray as one source of evidence for this construction. The expert construction utilized knowledge of anatomy, pathophysiology, disease conditions and knowledge of the mapping between these knowledge structures and x-ray film manifestations. The model of the patient was constructed, modified, and adjusted throughout the process of viewing.

(3) Experts were less dependent on and restricted by "classical frames." As with the pediatric cardiology work, novices were restricted by the classical forms of disease conditions and disease features which limited both what they attended to and how they interpreted what they saw. Experts were more flexible in interpretation - partly because of knowledge of more alternative variants of any disease process (more level 3 frames, disease variants, as in pediatric cardiology), but also because of better ability to use principles of pathophysiology and anatomy to adjust rules of interpretation to context.

(4) Some results from these studies provided a glimpse of the acquisition and development of diagnostic skill - from classicality to expert flexibility. There were several films where beginning residents and experts performed better than intermediate residents. Moreover, the difference in performance among these three groups varied systematically in ways that gave clues to the process of development. Films where these patterns were seen were ones which represented fairly classic film presentations of the medical conditions involved. A characterization of findings from these cases is as follows: Beginning residents diagnosed the cases correctly, using standard direct associational rules from film features to medical interpretations. Intermediate residents diagnosed the cases wrong, often giving diagnoses in terms of abnormal anatomical or pathophysiological structure. In the process, they tried to untangle the anatomical/physiological basis for the apparent features of the film and tried to account in this way for contextual characteristics of the case, e.g., interrelationships among film features. It was as though they were trying to understand the film in terms of trying to build a deep-structural model of the medical condition evidenced by the film. Expert interpretations were correct and again more direct, more like the novices. But, the contextual and deep issues that intermediates struggled with were incorporated and embedded within the interpretation rules that experts applied.

The process of development of skill appears to be something like the following: Novices learn a set of text-book medical conditions and rules of interpretation connecting film features to these interpretive models. As they diagnose more films, intermediates learn that these simple, direct rules at times lead to error, because of contextual factors, idiosyncracies of presentation in individual patients, interdependencies in film features whereby the significance of some features varies according to variations in others, etc. This leads to a period in which the student tries to understand the underlying basis of interpretive rules, in the principles of anatomy and pathophysiology responsible for the appearance of a particular film. This kind of deeper analysis is cognitively taxing and places high demands on available mental processing resources, sometimes leading to error. High levels of expertise are then achieved when the contextual considerations focused on by the intermediate are compiled (Anderson, 1982) within cognitively efficient direct associational rules of interpretation (productions). Understanding which the intermediate strives for through focused, deliberate effort and model building is then accounted for implicitly within the structure of the rules of interpretation that experts use. The next set of work to be discussed complements this interpretation.

## Clancey

Bill Clancey and his collaborators at Stanford have taken a somewhat different approach to studying expertise in clinical reasoning and have contributed substantial insights to what understanding in clinical reasoning might mean--both at the level of structural understanding (the underlying structure of clinically relevant knowledge and clinical problems) and in matters of clinical strategy.

Clancy works out of an artifical intelligence (AI) and "knowledge engineering" background, fields concerned with modeling expert knowledge and strategies so that computers can perform complex, "real-world" tasks (e.g., Feigerbaum, 1977). One of his original (starting in late 1970's) and continuing interests is in building a computer tutor for teaching clinical reasoning skills.

In attempting to build a clinical tutor, Clancey (1979) started with an intact, running expert system (MYCIN) that performs clinical diagnosis in the medical area of meningitis (Shortliffe, 1976). MYCIN works well, performing diagnosis in its area at a level comparable to medical specialists (Yu, et al., 1979). MYCIN functions by applying a large set of diagnostic rules (heuristic "rules of thumb") extracted from medical experts. Clancey, reasoning that he had in hand a set of expert performance rules, sufficient for accomplishing diagnosis at a high level, first thought that he might be able to teach these rules directly to students.

These direct expert rules proved unsatisfactory for teaching students, as students could not conceptualize what the rules were doing, and were forced to learn them largely by rote (Clancey, 1981/84). These limitations for teaching, along with developments from the psychology of problem solving, led Clancey (1983/84) to conduct an extensive investigation of the regularities and structure in MYCIN's "compiled" or performance-level rules, in order in guide the development of a new medical consultation program, NEOMYCIN.

A comparison of this approach to the radiology work is illuminating. The radiology work had studied the diagnostic process directly and suggested that experts compile (incorporate) various kinds of "deep" problem regularities (e.g. pathophysiological dynamics, types of chest classifications) into direct performance rules. Clancey, on the other hand, started with a set of expert heuristic performance rules (from MYCIN) and worked backward from the set of rules to proposing a set of cognitive structures that formed the basis of the problem structuring and strategical considerations implicit within the compiled rules.

An analysis of a MYCIN rule (taken directly from Clancey and Letsinger, 1981/84) illustrates some of the flavor of this endeavor:

## A Typical MYCIN Rule

If:    1) The infection is meningitis
        2. The subtype of meningitis is bacterial
        3) Only circumstantial evidence is available
        4) The patient is a least 17 years old
        5) The patient is an alcoholic
Then:      there is suggestive evidence that diplococcus-pneumoniae is an organism causing the meningitis.

"This rule is an example of "compiled expertise." We can list some of the individual steps of reasoning and knowledge sources out of which it is composed, unknown to MYCIN, but explicitly represented in NEOMYCIN:

--      Analysis of other rules shows that this rule (to determine the organism) is only invoked after it has been established that the patient has an infection. Thus, four major subgoals are established in this order: Is there an infection? Is it meningitis? Is it bacterial? Is it diplococcus-pneumoniae? Each of these subgoals hypothesizes a more specific cause of disease. Thus, the ordering of clauses constitutes a top-down refinement strategy. However, MYCIN does not know about this specialization hierarchy. It does not even know that diplococcuspneumoniae is a bacterium. Perhaps most seriously of all for meeting our teaching goals, MYCIN omits intermediate categories such as acute/chronic meningitis and "gram negative meningitis" that physicians find helpful. In NEOMYCIN these categories are represented explicitly in an etiological taxonomy by allowing parameters to be specializations of one another.

--      The clause about the patient's age prevents MYCIN from asking if a child is an alcoholic. MYCIN does not know that the ordering of these clauses is important, or what the relationship is. In NEOMYCIN these world relations are captured by separate "screening" rules.

--      When there is laboratory evidence (a culture with visible organisms), this rule does not apply (clause 3). However, a companion rule still allows the circumstantial evidence of alcoholism to be considered, but gives it less weight. This principle of considering circumstantial evidence even when there are hard, physical observations of the

cause, is not explicitly know to MYCIN. The principle is
compiled identically into 40 pairs of rules, rather than
being stated as a reasoning rule for combining hard and soft
evidence. NEOMYCIN has rules for reasoning about the evi-
dence it has collected, so connections between data and
hypotheses are separate from the contexts in which they will
be used."

While uch analyses provided some guidance as to the knowledge structures
(e.g , hierarchical organizations of causes) and strategical considerations
(e.g., refining more general to more direct causes) implicit in expert per-
formance, many issues could not be addressed—largely because of the nature of
MYCIN, and the clinical tasks it was designed to perform.

The task of MYCIN was to determine the specific cause of meningitis in a
patient, knowing that the patient had some form of meningitis. Hence it had
no need to deal with higher levels of problem structuring, e.g., forming a
differential diagnosis involving meningitis and other alternative possible
explanations of a case. Hence, for example, rules for configuring a differ-
ential did not appear.

To examine these other issues, and also to follow-up leads derived from
analyses of MYCIN's rules, Clancey has been involved in more direct studies of
expert clinical performance. However, under the assumption that much of
expert performance constitutes compiled behavior, incorporating implicit
structures, a guiding principle for these investigations is that expertise
cannot be discerned exclusively from observations of experts performing their
material task (e.g., diagnosis). Experts need to be interrupted and probed
for deeper understandings during the conduct of performance, and some of the
underlying bases for performance should be made more explicit in related but
tangential tasks, e.g., teaching diagnosis to students.

Hence the overall methodology is eclectic, involving rule analysis, inter-
views, diagnostic tasks, class-room observations, and formal computer model-
ing. Cycles of computer modeling lead to more directed observational studies
and refinement of the model, and so forth.

This work (Clancey, 1984) has suggested several dimensions of organiza-
tional structure in the medical knowledge base of the expert clinician that
contribute to case understanding and that provide for a stru cured
(meaningful) conception of diagnostic strategy:

1.  The expert's internal representational model of a "disease" is a com-
    plex structure, a "causal story" of how an illness comes to be in a
    patient, how it produces effects, and develops through time -- a
    "disease process frame." The "slots" or expected features in these
    frames can be highly relational, specifying temporal relations,
    causal relations in the development of illness, etc.

2.  Commonality in components of disease process frames  e.g., shared
    conditions o  exposure, similar elements of compromise in the host)
    constitutes a rich set of sources for configuring a differentia',

i.e., grouping case explanations, and for discriminating among diseases. Furthermore, abstract models of disease forms (e.g., infection as a story of exposure to and invasion by an external agent, disemination to a site, compromise in the host, etc.) serve as a means of grouping sets of diseases, eliminating whole sets of diseases (e.g., by excluding particular components of developmental phases), and for giving coherence to a set of patient findings, by integrating them into a causally connected structure.

3. The expert recognizes a rich vocabulary of clinical evidence types and their usefulness in the diagnostic process. These are abstract characterizations of types of evidence (e.g., compromised host factors: alcoholism, pregnancy, certain classes of medications; "Red--flag findings:" serious consequences of illness) and the functions they can generally serve in organizing the diagnostic process. For example, "red-flags" as a class are useful for establishing a differential. Compromised host factors are useful for expanding a differential by specifying unusual susceptibilities in the host.

Clancey's rich, causally oriented conception of diseases (and medical abnormality more generally, e.g., pathophysiological states), and the categorical organizations derived from these (dimensions for grouping, and causal disase forms built from these -- e.g., infectious processes), allows a model of the patient to be built. The model reflects a form of problem understanding by specifying at a detailed level how the components of the patient's problem are related (causally, temporally, etc.) among themselves.

The abstract conception of data types and their functions in diagnosis constitute an "intermediate level" vocabulary (between broad goals of diagnosis such as "form a differential and discriminate" and particular handling of individual data items) that provides meaning for diagnostic strategy and strategical acts. For example, an expert might examine "unusual exposure" factors (travel, big crowds) to establish a step in the infectious process scheme, or examine compromised host factors to broaden a current working differential

It should be noted that although Clancey's system construes case understanding in terms of relatively "deep" causally oriented disease process models and forms of disease types, the system does not focus on the even more basic underpinnings of these within functions and regularities of the human body (e.g., basic biochemical or physiological concerns). The system deals with aggregations and regularities of these more basic factors that the expert has found over time to be useful. Clancey speculates that the more basic biomedical dynamics may be necessary to consider when the expert is functioning outside his range of adaptation, i.e., in situations which might be novel or strange even for the expert. In this regard, Pople (1983) represents an AI system that deals with more fundamental causal factors more directly and proposes developmental (learning) mechanisms by which these are aggregated over time (compiled) into clinically efficient and useful structures more of the sort that Clancey proposes.

Patel

While Clancey started with expert performance rules in medicine and traced back, aided by empirical investigations, to underlying cognitive structures that give these rules meaning and organization, Patel and colleagues focus directly on what it means to understand or comprehend a clinical case.

The Patel work has its roots in the field of discourse processing, the field that studies how people comprehend textual or spoken material. This field treats comprehension as a process of building an internal cognitive model of the message contained in the written or spoken material. Construction of this model depends greatly on inference, processes of filling-in information not directly stated in the material or giving particular interpretations to the text material depending on one's prior knowledge.

The studies conducted by Patel and her colleagues are based on theories of comprehension. Two major approaches have contributed to such theories (cf., Collins, Brown, & Larkin, 1978). One approach, the text based approach, has focused on the textual material as a source of inference, emphasizing aspects of text structure that lead to inferences. Schema-based approaches emphasize the contribution of the comprehender's store of knowledge and prior experience as sources for filling-in and augmenting information contained in the text base.

Prior knowledge provides a kind of "guidance system" for regulating and coordinating inferences applied to the text. There have been two general approaches to specifying this guidance system. One emphasizes stereotypes and regularity -- and might be termed the fixed-frame approach. This view construes the guidance system as a collection of frames (scripts, schemas, etc.), mini world-models that specify the features of recurrent world events and situations (e.g., a disease and its typical clinical features). Under this view, constructing a model of the text is largely a matter of using text cues to engage relevant frames ("bottom-up," or text-to-knowledge processing) and using these structures, once activated, to fill-in or augment information provided ("top-down," or concept-driven processing). Even though the number of frames can be immense (especially in experts), can range in scope from large (models of whole situations) to small (models of pieces of situations), and can be composites of other frames, the approach still rests ultimately on known regularity. Another approach to the guidance system emphasizes generativity and novelty -- the generative frame approach. Under this view, the guidance system is a set of rules (e.g., a grammar) which specify the forms that frames can take, and criteria to which a comprehender's constructed model of a certain kind of text must conform (Collins, Brown, & Larkin, 1978). These rule systems can spawn a potentially infinite set of models.

The work of Patel and colleagues in medical comprehension can be characterized by the generative frame approach (Patel and Frederiksen, 1984a; Patel, 1984). Under their system, comprehending (understanding) a medical case is an interactive (bottom-up and top-down) process of constructing an interpretive frame (case model) for the case, that reflects both data/text properties and prior conceptual knowledge of the physician. The model is constructed under

the guidance of generative structures (grammars) that direct and constrain the form the emergent model can ta.e. This work focuses on the case or problem model [a focus shared recently especially by various artificial intelligence endeavors in medicine (e.g., Patil, Szolovits, & Schwartz, 1982)] that physicians create, and how it is assembled from data, knowledge, and interactions among these. The problem model can be thought of as a structure separate from the base of patient information and the physician's structures of knowledge, one that is particularly tailored for a case, utilizing these other structures and rules for building a model.

In its emphasis on the problem/case model, the Patel work relates to other lines of research in problem solving -- through the construct of a "problem representation." For example, one definition of a problem representation (Chi, Feltovich, Glaser, 1981) that har been given is that "a <u>problem representation</u> is a cognitive structure corresponding to a problem, constructed by the solver on the basis of his domain-related knowledge and its organization." A problem representation is the solver's internal model of the problem, its elements, relationships among them, and embellishments provided by the solver's knowledge. Greeno (1977) has proposed that the problem representation is a semantic network of entities and relationships among them, and that problem solving really amounts to developing such a structure for a problem - one in which <u>solution</u> elements are initially absent but are later incorporated as problem solving proceeds. The problem models that Patel and colleagues propose are constructed during comprehension are similar to some extent in that they are highly intricate, detailed semantic networks of entities and specific relationships among them which are structured in various ways into more complex organizations. Their studies are similar in orientation to the works of Kintsch and Greeno (in press), and they look at the relationship between comprehension and problem solving in medicine.

One of the strengths of the Patel work is that it has a uniform notational formalism for representing the semantic information in external events (e.g., a medical text), the structure of knowledge in human memory, and operations (relations) that act on and modify these structures to produce inferences. This formalism enables detailed analysis of the influences of text structure and internal knowledge in building a conceptual interpretation (comprehension) of medical material.

The representational formalism used in this research contains a system of "nodes" (entities) and relationships among them as developed by Frederiksen (1975). The most primitive units in this system are <u>concepts</u> (objects, actions, and attributes) which serve as the "nodes" (e.g., infection, fever) and a set of relationships (links) that can hold among concepts. The next higher level of organization is the proposition, a triple of concepts and relationships that constitutes the lowest level of organization that can have truth value (i.e., be true or false) when applied to a situation, e.g.:

All infections produce fever.

Propositions
1.0    CAU(produce)         [1.1], [Fever]:
       (link a)             (node x) (node y)
1.1    Infections           GEN,UNIV,(ALL);


Fever may last for weeks and months.

Propositions
2.0    Fever                TEM:_____,DUR:WEEKS OR MONTHS;
       (node Y)             MOD:QUAL(MAY), ASPCT:CONT(LAST)
2.1    OR-ALT               [weeks],[months]

Propositions can also function as nodes. Relationships among propositions
constitute yet higher order organizations. "Frames" in this system are com-
plex high order conceptual structures which are defined in terms of relation-
ships among propositions e.g.:

Conceptual Representation

```
                                          ┌──→  weeks
  ┌──────────────────────────────────┐    │  2.1 OR-ALT
  │              DURATION            2.1   └──→  months
┌─┴───────┐                      ┌────────────┐
│  FEVER  │  ←──── (produce) ────│  INFECTION │
└─────────┘          CAU         └────────────┘
    2.0              1.0               1.1
```

Frame generation grammars constitute a set of generative rules which bound the
forms that these frame ensembles can assume.

     A typical experimental paradigm used by Patel uses medically related
textual material (e.g., case descriptions or disease descriptions), and has
subjects read these materials and recall or summarize what they have read.
She also uses directed probes and questions to help determine the subjects'
knowledge base. Analysis is directed at comparing the semantic structure of
the presented text (formalized according to the scheme described above) with
the text produced in recall/summary. These analyses enable the experimenter
to make detailed inferences about the influence of text structure, the compre-
hender's knowledge structure, and interactions among these in the comprehen-
der's construction of a model for the text.

     The primary method for identifying the cognitive processes a subject uses
in constructing a text model is through techniques of "discourse analysis."
These include text level clausal analysis, propositional analysis, analysis of
inferences a subject produces, and analysis of the conceptual frames the sub-
ject employs. An illustration of such an analysis is traced below, using an
example segment of text.

The first step in the analysis is to separate the text into segments, corresponding to syntactic units, using Winograd's system of clausal analysis which is based on systemic grammar (Winograd, 1972).

Example Text:

"Examination of regional lymphatic drainage is important
because this may give clues to primary sites of malignancy."

The clausal analysis of this sentence identifies two clausal segments: a declarative major clause (DEC) and a secondary bound adjunct (BAJT). "Examination of the regional lymphatic drainage system is important" [Segment (1)] is the declarative major clause, and "because this may give clues to primary sites of malignancy" [Segment (2)] is the bound adjacent.

The next step involves the representation of discourse in terms of its propositional content and structure, the "propositional analysis." The propositional analysis is seminal in ultimately allowing one to see precisely how a subject's recall is related to text content and for specifying the inferences that a subject makes. Propositions are the basic sets of elements which represent the semantic content in text segments. The two most popular models for semantic analysis of propositions are those of Kintsch (1974) and Frederiksen (1975), and the latter forms the basis for the work of Patel and her colleagues. The basic assumption of these models is that the sentence is comprised of one or more propositions and that the pivot of each proposition is the verb.

As applied to the sample text, the propositional analysis results in the following propositional decomposition. Segment (1) contains two propositions (indicated by numbers):

| Predicate | Segments |
|-----------|----------|
| 2.0 EXAMINE | ACT.2.1-ATT:IMPORTANT; |
| 2.1 DRAINAGE | =ATT:REGIONAL, ATT:LYMPHATIC; |

Segment (2) contains the following six propositions:

| | | |
|---|---|---|
| 3.0 | COND | (3.1),(2.0); |
| 3.1 | CAU | (THIS),(3.2); |
| 3.2 | GIVE | THM:CLUES-MOD:QUAL(MAY); |
| 3.3 | CLUES | THM:(TO)(3.4); |
| 3.4 | MALIGNANCY | LOC:SITES(GEN,UNIV); |
| 3.5 | SITES(GEN,UNIV) | =ATT:PRIMARY; |

The propositions consist of a "predicate" and a series of labelled segments. A predicate may be an action (e.g. examine), an object (e.g. tract) or a relation which connects propositions (e.g. COND, a conditional relation). Arguments may be "case" relations, such as PAT, the patient of an action; ACT, the agent of an action; THM, the theme of a cognitive; and RSLT, the result of an action. Arguments can also be identified relations such as LOC, the location

of an object or action, and TENS, the tense, ASPCT, the aspect, or MOD, the modality. A grammar specifies the possible propositional structures that can occur.

By specifying, in this way, both the propositional structure of a text and the propositional structure of the "recall" that a subject provides after reading the text, one can determine which subject recalls are straightforward recapitulations of the text and which are modifications or inferences. Inferences are formally defined a₅ propositional transformations made by a set of inference rules.

The next level of analysis is frame analysis which requires analysis of abstract conceptual structures that are reflected in the organization among propositions contained in a text (Frederiksen et al., in press). A frame is a means of formalizing how the propositions are related to each other to form a more ᴬomplex conceptual structure. Patel's method of analysis of frame structures  begins with specifying the particular frame type to be analyzed (Patel & Frederiksen, 1984b). As a frame structure is constructed, propositions from the text that instantiate each sub-frame element are identified.

For our example text, the propositions which were identified represent elements of higher level frame structures. For example, propositions 2.0 through 3.5 present two procedures ("examining part of a body") and ("being given clues to disease site") and a conditional dependency clause between them ("clue₅ to the primary site of malignancy will be obtained from th₂ examination of the regional lymphatic drainage system"). These segments are part of a text which elaborates the procedure for examining a part of a body and consequences of particular procedures under malignant and nonmalignant disease conditions. Such a conceptual structure is a type of frame, in this instance a procedural frame.

Scoring a subject's recall of a text involves mapping a subject's recall segments (which themselves constitute a "text") against the propositions in the message base of the original passages. This involves identifying every item in the subject's recall that corresponds to the message base as defined for the original text (i.e., same propositions appear in the text and the recall protocol). Each recall segmen⁺ can then be analyzed to determine the types of inferences created by the reader (propositional transformations in the recalls). Frame use is measured by analyzing the distribution of propositions recalled that "instantiate" different specified frame elements. The distribution of information within a frame reflects the degree of complexity of the frame-construction processes of the reader.

Patel and colleagues have conducted numerous experiments on clinical comprehension using the basic methods which have been described, and a full discussion of these studies and their results is outside the scope of the present paper (the interested reader should consult Patel & Frederiksen, 1984b; Patel & Groen, 1984). What will be described here are some selected findings, particularly pertinent to the themes developed in the present paper, taken from a study which was focused on the role of physicians' prior knowledge in comprehension of medical material (Patel et al., 1983; Patel & Frederiksen, 1984b). Subjects in this experiment were entering students at McGill, medical students

after their basic science years but before any clinical experience, and physician internists. Materials were two texts, one on the clinical diagnosis of cancer, and the other on the clinical significance of fever. Prior knowledge was assessed by asking subjects to "free-recall" what they knew about the clinical significance of fever or the clinical diagnosis cancer <u>before</u> being presented with textual material on these topics. Subjects gave a one page written protocol on their prior knowledge, read the texts, and finally, gave recalls of the texts.

Three sets of results from this experiment are relevant to the themes developed in the present paper:

(1) Non-selectivity in intermediate students. Texts presented to subjects were separated into high and low relevance components. This was done prior to the experiment by asking six medical experts to read the texts and to eliminate material in order to reduce text size by one-half, while still retaining the essentials of the medical content. In this way, recall and inferences produced by subjects could be compared against the high and low importance components of the text.

Results of the experiment showed that intermediate students (2nd year) made more recalls and inferences from the <u>total</u> texts than either the beginning students or the experts, and that these were indiscriminate with regard to high or low importance information in the text base. This suggests that the knowledge of the intermediate student, while reasonably extensive, is not well enough consolidated to provide selectivity in constructing a mental model for medical material.

(2) Expert selectivity and medical science knowledge. Inspection of the high relevance components of the texts revealed that they contained a high density of propositions containing "causal" and "conditional" relations. These types of relationships are particularly pertinent to conveying the substantive biomedical and clinically important content of the texts, as illustrated by the following examples:

TEXT: 1. All infections caused by either viruses, fungi, rickettsiae or bacteriae cause fever.
2. Fever may last for weeks or months.


RECALL PROTOCOL OF A PHYSICIAN:

"If the infection lasts for several weeks, it is probably caused by a bacterial agent. Viral infections may only last for a few days in some self-limiting conditions."

When analyses of subjects' recall and inferences were conditioned or restricted to those components of the texts containing causal and conditional relations, it was found that both intermediate students and experts made substantially more inferences to these text components than did beginning students. Taken in conjunction with results described above, this suggests that intermediate students recognize the medically important components of a "case" and augment these through inference in constructing a case model; however, they have trouble separating the "wheat from the chaff," and by operating on

everything, probably place a high load on cognitive processing resources, which might be expected to lead to error and inefficiency in problem solving (e.g., Anderson, Farrell, & Sauers, 1983). Experts focus on the important medical and clinical science content in building a case model.

(3) The final set of results to be discussed involves the <u>kinds</u> of inferences experts make to the medical and clinical content of the texts. When these inferences were inspected it was found that they often led in the direction of establishing or supporting goals of clinical problem solving, e.g., establishing or differentiating possible diagnostic explanations for a case, as illustrated below:

TEXT: (includes two lower order procedures)

     1. (If) recent onset of hoarseness
        (then)--- look at vocal cords.
     2. (If), hemoptysis, persistent cough,
        (then)--- examine trachea bronchial tree.

RECALL PROTOCOL OF A PHYSICIAN:

"(If) hoarseness, persistent cough and hemoptysis, ---- then
check for malignancy in lungs. Furthermore, patients past
history will be important in making such a diagnosis."

The physician's inferences in this example involve the use of a super inference category from two lower order subcategories of information and the use of frame elaboration directed towards clinical diagnosis.

Furthermore, the nature of inferences that experts made seemed tacitly or implicitly to rely on an underlying basis in the biomedical sciences, as in the following set of inferences:

TEXT: 1. Regional lymphatic drainage gives clues to primary
        site of malignancy.
      2. Enlarged lymph glands in the axilla may indicate the
        presence of breast malignancy.

RECALL PROTOCOL OF A PHYSICIAN:

"If there is lump in the breast, then it is probably due to the
problems with the lymphatic drainage. Therefore, look for lumps
in axilla."

The basic knowledge the physician used to make the inference is that a block in the lymphatic drainage system would result in an enlarged lymph gland. The rationale or support for such an inference chain rests in knowledge of both the pathophysiological reactions to cancer in the breast and also knowledge of the anatomical pathways of the lymphatic system. The basis for such inferences is usually not stated directly as experts perform the tasks, but can be uncovered when the experts are required to explain the inferences they make.

As suggested by the work of Clancey, seeing the underlying support structure for clinical activity requires more than observation of experts involved in performing their clinical activiti:s. Additional types of probes and questioning are also required.

## Summary and Discussion

This paper has attempted to trace a trend in the study of clinical reasoning from a focus on the generic form of clinical reasoning, based largely on the more apparent actions of clinical performance (e.g., hypothesis formation and data-gathering activity), *     emphasis on the mental model that physicians create to provide under.    _ing of a patient's medical condition. This is the physician's model of the patient as a malfunctioning biological organism, where knowledge of the principles of the biomedical sciences might be expected to play a significant role.

In this regard, Feltovich represented a turn from an emphasis on the global form of clinical reasoning to an emphasis on the content of medical knowledge as it is involved in the clinical reasoning process. Studies of expert and novice diagnoses in pediatric cardiology indicated that experts organize their problem representation (model) for a case according to the major pathophysiological principle (e.g., right-sided volume overload) operating in the case, and at times discriminated diseases in the differential using "deep" pathophysiological causal relationships. Lesgold and colleagues, studying expert-novice differences in radiologic diagnosis, have suggested that experts build a richer model of the particular patient, capturing anatomical and pathophysiological dynamics relevant to a case. These investigations have also provided a glimpse of the development of medical skill. Beginning residents interpret films using direct, classic interpretation rules. Advanced residents attempt more deliberately to understand and model the anatomic and physiologic issues involved in a particular patient's film. Experts, once again, use direct performance rules - but rules in which the biomedical issues advanced residents attempt to understand are directly structurally embedded. Clancey, from an artificial intelligence perspective, has studied expert diagnostic performance rules directly, and has traced back to cognitive structures that give the rules their conceptual meaning. Included are basic-science oriented "disease process frames" which represent disease hypotheses in terms of their basic science underpinnings. Clancey has also begun to illuminate how such structures direct strategic thinking within the diagnostic process. Patel and colleagues, studying clinical understanding directly, have suggested that in modelling a particular clinical case, experts make inferences that are directed at supporting clinical actions, and that the basis for these inferences is often in the basic science knowledge relevant to a case. Complementing the other studies, they have also helped to clarify the nature of the developmental process in the acquisition of expertise. Basic science knowledge must not only be learned, but it must be consolidated and reformulated (cf., Feltovich, 1983) so that its significance can be recognized and represented for clinical problems, and so that it can direct achievement of professional clinical goals.

Understanding, in terms of building a dynamic, integrated model of the patient's biomedical condition, might be expected to play somewhat different roles in different circumstances encountered by the physician. The most straightforward and most commonly attributed role for such understanding is in circumstances of unusual, rare, or multi-abnormality conditions, where the physician cannot easily take advantage of experienced, systematic regularities in the case. In these circumstances, basic biomedical understanding provides a generative capacity for reasoning and inference. However, the present paper has attempted to outline the importance of basic science understanding even in the more routinized conditions of medical practice. Because reasoning from basic principles is not cognitively efficient, it is expected that experts, people with considerable clinical experience, will reformulate or "compile" their basic knowledge to take advantage of recurrent patterns of regularity (e.g., similar presenting complaints - Feltovich & Barrows, 1984). However, the studies reviewed suggest that even in these situations basic science understanding is retained, but is implicit in the knowledge structures (e.g., Clancey's "disease process frames") and inference rules that experts develop. In diagnostic situations such understanding shows through in such places as the configuration of potential case explanations and the relational nature of disease expectations. It is our speculation that biomedical understanding will have even more impact in the more manipulative aspects of clinical care, e.g., management and treatment of the patient. These aspects accentuate less the identification of abnormality and more the manipulation of a defective system to bring it as much as possible back to functional state. Detailed understanding of the malfunctional state would seem a necessary prerequisite that influences efficiency and success.

In addition to studies of case management as it relates to biomedical understanding, several other lines of investigation are suggested by the studies outlined in the paper. One is more detailed investigation of the longitudinal process of incorporating basic science knowledge into clinical performance, and study of the knowledge reformulations that occur. Another is the study of misconceptions of biomedical science that medical students and practitioners hold and their influences on clinical performance, as this might clarify the role of more correct conceptions. Yet another is the comparison of conditions of training for the basic sciences to investigate how basic science and clinical instruction can best be blended to promote ultimate practical competence.

References

Anderson, J. R.  Acquisition of cognitive skill.  *Psychological Review*,
   1982, 89, 369-406.

Anderson, J. R.,  Farrell, R., & Sauers, R.  *Learning to program in LISP*.
   ONR Tech. Rept. #83-1, ONR Contract No. N00014-81-C-0335.  Pittsburgh,
   PA:  Carnegie Mellon Univ., Sept., 1983.

Barrows, H. S., Feightner, J. W., Neufeld, V. R., & Norman, G. R.  *Analysis
   of the clinical methods of medical students and physicians*.  Final
   report, Ontario Department of Health Grants (ODH-PR-273 and ODH-DM-226),
   Hamilton, Ontario, Canada:  McMaster University, March, 1978.

Barrows, H. S., Norman, G. R. Neufeld, V. R., and Feightner, J. W.  The
   clinical reasoning of randomly selected physicians in general medical
   practice.  *Clinical and Investigative Medicine*, 1982, 5, 49-55.

Bobrow, D. G., & Norman, D. A.  Some principles of memory schemata.  In
   D. G. Bobrow & A. Collins (Eds.), *Representation and understanding:
   Studies in cognitive science*.  New York:  Academic Press, 1975.

Chase, W. G. & Chi, M. T. H.  Cognitive Skill:  Implications for spatial
   skill in large-scale environments.  In J. Harvey (Ed.), *Cognition,
   social behavior, and the environment*.  Potomac MD:  Lawrence Erlbaum
   Associates, 1980.

Chase, W. G., & Simon, H. A.  Perception in chess.  *Cognitive Psychology*,
   1973, 1, 55-81.

Chi, M. T. H., Feltovich, P. J., & Glaser, R.  Representation of physics
   knowledge by Experts and Novices.  *Cognitive Science*, 1981, 5, 121-152.

Clancey, W. J.  Acquiring, representing, and evaluating a competence model
   of diagnostic strategy.  Prepublication document, Stanford University,
   1984.

Clancey, W. J.  Tutoring rules for guiding a case - method dialogue.
   *International Journal of Man-Machine Studies*, 1979, 11, 25-49.

Clancey, W. J.  *Methodology for Building an intelligent tutoring system*.
   Technical Report STAN-CS-81-894, HPP-81-18, Stanford University, 1981.
   (Also in *Methods and tactics in cognitive science*, Kintsch, Polson and
   Miller (Eds.), Lawrence Erlbaum Associates, Hillsdale, N.J., 1984.

Clancey, W. J.  The epistemology of a rule-based expert system:  A framework
   for explanation.  *Artificial Intelligence*, 1983, 20(3), 215-251.  (Also
   to appear in Buchanan and Shortliffe (Eds.), *Rule-based expert systems:
   The MYCIN experiments of the Stanford Heuristic Programming Project*.
   Reading, MA:  Addison-Wesley, 1984.

Clan. ey, W. J. and Letsinger, R.  Neomycin:  Reconfiguring a rule-based
    expert system for application to teaching, in Proceedings of the Seventh
    International Joint Conference on Artificial Intelligence, pages
    829-836, 1981.  (Revised version in Clancey and Shortliffe (Eds.),
    Readings in medical artificial intelligence:  The first decade,
    Addison-Wesley, 1984).

Collins, A., Brown, J. S., & Larkin, K. M.  Inference in Text
    Understanding.  In R. Spiro, R. Bruce, & W. Brewer (eds), Theoretical
    Issues in Reading Comprehension.  Hillsdale, N.J.:  Erlbaum, 1978.

deKleer, J.  Multiple representations of knowledge in a mechanics problems
    solver.  Proc. of the 5th Int. Conf. on Artifical Intelligence.
    Cambridge, MA:  MIT Press, 1977.

Elstein, A. S., Shulman, L. S. , & Sprafka, S. A.  Medical problem solving.
    Cambridge, Mass.:  Harvard University Press, 1978.

Feigenbaum, E. A.  The art of artificial intelligence:  I.  Themes and case
    studies of knowledge engineering.  In Proceeding of the Fifth
    International Joint Conference on Artificial Intelligence, August, 1977.

Feltovich, P. J.  Expertise:  Reorganizing and Refining Knowledge for Use.
    Professions Education Researcher Notes, 1983, 4(3), 5-9.

Feltovich, P. J. and Barrows, H.S.  Issues of generality in medical problem
    solving.  In H. G. Schmidt & M. L. DeVolder (Eds.), Tutorials in problem
    based learning.  Assen, Holland:  Van Gorcum, 1984.

Feltovich, P. J.  Knowle_ge based components of expertise in medical
    diagnosis.  Doctoral Dissertation, University of Minnesota, 1981.
    (Also, Feltovich, P. J., Johnson, P. E., Moller, J. H., and Swanson,
    D. B.  LCS:  the role and development of medical knowledge in diagnostic
    expertise.  In W. J. Clancey & E. H. Shortliffe (Eds.), Readings in
    medical artificial intelligence:  The first decade.  Reading, MA:
    Addison Wesley, 1984.

Frederiksen, C. H.  Representing logical and semantic structure of knowledge
    acquired from discourse.  Cognitive Psychology, 1975, 7:371-458.

Frederiksen, C.H., Frederiksen, J.D., Bracewell, R.J.  Discourse analysis of
    children's text production.  To appear in A. Matsuhasi (Ed.), Writing in
    real time.  Longman:  London, in press.

Greeno, J. G.  Process of understanding in problem solving.  In N. J.
    Castellan, D. B. Pisoni, & G. R. Potts (Eds.), Cognitive theory
    (Vol 2).  Hillsdale, N.J.:  Lawrence Erlbaum Associates, 1977.

Kintsch, W.  The representation of meaning in memory.  Hillsdale, N.J.:
    Lawrence Erlbaum Associates, 1974.

Kintsch, W. and Greeno, J.G.  Understanding and solving word arithmetic
    problems.  Psychological Review (in press).

Lesgold, A. M. Acquiring expertise. In J.R. Anderson & S.M. Kosslyn (Eds.), Tutorials in learning and memory: Essays in honor of Gordon Bower. San Francisco: W.H. Freeman, 1984.

Lesgold, A. M., Feltovich, P. J., Glaser, R., & Wang, Y. The acquisition of perceptual diagnostic skill in radiology. Tech. Report No. PDS-1, ONR Contract No. N00014-79-C-0215. Pittsburgh, PA:LRDC, University of Pittsburgh, 1981.

Lesgold, A. M., Rubinson, H., Feltovich, P. J., Glaser, R., & Klopfer, D. Expertise in a complex skill: Diagnosing x-ray pictures. To appear in M. T. H. Chi, R. Glaser, and M. Farr (Eds.), The nature of expertise, forthcoming.

McDermott, J. & Larkin, J. H. Re-representing text book physics problems. Proceedings of the 2nd National Conference of the Canadian Society for Computational Studies of Intelligence. Toronto: Univ. Toronto Press, 1978.

Minsky, M. A. A framework for representing knowledge. In P. H. Winston, (Ed.). The psychology of computer vision. New York: McGraw-Hill, 1975.

Patel, V.L. Expert diagnostic reasoning: an interactive process of case comprehension. Professions Education Researcher Notes, 5(3):7-11, 1984.

Patel, V. L. & Frederiksen, C. H. Cognitive processes in comprehension and knowledge acquisition by medical students and physicians. In H. G. Schmidt and M. L. DeVolder (Eds.). Tutorials in problem based learning. Assen, Holland: Van Gorcum, 1984a.

Patel, V.L. and Frederiksen, C.H. Role of prior knowledge in expert/novice differences in comprehension of medical texts. Report No. CME84-5, Centre for Medical Education, McGill University, Montreal, 1984b. (Paper presented at AERA meeting, 1984).

Patel, V. L., Frederiksen, C. H., & Groen, G. J. Relationship between comprehension and medical problem solving. Paper presented as part of a symposium on "The essence of clinical competence: Psychological studies of expert reasoning in medicine" at the Research in Medical Education Conference of the Association of American Medical Colleges, Washington, D.C., 1983.

Patel, V.L. and Groen, G.J. Knowledge based solution strategies in medical diagnosis. Report No. CME84-7. Centre for Medical Education, McGill University, 1984. (Paper presented at Cognitive Science Society Meeting Boulder, Colorado, 1984).

Patil, R. S., Szolovits, P., Schwartz, W. B.: Modeling knowledge of the patient in acid-base and electrolyte disorders, in Szolovits, P. (Ed.): Artificial Intelligence in Medicine. Boulder, Colorado, Westview Press 1982, 191-226.

Pople, H. E.: Heuristic methods for using structure on ill-structured problems: The structuring of medical diagnostics, in Szolovits, P. (Ed.): Artificial Intelligence in Medicine. Boulder, Colorado, Westview Press, 1982, 119-190.

Sacerdoti, E. D. Planning in a hierarchy of abstraction spaces. Artificial Intelligence. 1974, 5, 115-135.

Sacerdoti, E. D. A structure of plans for behavior. New York: American Elsevier, 1977.

Simon, H. A., & Chase, W. G. Skill in chess. American Scientist, 1973, 61, 394-403.

Swanson, D. B. Computer simulation of expert problem solving in medical diagnosis. Ph.D. Thesis, University of Minnesota, 1978.

Swanson, D. B., Feltovich, P. J., & Johnson, P. E. Psychological analysis of physician expertise: Implications for the design of decision support systems. In C. B. Shires and H. Wolfe (Eds.), MEDINFO 77, North Holland, 1977.

Winograd, T. Frame representation and the declarative/procedural controversy. In D. G. Bobrow & A. Collins (Eds.). Representation and Understanding. New York: Academic Press, 1975.

Winograd, T. Understanding natural language. Edinburgh: Edinburgh University Press, 1972.

Wortman, P. M. Medical diagnosis: An information processing approach. Computers and Biomedical Research, 1972, 5, 315-328.

Wortman, P. M., & Greenberg, L. D. Coding, recoding, and decoding of hierarchical information in long-term memory. Journal of Verbal Learning and Verbal Behavior, 1971, 10, 234-243.

Yu, V. L., Fagan, L. M., Wraith, S. M., Clancey, W. J., Scott, A. C., Hanigan, J. F., Blum, R. L., Buchanan, B. G., & Cohen, S. N. Antimicrobial selection by computer: A blinded evaluation by infectious disease experts. JAMA, 1979. 242, 1279-1282.

J800m/23-26