Issues of generality in medical problem solving

Paul J. Feltovich & Howard S. Barrows
Southern Illinois University
U.S.A.

Despite much effort at trying to identify them, few general, cross-situational aspects of clinical reasoning have been discovered beyond hypothetico-deduction and the early generation of clinical hypotheses. The importance of good hypothesis sets leads to questions about the nature of problem representation in medicine. Results from studies of problem representation in fields other than medicine suggest that clinical reasoning, at levels more specific than hypothetico-deduction, is not likely to be generic and uniform, but is likely to be conditioned on classes of medical problems and problem components that practitioners recognize and include in their mental problem representations. A way of characterizing medical problem representation, involving illness scripts, is proposed for adressing the issue of conditional generality. The illness script is an integrated model of medical abnormalities, specifying illness in terms of enabling conditions, faults, and their consequences. It is proposed that areas of regularity in clinical reasoning may be associated with components of this script.

The search for components of medical clinical reasoning which are consistent across situations and problems has met with mixed success and has yielded a rather confusing picture. On the one hand, some aspects of clinical reasoning are clearly general and cross-situational. Studies have contributed to a consensus about the broad form of clinical reasoning (Barrows, Feightner, Neufeld & Norman, 1978; Barrows, Norman, Neufeld & Feightner, 1982; Elstein, Shulman & Sprafka, 1978; Neufeld, Norman, Feightner & Barrows, 1981). Cues in patient data (signs, symptoms, etc.) suggest diagnostic hypotheses which are, in turn, projected toward and compared with subsequent data of a case. This basic hypothetico-deductive reasoning method (generate and test hypotheses) is shared by different specialties of physicians and by experienced, inexperienced, diagnostically successful, and unsuccessful physicians alike. As well as being general to types and characteristics of physicians, the hypothesis testing method is applied to different kinds of medical cases, emphasizing different

domains of medical subject-matter content.
However, more fine-grained analyses of the clinical reasoning process have produced less uniform results. These analyses have focussed on three basic kinds of components of the general hypothesis testing process.

Hypothesis-related parameters

Investigations of the parameters of hypothesis testing are concerned with the timing and number of hypotheses created during a diagnostic encounter. These parameters include such elements as time or percentage of patient data items to creation of first hypotheses, total number of hypotheses considered, number of hypotheses actively considered at any one time, etc. Such characteristics of reasoning have been found to be similar across different specialties of physicians and clinicians with different levels of experience (Barrows et al., 1982; Elstein et al., 1978; Neufeld et al., 1981). While hypothesis related parameters are common to different kinds of physicians, their importance to the clinical reasoning process is unclear. For example, measures of numbers of hypotheses, total generated during a case, or number actively considered at different case points, are not predictive of the quality of outcomes of clinical reasoning, that is, quality of final diagnosis or management plan. Two measures of timing in hypotheses generation, time to first hypothesis and time to correct hypothesis, have been found to be moderately predictive of outcome quality (Barrows et al., 1978, time to correct hypothesis; Barrows et al., 1982, both measures), but investigations showing neither relationship also exist (e.g., Neufeld et al., 1981). Furthermore, even though some hypotheses are generally proposed early in a clinical encounter and the number of hypotheses considered falls within a fairly small range (approximatly 3-7), within these bounds the timing and number of hypotheses does show variability across different cases (Barrows et al., 1978, 1982; Elstein et al., 1978).

Data-gathering

Components of the reasoning process associated with data-gathering to inquire about and test hypotheses show an even more muddled pattern in clinical use. Despite particular variations within different studies, measures used to assess data-gathering characteristics have centered around the general themes of thoroughness and efficiency in information gathering, and activeness and accuracy in interpretation of the information collected (e.g., Elstein et al., 1978; Neufeld et al., 1981).
Measures of thoroughness in data-gathering are usually defined with regard to a set of criterial patient data items, items judged by some standard (e.g., a panel of experts) to be critical for successful solution of a patient case. Thoroughness refers

to the acquisition of this criterial data items during a patient encounter. Efficiency refers to some ratio of criterial to non-criterial items elicited by the clinician, and attempts to measure the extent to which the clinician is focusing on important data items to the exclusion of others. Activeness involves the extent to which data collected are evaluated in relationship to hypotheses being considered, in order to test the appropriateness of the hypotheses. Accuracy in interpretation attempts to determine if a clinician interprets data items appropriately as confirmatory or disconfirmatory of particular hypotheses.

The value of data-gathering measures for revealing important aspects of clinical reasoning or in addressing the issue of generality in reasoning is dubious. They are general to practitioners in that such data-gathering activities are common to and do not discriminate among different types of clinicians, clinicians with different levels of experience, or those of different peer-judged proficiency (e.g., Barrows et al., 1982; Elstein et al., 1978; Neufeld et al., 1981). However, such measures appear to hold little importance in the quality of the reasoning process. With the exception of "accuracy of interpretation" as reported by Elstein et al., (1978) no other data-gathering measures of the kind discussed have been shown to influence the quality of clinical outcomes, and the importance of accuracy of interpretation itself must be questioned (Barrows et al., 1978, 1982; Neufeld et al., 1981, Norman et al., 1982). Furthermore, there is indication that traditional data-gathering measures are unreliable, in the sense that the same clinician does not performsimilarly on these measures during two presentations of essentially the same case (Norman et al., 1982). Finally, case-dependency in data-gathering activities has been indicated, with performance on these activities varying substantially across different clinical cases (Barrows et al., 1978; Elstein et al., 1978).

Quality of hypotheses

By far the best indicator of competent clinical reasoning and best clue to the nature of the reasoning process appears to be the hypothesis sets that clinicians create during the course of a case to coordinate and direct the process of clinical inquiry. The quality (as judged by expert standards) of the hypothesis sets has been shown to increase systematically with increasing clinical experience and training (Barrows et al., 1978; Feltovich, Johnson, Moller & Swanson, 1983; Neufeld et al., 1981). Quality of hypothesis sets is also the strongest indicator of quality in the clinical outcomes of reasoning such as final diagnosis, management, and therapy (Barrows et al., 1978, 1982; Neufeld et al., 1981; Norman et al., 1982). In addition, quality of hypotheses has been found to be considerably more consistent over multiple presentations of the same case than are other measures of reasoning such as data-gathering activities (Norman et al., 1982). Yet, despite its validity in discriminating experi-

ence levels and in predicting successful outcome, and its apparent reliability, the cross-situational nature of hypothesis quality is still an issue. While some investigations have found no correlations between scores of hypothesis quality across different cases (Barrows et al., 1978; Williams, Vu, Barrows & Verhulst, 1984), others have found more consistency across cases (Norman et al., 1982).

As a result of the investigations of clinical reasoning that have been discussed, an overall picture of the reasoning process, its critical components, and its points of generality can be ventured. At a superficial level of description, the process of clinical reasoning is nearly universal, and consists of the hypothetico-deductive method of generating and testing hypotheses. Furthermore, clinicians generate hypotheses early in a clinical encounter to coordinate and guide their activity, and the earliness with which they construct a framework of hypotheses, especially a good one, is central in determining successful reasoning. Beyond this set of facts, aspects of the clinical reasoning process are labile. In particular, it appears extremely difficult to determine necessary and critical data-gathering maneuvers that reveal anything important about the clinical reasoning process, although accuracy of interpretation of findings in relation to hypotheses may be indicative of successful performance.

Although some consistent characteristics of the clinical reasoning process can be identified, they are not particularly satisfying for understanding the process of reasoning or useful for communicating it when training future clinicians. They are deficient for understanding in that they leave critical components of reasoning unexplained. For example, how are good hypotheses generated, and what is the nature of a good hypothesis set? They are deficient for training for much the same reason. It is not greatly helpful to tell novice diagnosticians that they must generate and test hypotheses and, most importantly, generate good hypotheses early, without some better specification of the mechanics involved.

Furthermore, the consistent characteristics of clinical reasoning which have been identified, beyond hypothetico-deduction, do not point to a singular, homogeneous process to be taught or understood. The criticality of a good conceptualization of the problem, as reflected in high-quality hypothesis sets, suggests the importance of the clinician's medical knowledge-base; knowledge-base contributes to knowing which abnormalities and processes of abnormality share common underlying characteristics andmay be confused with each other clinically (Feltovich et al., 1983). the issue of accuracy of interpretation of clinical data with respect to hypotheses also implicates the role of the medical knowledge base, since it specifies the relationship between processes of abnormality and clinical manifestations to be expected or not expected under their influence. To the extent knowledge-based factors are critical in reasoning, it might be expected that the quality of reasoning might be uneven across particular medical content areas, reflecting the idiosyncratic training and

exposure history of the clinician. Troublesome but conflicting findings associated with case-specificity in clinical reasoning contribute to this interpretation; whether or not inter-case generality in reasoning is found within particular studies, probably depends on the medical similarity of the cases involved. How can we reconcile notions of knowledge-base importance, case and content dependence, and the critical nature of hypothesis sets with any useful conception of generality in clinical reasoning? The importance of a strong conceptualization of the problem, as reflected in quality hypothesis sets, suggests that one possible answer might lie in the construct of 'problem representation' which, in studies of problem-solving outside the field of medicine, has been shown to greatly influence reasoningand problem solution and to differ considerably between experts and novices - much as hypothesis sets do in medicine. Problem representation and its influences in problem solving will be discussed next. Possible implications of problem representation for the investigation of generality in clinical reasoning will then be addressed in the remainder of the paper.

Problem representation

Studies of problem solving and expert-novice differences in problem solving in fields other than medicine have pointed to the importance of an individual's problem representation for guiding reasoning and determining successful problem solution. What is a problem representation? One characterization which has been given defines a problem representation as "a cognitive structure corresponding to a problem, (which is) constructed by a solver on the basis of his domain-related knowledge and its organization" (Chi, Feltovich & Glaser, 1981). A problem representation is the solver's internal model of the problem, containing the solver's conception of problem elements, their relationships to each other, the goals of problem solving, etc. It is well known that differences in problem representation among problem solvers greatly influence the ease or difficulty with which a problem can be solved (e.g., Hayes & Simon, 1976; Newell & Simon, 1972). Furthermore, the problem representations of experts and novices within given fields have been shown to be characteristically different. Expert representations contain <u>abstractions</u> of the problem setting that capture the deeper commonalities of superficially disparate events. For example, when viewing a chess-board, during the course of play, a chess expert aggregates particular pieces of the board into larger recognizable piece configurations reflecting areas of strategic vulnerability and strength, positions supporting maneuvers of attack and defense (Chase and Simon, 1973; de Groot, 1965). Similarly, when electronics experts view an electronic circuit board, they create a model of the board that organizes elementary circuit components into the major structural units (e.g., rectifiers, amplifiers) which determine the global function of the circuit (Egan & Schwartz, 1979). Nov-

ices, in contrast, represent elementary circuit components by spatial proximity, independently of their integrated electrical function.

These results indicate that the abstractions in expert representation are task and goal-oriented. Further indication of this goal-orientation comes from studies of the games of baseball and bridge. When baseball aficionados observe portions of a baseball game, their representation of the event is structured by major goal-related sequences of the game such as advancing runners, scoring runs, and preventing scoring. Novice representation involves less integral components, e.g., the weather and crowd mood, and does not capture goal-directed sequential activity well (Spilich, Vesonder, Chiesi & Voss, 1979). Expert bridge players organize the cards in their own and opponents' hands by suit, across hands - an organization useful for playing the game of bridge. Novices organize cards by order of rank within hands, an organization less supportive of useful strategy within the game (e.g., Charness, 1979).

Finally, studies of problem solving in physics suggest the importance in problem representation of fundamental laws and principles of a field. Expert physics problem solvers represent physics problems as instances of the major physics laws (e.g., conservation of energy) applicable to the physical situation, and in doing so recode and transform overt problem features into more abstract terms (e.g., "before and after situation", "well defined initial conditions"), so that their consistency with proposed laws and principles can be tested (e.g., Chi, Glaser & Rees, 1982; Chi, Feltovich & Glaser, 1981). Novice problem representation in physics is more literal, fragmented, and tied to overt features of a problem (e.g., representations of problems as 'spring' or 'pulley' problems because of the physical presence of these objects in the problem description).

The discussion so far has suggested that in a number of different problem areas, expert problem representation develops in the direction of increasing value for accomplishing the tasks of a field. Expert representations are useful for playing chess or bridge, for solving physics problems, etc. For example, when a chess expert characterizes a game-situation abstractly with regard to its elements of vulnerability and strength, the appropriate moves from the position are greatly constrained. When the expert physicist characterizes a physical situation abstractly as an instance of a physics law, the law itself substantially directs the form and application of equations that will be used in solution (Chi et al., 1981, Larkin, 1978). Novices also impose organization on problems, but the nature of their representations (e.g., proximity in electronic circuits, literal objects in physics) may be expected not to constrain and guide subsequent solutions as well.

One way to conceptualize the influence of problem representation on reasoning and problem solving is with regard to a distinction between 'weak' and 'strong' methods of problem solution, originally proposed by Newell (Newell, 1969, 1973). Weak methods

of problem solution are general methods that can be made to apply almost universally; the conditions which must be met in a problem for the methods to be legitimately applied are non-stringent. Examples of weak methods are 'means-ends analysis' (comparing one's current progress in achieving a goal to the desired goal and trying to close the difference) and 'generate and test' a kind of hypothesis testing in which possible solutions to a problem are proposed and then tested for their appropriateness. Although weak methods are widely applicable, they offer little guarantee of success in any particular application; hence their characterization as 'weak'. In contrast, strong methods are those that have more restrictive conditions which must be met for their applicability and, hence, are tailored to particular closed environments. Within their domain of applicability, however, they are highly effective in achieving solution. An example of a strong method is the mathematical procedure for solving quadratic equations. It applies only to an abstract class of equations, but within equations of specified form, successful solution is assured.

The studies of problem solving reviewed in this section suggest that the development of competent problem solving ability is associated with the development of useful strong methods. Experts represent problems and problem features in terms of abstract categories and partitions that define the bounds of directed problem solving procedures. Novices also attempt to partition problems within classes, but do so in less useful ways (e.g., 'pulley' problems in physics).

There are several implications of this characterization of problem solving for the attempt to find areas of generality in medical clinical reasoning and problem solving. First, it appears that the weak method of clinical reasoning is hypothetico-deduction, and that as a weak method it is applied nearly universally. Second, with the exclusion of hypothetico-deduction, other characterizations of generality are likely to be conditional, that is, to be found only in relation to the classes of problems and problem components individuals create with experience and recognize as pertaining within a setting. And finally, in order to characterize areas of conditional regularity in medical problem solving, we will have to determine the ways that medical problem solvers of different kinds represent clinical problems. It is necessary to determine the abstract partitions of the space of medical problems that clinicians utilize in representation, so that regularity of reasoning within these classes can be identified.

Problem representation in Medicine

Little is known about the ways that medical clinicians represent problems, or the associated categories and abstractions they use in characterizing problems or problem components.

However, if the studies of problem representation discussed

above can be taken as a guide, there are characteristics that at least expert representation might be expected to show. First, we would expect expert problem representation to capture commonality across different problems at some level of abstraction. Second, we would expect expert problem representation to involve basic 'laws' and principles of medicine; candidates are systematic relationships provided by the basis sciences in medicine, for example, anatomy and pathophysiology. Third, there should be somekind of problem feature consolidation, whereby the evidence and patient data provided by a case are transformed and aggregated into different types, which serve different roles and purposes. Finally, expert problem representation should be useful for serving the diverse goals of practice, including diagnosis, but also management, prognosis, therapy and so forth.

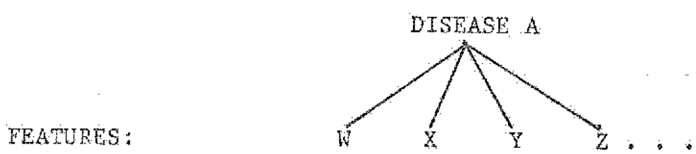

Figure 1: Templates and slot matching

Although problem representations has not been addressed directly in medicine, the practices of researchers, in terms of the operations they measure in studying clinical reasoning and the standards they use for defining competent performance, reflect an implicit conception of the manner in which clinical problems are represented.

This conception of problem representation might be characterized as "templates and slot matching" (Figure 1). The clinician is viewed as structuring a problem in relation to disease or abnormality forms, each of which specifies associated problem features (e.g., patient data) which should occur in the presence of the condition. The templates are generated as hypotheses, and further clinical data-gathering is directed at filling in and matching related feature slots. When researchers create matrices crossing disease conditions with problem features, solicit expert judgment as to the confirmatory or disconfirmatory relationship between features and conditions, and define criteria for quality of feature interpretation and information gathering with regard to these standards, then they are assuming a problem representation such as templates with slot filling (e.g., Elstein et al., 1978). It is also the kind of representation which has been used explicitly in early attempts at modeling clinical reasoning on computers (Pople, 1977; Miller, Pople & Myers, 1982).

What is wrong with this conception of problem representation in medicine? There are four areas in which it seems particularly limited. First, it does not capture well the interdependency in problem features and data items. Patient data items are oftentimes interrelated and partially redundant, with different items all bearing on the same characteristic of the patient's condi-

tion (e.g., history items of frequent upper respiratory infections, certain heart sounds, and suggestive x-ray findings all indicating increased pulmonary blood flow). These interrelationships need to be captured, since clinicians can often infer some findings from others, and do not need to observe all relevant evidence in order to reach correct conclusions. Second, it does not represent causal relationships among items. All features are treated homogeneously, in situations where causal subordination and superordination may exist. When causal relationships among features are not represented, the total picture of expected features cannot be adjusted for extreme or unusual values of individual features. Third, there is no distinction among different problem features and patient data items regarding their different types and functions in the clinical reasoning process. And, finally, the template conception of problem representation does not capture the dynamic nature of illness as a process that is somehow acquired, and develops and changes with time.

Some examples of the kinds of relationships among features that need to be captured in problem representation can be illustrated by an example suggested by research by the first author and colleagues in radiological diagnosis (Lesgold, Feltovich, Glaser & Wang, 1981). Figure 2 shows a disease template and associated features that might be proposed for the condition of atelectasis, that is, the collapse of part of a lung. The template specifies features that might be expected to appear on the x-ray (e.g., mediastinal shift, elevated diaphragms) or as part of the patient's history (e.g., chronic pulmonary infections). One can envision a template-matching scheme for diagnosis that generates atelectasis as an hypothesis, based on some subset of these features, and proceeds to match against the rest.

However, as indicated in Figure 2, there is a great deal of structure in the specified template which influences how it will be applied and interpreted. There is structure at the level of the case data themselves. The x-ray items are 'illness manifestations', that is, they are consequences of the condition in the patient, while the historical items are 'predisposing factors', conditions which incline an individual toward acquisition of the abnormality. These are very different kinds of items and might be expected to function differently in the diagnostic process. Furthermore, there is structure among the manifestations appearing on x-ray. These features are interrelated by principles of anatomy, biophysics, etc, so that, for example, the degree of expected elevation in the diaphragm and the amount of hilar displacement can be adjusted to degrees of compensatory hyperinflation in the non-collapsed lobes of the lung. The particular x-ray manifestations of any one case of collapse depend on a subtle interaction of such rule-governed factors. Finally, when all of the features specified in the template are taken together, different 'stories' or scenarios of the illness and its development can be produced. For example, a history of alcoholism is a predisposing factor because alcoholics sometimes inhale their own oral and nasal mucous, leading to obstruction of air-passages

in the lung and subsequent collapse. Heavy smoking leads to emphysema, overinflation of lung lobes, which can compress and collapse other lobes. Long term, recurrent pulmonary infections and inflammation can cause severe scarring and resulting contractions of the lung.

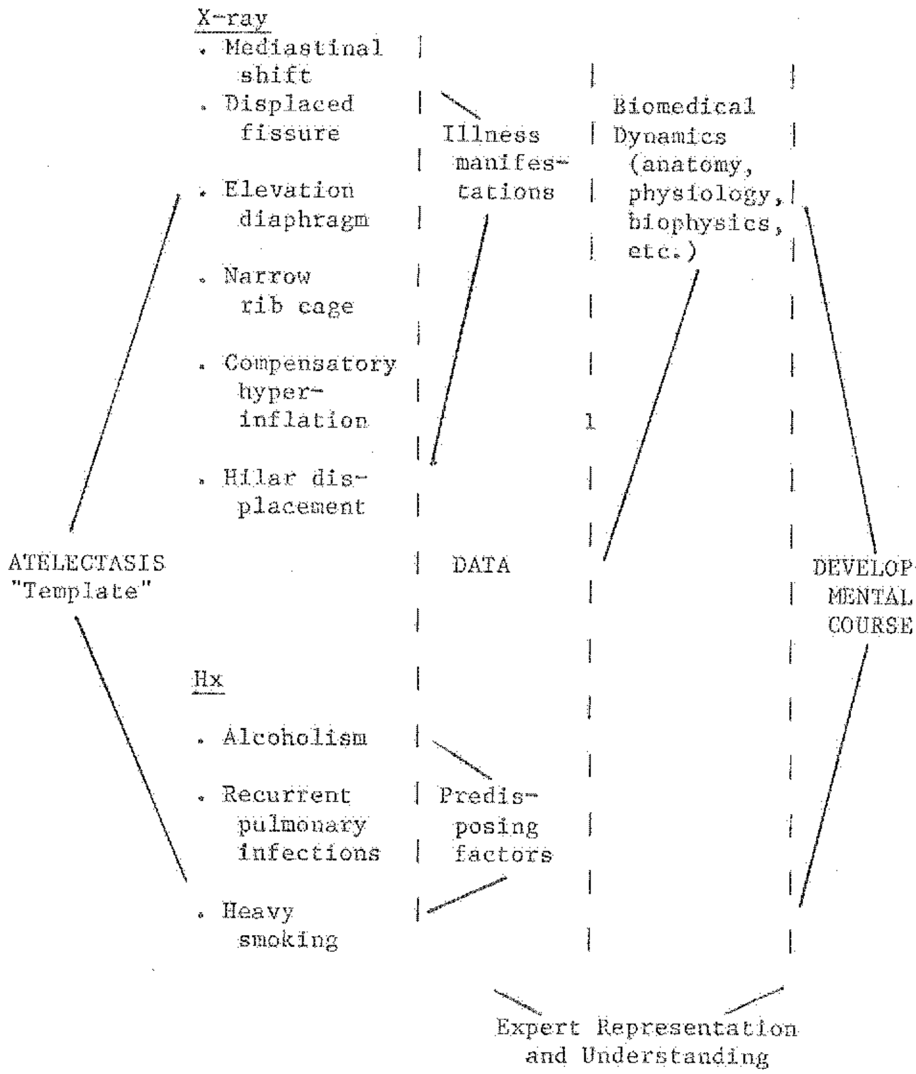

Figure 2: Examples of structure within a template, atelectasis

All of these levels of structure exist in a template of an illness and its associated clinical features, and may be candidates for the kinds of organizations competent practitioners recognize and utilize in representing a clinical problem. In turn, they are also candidates for meaningful dimensions on which clinicians can partition sets of problems and establish related consistent lines

of clinical inquiry. However, investigation of this possibility will require new and richer ways of thinking about problem representation in medicine.

```
Illness script ---> enabling conditions, fault, consequences

   enabling conditions ---> predisposing factors, boundary
                            conditions, hereditary factors, etc.

           predisposing factors ---> compromised host factors,
                                     travel, drugs, etc.

           boundary conditions ---> age, sex, etc.

   faults ---> invasion of tissue by pathogenic organism
                 inadequate nutrient supply,
                 inability of tissue to survive/thrive, etc.

   consequences ---> signs, symptoms
           symptoms ---> initial symptoms, symptom pattern...
```

Basic sciences (anatomy physiology, etc.) are constraints on how the pieces can be put together.

Figure 3: Illness script

We have recently begun to explore a device for studying clinical problem representation which we call an 'illness script' (Figure 3). This is a model for representing medical illness or abnormality which proposes general components of illness and their interrelationships. (In our thinking about the illness script, we are indebted to Clancey, in press, which serves as a foundation.)

At its most general level of description, the script proposes that an illness can be characterized by three component parts - enabling conditions, faults and a set of consequences. These are abstract categories of illness features, and each can be further specified into more specific categories. Enabling conditions are illness features associated with the acquisition of illness. They include more specific categories such as predisposing factors (e.g., compromised host factors, unusual travel, hereditary factors), which can make an individual more susceptible than usual to illness in general or to particular illnesses. Faults are the major real malfunctions in illness, and are characterized abstractly as major different types, e.g., invasion of tissue by pathogenic organisms (a site of infection), inadequate nutrient supply (metabolic disorder), or inability of tissue to survive/ thrive (degeneration). Consequences are the secondary conse-

quences of faults within the organism, and generally comprise different types of signs and symptoms.

In the system of clinical reasoning we envision, a clinician attempts to represent and understand a patient problem by constructing an integrated script or scenario for how the patient's condition came to be, its major points of malfunction, and its subsequent associated consequences. We also propose that basic science knowledge (physiology, etc.) plays an integral role in the construction of such scripts, by constraining and guiding the manner in which script components can be structured together. In this system, it is the illness script that holds preeminence, and other clinical activities such as hypothesis generation and information gathering must be understood in relation to script development.

We have only begun studies to explore the nature of medical problem representation and to determine if the illness script is a useful conceptualization for how problems are represented, but we can show an example of the kind of results we might find. The example reported is from a study in which we are using problem syntheses and summaries to investigate the structure of clinician's problem representations. Physicians take histories from and conduct physical examinations on patient-actors (individuals trained to mimic the clinical presentation of real patients), and are then asked to give summaries of the patient data they have collected and a brief synopsis of the patient's underlying biomedical condition as they understand it.

Q. Please review all patient data you consider to be important to the case.

A. This is a thirty-year-old, married woman, who has been in good health, who has had two episodes of (uh) shortness of breath in the last week. These were characterized by hyperventilation, and subsequently with chest discomfort and some numbness and tingling. There has been a modest amount of stress in her life recently. (uh) There appears to be nothing else in her history which suggests any underlying pulmonary disease. (uh) (pause) She's not taking any birth control pills, she's not a smoker, she has no evidence of (um) thromboembolism. (um) (pause) She's generally in good health. Okay?

Q. Please summarize important aspects of the patient's biomedical condition: for example, anatomy of pathophysiology.

A. (uh) I think she has episodes of hyperventilation (uh) which appear to be (uh) initiated by stress, (um) manifested by (uh) (pause) tachypnea which results in (uh) excess CO_2 (uh) loss (uh) and results in some peripheral (uh) (pause) neurologic symptoms. Some Chest pain.

Figure 4: Examples of physician summaries

Figure 4 gives these summaries given by a cardiologist for a young woman complaining of 'difficult breathing'. When he is asked to summarize the case data (first summary), the structure in the physician's response is not clear (at least to a non-medical eye). However when these data are compared with the second summary, the biomedical condition, a distinct pattern emerges. The second response is a high-level ouline of an illness script which can be characterized as follows: Because of stressful situational factors in the patient's life (enabling condition), the patient becomes anxious and breathes strangely, blowing off CO_2 (the fault). The loss of CO_2 results in neurological symptoms and chest pain (consequences). The correspondence of the script to more elemental items given in the first part of the first summary (e.g., numbness and tingling) can be seen. Furthermore, rudiments of another, alternatively considered, script can be seen in the first summary. To a medically knowledgable observer, this script involves the taking of birth control pills leading to the formation of thromboembolism in a leg, migrating to the lungs, and causing the patient's difficulty in breathing.

Of course, more work needs to be dome before the role and structure of illness scripts in clinical problem representation can be verified and specified. However, this way of thinking about problem representation has several implications for the search for conditional generality in clinical reasoning. The first is with regard to good hypotheses and the partitioning of the problem space. The illness script specifies abstract categories of problem features (e.g., enabling conditions, faults), occurring across different illnesses. Medical conditions which are similar on any of these problem dimensions might constitute a problem category in the problem representation of the physician, and serve the formation of a good set of hypotheses. Other implications are with regard to clinical 'strategy' and data-gathering. First, it is within problem categories based on similarity in script components that we might expect to see conditional regularity in reasoning. For example, if illnesses in a category are defined by common faults and consequences, we would expect inquiry to focus on establishing differences in boundary conditions of various kinds. Second, it is an integrated script that needs to be constructed for a patient, and there is leeway in the <u>particular</u> data items which may be used to complete it. For example, in questioning the patient, the clinician in our example asked numerous stress-related questions concerning husband, family, work, etc. At the level of the problem representation, it was the <u>stress</u> that needed to be verified, and particular questions and questioning maneuvers could vary. Finally, the abstract representation of different types of illness features and feature roles within a script may support another level of strategy, not conditioned on the particular illness sets being considered as hypotheses, but, rather, directed at the control and management of the global process of hypothetico-deduction. For example,when physicians judge that they are becoming overly focused within

diagnosis, they may choose to do a general 'scan' for unusual predisposing factors, e.g., compromised host conditions, unusual drugs, or travel, etc. This level of strategy would also be conditional, but on specified aspects of the reasoning process (e.g., 'premature closure'), rather than on hypothesis sets or other elements of the content of reasoning (cf. Clancey, in press). Our research efforts directed at clinical problem representation and the role of scripts should help to clarify some of these implications.

If we can characterize classes of problems, problem components, and management episodes that clinicians represent for a problem, then associated consistent lines of reasoning (strong methods) and their necessary components can also be investigated. This will contribute to our understanding of the regularities of clinical reasoning at levels more detailed than general hypothesis testing. Training can then be directed at teaching novices to recognize these classes, and to engage related, more specific lines of clinical inquiry.

REFERENCES

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. & Feightner, J.W. The clinical reasoning of randomly selected physicians in general medical practice. Clinical and Investigative Medicine, 1982, 5, 49-55.

Charness, N. Components of skill in bridge. Canadian Journal of Psychology, 1979, 33, 1-16.

Chase, W.G., & Simon, H.A. Perception in chess. Cognitive Psychology, 1973, 4, 55-81.

Chi, M.T.H., Feltovich, P.J. & Glaser, R. Categorization and representation of physics problems by experts and novices. Cognitive Science, 1981, 5, 121-152.

Chi, M.T.H., Glaser, R. & Rees, E. Expertise in problem solving. In R. J. Sternberg (Ed.), Advances in the psychology of human intelligence (Vol 1). Hillsdale, NJ: Lawrence Erlbaum, 1982.

Clancey, W.J. The epistemology of a rule-based expert system: A framework for explanation. Artificial Intelligence, in press.

De Groot, A.D. Thought and choice in chess. New York: Basic Books, 1965.

Egan, D.E & Schwartz, B.J. Chunking in recall of symbolic drawings. Memory and Cognition, 1979, 7, 149-158.

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

Feltovich, P.J., Johnson, P.E., Moller, J.H. & Swanson, D.B.

LCS: The role and development of medical knowledge in diagnostic expertise. In W.J. Clancey & E.H. Shortliffe (Eds.), Readings in artificial intelligence in medicine: The first decade. Reading, MA: Addison Wesley, 1983.

Hayes, J.R. & Simon, H.A. The understanding process: Problem isomorphs. Cognitive Psychology, 1976, 8, 165-190.

Larkin, J.H. Skilled problem solving in physics: A hierarchical planning model. Unpublished manuscript, Univ. California at Berkeley, 1978.

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.

Miller, R.A., Pople, H.E. & Myers, J.D. INTERNIST-I, An experimental computer-based diagnostic consultant for general internal medicine. New England Journal of Medicine, 1982, 307, 468-476.

Neufeld, V.R., Norman, G.R., Feightner, J.W. & Barrows, H.S. Clinical problem-solving by medical students: A cross-sectional and longitudinal analysis. Medical Education, 1981, 15, 315-322.

Newell, A. Artificial intelligence and the concept of mind. In R.C. Shank & K.M. Colby (Eds.), Computer models of language and thought. San Francisco: W.H. Freeman, 1973.

Newell, A. Heuristic programming: Ill-structured problems. In J. Aronofsky (Ed.), Progress in operations research (Vol 3). New York: Wiley, 1969.

Newell, A. & Simon, H. Human problem solving. Englewood Clifs, NJ: Prentice-Hall, 1972.

Norman, G.R., Tugwell, P. & Feightner, J.W. A comparison of resident performance on real and simulated patients. Journal of Medical Education, 1982, 57, 708-715.

Pople, H.E. The formation of composite Hypotheses in diagnostic problem solving: An exercise in synthetic reasoning. Proc. Fifth International Joint Conference on Artificial Intelligence, 1977.

Spilich, G.J., Vesonder, G.T., Chiesi, H.L. & Voss, J.F. Text processing of domain-related information for individuals with high and low domain knowledge. Journal of Verbal Learning and Verbal Behavior, 1979, 18, 275-290.

Williams, R.G., Vu, N.V., Barrows, H.S. & Verhulst, S. Profile of the clinical reasoning test (CRT): An objectively scorable measure of clinical reasoning skills and proficiency in using medical knowledge. In H.G. Schmidt & M.L. de Volder (Eds.) Tutorials in problem-based learning. Assen: Van Gorcum, 1984.