

PREDICTIVE VERSUS DIAGNOSTIC REASONING IN THE APPLICATION OF BIOMEDICAL KNOWLEDGE

VIMLA L. PATEL
MCGILL UNIVERSITY

DAVID A. EVANS
CARNEGIE-MELLON UNIVERSITY

ANOOP CHAWLA
MCGILL UNIVERSITY

ABSTRACT: Clinical problem solving involves both diagnostic and predictive reasoning. Diagnostic reasoning is characterized by inference from observations to hypotheses; predictive reasoning, by inference from hypotheses to observations. We investigate the use of such strategies by medical students at three levels of training in explaining the underlying pathophysiology of a clinical case. Our results show that without a sound, pre-existing disease classification, the use of basic biomedical knowledge interferes with diagnostic reasoning; however, with sound classification, biomedical knowledge facilitates both diagnostic and predictive reasoning.

INTRODUCTION:

BASIC SCIENCE VERSUS CLINICAL KNOWLEDGE

It is generally believed that knowledge of basic science is required for competence in the practice of clinical medicine. Yet the precise role of basic science in *diagnostic reasoning* has remained controversial (Clancey, in press; Patil, Szolovitz & Schwartz, 1984). Recent evidence suggests that basic science and clinical problem solving may represent domains of knowledge with very limited overlap, as judged by the inability of students and clinicians to integrate basic science knowledge in the context of diagnostic explanation tasks (Patel & Groen, 1986; Patel, Arocha & Groen, 1986). Indeed, some investigators have proposed that the uses of clinical and basic science knowledge involve *qualitatively* different types of reasoning (Patel, Groen & Scott, in press; Evans, Gadd & Pople, forthcoming).

One source of evidence on the relation between basic science and clinical knowledge derives from analysis of pathophysiological explanations offered by clinicians to describe medical problems (e.g., Patil & Szolovitz, 1981). In a recent study, Patel and Groen (1986) found that expert cardiologists who accurately diagnosed a medical case manifested a pattern of inference in which clinical findings were progressively eliminated by subsumption under a diagnostic hypothesis. This pattern of monotonic reduction of uncertainty was termed *forward chaining*. Such subjects generated very few basic science descriptions in explaining the underlying pathophysiology of the problem. By contrast, cardiologists who failed to diagnose the problem generated more intermediate steps, including more basic science descriptions. While they, too, gave evidence of forward chaining, their performance was characterized by episodes of *backward chaining*, in which secondary hypotheses were produced, effectively *increasing* the number of variables that required explanations. Since the hypotheses that led to greater uncertainty were derived from basic science knowledge, it is tempting to conclude that clinical and scientific knowledge are not well integrated for problem-solving tasks.

What explanations can we offer to account for such phenomena? In routine medical problem solving, the most expeditious path towards a diagnosis leads through a chain of clinical associations, resulting in progressive elaboration of constraints (specific diagnostic hypotheses) and incremental reduction of the problem space (the clinical findings that must be explained). When a patient presents with a difficult, multi-system problem that does not conform neatly to a familiar pattern, the physician must attempt to partition the problem space by grouping findings according to their interrelations. It is in this context that the physician may rely on knowledge of

basic science (Joseph & Patel, 1986). Similarly, a clinician who lacks the domain-specific knowledge or necessary clinical experience to reason diagnostically *via* clinical associations could be expected to resort to pathophysiological models of systemic disease processes to predict associations of findings.

Given the different demands of clinical practice and basic biomedical research, it is reasonable to assume that experts will show preferences in the kinds of knowledge they employ in solving a case. Patel, Arocha and Groen (1986) compared endocrine researchers and endocrine clinicians presented with pathophysiological explanation tasks. The researchers generated very detailed basic science explanations while the practitioners used principally clinical inferences to explain the underlying problem. The authors concluded that the basis of explanation was a function of the daily tasks the specialists are required to perform. Practitioners typically see a great many patients in a short period of time and are required to classify problems in order to recommend treatment. By contrast, researchers typically must attempt to understand and elucidate the mechanisms of specific biomedical phenomena.

Medical students, unlike practicing physicians or researchers, have only partially developed knowledge of basic science and clinical phenomena. In most medical schools, students spend their first two years taking basic science courses such as anatomy, biochemistry and physiology, and only in their second two years begin to focus on problems in clinical medicine. It is reasonable to suppose that, in problem-solving tasks, students' use of basic science knowledge will depend on the level of medical training, task demand, and the degree of uncertainty inherent in the problem. This paper reports on one investigation of this intuition.

HYPOTHESES ON PROBLEM SOLVING

An obvious hypothesis is that efficient causal reasoning about a disease process requires management of cognitive load -- maximizing the constraints on the problem space and minimizing the number of variables (uncertainty) that must be held in memory. Given that the task is to associate clinical manifestations to a disease process, there are *a priori* two ways in which causal networks can be built, *viz.*, by reasoning *predictively* from hypotheses to manifestations, or *diagnostically* from manifestations to hypotheses. In the sense that predictive reasoning involves the use of generalizations associated with diagnoses to identify candidates' specific details, in a case it can be regarded as *deductive*. Similarly, in the sense that diagnostic reasoning involves the use of particular details to suggest the appropriate generalized "diagnoses", it can be regarded as *inductive*.

In predictive reasoning, the problem space is controlled because one entertains a hypothesis of sufficient power and generalization to account for many possible manifestations. Uncertainty is controlled because inference is limited to what is entailed by the hypothesis. For example, if a physician *assumes* that a patient has bacterial endocarditis, he or she can know that an infectious process is involved and, hence, can predict (and account for) a finding of fever. Clearly, for predictive reasoning to be effective in problem solving, knowledge of diagnoses and their associated manifestations is important.

In diagnostic reasoning, it is much more difficult to constrain the problem space and eliminate uncertainty. For example, one could link fever with an infectious process, diagnostically, but it would be impossible to say which infectious process was involved without considering combinations of other findings. Furthermore, fever is not caused by infectious processes, but occurs in inflammatory disorders and certain cancers. In reasoning from specific manifestations to possible hypotheses, one introduces numerous alternatives that must be reconciled against one another. A physician who cannot quickly classify findings according to the most likely diagnoses that would account for them will be overwhelmed with the problem of managing information and inference.

We would suggest that physicians, in fact, use relatively simple classification schemata to transform the problem space of individual findings to one of a small number of diagnostic hypotheses and clustered findings. Without clinical knowledge as a basis for classification, knowledge of physiological mechanisms and scientific principles will be used to drive inferences and associate observations. Our characterization of reasoning in the domain of medicine, thus,

underscores the epistemological complexity of the task. This suggests that the management of uncertainty will be the principal component of effective performance. It is, therefore, important to consider structural problems that arise in reasoning under uncertainty. In particular, Henrion (1986) has identified certain phenomena associated with the representation of causal knowledge that lead to *propagation* of uncertainty in inference. We focus on two problems discussed by Henrion, *dependency effects* and *cyclical inferencing*.

Dependency effects can be illustrated with the aid of Figure 1, below. Under the first representation of the relations of observations to hypotheses, on the left, evidence from three different sources A, B, and C together strengthen the hypothesis, H_1 . However, if A, B, and C are dependent on some other source, D, as represented on the right, the hypothesis is not strengthened. One method for avoiding this, suggested by Henrion, is to insure independence of evidence.




FIGURE 1: DEPENDENCY EFFECTS

Cyclical inferencing occurs when a hypothesis, based on certain specific evidence, is used to account for phenomena, including the evidence that originally gave rise to the hypothesis. Since reasoning in the medical domain involves both predictive and diagnostic strategies - i.e., reasoning from hypothesis to disease manifestations and *vice versa* - there is the danger that, without a means of keeping evidence from predictive and diagnostic sources separate, cyclical inferencing can occur. For example, consider the following reasoning, given schematically in Figure 2: *Intravenous drug use can lead to infection, which leads to fever. In a patient with fever and heart murmur, there is a high possibility of bacterial endocarditis. Bacterial endocarditis is associated with intravenous drug users, hence accounts for drug use.* The resulting network leads to propagation of uncertainty.





FIGURE 2: CYCLICAL INFERENCE

Henrion argues that this problem arises as a result of failure to maintain separation between the flow of predictive and diagnostic reasoning. We are particularly interested in the interaction of basic biomedical and clinical knowledge with such phenomena. Our approach has been to analyze the protocols of medical students at different levels of training who were asked to perform a diagnostic task. Our hypothesis is that clinical training provides knowledge of appropriate disease classification that enables efficient predictive reasoning, leading to explanations that manifest causal coherence. In terms of reasoning strategies, we hypothesize that classification-guided predictive reasoning will be preferred to diagnostic reasoning; and that in the absence of well-developed classification schemata, some combination of diagnostic reasoning and predictive reasoning from *naive* or *for-the-nonce* classification will be used. *Ad hoc* classification is facilitated, we suggest, by application of basic scientific knowledge edge; and in such instances, we expect to see a compounding of problems of dependency effects and cyclical inferencing. More generally, we hypothesize that any use of *induction* without an adequate basis for *deduction* - such as afforded by well-developed classification schemata - will lead to inefficient, incomplete, and incoherent diagnostic problem solving.

EXPERIMENTAL METHOD

A total of 24 subjects were selected from three levels of medical school training. Level 1 included six students just entering their first year of medical school. Level 2 included six second-year medical students who had completed all basic medical sciences, but had had no clinical work. Level 3 included twelve final-year medical students three months before graduation.

PROCEDURE: The following empirical paradigm was used: We (1) presented the subject with a clinical problem text; (2) obtained a summary protocol of the problem; (3) asked the subject to explain the underlying pathophysiology; (4) asked the subject for a diagnosis; and (5) presented the subject with three basic science texts and asked for an explanation of the clinical problem.

ANALYSIS: We use a combination of propositional and protocol analysis. Our first step is to separate the texts into segments, corresponding to syntactic units, using Winograd's system of clausal analysis (Winograd, 1972). The next step involves the representation of clauses in terms of their propositional content and structure, using the techniques of propositional analysis based on Frederiksen (1975). A detailed analysis of a clinical case identical to that used here is given in Patel & Groen (1986) and Patel & Frederiksen (1984). At a second level of analysis, propositions are linked to form a higher-ordered relational structure (frames). Such structures are representable as the causal networks discussed in the literature on medical artificial intelligence. The links in the causal networks can be converted into production rules. The directionality of such rules is important since causal rules generally lead away from a diagnosis (in a predictive direction) and only conditional rules lead toward one (in a diagnostic direction). Our analysis also uses the notion of *reference frame* of a disease, which aids us in identifying where basic science knowledge is used in reasoning in either direction.

COGNITIVE TASK: The tasks require each subject to read the case and form a representation of the problem. Subsequently they are asked to summarize the problem. This provides insight into the relations among clinical findings that they have considered. As subjects are also asked to explain the underlying pathophysiology of the problem, it is possible to assess the degree to which they can coherently account for the disease process. There are several levels of abstraction from which they can approach this task. For example, they can describe the disease process entirely in terms of clinical or anatomical features (as a static process) or they can explore the perturbations in the physiological processes that result in the patient's presenting condition (a dynamic process). After presenting a pathophysiological account, they are required to make a diagnosis. The text of the clinical problem, acute bacterial endocarditis, is given in Figure 3 and the components of the reference frames are given in Figure 4.

THIS 27-YEAR OLD UNEMPLOYED MALE WAS ADMITTED TO THE EMERGENCY ROOM WITH THE COMPLAINT OF SHAKING CHILLS AND FEVER OF FOUR DAYS DURATION. HE TOOK HIS OWN TEMPERATURE AND WAS RECORDED AT 40°C ON THE MORNING OF ADMISSION. THE FEVER AND CHILLS WERE ACCOMPANIED BY SWEATING AND A FEELING OF PROSTRATION. HE ALSO COMPLAINED OF SHORTNESS OF BREATH WHEN HE TRIED TO CLIMB THE TWO FLIGHTS OF STAIRS IN HIS APARTMENT. FUNCTIONAL ENQUIRY REVEALED A TRANSIENT LOSS OF VISION IN HIS RIGHT EYE WHICH LASTED APPROXIMATELY 45" ON THE DAY BEFORE HIS ADMISSION TO THE EMERGENCY WARD.

PHYSICAL EXAMINATION REVEALED A TOXIC LOOKING YOUNG MAN WHO WAS HAVING A RIGOR. HIS TEMPERATURE WAS 41°C, PULSE 120, BP 110/40. MUCUS MEMBRANES WERE PINK. EXAMINATION OF HIS LIMBS SHOWED PUNCTURE WOUNDS IN HIS LEFT ANTECUBITAL FOSSA. THE PATIENT VOLUNTEERED THAT HE HAD BEEN BITTEN BY A CAT AT A FRIEND'S HOUSE ABOUT A WEEK BEFORE ADMISSION. THERE WERE NO OTHER SKIN FINDINGS. EXAMINATION OF THE CARDIOVASCULAR SYSTEM SHOWED NO JUGULAR VENOUS DISTENSION, PULSE WAS 120 PER MINUTE, REGULAR, EQUAL AN DISYNCHRONOUS. THE PULSE WAS ALSO NOTED TO BE COLLAPSING. THE APEX BEAT WAS NOT DISPLACED. AUSCULTATION OF HIS HEART REVEALED A 2/6 EARLY DIASTOLIC MURMUR IN THE AORTIC AREA AND FUNDOSCOPY REVEALED A FLAME-SHAPED HEMORRHAGE IN THE LEFT EYE. THERE WAS NO SPLENOMEGALY. URINALYSIS SHOWED NUMEROUS RED CELLS BUT THERE WERE NO RED CELL CASTS.

FIGURE 3: BACTERIAL ENDOCARDITIS TEXT




FIGURE 4: Components of Canonical Frames for Acute Infectious Endocarditis with Aortic Insufficiency

CAUS Causal Relation
 COND Conditional Relation
 PROX Proximity Relation

In the second part of the experiment the subjects are asked to read three basic science texts, which they must subsequently relate to the clinical case. In effect, the subjects are provided with the opportunity at this stage to update their pathophysiological explanations using the basic science principles, abstractions, and concepts presented in the texts. There are several ways in which their explanations can change. The subjects can choose to elaborate and embellish previous characterizations of the disease process with additional detail; they may choose to reject earlier characterizations on the basis of the new information; or they may offer new pathophysiological descriptions, entirely orthogonal to their previous descriptions, resulting from a complete shift in focus. The nature of their revised representations is a function of several factors including (a) their ability to form a coherent representation of the clinical text; (b) their comprehension of the basic science material; and (c) their ability to synthesize and integrate the relevant information from the basic science texts into the context of the clinical problem.

RESULTS AND DISCUSSIONS

The performance of subjects at each level was distinct and quite uniform in terms of diagnostic and predictive reasoning using basic science knowledge in the explanation task. Students at level 1 were rated *poor* on both diagnostic and predictive reasoning; students at level 2 were rated *fair* on both; and students at level 3 were rated *good* on both. Thus, though all groups of subjects showed evidence of both types of reasoning in explaining the patient's underlying problem, reasoning improved in accuracy and coherency with level of training.

DIAGNOSTIC REASONING: It appears that reasoning in a diagnostic direction is frustrated by the application of basic science knowledge *unless* there is a strong classification of hypotheses (diagnoses). The final-year medical students, who have had some clinical experience, have partial classification schemata which assist them both in selecting an initial, accurate diagnostic hypothesis and in predicting selectively, the variables to be tested to confirm the selected hypothesis. This is illustrated in Table 1 by an analysis of a protocol, taken from a final-year student, showing the explanations produced before and after receiving basic science information. Figures 5 and 6 give the causal networks generated from these protocols. There is no evidence of cyclical inferencing or dependency effects in pre-basic science protocols but they are present to a limited extent in post-basic science protocols.

TABLE 1:

PATHOPHYSICAL EXPLANATION PROTOCOLS OF BACTERIAL; ENDOCARDITIS
BY A FINAL YEAR MEDICAL STUDENT, WITH AND WITHOUT THE
BASIC SCIENCE TEXT PROVIDED.

<u>CLINICAL TEXT ONLY</u>	<u>WITH CLINICAL AND BASIC SCIENCE TEXTS</u>
<p>Bacteremia resulting in generalized symptoms of illness, i.e., fatigue, chills, fever, sweating. These bacteria accumulated in aortic valve with resultant vegetations and murmur due to auscultation of valve. Emboli from these vegetations travel via carotid on left to left retinal artery resulting in temporary blindness and flame hemorrhages due to the occluded blood supply.</p>	<ol style="list-style-type: none"> 1) Bacteria introduced into bloodstream perhaps via vena-puncture. Subsequently release endotoxins which act on bone marrow cells to produce pyrogens. These pyrogens act on the hypothalamus (pre-after) to result in fever. Chills occur due to need to increase body temperature as indicated by the hypothalamus. 2) This bacteremia results in vegetations occurring in aortic valve. The diastolic murmur is due to aortic regurgitation as these vegetations alter the closing of the valve. The murmur radiates along the blood flow tract of the aortic valve which is the left sternal border, where the left ventricle is. 3) Micro emboli are released from these vegetations and flow in the arterial circulation and eventually occlude capillaries going to the retina resulting in transient blindness. This increased capillary pressure due to occlusion resulted in retinal hemorrhage and thus are seen on fundoscopic examination.

It is clear that the subject inductively decides that the patient has bacteremia and the bacteremia has affected the aortic heart valve (Figure 5). In order to account for these two facets of the overall disease, certain predictions have to be made. The subject confirms that it is bacteremia because the patient has fatigue, chills, fever, and sweating, which are the findings one would expect. In order to confirm aortic valve involvement, the subject explains heart murmur and subsequent hemorrhage and blindness. The classification of the disease presumably directed his or here selection of particular variables (findings), whose presence can be taken as confirmatory evidence that the prediction was valid.

After the basic science text has been read, the subject once again uses a classification schema to provide the necessary basic science knowledge required to explain the findings (Figure 6). Here, a deeper account of the disease process is offered. There is no change in either the diagnostic explanation or in the selection of predictions in confirming the diagnosis. Basic scientific information is used to provide detailed physiological mechanisms that can account for the existence of the selected variables identified earlier. There is an effect of global coherence in both explanation protocols, due in part to the rich connectedness of details. Both causal networks resulted in only one diagnostic component each, corresponding to the diagnosis, and produced no disjoint components.

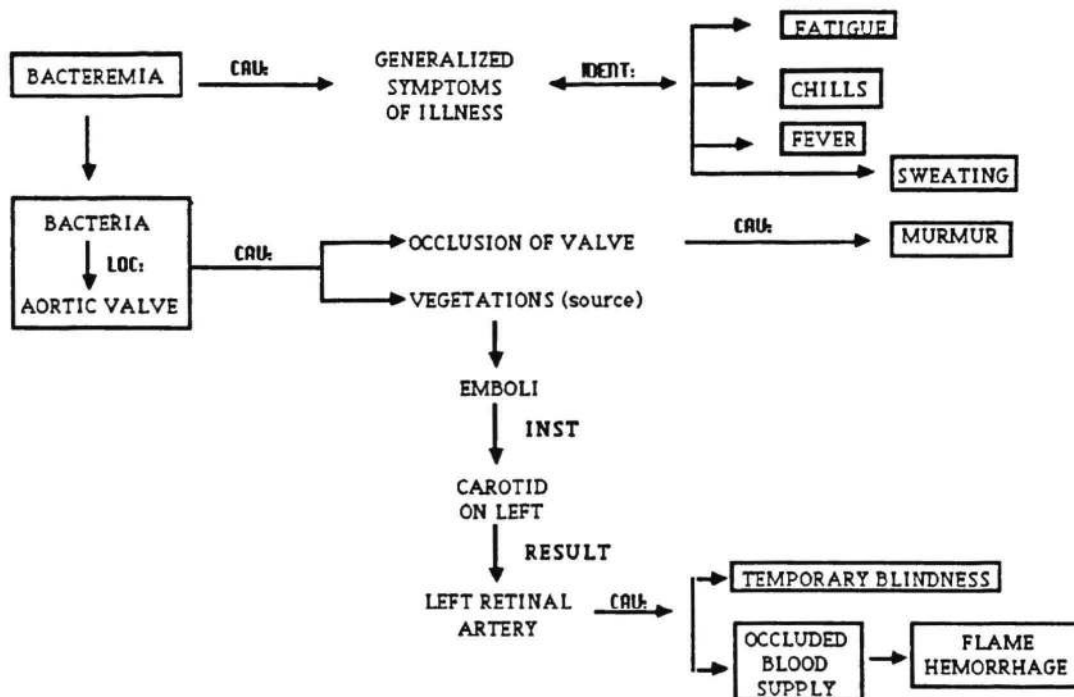


FIGURE 5: CAUSAL NETWORK OF PATHOPHYSIOLOGICAL EXPLANATION OF ENDOCARDITIS BY A FINAL YEAR MEDICAL STUDENT (clinical text only)

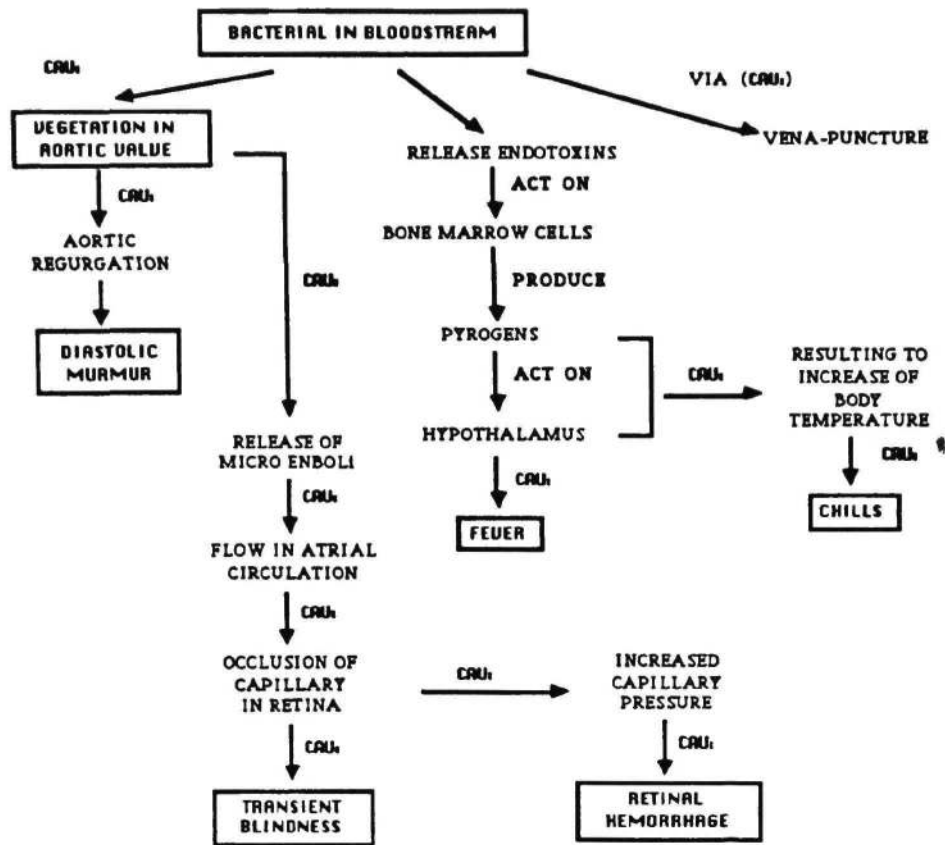


FIGURE 6: CAUSAL NETWORK OF PATHOPHYSIOLOGICAL EXPLANATION OF ENDOCARDITIS BY A FINAL YEAR MEDICAL STUDENT
(clinical and basic science text)

The second-year medical students, who do not have a clinical basis for disease classification, are not able to use basic science knowledge adequately in diagnostic reasoning. However, these students do have good taxonomies of some for the *components* of disease processes, such as *general infection*, which help focus hypotheses and provide local coherence in the organization of findings. This is illustrated by an explanation from a second-year student, as analyzed in the causal network in Figure 7. The analysis shows that diagnostic reasoning is facilitated by the identification of the infection component of bacterial endocarditis. Further, the subject rules out other sources of infection, *viz.*, common causes and tropical fever, and thereby increases the certainty of infection from cat bite. The subject does not use basic science information any further to reason diagnostically; rather, moves backward to explain the findings, to confirm the generated hypothesis. When more basic science information is provided, to confirm the generated hypothesis. When more basic science information is provided, it is used by the student to rule out the possibility of immune response as a differential diagnosis (Figure 8). Besides that, the additional basic science knowledge leads to a reduction in global coherence (6 hypothetical components appeared where there previously was only 1 candidate diagnosis), without a loss of local coherence. Furthermore, some use of basic science knowledge is wholly inaccurate. For example, decreases in the transfer of energy from systole to diastole do not cause hypotension.

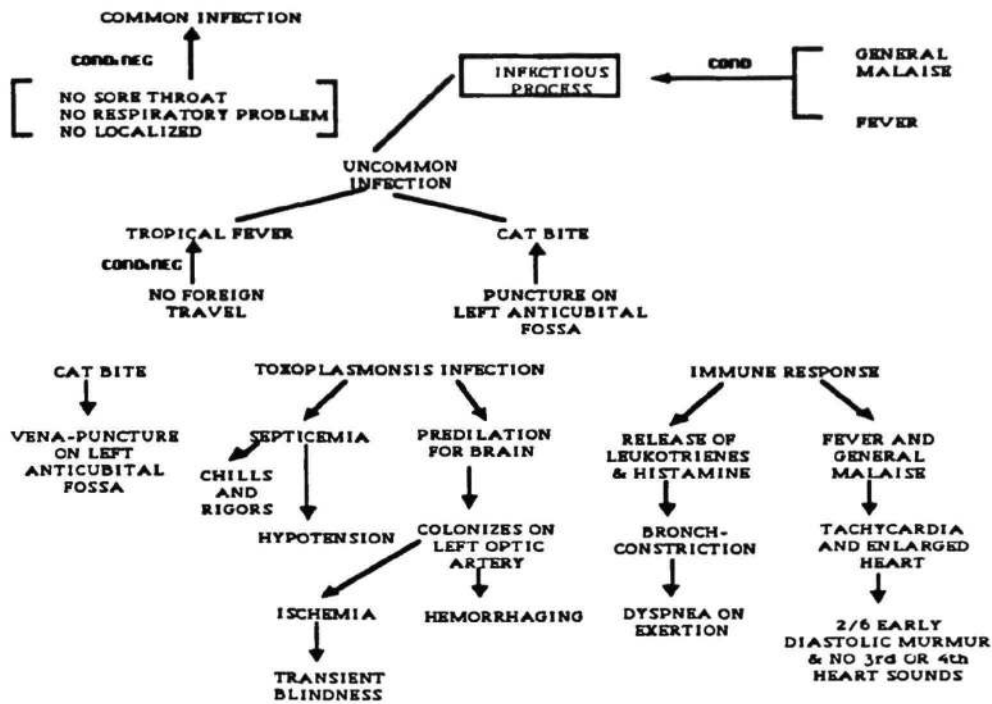


FIGURE 7: CAUSAL NETWORK OF PATHOPHYSIOLOGICAL EXPLANATION BY A SECOND YEAR MEDICAL STUDENT (clinical text only)




FIGURE 8: CAUSAL NETWORK OF PATHOPHYSIOLOGICAL EXPLANATION BY A SECOND YEAR MEDICAL STUDENT (Clinical & Basic Science Text)

The first-year students, who have only naive taxonomies of disease phenomena, choose a general hypothesis about the problem, which they are prepared to understand. For example, one subject diagnoses the problem as *mild shock*, then classifies the symptoms of this condition as increased heart rate, low diastolic pressure, low systolic pressure, and poor venous return (Figure 9). The balance of the student's effort is spend reflecting on connections among individual findings, to provide some coherence to the text. All subsequent explanations support blood loss and hemorrhaging, which account for the shaking chills. The sample explanations have only modest local coherence with very little global coherence (3 components remain unresolved). After the first-year students have read the basic science text, global coherence diminishes (6 unresolved components appear); though local coherence, which is still poor, improves over the first explanation. As can be seen from Figure 10, the additional basic science knowledge does not provide extra understanding affecting diagnostic reasoning.

PREDICTIVE REASONING: We might conjecture that reasoning in predictive directions would be facilitated by the application of basic science knowledge. In fact, final-year medical students, who have knowledge of relevant basic science and some disease classification schemata, have no problem in using their knowledge to predict selected variables to confirm the hypothesis. As the causal network in Figure 6 shows, additional basic science knowledge is used to add coherence to the clinical model which in turn provides the basis of the disease classification. There is a clear evidence of predictive reasoning using basic science knowledge to support the hypothesis that the patient was suffering from bacterial endocarditis.

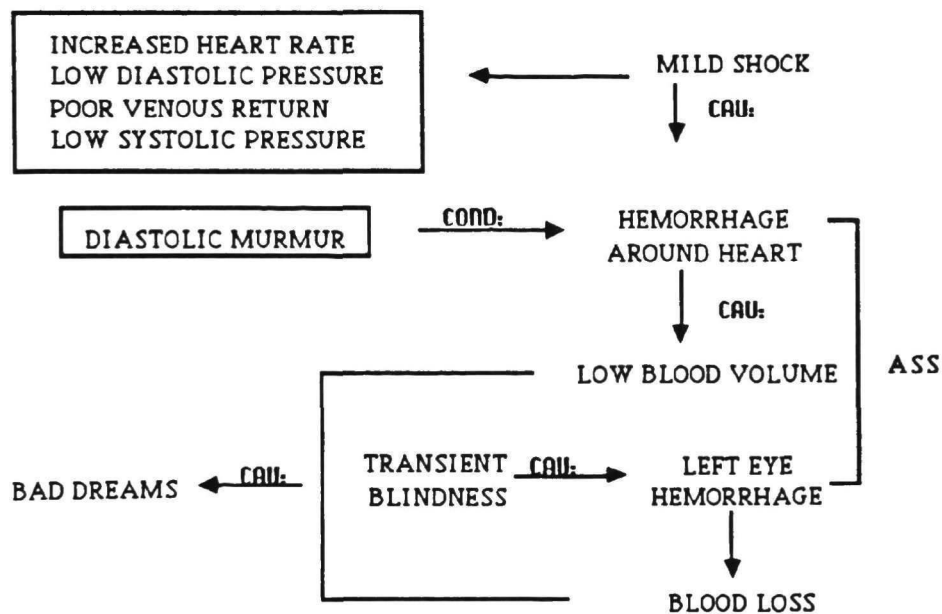


FIGURE 9: CAUSAL NETWORK OF PATHOPHYSIOLOGICAL ENDOCARDITIS BY A FIRST YEAR STUDENT (clinical text)

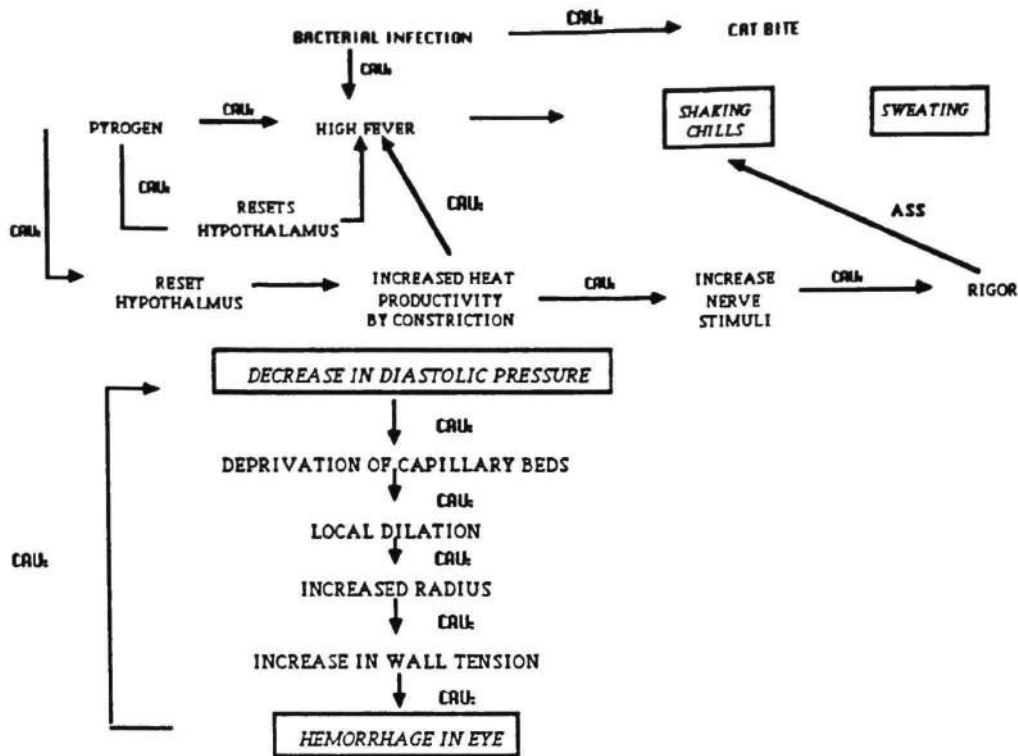


FIGURE 10: CAUSAL NETWORK OF PATHOPHYSIOLOGICAL EXPLANATION OF ENDOCARDITIS BY A FIRST YEAR MEDICAL STUDENT (clinical and basic science text)

Basic science knowledge did not adequately assist the second-year students in predictive reasoning (Figure 8). It provided coherence to the local infection component of the disease, only to the extent that chills, rigor, and paleness could be predicted. The rest of the protocol shows signs of predictive reasoning using basic science knowledge, but without leading to the selection of relevant predictions.

The first-year students' protocols show poor predictive reasoning throughout. In the example case, the variables selected for confirmation were based on the diagnostic hypothesis (mild shock), which was grossly inaccurate. This hypothesis guided the selection of predictors such as blood loss and hemorrhaging. Although the use of knowledge irrelevant to the diagnosis provides some coherence to the text, it does not contribute to an explanation of the basic underlying pathophysiology of the problem. With additional basic science knowledge there is evidence of one variable predicting another, cyclically. This is seen when decreased diastolic pressure accounts for hemorrhage in the eye and *vice versa* (Figure 10). Further evidence of cyclical inference appears when fever is used to explain bacterial infection. Chills, sweating, rigor, and pink mucus membrane are caused by fever and thus provide evidence for fever. High fever is said to be caused by pyrogens which account for all the findings - fever, chills, sweating, and rigor.

The pattern of results described in detail for one subject at each level of medical training was also true for the other subjects at that level. Nine out of twelve students in their final-year of training conformed to the pattern of diagnostic and predictive reasoning described. All six of the second-year students protocols reflected the general phenomena described for one subject. However, there was somewhat greater variance seen in the protocols of the first-year medical students where, although four out of six subject protocols showed the described pattern of results, there were differences in their specific prior knowledge and in interpreting the protocols. Two subjects in this year did not follow the described pattern of results for the first-year students because of their unusual background of having a degree of Ph.D. and nursing.

CONCLUSION

For clinical problem solving, basic science knowledge serves as a powerful resource for bridging gaps during deductive reasoning; but is too powerful to be used in inductive reasoning. Inductive reasoning is successful only in the context of a well-differentiated taxonomy of goals, in particular, a sound classification schema for diagnoses and their facets. The ability to organize observations based on relevant clinical associations provides sufficient structuring of a diagnostic-problem space to exploit the details of basic science knowledge most efficiently. Once partitioned, the problem space can be constrained further by either diagnostic reasoning or predictive reasoning, with or without the addition of basic science knowledge. In sum, we find that basic biomedical knowledge interferes with diagnostic problem solving, unless there is a basis for predictive reasoning.

ACKNOWLEDGEMENTS

We wish to thank David Kaufman for his thoughtful comments during many discussions of the data in our studies. We would also like to thank Susan Young for typing the manuscript and Armar A. Archbold for his technical assistance with the figures. The authors alone, of course, are solely responsible for any errors or misrepresentations in this paper. The research reported here has been supported, in part, by grants from the Josiah Macy, Jr. Foundation to Patel and Evans.

REFERENCES

- Clancey, W.J. (in press) Acquiring, representing, and evaluating a competence model of diagnostic strategy. In: M.T.H. Chi, R. Glaser & M. Farr (Eds.), *Nature of Expertise*. Hillsdale, NJ: Lawrence Erlbaum.
- Evans, D.A., Gadd, C.S. & Pople, H.E., Jr. (in press) Managing coherence and context in medical problem-solving discourse. In: D. Evans & V. Patel (Eds.), *Cognitive Science in Medicine*. Cambridge, MA: MIT Press
- Frederiksen, C.H. (1975) Representing logical and semantic structure of knowledge acquired from discourse. *Cognitive Psychology*, 7, 371-45.
- Henrion, M. (1986) Uncertainty in artificial intelligence: Is probability epistemologically and heuristically adequate? In: J. Mumpower & O. Renn (Eds.), *Expert Systems and Expert Judgement*, Proceedings of the NATO Advanced Research Workshop, in Porto, Portugal, August, 1986.
- Joseph, G.-M. & Patel, V.L. (1986) Domain specificity and expertise in medicine. *Proceedings of the Cognitive Science Society Meeting, Amherst, MA*, August 1986. Hillsdale, NJ: Lawrence Erlbaum . 331-345.
- Patel, V.L. & Frederiksen, C.H. (1984) Cognitive processes in comprehension and knowledge acquisition by medical students and physicians. In: H. Schmidt & H.G. deVolder (Eds.), *Tutorials in Problem Based Learning*. Assen, Holland: Van Gorcum.
- Patel, V.L. & Groen, G.J. (1986) Knowledge-based solution strategies in medical reasoning. *Cognitive Science*, 10, 91-116.
- Patel, V.L., Groen, G.J. & Scott, H. (in press) The uses of biomedical knowledge in clinical reasoning. In: D. Evans & V. Patel (Eds.), *Cognitive Science in Medicine*. Cambridge, MA: MIT Press.

- Patel, V.L., Arocha, J.F. & Groen, G.J. (1986) Strategy selection and degree of expertise in medical reasoning. *Proceedings of the Cognitive Science Society Meeting, Amherst, MA, August 1986*. Hillsdale, NJ: Lawrence Erlbaum. 780-93.
- Patil, R.S. & Szolovitz, P. (1981) Causal understanding of patient illness in medical diagnosis. *Proceedings of the 7th IJCAI*. Vancouver, B.C., Canada.
- Patil, R.S., Szolovitz, P. & Schwartz, W. (1984) Causal understanding of patient illness in medical diagnosis. In: W.J. Clancey & H.E. Shortliffe (Eds.), *Readings in Medical Artificial Intelligence*. Reading, MA: Addison-Wesley Publishing Company.
- Winograd, T. (1972) *Understanding Natural Language*. New York: Academic Press.