

THE NATURE OF EXPERTISE IN THE CLINICAL INTERVIEW: INTERACTIVE MEDICAL PROBLEM SOLVING

David R. Kaufman and Vimla L. Patel
Centre for Medical Education
McGill University

The doctor patient interview is an interactive, goal-directed process; diagnosis being one of the primary goals. The clinical interview is perhaps the most important source of medical data that a physician has access to. Although diagnoses are most often made during the history-taking (Kassirer & Gorry, 1978), it is one of the least understood aspects of clinical cognition. In medical problem solving there have been comparatively few studies that have examined the acquisition and representation of clinical information. Elstein, Shulman and Sprafka (1978) used doctor-patient simulations to evaluate the performance of physicians. Barrows, et al., (1978) extended the research to evaluate the nature of clinical expertise along a developmental continuum, from novice medical student to expert physician. The principal goal of these studies was to characterize the general aspects of clinical performance. The emphasis was on process rather than content or knowledge. Their findings did not discriminate between student and expert physician except for the accuracy of diagnostic hypotheses.

Recent investigations in medical problem solving have shifted from an emphasis on global aspects of clinical reasoning to a focus on the nature and content of medical knowledge used to solve a problem. Feltovich and colleagues investigated knowledge-based differences between expert and novice in the domain of pediatric cardiology (Feltovich et al., 1984). The results suggested that there are systematic differences in the structure of clinicians' knowledge-base, corresponding to certain levels of expertise that direct them towards particular inferences in generating diagnostic hypotheses. Patel and Groen (1986) demonstrated that expert physicians who accurately diagnosed a clinical case, developed explanations of the disease process that could be accounted for in terms of a forward-chaining strategy, moving from propositions in the stimulus text to conditions that suggested a component of the diagnosis. Expert physicians who misdiagnosed the case used a backward-chaining strategy, characterized by the posting of a potential cause and the subsequent generation of confirmatory evidence.

These studies provide converging evidence that it is primarily the differences in a clinician's knowledge base that allows them to effectively represent a problem and generate an appropriate solution. This would suggest that any characterization of expertise in the context of the clinical interview cannot be made independent of the knowledge necessary to solve the problem. The research presented here represents an attempt to investigate the clinical interview as a problem-solving situation. In particular, this study investigates expert-novice differences in the acquisition, representation and utilization of patient information in the course of the clinical interview.

THEORETICAL FRAMEWORK

It was determined that to effectively investigate this area we needed a coherent epistemological framework— identifying appropriate units of knowledge that correspond to medical problem solving and a characterization of the task environment and cognitive demands of the clinical interview. The overall framework is described in greater detail in Kaufman (1987) and Patel, Evans, and Kaufman (in press).

The epistemological framework is adopted from a model proposed by Evans, Gadd and Pople (in press) who distinguish six levels at which clinical knowledge stratifies. Four of these levels are of interest to this research. **Observations** are units of information that are recognized as potentially relevant in the problem-solving context. They do not constitute clinically useful facts and many observations will be disregarded when formulating diagnostic hypotheses. **Findings** are comprised of observations that are anchored in the diagnostic context and have potential clinical significance. **Facets** are clusters of findings that are suggestive of prediagnostic interpretations. Facets reflect general pathological descriptions such as *aortic insufficiency* or categorical descriptions such as *endocrine problem*. **Diagnosis** is the level of classification that subsumes and explains all levels beneath it. The diagnostic label serves to summarize and organize the patient findings into a coherent and functional unit that facilitates therapeutic and management decisions.

Facets and diagnoses define the context and provide a framework in which observations and findings can be interpreted. This results in the instantiation of default expectations that guide the physician's subsequent questions. Observations form the minimal unit of the problem space. Observations

convey the propositional information that is subjected to inferential operations and results in the construction of higher order conceptual structures.

Medical problem solving has been described as an ill-structured task (Johnson, 1983). Ill-structured problems are ones in which the initial states, the goal state and the necessary constraints, are unknown at the beginning of the problem-solving process. In solving a patient's problem, the initial state, the goal state, and the elements of the problem space are undefined prior to the interview. The interim subgoal is to formulate a representation of the patient's presenting complaint. This as well as patient variables become the initial knowledge state. The clinician's representation of this initial state will in part determine subsequent inquiry. The permissible operators are essentially the range of questions asked by the clinician to elicit specific information. The application of these operators can be guided towards specific inquiry, such as evaluating a hypotheses or oriented towards general information gathering, such as assessing the patient's past medical history.

Data-gathering strategies can be guided by a goal to elaborate the general parameters of a problem-space; that is, to derive findings and meaningful relations among findings in the patient's history. This process is principally bottom-up, generating facets and differential diagnoses from the elements in a problem space. This is referred to as *reasoning diagnostically* (Patel, Evans & Chawla, 1987). Once the problem representation is sufficiently well elaborated, certain expectations are built up. The inquiry is directed towards eliciting specific findings that correspond to the interim interpretations (facet hypotheses) and diagnostic hypotheses that the physician has generated. This top-down process is referred to as *reasoning predictively*. It is expected that experienced physicians have a greater capacity to reason in the predictive mode, because they are more able to recognize patterns of elements in a problem that are associated with certain diagnoses.

EXPERIMENTAL METHODS

Fifteen volunteer subjects were selected: 5 endocrinologists (experts), 5 physicians in their third year of residency programs (intermediates), and 5 final year medical students (novices) who had completed their clinical training. Each subject was given 15 minutes to take the patient's history. The patient was a paid volunteer outpatient. Each session was videotaped and the dialogue was subsequently transcribed. The patient was a 22 year old oriental male who presented with two episodes of severe muscle weakness and other manifestations. This problem had been diagnosed as hypokalemic periodic paralysis associated with thyrotoxicosis, an uncommon disorder of the thyroid gland, involving episodes of paralysis associated with a marked fall in serum potassium. This problem can be decomposed into two components, *hyperthyroidism* and *hypokalemia*.

Diagnostic accuracy was evaluated according to whether the two diagnostic components, hyperthyroidism and hypokalemic periodic paralysis were present in the subjects' differential diagnoses. The transcribed dialogue was segmented into units of doctor-patient exchanges, which are physician-question/patient-response pairings. The method used to characterize the doctor-patient dialogue incorporated a pragmatic analysis of the discourse (described in Patel, Evans & Kaufman, in press), whose primary function is to characterize the information in focus during the course of the interview. This analysis is designed to evaluate how a clinician manipulates the context of the interview, to elicit observations, conclude findings, and derive meaning from higher order relations in the data. It also permits us to evaluate data-gathering strategies, as evidenced by sequences of questions elicited and findings concluded during the course of the interview.

Observations and findings are the main units of data used in this study to characterize the representation that the clinician is constructing. Observations in this analysis are units of information expressed by the patient or elicited by the physician in the course of the interview. The criterion for coding an observation is that it contributes unique information to the interview and is accepted by the physician.

Findings reflect a decision made by a physician concluding that a particular array of observational data contains a significant cue or cues that need to be accounted for in the diagnostic context. Findings are multi-faceted concepts that have certain well-formedness conditions. Certain propositional parameters such as location, duration, and sensation quality, need to be satisfied to conclude a specific finding. Findings are specific to particular individuals with unique presenting problems. However, it is possible to specify a relatively comprehensive set of generic findings. Towards this end, we appeal to a highly developed source of medical knowledge. Internist-I is an expert system for computer-assisted diagnosis in internal medicine (Miller, et al., 1984). It has an extensive knowledge-base incorporating over 550 diagnostic classifications and more than 4000 findings. Findings are determined to be concluded if they correspond in form to internist findings. This method is explicated in Patel, Evans, and Kaufman (in press).

Findings can be either negative or positive. A positive finding is a determination by the subject that a finding is characteristic of the patient's condition. A negative finding is a result of a conclusion that a particular finding is not indicative of the patient's problem. Negative findings are particularly useful in discriminating between competing diagnostic hypotheses. However, it is reasonable to speculate that the more the subjects elicit negative findings, the more likely the subject has not been able to develop an adequate representation of the problem.

A reference model was created, with the assistance of an expert endocrinologist, to identify and classify findings with respect to their importance in this particular case. The model identified the relevant findings in this case. Relevant findings are suggestive of this diagnosis and are contributory to the identification of the problem. In total, there were 16 findings designated as relevant.

The variables used in this study include the accuracy of diagnosis, the number of question/answer exchanges, the number of observations elicited, the number of total findings acquired, the number of positive and negative findings elicited and the percentage of relevant findings elicited during the interview. In addition, to these variables we can derive three measures of efficiency: a) the number of observations per question/answer exchange expressed as a percentage. This is designed to assess the efficacy of a subjects' ability to elicit potentially relevant information; b) the number of findings per question/answer exchange expressed as a percentage. This is a more specific index of a clinicians' ability to recognize significant information in the patients statements and to focus probes directly on the problem that lead to the generation of specific findings; c) the number of positive findings over question/answer exchanges as a percentage. This measure assesses the ability of a subject to develop a problem representation, anticipate candidate findings and reason predictively.

Three exploratory multivariate and two univariate analyses of variance were performed. The level of expertise was the independent variable for each analysis. The first multivariate analysis included question/answer exchanges, total number of observations and total number of findings as the dependent variables. In the second multivariate analysis, positive and negative findings were the dependent variables. In the third multivariate analysis, the first two efficiency measures, observations divided by exchanges and findings divided by exchanges were the dependent variables. The first univariate analysis used relevant findings as the dependent measure and the second univariate analyses used the third efficiency measure, positive findings over question/answer exchanges as the dependent measure.

RESULTS AND DISCUSSION

DIAGNOSTIC ACCURACY: Four out of five expert physicians accurately identified both components of the diagnosis. The other expert identified only the thyroid problem. One of the residents accurately identified the problem, while two others identified the more general aspect, hyperthyroidism. One resident recognized the fact that the patient was suffering from periodic paralysis, but failed to identify hypokalemia as the causal agent. One subject did not recognize either component. No novice was able to completely diagnose the problem. However, three of the novice medical students recognized the thyroid aspect.

INFORMATION ACQUISITION: A summary of the results pertaining to question/answer exchanges, findings and observations are presented in Table 1. The first multivariate analysis with question/answer exchanges, observations, and findings as dependent variables yielded a significant group (level of expertise) effect ($F[2,12] = 4.64, p < .05$). There was also a significant group by measure interaction effect ($F[4,20] = 3.43, p < .05$). The endocrinologists had on average the shortest interviews as judged by the number of question/answer exchanges, followed by the residents with the medical students on average conducting the longest interviews. As the interaction effect would suggest, the pattern changes somewhat for the findings and observations. The residents elicited the greatest number of findings, followed by the medical students and the endocrinologists. In each of these 3 measures the endocrinologists generated significantly lower frequencies than either of the other 2 groups.

The second multivariate analysis, which compared the positive and negative findings generated, yielded no significant differences. It is assumed that positive findings have a bearing on evaluating or confirming working hypotheses, while negative findings are useful in discriminating between alternative hypotheses. In examining the group means, it becomes apparent that a pattern is emerging. As mentioned previously, the experts elicited fewer observations and generated fewer findings than the other groups, which on the surface would seem puzzling. However, the experts generated a comparable number of positive findings (14.4) to the residents (14.6), which was on average more than the students generated (13). This would suggest that the experts are more focused in that they need fewer questions to generate significant information and to conclude findings. Generally, the resident group generated almost as many negative

findings as positive findings. The comparably large number of negative findings are suggestive of the fact that this group had difficulty accounting for the myriad of problems the patient was presenting with and could not develop a coherent problem representation. The generation of numerous negative findings would also indicate that these subjects were considering multiple hypotheses and were attempting to rule-out as many as possible. The medical students, on average, elicited somewhat fewer negative findings than did the residents. This may suggest that those subjects who were able to identify the thyroid aspect were content that they had correctly diagnosed the problem. The residents recognized that there were many discrepant findings that could not be accounted for by hyperthyroidism and pursued alternative explanations. The experts develop a coherent representation of the distinctive pattern of findings and elicit fewer negative findings because they do not perceive as great a need to rule out competing diagnoses.

TABLE 1: Categories of information acquisition by level of expertise.

CATEGORY	EXPERTS		RESIDENTS		NOVICES	
	MEAN	S.D.	MEAN	S.D.	MEAN	S.D.
Question/Answer Exchanges	77.4	31.9	119.4	19.1	122.8	20.1
Total Observations	41.6	19.2	61	7.3	56.8	9.9
Total Findings	20.8	2.9	27.4	4.5	23	4.8
Positive Findings	14.4	2.5	14.6	3.1	13	3.8
Negative Findings	6.4	3.6	12.8	5.6	10	4.9
Percentage of Relevant Findings	72	13	60	19.5	45	8.2

The efficiency measures are supportive of the findings suggesting a developmental continuum in the selectivity of information. The multivariate analysis with the first two efficiency measures, observations divided by exchanges and findings divided by exchanges, as dependent variables, resulted in significant group differences ($F[2,12] = 8.69, p < .01$). The expert endocrinologists had on average the highest efficiency scores on both measures followed by the residents and students. The differences are most salient on the second measure, the percentage of findings generated per question/answer exchange. This indicates that the experts are more efficient in identifying cues in the patient data and ask the appropriate questions which lead to the generation of findings. The medical students are the least efficient in focusing probes on aspects of the problem that lead to the generation of findings. Since they do not have much clinical experience, they may not be as capable of recognizing the specific contexts where findings are manifested. The univariate analysis of variance on the third measure of positive findings per question/answer exchange also resulted in highly significant group differences ($F[2,12] = 7.28, p < .01$). The expert group had a significantly higher group mean (18.6) on this measure than either the residents (12.2) or the students (10.6). This would suggest that the expert group tend to recognize aspects of the problem early on and attempt to pursue hypotheses that suggest candidate findings that lead to a component of the diagnostic solution. While the other two groups have greater difficulty in isolating the problem and are less capable of anticipating potential findings.

The differences that have emerged thus far are suggestive of differences related to the selective acquisition and processing of information. The results pertaining to the percentage of relevant findings generated are in the expected direction. The univariate analysis comparing the relevant findings elicited yielded a significant group effect ($F[2,12] = 4.46, p < .05$). The endocrinologists generated 72% of the available relevant findings, followed by the residents with 60% and the medical students who were only able to elicit 45% of the relevant findings.

To accurately diagnose the case, it is essential to elucidate the most significant findings in the case. Perhaps the most telling index of the selective acquisition is the percentage of relevant findings over the total number of findings generated. This reflects an ability to isolate and elaborate upon the most fundamental aspects of the clinical problem. It is this statistic that reveals the most striking differences

between groups. Fifty percent of all findings elicited by the expert physicians resulted in the generation of a relevant finding. This was considerably more than the other two groups, with the residents yielding 33% and the medical students yielding 31% of the total findings they generated as relevant. This indicates that the endocrinologists are able to recognize patterns in the data and focus on the pertinent information leading to components of the diagnosis. The residents generally can identify much of the relevant information. However, they do not possess the same domain specific knowledge and are consequently less selective in the information they elicit. The students do not have the knowledge nor the clinical experience to generate the relevant information that has to be accounted for.

PROBLEM REPRESENTATION STRATEGIES: Thus far, it is at the level of findings that the most salient differences have emerged. This section attempts to elucidate the broader spectrum of context and describe the nature and sequence of information in focus. This accounting provides a framework to discuss differences in the hierarchical structuring of information and the use of strategies in the construction and modification of a patient problem representation. This analysis is primarily descriptive with illustrations from individual protocols provided to highlight certain points. In the context of the clinical interview, the problem space is not defined a priori, it is constructed in the course of the evolving interaction. We can assume that there are multiple pathways to the goal state of diagnosis, which vary greatly in their efficiency and efficacy. It is proposed that through examining the sequence and nature of information in focus, we can characterize the strategic construction of a problem representation.

Predictive reasoning involves the generation of an initial problem formulation and the subsequent pursuit of specific candidate findings, driven by a hypotheses or small subset of related hypotheses. It is analogous to a forward-chain of reasoning. It is a very powerful strategy with an associated degree of risk. It has the effect of controlling uncertainty, by delimiting the number of variables that must be accounted for in the problem resolution process (Patel, Evans & Chawla, 1987). It is an efficient strategy that minimizes cognitive load and maximizes the allocation of resources towards the assessment of a particular hypotheses or a small subset of related hypotheses. The associated risk is that the scope of probing for information will be too narrow and essential elements may be overlooked.

An excerpt from the interview of expert 2 is presented in Table 2. At this point in the interview this subject had documented the patient's presenting complaint. In question/answer exchange 20, she focuses on one of the cardinal signs of hyperthyroidism, *heat intolerance*. She provides an exemplar, describing a circumstance in which the patient may be most cognizant of this problem. The patient at this point rejects that characterization. She continues to pursue this line of reasoning, eliciting several of the critical findings characteristic of hyperthyroidism. After establishing that the patient had been eating more (exchange 24), she attempts to link the thyroid component with the episodes of muscle paralysis. In exchanges 25 and 26, she focuses on the particular foods that the patient had been indulging in. Carbohydrates predispose patients to attacks of hypokalemic paralysis. After the patient agrees with this description, the physician subsequently attempts to determine whether the patient recalls eating a carbohydrate meal prior to the second muscle attack. This protocol documents the efficacy of predictive reasoning in generating findings from facet-level hypotheses and in integrating facets in the formulation of diagnostic hypotheses.

TABLE 2: DOCTOR-PATIENT DIALOGUE EXPERT 2

20. D. No. Had you noticed that you were minding the heat any more than other people you live with, or other people in the room?
P. Umm, no
21. D. No. Did you notice any tremor of your hands?
P. Well, yeah.
22. D. How long before?
P. Four months.
23. D. Four months. Did you notice any weight loss?
P. Around ten pounds.
24. D. Ten pounds. Were you eating more?
P. Yes.
25. D. Yes. What kinds of things were you eating?
P. Uhh, meat; beef, well, fish, sometimes ribs.
27. D. Well, sweet things, or macaroni, noodles, kind of carbohydrate things?
P. Well, yeah.
30. D. Do you recall if you were eating anything unusual just before this attack happened?

P. Umm, No.

An excerpt from a residents (R3) is presented in Table 3. In this situation, the interview had just begun and the subject was attempting to characterize the patient's presenting complaint. In the excerpt the topic in focus is leg muscle weakness. The resident is attempting to characterize the problem in terms of a circumstance that he could recognize as being exemplary of a medical context he could explain. Despite the fact that the interview had just started, the subject makes no attempt to elaborate on the patient's complaint. He presents the patient with a series of probes in the form of binary response choices (e.g. yes or no), using exemplars depicting circumstances in which the weakness may manifest itself.

In the course of this exchange, he fails to pick up on certain critical pieces of information. In question/answer exchange 6, for instance, the subject focuses on the periodic pattern of the weakness. The patient responds that it has only happened twice, a critical piece of information for determining that this is *periodic* as opposed to *generalized weakness*. The clinician continues to attempt to characterize the nature of the weakness, under the assumption that this is a problem of *chronic muscle weakness*. He provides the patient with potential situations in which the weakness may manifest itself, demonstrating that he has a causal model of different leg muscle problems and is attempting to instantiate the corresponding findings. In the course of pursuing these hypotheses, he has overlooked three essential findings, including the aforementioned finding. In exchange 9, the patient informs the physician that it just happened suddenly. From his subsequent probes, it appears that he does not conclude *weakness of sudden onset*, but rather infers that the weakness is generalized, but intermittent in severity. The follow-up questions, only a sample of which are included here, focus on potential palliative and exacerbating factors that impact on the weakness.

In this stretch of discourse, it is apparent that he has misrepresented the patient's complaint. The initial formulation of a problem representation constrains the search space and determines the probes that follow it. This was clearly the case in this protocol. This subject was unable to characterize any aspect of the patient problem, primarily for the reason that he was working under the assumption that this was a problem of *chronic muscle weakness*, and never triggered the appropriate facet of *acute periodic weakness*. As a result, the subject introduced greater uncertainty into the problem space, by generating multiple competing facet-level hypotheses, none of which adequately characterized the patient's condition. This is apparent throughout the protocol and even in the diagnoses listed in his differential.

TABLE 3: DOCTOR-PATIENT DIALOGUE RESIDENT 3

3. D. What's the problem?
P. I have muscle problems.
4. D. You have muscle problems? In your legs? What do you have, muscle weakness, or?
P. Weakness, no pain.
5. D. No pain at all. How long have you had the weakness?
P. Three months.
6. D. Three months. Is it there all the time, or does it come and go?
P. It's coming and going. Its happened twice.
7. D. Ok. Do you find let's ..say...Do you have the weakness when you wake up in the morning?
P. No.
8. D. Do you have to do a certain amount of exercise before you get weak?
P. Ah, no.
9. D. So, let's say...let's say you could walk down the street and as soon as you started walking you could notice that there is a weakness. Would you say that you walk a distance first and then you get weak?
P. Ahh... not really, it just suddenly felt weak.
10. D. And do you have periods let's say where you feel weak and then you feel normal again and then you are weak? Like that?
P. Yeah.
12. D. Ok. Is there anything you know that makes it better? You feel less weak ...or.
P. No.

The primary differences between the experts and the residents protocols are reflected in the ability to formulate an initial representation of the patient problem. The expert effectively characterizes the patient's presenting complaint and subsequently pursues a high-yield knowledge-based strategy to confirm a diagnostic hypothesis. The resident fails to elaborate the patient's presenting complaint and as a consequence, pursues general data-gathering strategies that continuously widen the scope of diagnostic

candidates without any forward movement towards the solution state. The differences can be attributed to domain-specific knowledge and domain-specific experience with patient's presenting with similar histories.

As expected, there were differences in terms of the accuracy of diagnosis, with the experts generating the most accurate diagnoses, followed by the residents and medical students, respectively. The most global measures of subjects' performance such as the total number of observations and findings generated are neither adequate indices of expertise nor meaningful indicators of successful performance. Novices and residents tended to elicit more total observations and findings than the experts did. It appears that the subexpert groups needed to elicit more observations to generate findings and more findings to derive facets. However, it is the nature of the findings and not the number of findings that is associated with successful performance. Findings need to be coordinated and organized under the relevant facets, which in turn are extrapolated to the appropriate diagnostic context. There were several instances where elements of a significant finding were elicited but because of inappropriate background assumptions were misconstrued.

The study demonstrated that experts have a greater capacity to elicit the appropriate information in context and elaborate the constraints of a problem-representation. Given a complex problem, intermediate subjects experience difficulty in integrating information from multiple sources. They tend to pursue hypothesis-driven strategies that lead them down multiple search paths. Novice medical students experience greater difficulty at the local level—deriving findings from the patient's observations.

In general, the results of this study suggest that the selective acquisition and efficient utilization of information are hallmarks of expertise. This finding is consonant with results of other problem-solving studies in medicine and in other domains that indicate that the rapid, systematic application of knowledge to problems and the use of effective knowledge-based strategies leading to accurate solutions distinguish performance at different levels of expertise and individuals varying in their competency.

REFERENCES

- Barrows, H.S., Feightner, J. W., Neufeld, V. R. & Norman G. R. (1978) Analysis of the Clinical Methods of Medical Students and Physicians. Final Report, Ontario Department of Health, Hamilton, Ontario.
- Elstein, A. S., Shulman, L. S., & Sprafka, S. A. (1978). Medical problem solving: An analysis of clinical reasoning. Cambridge, MA: Harvard University Press.
- Evans, D. A., Gadd, C. S., & Pople, H. E. (in press). Managing coherence and context in medical problem-solving discourse. In Evans, D. A. & Patel, V. L.(Eds.) Cognitive Science in Medicine. Boston: MIT Press.
- Feltovich, P. J., Johnson, P. E., Moller, J. H., & Swanson, D. B. (1984). LCS: The role and development of medical knowledge in diagnostic expertise. In W. J. Clancey & E. H., Shortliffe (Eds.), Readings in medical artificial intelligence. Reading, MA.: Addison-Wesley.
- Johnson, P. E. (1983). What kind of expert should a system be? The Journal of Medicine and Philosophy, 8, 77-97.
- Kaufman, D. R., (1987). Representation and utilization of information during the clinical interview in medicine. Unpublished M.A. Thesis, McGill University, Montreal, Quebec.
- Miller, R. A., Pople, H. E. & Myers, J. D. (1984). Internist-I, an experimental computer-based diagnostic for general internal medicine. In W. J. Clancey & E. H., Shortliffe (Eds.), Readings in medical artificial intelligence. Reading, MA.: Addison-Wesley.
- Patel, V. L., Evans, D. A., & Chawla, A. S. (1987). Predictive versus diagnostic reasoning in the application of biomedical knowledge. In Proceedings of the Ninth Annual Meeting of the Cognitive Science Society, Seattle, WA.
- Patel, V. L., Evans, D. A., & Kaufman, D. R. (in press). Cognitive framework for doctor-patient communication. In Evans, D. A. & Patel, V. L.(Eds.) Cognitive Science in Medicine. Boston: MIT Press.
- Patel, V. L., & Groen, G. J. (1986). Knowledge based solution strategies in medical reasoning. Cognitive Science, 10, 91-116.

ACKNOWLEDGEMENTS: I would like to thank all the volunteers who participated in the study. This research was supported in part by a Natural Science and Engineering Council of Canada grant (# 82598) to Vimla Patel. I wish to express my appreciation to Dr. Y. Patel who assisted in the medical aspects of the study and Anoop Chawla, Stephen Chase, José Arocha, and Aldo Braccio for their editorial assistance.