Title
Hypothesis Generation and the Coordination of Theory and Evidence in Medical Diagnostic Reasoning

Journal
Proceedings of the Annual Meeting of the Cognitive Science Society, 13(0)

Authors
Arocha, Jose F.
Patel, Vimla L.

Publication Date
1991

Peer reviewed

Hypothesis Generation and the Coordination of Theory and Evidence in Medical Diagnostic Reasoning

José F. Arocha and Vimla L. Patel

Cognitive Studies in Medicine
Centre for Medical Education
3655 Drummond Street, Room 529
McGill University
Montreal, Quebec, Canada
H3G 1Y6
AX53@MUSICA.McGill.CA

Abstract

This paper investigates the process of hypothesis generation and the coordination of hypothesis and evidence in medical diagnostic tasks. Two issues are addressed: the generation of hypothesis and the directionality of reasoning. Two problems whose initial presentation suggested an initial hypothesis were presented to subjects with different degrees of expertise in clinical medicine. When faced with contradictory evidence against the initial hypothesis, 1) early novices either modified the initial hypothesis, or ignored, or reinterpreted the cues in the problem to fit the hypothesis; 2) intermediate novices generated concurrent hypotheses to account for different sets of data; and 3) advanced novices generated several initial hypotheses and subsequently narrowed the hypothesis space by generating a single coherent diagnostic hypothesis. All subjects, used a mixture of forward reasoning and backward reasoning. A more forward-directed reasoning was related to diagnostic accuracy. These results on diagnostic reasoning are discussed in relation to findings on scientific reasoning.

The study of diagnostic reasoning has had a long history in various fields, especially in troubleshooting and medicine. In medicine, this study has had a number of different forms. One is the normative study of diagnostic reasoning based either on the rules of deductive logic or on the experiential knowledge and intuition of physicians (Feinstein, 1967). Another is the psychometric study of the factors affecting physicians' diagnostic skills (Rimoldi, 1961). Still, another form is the comparison with the processes assumed to be characteristic of scientific reasoning. The cognitive tradition in the study of medical reasoning, which started with the publication of the book by Elstein, Shulman, and Sprafka (1978), gave support for the view that diagnosticians reasoned in way similar to the way that scientists are assumed to reason; that is, by applying hypothetico-deductive methods. This view was later challenged (Groen & Patel, 1985) given arguments from philosophy of science and results from studies in cognitive psychology. This led to the conceptualization of reasoning as analogical, rather than deductive. Philosophical argument as well as cognitive science research supports this position.

This work was supported by a grant from the Medical Research Council of Canada (MH004) to Vimla Patel and by a fellowship award from the Royal Victoria Hospital Research Institute (McGill University) to José Arocha.

Criticism of the work of Elstein, Shulman, and Sprafka (1978) was made because they seemed to have proposed hypothetico-deductive reasoning as psychologically real, rather than as an intellectual tool. In actuality, people do not think logically; rather, they think by means of "retrieving" analogous cases from prior experience. Furthermore, as it is now widely known, people make logical mistakes in their reasoning; and obviously, scientists or physicians are no exception. However, this does not mean that they do not *use* hypothetico-deductive methods. They may use it in circumstances where it is required.

Hypothesis Generation in Scientific and Diagnostic Reasoning

Several attempts have been made by researchers working within different theoretical approaches to investigate the reasoning processes in scientific discovery tasks: Langley, Simon, Bradshaw, and Zytkow (1987) in studies and simulation models of how scientists solve scientific problems; Dunbar (1989), Dunbar and Schunn (1990), Klahr and Dunbar (1989), Kuhn (1989), who have worked on simulated scientific problems; and Tweeney and Hoffner (1987) in their reconstruction of the scientific reasoning from Faraday's diaries.

Kuhn (1989) addressed the issue of the conceptual understanding and its relationship to processes and strategies in scientific problem solving. She argued for a developmental shift in the processes used in coordinating theory and evidence from childhood and adulthood. Children, more than adults, fail to differentiate evidence and theory; these become integrated, undifferentiated, in a single whole. Adults, contrarily, are able to differentiate between the two and can therefore accommodate or revise their theories when faced with inconsistent evidence, something that children are unable to do. Kuhn (1989) has shown that subjects tend to ignore evidence that contradicts the theories they believe to be true. She explains this by hypothesizing that subjects fail to separate their beliefs about the world from the evidence that support them. This separation is, in part, determined by the type of evidence the subjects have to coordinate. The weaker the evidence the more difficult it is to achieve such a separation. The difficulty of problem solving increases with the the presence of weak evidence. The process of solution with weak evidence requires combining the pieces of evidence such that they add up to the equivalent value of a strong evidence. Strong evidence more directly suggests a solution.

A theoretical distinction has been made by Klahr and Dunbar (1988) in their work on scientific discovery: working on a science domain involves a search on, at least, two spaces, the space of the *hypotheses* and the space of the *experiments*. They have found that subjects' doing scientific discoveries do a search on these two spaces to solve problems. In a recent study, Dunbar and Schunn (1990) presented two problems with the same underlying structure to subjects in a scientific discovery task. The authors concluded that in scientific discovery, 1) subjects search for a hypothesis that can account for a previous instance which is similar to the present problem; 2) the hypothesis searched is then tested; 3) if this test fails, then the subject searches for an alternative hypothesis. This search is for an underlying mechanism and if a mechanism is primed then it is selected ; 4) if no hypothesis is found, then the subject *adjusts the present hypothesis to conform to the data.*

In a medical context, the process of hypothesis generation in the diagnostic process has been studied by Joseph and Patel (1990). They compared low-knowledge and high-knowledge subjects in the process of evaluating a clinical case in endocrinology. Low-knowledge subjects were physicians who had general knowledge of the domain but lacked specialized knowledge; high-knowledge subjects were specialists in endocrinology. The results showed that it is not the process of generating hypotheses what distinguishes low and high knowledge subjects. It is the process of hypothesis evaluation. High-knowledge subjects evaluated their hypotheses by searching for confirmation and refining their initial hypothesis. Low-knowledge subjects generated alternative hypotheses after producing the correct hypothesis, even though the data could not discriminate among them. Joseph and Patel (1990) concluded that diagnostic reasoning can be described as a two stage process: a hypothesis generation phase, which is dependent on general domain knowledge; and a process of evaluating the hypotheses thus generated. This stage is dependent on the specific domain knowledge possessed by the subjects.

Directionality of Diagnostic Reasoning

At a more detailed level of analysis, the processes involved in diagnostic reasoning have been shown to make use of two types of inferences. These have been termed *forward* and *backward*. The term forward reasoning refers to an inference pattern that acts on evidence and generates a hypothesis from it. Research has shown that forward reasoning is characteristic of successful expert performance in some domains, and that it is used when subjects possess the relevant knowledge to solve the problem. It is highly efficient and fast. Backward reasoning is a pattern of inference in which a hypothesis serves to generate the evidence. Its direction, then, is from hypothesis to data. An alternative, but related definition is that it goes from the goal to the problem data. In such a instance, it may describe the directionality of a procedure. In either case, there is first the generation of a hypothesis and then the checking of the evidence against the hypothesis. Backward inference is used mainly when there is lack of sufficient knowledge to reach the solution from the data alone. It is slow and makes heavy demands on working memory. It is most characteristic of the novice and it is frequently tied to unsuccessful performance. However, its use is a function of both the problem solver and the problem to be solved. Nonetheless, it has been found to characterize the problem solving of experts in various domains (Larkin, McDermott, Simon, & Simon, 1980; see Anderson (1987) for a curious situation in the domain of programming).

Although forward reasoning is frequently associated with expert performance, there is evidence that it is also used by novices. Recent research (Patel, Groen, & Norman, in press) has provided evidence that the directional pattern of reasoning may be linked to curricular factors in novices. Patel, Groen, and Norman carried out two studies which they compared the diagnostic reasoning process of students who were trained in two different curricula. One group (the PBL group) was trained in a problem solving environment, which emphasizes the utilization of a hypothetico-deductive approach to medical diagnosis, and where the acquisition of basic science is concurrent with clinical practice. The other (the CC group) was trained in a traditional medical curriculum, where most of basic science is taught independently of, and prior to, clinical training. The results showed that whereas the PBL group relied more on backward reasoning and produced more elaborations, the CC group relied more on the use of forward reasoning.

This paper presents some results from a larger research project (Arocha, 1991) on the process of problem solving by novice diagnosticians. The results reported here regard the process by which medical students, trained in a traditional curriculum, generate and evaluate diagnostic hypothesis.

Method

Subjects: The subjects were second, third, and fourth year medical students at McGill University (4 subjects per group) and one cardiology resident. Second year students had just started an introduction to clinical medicine; they had no experience in the hospital and, therefore, their knowledge of disease comes principally from lectures and textbooks. Third year students were in their first year of clinical training in the hospitals. Fourth year students had just completed their medical. degree requirements with two years of biomedical sciences and two years of clinical training, and were preparing for their provincial and federal certification exams.

Materials: Two real-life clinical problems in cardiology selected were modified in such a way that the relevant information was presented later in the case. Both cases were similar in their initial presentation, but differ in subsequent information. The cases were divided into three segments:

(1) presenting complaint; (2) past history; and (3) results of the physical examination.

Procedure: The subjects were all tested individually in a single 90 minute session. Their task was to explain each case as completely as possible. The cases were presented on three consecutive cards, each containing one segment of information. The subjects were instructed to think out loud while they read the segments. They were asked to make a tentative diagnosis after each segment. Then, they followed to the next card and repeated the procedure for the second and third cards.

Analysis: The analysis consisted of representing the cases and protocols in terms of propositions (Frederiksen, 1975; van Dijk & Kintsch, 1983). From these propositions, a reference frame for each case was constructed as a semantic network, which has been successfully used in cognitive research in medicine by Patel and Groen (1986). The reference frames so developed served as a comparison for the assessment of individual frames constructed from the subjects' protocols. The network representations contained mostly conditional and causal relationships. Other relationships were attributive, locative, and temporal. Coordination of theory of evidence was assessed by confirmation (i.e., a positive cue serves to generate a hypothesis) and disconfirmation (a positive or negative cue serves to rule out a hypothesis) strategies. Directionality of reasoning was assessed by the use of backward (i.e., hypothesis is matched to cue) and forward (cue suggests hypothesis) inferences.

Results

A summary of the results for all groups of subjects is presented first. The results include accuracy of the diagnosis, the hypothesis generation and the evaluation strategies the subjects used, and the directionality of their reasoning. A detailed analysis of two subjects' schematic representations follows. This allows us to see the process by which these subjects dealt with the evidence for or against the diagnostic hypotheses in the cases and describe the processes they used to solve the problems, specially the strategies they followed and the directionality of their reasoning.

Diagnostic Accuracy: Five subjects reached the correct diagnostic hypothesis: Two second year students, two fourth year students and the resident. Only the resident was successful in both problems. No third year student generated the correct solution. Although there was little success in reaching accurate diagnosis in most cases, the subjects were able to select all the important cues in the case for interpretation.

Hypothesis Generation: The number of hypotheses generated varied across groups; specially, those generated at the first segment, with the second year students generating the fewer diagnostic hypotheses ($\bar{x}$= 1.5) and third and fourth year students generating a larger number ($\bar{x}$= 3 and $\bar{x}$=3.4, respectively). The resident generated the same number of diagnostic hypotheses for both cases (n=5). In terms of the quality of the diagnoses, all subjects generated the initially suggested hypothesis at the first segment.

Evaluation Strategies: The most widely used evaluation strategy was confirmation. A confirmation strategy was one which linked a positive cue to a hypothesis as support for the hypothesis. Seventy one percent of the cues by second year students were of this kind; 88% were confirming relations for both third year and fourth year students. The resident generated 92% confirming relations. All second year students kept their initial hypothesis even when contradictory data were presented. They ignored contradictory evidence or modified slightly the initial hypothesis to fit the data. Third and fourth year students used concurrent hypotheses to account for different cues in the case. These were of two kinds: cardiac hypotheses to account for the initial cues and pleuro-pulmonary to account for cues that did not fit the initial diagnosis. The expert subject generated more diagnostic hypothesis, but was able to produce a single hypothesis in the third segment by selecting the one that accounted for more of the data.

Directionality of Reasoning: Most reasoning inferences were forward inferences. In the cases in which the correct diagnoses were reached, 4 subjects used forward inferences to conclude the final diagnosis. All these corresponded to case 2 (AD45). The resident also used backward inference in reaching the solution to case 1 (VP45), when provided with contradictory evidence.

Detailed Analysis of Two Protocols: This section presents a comparison of the solution processes given by a second year student and the resident to case 1 (VP45) . Figure 1 presents the schematic representation of the problem solving process generated from the protocol of subject 2.1. The subject starts by generating the hypothesis of myocardial infarction, given the cues of *severe central crushing chest pain* and *chest pain of four hours*. The subject adds support for this hypothesis by noting that the age of the patient is in keeping with a related affection: coronary artery disease. A second hypothesis, which the student disconfirms, is rheumatic fever (streptoccocal infection), based on the finding of *chest pain of 4 hours*. The negative findings in the case are used to disconfirm pulmonary edema secondary to left ventricular failure, which may result from myocardial infarction. Lastly, the subject concludes that the *cough* is irrelevant to the patient's main condition. The subject bases her conclusion on the negative findings of *no hemoptysis* and *no sputum*, which are associated with cough in myocardial infarction.

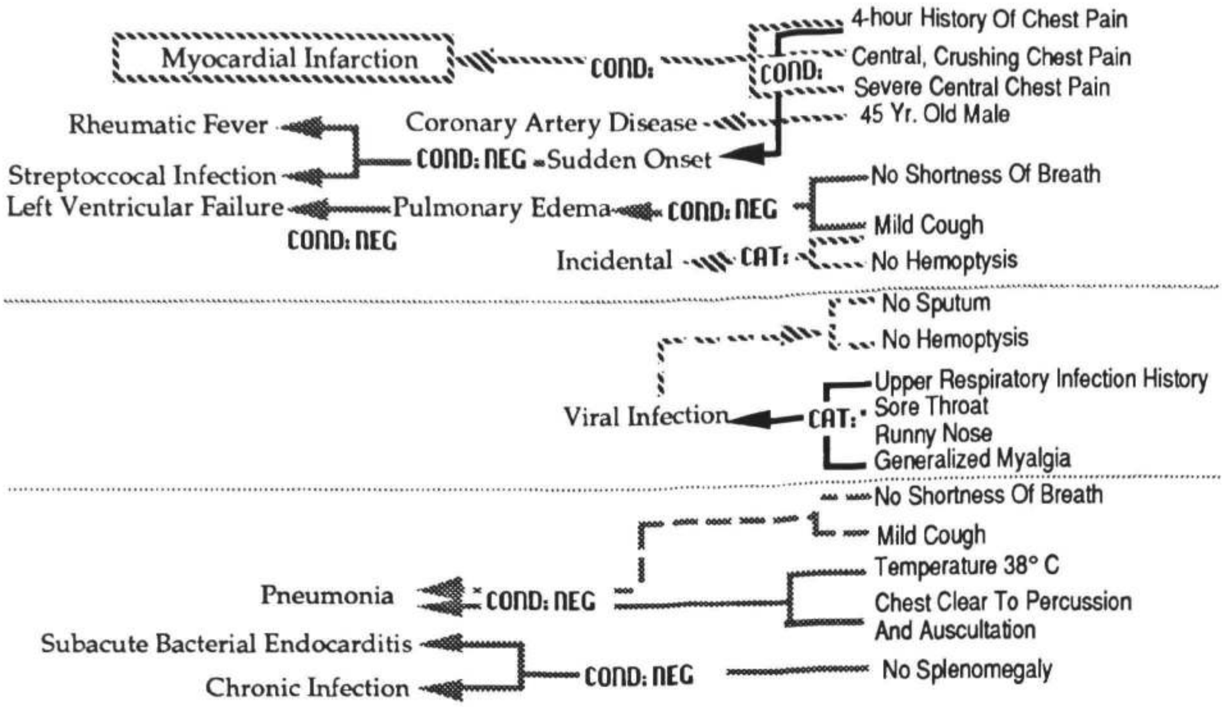

Figure 1. Schematic representation of the problem solving process of case 1 (VP45) by the a novice. A ▬▬ represents a confirming relation towards the correct diagnosis; a ⋯⋯ represents a confirming relation towards and alternative diagnosis; a ▬▬ represents a disconfirming relation. A dashed line represents a relation linking a cue in the current segment to either a hypothesis or another cue in previous segments. COND: represents a conditional relation; CAU: represents a causal relation; NEG represents a negation; CAT: represents a category relation. A box indicates the final diagnosi reached by the subject.

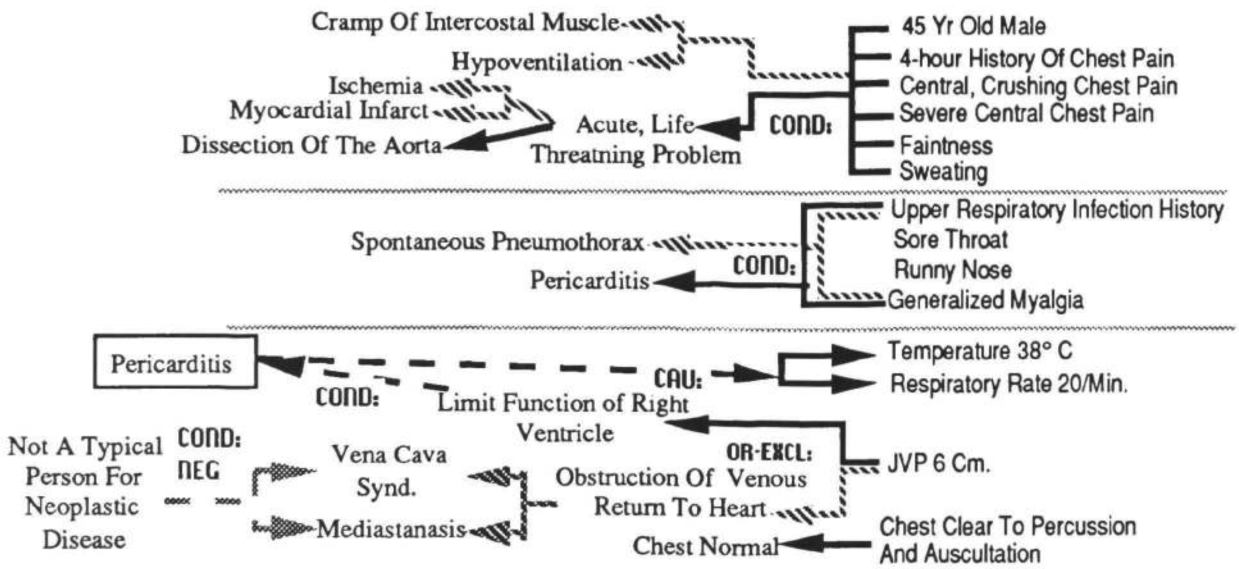

Figure 2. Schematic representation of the problem solving process of case 1 (VP45) by the resident. A ▬▬ represents a confirming relation towards the correct diagnosis; a ⋯⋯ represents a confirming relation towards and alternative diagnosis; a ▬▬ represents a disconfirming relation. A dashed line represents a relation linking a cue in the current segment to either a hypothesis or another cue in previous segments. COND: represents a conditional relation; CAU: represents a causal relation; NEG represents a negation; OR-EXCL: represents an exclusive disjunctive relation. A box indicates the final diagnosis reached by the subject.

Thus, up to this point, this second year student has generated a single diagnosis, that of myocardial infarction, with one modification: no left-ventricular failure. The rest of the hypotheses she generates are disconfirmed: they are ruled out by the evidence. The second segment introduces the findings of *history of upper respiratory infection*, *sore throat*, *runny nose* and *generalized myalgia*, which the subject correctly interprets as a viral infection. However, she uses these findings to conclude that is not a streptoccocal infection, which is related to the hypothesis, disconfirmed in the first segment, of rheumatic fever (streptoccocal infection precedes the appearance of rheumatic fever). This subject then concludes that the infection is unrelated to the chest pain. Thus, the subject at this segment only rules out the relevance of the infection while maintaining the diagnosis of myocardial infarction. The third segment introduces 7 new findings of which the subject selects 3 for comment. These are *temperature 38°C*, *chest clear to percussion and auscultation*, and *no splenomegaly*. The same strategy she used in the previous segment is used here: she uses the findings to disconfirm alternative diagnoses. This time, pneumonia and subacute bacterial endocarditis are ruled out: Pneumonia involves congested chest and subacute bacterial endocarditis involves enlarged spleen. In summary, this subject generated the initial diagnosis of myocardial infarction and maintained it in the presence of contradictory evidence. Further information, which go against such initial diagnosis served to modify the initial hypothesis without rejecting it (i.e., right ventricular rather than left ventricular failure). Other information is either acknowledged, but interpreted as unrelated to the episode of chest pain, or ignored.

Consider now the solution given by the resident to the same case. First, the resident considers a larger number of diagnoses: from a cramp of the intercostal muscles to hypoventilation to myocardial infarction or dissection of the aorta, to pericarditis. Thus, instead of assuming that the central chest pain corresponds to myocardial infarction, the resident acknowledges that this presentation may be produced by other diseases. The subject employs a breadth-first search by considering the possibility of different alternative diagnoses of varied origin, quality, and seriousness. The second segment suggest to the resident the possibility of an infection, but the main focus is on ruling out any serious disease. The infection does not strongly suggest any particular diagnosis. However, he pays more attention to diagnoses that may relate an infectious process and chest pain: pneumothorax or pericarditis At the third segment, the subject focuses on the findings of *fever*, *jugular venous distension* and *chest clear to percussion and auscultation*. To account for the case, then the resident looks for hypothesis that would compatible with the pattern of chest pain, the jugular venous distension, and the clearness of the chest. He generates two intermediate, non-diagnostic, hypotheses that can produce jugular venous distension and chest pain: one is diseases that involve an obstruction of the venous return to the heart (*vena cava syndrome*, *mediastinitis*) the other is diseases that limit of the diastolic function of the right ventricle (*pericarditis*). For the former, the subject finds no evidence in the case presentation. He apparently matches the features that identify a typical person with neoplastic disease with the features of the present case, from which he discards the hypothesis. The other type of underlying process that may cause such a pattern is one that limits the diastolic function of the right ventricle, such as the case of pericardits. Therefore, pericarditis must be the correct diagnosis. The resident does not consider the possibility of myocardial infarction because the *chest is clear*, the patient does not have pulmonary edema, which he would have if he had an myocardial infarction and because the finding *respiratory rate of 20* is not consistent with heart failure.

In summary, the strategy employed by this subject differs substantially from the one used by the first subject. First, he starts by generating several initial diagnoses, even though the case strongly suggests only one. The subject uses subsequent information to narrow the possible hypotheses: in the second segment, by focusing on diseases of infectious character; and in the third segment, by trying to fit all the main case cues into a single hypothesis that covers them all. He uses a form of backward inference which serves to generate the correct diagnostic hypothesis. By postulating two possible causal mechanisms, either an obstruction of venous blood flow to the heart or a limiting of the diastolic function of the right ventricle, and matching the pattern of cues to what would be produced by the two such processes, the subject concludes the correct diagnosis. Such use of bacward inference requires a vast background of biomedical knowledge to be successful (cf. Patel, Groen, Arocha, 1990).

Discussion

First, the results show that degree of training does not have an important effect on what findings are chosen as being relevant; that is, all subjects focused their attention on the most important findings. This is consistent with findings in previous research (Braccio, 1988). Also, it has been found that novices are frequently able to recall the important findings in a case but are, nonetheless, unable to use those findings correctly in the solution of the case (Patel & Frederiksen, 1984). Second, second year students seemed to have been strongly influenced by the initial hypothesis and were unable to change their diagnosis in light of contradictory evidence. More experienced students seem to have recognized the importance of contradictory evidence but were also unable to override their initial hypothesis; instead, they generated concurrent hypotheses to explain sets of cues. The resident was the only subject who changed the initial diagnosis as contradictory evidence was presented and who attempted to provide a single coherent hypothesis to account for all the case cues. These results are compatible with those obtained in research on the scientific discovery process (Dunbar, 1989; Dunbar & Schunn, 1990; Klahr & Dunbar, 1988). However, there may be specific differences between levels of expertise, as the present results suggest. As expertise with a domain increases, subjects seem to be more aware of the role of providing coherence to the problem interpretation. The three general processes reported here, namely, 1) generating a major hypothesis and either reinterpreting or ignoring contradictory data; 2) generating concurrent hypotheses to account for different sets of cues; and 3) generating various initial hypotheses and then narrowing the hypothesis space and interpreting the data in terms of a single hypothesis by imposing coherence, may be three general strategies linked to the development of expertise. Further investigation is required to validate this conclusion.

The two schematic representations presented on the problem solving process of two subjects serve to highlight some points. Previous research has found the use of forward inferences linked to successful performance of expert subjects. In this paper, as it has been reported elsewhere (Patel, Groen, & Norman, in press), the use of forward reasoning was the most consistently used type of inference. However, it consisted of a "shallow" form of forward inference as it was not supported by knowledge of the underlying mechanism. Backward reasoning was used in a few occasions mainly by third year students, who generated the longest protocols and gave the largest number of explanations for the cases, and by the resident.

The case of the resident shows an instance where the use of a backward inference is successfully used to generate the correct diagnosis. This case is also interesting because it presents the use of biomedical knowledge to reach the correct diagnosis. It has been argued that biomedical knowledge forms a separate "world" from the clinical knowledge. This interpretation is, to some extent, supported in the present study. Recall that second year students had no clinical training; their only knowledge of clinical medicine was through lectures, but they had 2 years of basic science teaching. Even these subjects, who one may expect to use their knowledge of basic science to solve the clinical problem, did not do so (cf. Patel, Groen & Norman 1990). On the contrary, they relied almost exclusively on the association between clinical cues and diagnostic categories. The resident, however, was the subject with longest clinical training. Nonetheless, he used knowledge of basic science to reach the diagnosis.

References

Anderson, J.R. (1987). Skill acquisition: Compilation of weak–method problem solutions. *Psychological Review, 94*, 192–210.

Arocha, J.F. 1991. Clinical Case Similarity and Diagnostic Reasoning in Medicine. Unpublished Ph.D. diss. Department of Educational Psychology and Centre for Medical Education., McGill Univ.

Braccio, A. 1988. On-line Analysis of Novice Problem Solving on Medicine. M.A. thesis. Dept. of Educational Psychology, McGill Univ.

van Dijk, T. and Kintsch, W. 1983. *Strategies of Discourse Comprehension*. New York: Academic Press.

Dunbar, K. 1989. Scientific Reasoning Strategies in a Simulated Molecular Genetics Environment. In Proceedings of The 11th Annual Conference of the Cognitive Science Society. Hillsdale, NJ: Lawrence Erlbaum.

Dunbar, K. and Schunn, C.D. 1990. The Temporal Nature of Scientific Discovery: The Roles of Priming and Analogy. Proceedings of The 12th Annual Conference of the Cognitive Science Society. Hillsdale, NJ: Lawrence Erlbaum

Elstein, A.S., Shulman, L.S., and Sprafka, S.A. 1978. *Medical Problem Solving: An Analysis of Clinical Reasoning*. Cambridge, MA: Harvard University Press.

Evans, D.A. and Gadd, C.S 1989. Managing Coherence and Context in Medical Problem-solving Discourse. In D.A. Evans and V.L. Patel , eds. *Cognitive Science in Medicine*. Cambridge, MA: MIT Press.

Feinstein, A.R. 1967. *Clinical Judgement*. Baltimore: The Wilkins and Williams Company.

Frederiksen, C. H. 1975. Representing Logical and Semantic Structure of Knowledge Acquired From Discourse. *Cognitive Psychology, 7*, 371-458.

Groen, G.J., and Patel, V.L. 1985. Medical Problem Solving and Cognitive Psychology: Some Questionable Assumptions. *Medical Education, 19* , 95-100.

Joseph, G.-M., and Patel, V. L. 1990. Domain Knowledge and Hypothesis Generation in Diagnostic Reasoning. *Medical Decision Making, 10*, 31-46.

Klahr, D. and Dunbar, K. 1988. Dual Space Search During Scientific Reasoning. *Cognitive Science, 12*, 1-55.

Kuhn, D. 1989. Children and Adults as Intuitive Scientists. *Psychological Review, 96*, 674-689.

Langley, P., Simon, H., Bradshaw, G.L., and Zytkow, J.M. 1987. *Scientific Discovery: Computational Explorations of the Creative Process*. Cambridge, MA: MIT Press.

Larkin, J.H., McDermott, J., Simon, D.P., and Simon, H.A. 1980. Expert and Novice Performance in Solving Physics Problems. *Science, 208*, 1335-1342.

Patel, V.L., and Groen, G.J. 1986. Knowledge-based Solution Strategies in Medical Reasoning. *Cognitive Science, 10*, 91-116.

Patel, V.L. and Frederiksen, C.H. 1984. Cognitive Processes in Comprehension and Knowledge Acquisition by Medical Students and Physicians. In. H.G. Schmidt, and M.C. DeVolde, eds. *Tutorials in Problem-Based Learning*. Assen, Holland: Van Gorcum.

Patel, V.L., Groen, G.J., and Arocha, J.F. 1990. Medical Expertise as a Function of Task Difficulty. *Memory & Cognition, 18* (4), 394-406.

Patel, V.L., Groen, G.J., and Norman, G.R. 1991. Two Modes of Thought: A Comparison of Effects of Conventional and Problem-based Medical Curricula. *Academic Medicine*, July.

Rimoldi, H.J.A. 1961. The Test of Diagnostic Skills. *Journal of Medical Education, 36*, 73-79.

Tweeney, R.D. and Hoffner, C.E. 1987. Understanding the Microstructure of Science: An Example. In Proceedings of The Ninth Annual Conference of The Cognitive Science Society. Hillsdale, NJ: Lawrence Erlbaum.