

PROBLEM REPRESENTATION AND HYPOTHESIS GENERATION IN DIAGNOSTIC REASONING

**GUY-MARIE JOSEPH
VIMLA L. PATEL**

MCGILL UNIVERSITY

ABSTRACT

In this paper we examine the role of domain knowledge in the process of hypothesis generation and problem representation during diagnostic reasoning. An on-line task environment and the combination of discourse and protocol analysis techniques were used to test the differences between two groups of experts solving a clinical problem. The groups consisted of high domain-knowledge subjects (HDK) -endocrinologists- and low domain-knowledge (LDK) subjects cardiologists-. The results show that HDK subjects used a more efficient process of diagnostic reasoning as generated a more coherent representation of the problem. A two-stage model describing the process of hypothesis generation was proposed to explain the differences in the process of hypothesis generation.

The most frequently investigated aspect of medical problem solving is diagnostic reasoning. A central question for researchers in fields such as cognitive science, medical education, medical problem solving and decision making, is how do expert physicians go about making a diagnosis? This interest has been motivated by both theoretical and practical concerns.

Theoretically diagnosis is viewed as representative of the problem-solving processes wherein the problem solver is required to examine, evaluate and select information in order to generate an accurate solution. According to Schwartz and Griffin (1986) diagnosis is a subset of the more general skill of classification (i.e., assigning entities to different classes or categories) of diseases. Research in clinical reasoning has played a significant role in developing psychological theories of human problem solving (Wortman, 1972). In practical terms, this is directly applicable to the teaching of medicine (Wortman, 1972).

Diagnostic reasoning has been studied from two perspectives. On the one hand, computational/mechanistic models have been used to describe describe the reasoning processes of physicians while solving a problem (Johnson, 1983; Shortliffe, Buchanan, & Feigenbaum, 1979). The main goal of the computational approach is to develop computer programs that can perform (diagnose problems) as accurately as human experts. The first step in developing such programs consists in

collecting think aloud protocols from experts engaged in simulated clinical tasks (e.g., the diagnosis of a patient's problem), and analyzing these transcripts to formulate models of the solution strategies and reasoning processes used. These models are then represented in a computer program (simulations) that, when executed, produces behavior that can be compared with that observed in physicians (Kassirer et al., 1982).

This is considered by many (e.g., Kassirer et al., 1978; Pople, 1982) to be a promising approach for the development of knowledge-base components in expert systems, and can contribute toward the development of theories of knowledge representation (Brachman & Levesque, 1985). The main weakness of this this approach is a lack of theories and methods for dealing with the problem of natural language in protocol analysis. Central to building a more complete account of problem solving in complex, knowledge-rich domains is the need for specific information about both how the subject's representation of the problem changes over time and how his domain knowledge is used to construct and modify these representations. "Think-aloud" protocols (e.g. Greeno & Simon, 1985; Kassirer et al., 1982) provide rich, complex data that are approximately concurrent with the subject's reasoning and therefore provide information about the subject's changing representation of the problem. However, protocol analysis methods

(e.g., Ericsson & Simon, 1984) have been limited in their success at providing more than global information about subjects' processing (cf. Joseph, 1987).

From a second perspective, the psychological research on medical diagnosis has treated the physician as a processor of information. These investigations were influenced by then current models of general problem-solving ability which are summarized in Newell & Simon, (1972). Elstein and colleagues (1978) were among the first researchers to apply the information processing approach to study diagnostic reasoning. They proposed the 'hypothetico-deductive' model as the general and logical model of expert reasoning.

"Diagnosis Diagnostic problems are solved through a process of hypothesis generation and verification. Hypotheses are consistently generated early in a workup when only a very limited data base has been obtained. While any early information may be revised or discarded if subsequent data fail to confirm it, there is a high probability that at least some of the formulations of experienced physicians will be correct. Hypotheses serve as organizing rubrics in working memory. They help to overcome limitations of memory capacity and serve to narrow the size of the problem space that must be searched for solution. Since it would be impossible to conduct an efficient inquiry without some hypothetical goal that would tell the inquirer when to stop [control structures], hypotheses serve to transform an open medical problem (What is the patient's illness?) into a set of closed problems that are much easier to solve (Is the illness X? or Y? or Z?) Elstein et al., 1978, .176".

The early generation of diagnostic hypotheses is considered as the main strategy that physicians used to reduce search in the problem space that is most likely to yield the accurate solution. "In sum, we may propose that a set of problem formulations defines the dimensions of the fundamental problem space in which a physician's search for a diagnosis is conducted" (Elstein et al., 1978, p.176).

The basic assumptions of the use of the hypothetico-deductive method of reasoning used by expert physicians in clinical reasoning have increasingly come into question (1983; Groen & Patel, 1985; and Kassirer, et al., 1982). The assertion that expert diagnostic reasoning is characterized by a hypothetico-deductive method is questionable in view of the fact that the research in other domains have demonstrated this to be a 'weak' and inefficient method of problem solving, more characteristic of novice rather than expert performance (e.g.,

Feltovich & Barrows, 1984; and Simon & Simon, 1978). Weak methods are very general, in the sense that they can be used in almost any domain. Under realistic time constraints, however, they tend to yield either the wrong answer or no answer at all. In contrast strong methods almost always yield the correct answer, but are only applicable in a limited domain. The main problem however, is that the use of diagnostic hypotheses is not necessarily a reliable indicator of diagnostic expertise or the quality of diagnostic reasoning. In addition, further investigation on hypothesis generation parameters (e.g., timing of hypothesis generation, number of hypotheses generated), and hypothesis quality either yield inconsistent cross-situational results or are insufficient in discriminating expertise (Feltovich & Barrows, 1984).

Studies in problem solving and expertise suggest that the diagnostic hypotheses generated for a medical case constitute one component of an expert physician's representation of a patient's condition (Hasserbrock & Prietula, 1986). An expert representation should be an abstraction of commonalities existing in a patient's overall symptoms. Knowledge of pathophysiological processes provides experts with another level of abstraction. The appropriateness of the disease hypotheses being considered for a given case is likely to change if a physician shifts to a different underlying pathophysiological condition.

In the present study we are primarily interested in testing the role of domain knowledge in the process of hypothesis generation and problem representation in diagnostic reasoning. The use of an on-line task environment and the combination of discourse and protocol analysis will be used to test the differences in problem representation and hypothesis generation while solving the problem (Joseph & Patel, 1986a). Experts with specific domain-knowledge are expected to use a more efficient process of hypothesis generation as well as a more coherent representation of the problem.

METHOD

Materials: The stimulus was a text based on a real patient in a Montreal hospital and modified by an endocrinologist for the purposes of this study. The clinical findings for the case were assembled in a typed "patient file" format. That is, the clinical information

was arranged in the typical order of medical history, findings from physical examination, and x-ray and laboratory test results.

Subjects: Nine senior physicians associated with the Faculty of Medicine at McGill University volunteered as subjects for the study. Given that the task was the diagnosis of a patient with an endocrine disorder, the high domain knowledge (HDK) group consisted of four endocrinologists and the low domain knowledge (LDK) group consisted of five cardiologists. The physicians were all practicing physicians with from five to ten years of experience (i.e., practice) in their respective fields.

Clinical problem The endocrinology problem (Table 1) describes the case of an elderly lady who was brought to the emergency room by her daughter suffering from severe hypothyroidism. Prior to her admission into the emergency room, the patient consulted a physician complaining of difficulty in speaking and throat irritation. His diagnosis was chronic laryngitis, for which he prescribed a potassium iodide mixture as an expectorant. When taken by a healthy person, potassium iodide is quite harmless. However, when taken by a hypothyroid patient, it precipitates an acute hypothyroid crisis which leads to myxedema.

The accurate diagnosis of the case is: Hashimoto's hypothyroidism precipitated to myxedema pre-coma by the potassium iodide mixture. The diagnosis can be divided into three subcomponents varying in specificity, from general to specific. The first subcomponent is hypothyroidism. The second one is myxedema pre-coma. The third subcomponent is an autoimmune condition called Hashimoto's thyroiditis, which is the cause of the hypothyroidism.

Table 1. Endocrine problem.

<u>Medical History</u>
S1. A 63 year old woman with a one-week history of increasing drowsiness and shortness of breath was brought to the emergency room by her daughter.
S2. The patient had not been well for over a year.
S3. She complained of feeling tired all the time, had a loss of appetite, a 30 lb. weight gain and constipation.
S4. A month later she had been diagnosed as having "chronic laryngitis" and was prescribed a potassium iodide mixture as an expectorant.
<u>Physical Examination</u>
S5. Physical examination revealed a pale, drowsy, obese lady with marked periorbital edema.
S6. She had difficulty speaking, and when she did speak her voice was noted to be slow and hoarse.
S7. There were patches of vitiligo over both her legs.
S8. Her skin felt rough and scaly.

- S9. Her body temperature was 36 deg. C.
S10. Pulse was 60/minute and regular.
S11. B. P. was 160/95.
S12. Examination of her neck revealed no jugular venous distention.
S13. The thyroid gland was enlarged to approximately twice the normal size.
S14. It felt firm and irregular.
S15. There was grade 1 galactorrhea.
S16. The apex beat could not be palpated.
S17. Chest examination showed decreased movements bilaterally and dullness to percussion.
S18. There was no splenomegaly.

Results From Neurology, X-Ray and Laboratory Tests

- S19. Neurological testing revealed symmetrical and normal tendon reflexes but, with a delayed relaxation phase.
S20. Urinalysis was normal.
S21. Chest X-ray showed large pleural effusions bilaterally.
S22. ECG revealed sinus bradycardia, low voltage complexes and non-specific T-wave flattening.
S23. Routine biochemistry (SMA=16) showed Na=125, K=3.8, BUN=8 mg/100ml.
S24. Arterial blood gases $P_{O_2}=50$ mm Hg, $PCO_2=60$ mm Hg.
S25. The patient was admitted to the intensive care unit for further management

Procedure: An on-line task environment (Joseph, 1987) was used to present the stimulus material to individual subjects one segment at a time on a microcomputer. The order of presentation of text segments was controlled by the experimenter. However, the rate of presentation of each segment was controlled by the subjects. Segments consisted of one or two sentences each. After the presentation of each segment subjects were asked to give a verbal report of their interpretation of the sentence with respect to a possible diagnosis. After the presentation of the entire case, subjects were asked to summarize the case and then to provide a diagnosis. Subjects' interpretations were tape recorded and later transcribed verbatim.

ANALYSES

The methods of analysis used in this study include the use of techniques from both discourse and protocol analyses as well as as the use of quantitative and qualitative measures. Three analyses were carried out on each subject's protocols.

- 1) Segmentation.
- 2) Hypothesis generation.
- 3) Problem representation.

Techniques were taken from discourse analysis for application in analysis 1

(functional-syntactic analysis), and analysis 3 (representing conceptual structures as frames or semantic networks). Standard techniques of protocol analysis (e.g., tabulating specific kinds of information in the protocols) were used in analysis 2.

Segmentation

At the first level of analysis, subjects' protocols were transcribed and divided into syntactic units or 'segments' (Dillinger, 1984). This division facilitates identification and further analysis of the parts of the text and of the protocols. This method is based on Winograd's system of clausal analysis (Winograd, 1983) and derived from the systemic grammar of Halliday (1967).

Process of hypothesis generation

After the segmentation, subjects' protocols were analyzed and coded for the generation of diagnostic hypotheses. In clinical reasoning the term 'diagnostic hypothesis' refers to any ideas, diagnoses, or guesses that label or propose explanations which will guide investigation of the patient's problem (Barrows, & Tamblyn, 1980). These hypotheses can refer to syndromes, specific disease entities, disorders, pathophysiological processes, and anatomical or biochemical disturbances.

<p>Segment #3: She complained of feeling tired all the time, had a loss of appetite, a 30 pound weight gain and constipation.</p> <p>Subject's Comment: She had a loss of appetite, a 30 pound weight gain, and constipation. OK, right now you are wondering whether she has got hypothyroidism when you are looking at this.</p>

Hypotheses were coded only the first time they were generated. In the above example **hypothyroidism** was coded as a diagnostic hypothesis generated. Hypotheses which were repeated were not coded twice.

<p>Segment 7: There were patches of vitiligo over both her legs:</p> <p>Subject's Comment: Vitiligo over the legs, It fits with hypothyroidism, that you have one autoimmune disease that you get another autoimmune disease Hashimoto's thyroiditis and vitiligo.</p>

In this segment only **Hashimoto's thyroiditis** was coded as a diagnostic hypothesis generated. Hypothyroidism was not coded as a diagnostic hypothesis generated because it was previously generated, and the

autoimmune disease was considered to be part of the Hashimoto's thyroiditis. When a subject ruled out the possibility of an hypothesis generated previously during the presentation of the case, it was also coded as an hypothesis generated.

Problem representation

The objective of this analysis is to generate data about the subjects' changing representation of the problem, from which it will be possible to infer some characteristics about the processes they used for constructing a problem representation and about the solution strategies they used. To do this it was necessary to develop two models: the first was a general model of the knowledge necessary for generating an accurate solution to the problem, represented as a causal network and referred to here as the "reference" or "canonical" model (Joseph, 1987; Patel & Groen, 1986). The second was a specific model of the problem (i.e., a subset of the general model) generated by each subject at each stage of solving it. This technique has proven useful in studying knowledge-based differences between experts and novices (Patel & Groen, 1986), and experts with different levels of expertise (Joseph & Patel, 1986a,b) in diagnostic reasoning.

RESULTS AND DISCUSSION

The Process of Hypothesis Generation. Analysis of the time course of the production of diagnostic hypotheses focussed on differences between HDK and LDK subjects in: a) the production of the accurate diagnostic subcomponents, b) the pattern of hypothesis generation and confirmation, and finally c) the number of diagnostic hypotheses generated by each group.

First, the HDK subjects are expected to generate the accurate diagnostic subcomponents before the LDK subjects. Second, the pattern of hypothesis generation is expected to reflect two distinct stages. The first stage consists of the generation of hypotheses, followed by a stage of hypothesis confirmation and/or ruling out inaccurate hypotheses. Third, the HDK subjects are expected to generate fewer diagnostic hypotheses than the LDK subjects. This is primarily due to the differences in domain-specific knowledge of the HDK subjects.

The overall number of diagnostic hypotheses produced by the HDK and LDK subjects is presented in Figure 1. It shows

clear differences in: the pattern of hypothesis generation, the number of diagnostic hypotheses generated by the subject, and the time taken for the production of the accurate diagnostic subcomponents.

Before the presentation of Segment 7 (by which time the HDK subjects have generated all three accurate diagnostic subcomponents) the HDK subjects generated more diagnostic hypotheses than the LDK subjects. After the presentation of Segment 7 the HDKs generated fewer diagnostic hypotheses than the LDK subjects. On the other hand, the LDK subjects show a rather different pattern of hypothesis generation both before and after the presentation of Segment 7 (Figure 1). Specifically, the LDK subjects show little difference in their pattern of hypothesis generation before and after the complete diagnosis is reached.

7 (Segments 1 to 7). The letters A, B and C in Figure 1 indicate the mean segment number where the accurate diagnostic subcomponents were generated. As expected the HDK subjects generated the accurate diagnostic subcomponents earlier than the LDK subjects. Before the presentation of Segment 7 the HDK subjects generated more hypotheses than the LDK subjects. Both groups generated the first two subcomponents of the accurate diagnosis (Hypothyroidism and myxedema). However there was a time difference (i.e., number of segments) between the generation of these diagnoses in the two groups. The difference between the production of the first diagnostic subcomponent was of 1 segment (Figure 1: Diff A). The difference for the generation of the second subcomponent was of 2 segments (Figure 1: Diff B).

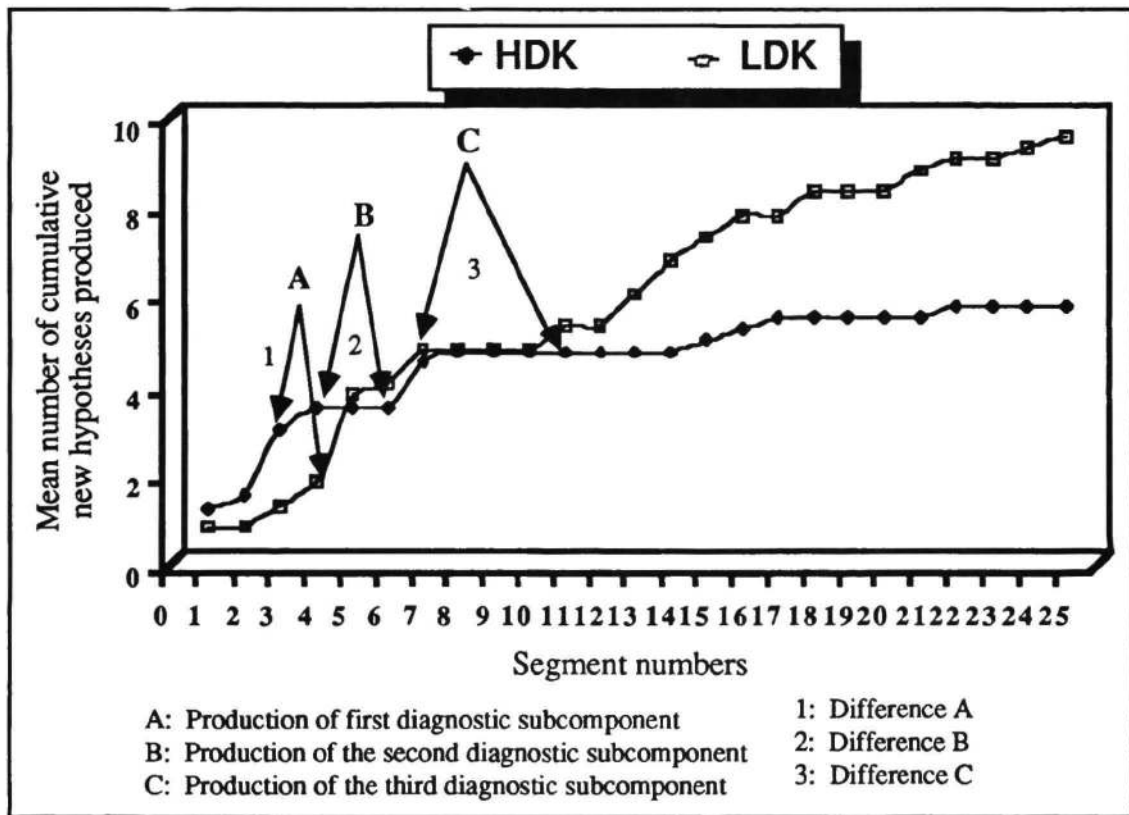


Figure 1. Cumulative number of diagnostic hypotheses generated by HDK and LDK subjects.

Hypothesis generation after Segment 7 (Segments 7 to 25): The pattern of hypothesis generation from the presentation of Segment 7 to the end of the case is also presented in Figure 1. Once again, as expected, the HDK subjects generated the third subcomponent of the accurate diagnosis earlier than the LDK subjects. The difference between the two

groups in the production of the third diagnostic subcomponent was of 4 segments (Figure 1: Diff C). Furthermore, after the production of the accurate diagnosis the HDK subjects generated very few new diagnostic hypotheses. The HDK subjects ruled out some of the hypotheses generated earlier, used the findings

from the physical examination to confirm the diagnosis, and determined secondary problems.

The analysis of the time course of hypothesis generation suggests that the process of hypothesis generation differs across groups. The HDK subjects seem to distinguish two phases in their process of hypothesis generation, whereas the LDK subjects do not. The first phase consists in hypothesis generation and the second in diagnostic confirmation.

In the first stage (before the presentation of Segment 7) the HDK subjects use the information from the medical history and a few findings from the physical examination to generate the accurate diagnostic subcomponents. This explains the finding that HDK subjects generate more diagnostic hypotheses before the presentation of Segment 7. The rapidity with which the initial hypotheses are generated constitutes the most striking feature of the behavior of experienced clinicians. Often with only the age, sex and present complaint of the patient, the clinician unhesitatingly selects a single working hypothesis.

The early generation of accurate diagnostic hypotheses is also important in determining the accurate diagnosis. Work by Barrows et al. (1978) has shown that the earlier a good hypothesis set is created, the more predictive it is of the quality of the diagnosis. Differences in the time course of the generation of the accurate diagnostic subcomponents seem consistent with the findings of speed difference found in the chess-playing ability of masters vs beginners (Chase & Simon, 1973) as well as the findings from studies contrasting experts and novices solving physics problems (e.g., Larkin et al, 1980; Simon & Simon, 1978). This difference is explained by the difference in domain knowledge between the HDK and LDK subjects.

The hypothesis generation process is interpreted as an indication of subjects' organization of the information provided into a more or less coherent problem representation. The results discussed so far indicate considerable differences in terms of the processes used by the two groups in generating diagnostic hypotheses and the accurate diagnostic subcomponents. This poses the question of the generative difference in how the two groups construct their problem representation over time. In the final set of analyses, the information selected from the input and the hypotheses adduced to structure

that information are examined over the first six segments, i.e., the hypothesis generation phase.

Problem representation

The objective of this analysis is to generate data about the subjects' changing representation of the problem, from which it will be possible to infer some characteristics about the processes they used for constructing the problem representation and about the solution strategies they used. The method of analysis developed for this study required the use of a general model of the knowledge necessary for generating an accurate solution to the problem, represented as a causal network and referred to here as the "reference" or "canonical" model (see Patel & Groen, 1986). Using the reference model as a template, the subject's current representation of the problem was generated after each sentence presented. The subset of the reference model that the subject used was highlighted to represent the knowledge that he used to construct his current interpretation of the problem, thus yielding problem representation after sentence 1, after sentence 2, etc. for each subject.

Reference model. The reference model (Figure 2) depicts, based on techniques developed in discourse analysis and artificial intelligence (see Frederiksen, 1986; Brachman & Levesque, 1985), the conceptual relations between the cues given in the case description, some of the underlying pathophysiological processes and the components of the correct diagnosis. It was developed from protocols and interviews with two expert endocrinologists and several other physicians as well as from standard textbooks (e.g., Isselbacher et al., 1980), and is taken to represent the knowledge required for generating an expert diagnosis.

The nodes represent cues (critical in rectangles; relevant in rounded rectangles), the components of the diagnosis (filled ovals), and some of the pathophysiology linking them (text in italics). The links in the reference model represent relations between the nodes; most are causal relations (CAU) with the exception of an occasional association relation (AND), category relation (CAT), or conditional dependency relation (COND). (These relations are defined in Frederiksen, 1975.) The nodes and links are arranged to reflect the chain of ramifications which tie the disease to its manifestations: the central determinants of the patient's condition (in the middle of the diagram) cause particular conditions which eventually lead to the observed signs,

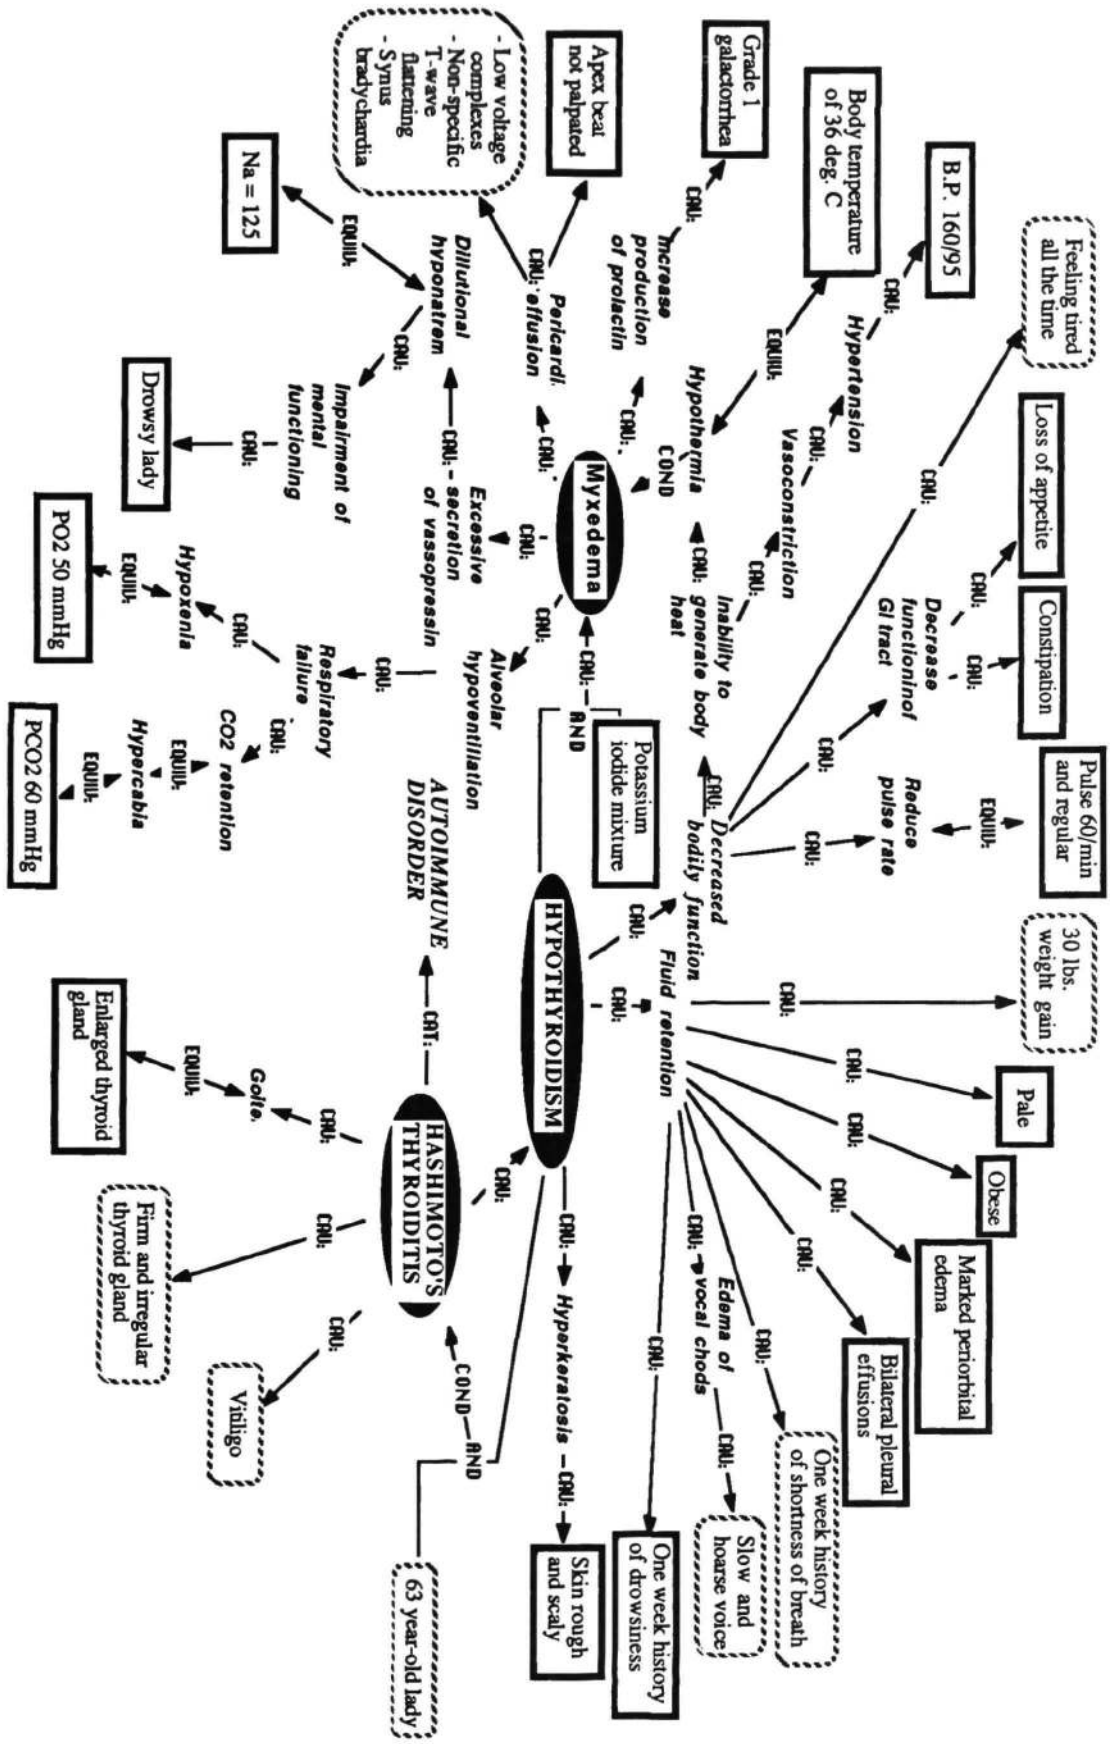


Figure 2. Reference model for the endocrinology problem.

- Critical Cues**
- Relevant Cues**
- Non-Relevant Cues**
- Accurate Diagnostic Subcomponents**
- CNLI: Causal relation**
- COND: Conditional relation**
- AND: Associative relation**
- EQUIL: Equivalence relation**
- CRT: Category relation**

symptoms or laboratory results (at the periphery of the diagram). The links also represent the rules necessary to generate an accurate solution to the diagnostic problem (Patel & Groen, 1986). The reference model is used here to compare and contrast the representations of the problem which subjects generate after the presentation of each text sentence.

The HDK subjects are expected to generate a representation that is better organized and structured than the LDK subjects. Thus far the results have indicated that the most important difference between the two groups is a better organization of the information which allows a more efficient pattern of hypothesis generation and confirmation by the HDK subjects. In addition, the representation generated by the HDK subjects should be more focussed than that of the LDK subjects, primarily due to differences in domain-specific knowledge between the two groups of subjects. The results of the detailed analysis of the sample protocols will precede the results of the overall analysis of the generation of the problem representations. A detailed analysis of two sample protocols will be described in the next section.

DETAILED ANALYSIS OF TWO SAMPLE PROTOCOLS

The mapping of the subjects' representations on the reference model will be discussed for Segment nos. 3, 4, and 7. These three segments are chosen because they correspond to the segments in which the HDK subjects generated the accurate diagnostic subcomponents. The mapping consists in doing an overlay of the subjects' representation onto the reference model (Figure 2). The different types of text segments are in boxes of different types and the diagnostic hypotheses generated by the subjects are surrounded by dark circles. The pathophysiological mechanisms used by the subjects are in italics. A more detailed description of the method was given in Joseph and Patel (1986a,b). This method was originally used by Patel & Groen (1986) to study solution strategies of experts and novices in diagnostic reasoning in medicine.

The mapping of the HDK subject is expected to be more closely related to the reference model. His representation will be more organized (contain more links) and more focused (fewer types of diagnostic hypotheses generated). The HDK subject is also expected to elaborate more on the findings about the

patient, i.e., generate more links from diagnosis to pathophysiological mechanisms and findings, than the LDK subject. The representation of the HDK subject is also expected to be more focused than that of the LDK subjects, due to the two-stage pattern of hypothesis generation of the HDK subjects. In addition, in the second stage the HDK subjects generated very few additional or new diagnostic hypotheses.

SEGMENT #3

The mapping of the HDK and LDK subject's representation for Segment 3 onto the reference model is illustrated in Figure 3. The findings presented in that segment activated the generation of the general diagnostic subcomponent (Hypothyroidism). After that the subject elaborated on the general condition of the patient's state, generating causal links from pathophysiological processes to other text cues (e.g., respiratory distress causing the one week history of shortness of breath and the fact that hypothyroidism caused some fluid retention which in turn caused the patient's weight gain).

The organizational pattern of the LDK's representation is also consistent with the previous segments. In this case, however, it is the evaluation of individual findings which leads to the generation of more diagnostic hypotheses (psychiatric problem, depression, and decreased glandular function). Nonetheless, the LDK subject generated a general description of the first diagnostic subcomponent (endocrine disorder or one of decreased glandular function).

The HDK subjects seem to organize the information in a way that limits the generation of multiple hypotheses, whereas the LDK subject is not able to use such efficient patterns of organization of information to limit the generation of new hypotheses. Both subjects generate the first subcomponent of the diagnosis.

SEGMENT #4

The mapping of the HDK and LDK subject's representation for this segment onto the reference model is illustrated in Figure 4. At the presentation of this segment the HDK subject generated the second subcomponent of the accurate diagnosis (myxedema), and used the same links as those illustrated in the reference model. The LDK subject also generated the second subcomponent of the diagnosis. However, as was the case in the preceding segment (#3) the subject also provided a more general explanation for myxedema: "advanced hypothyroidism".

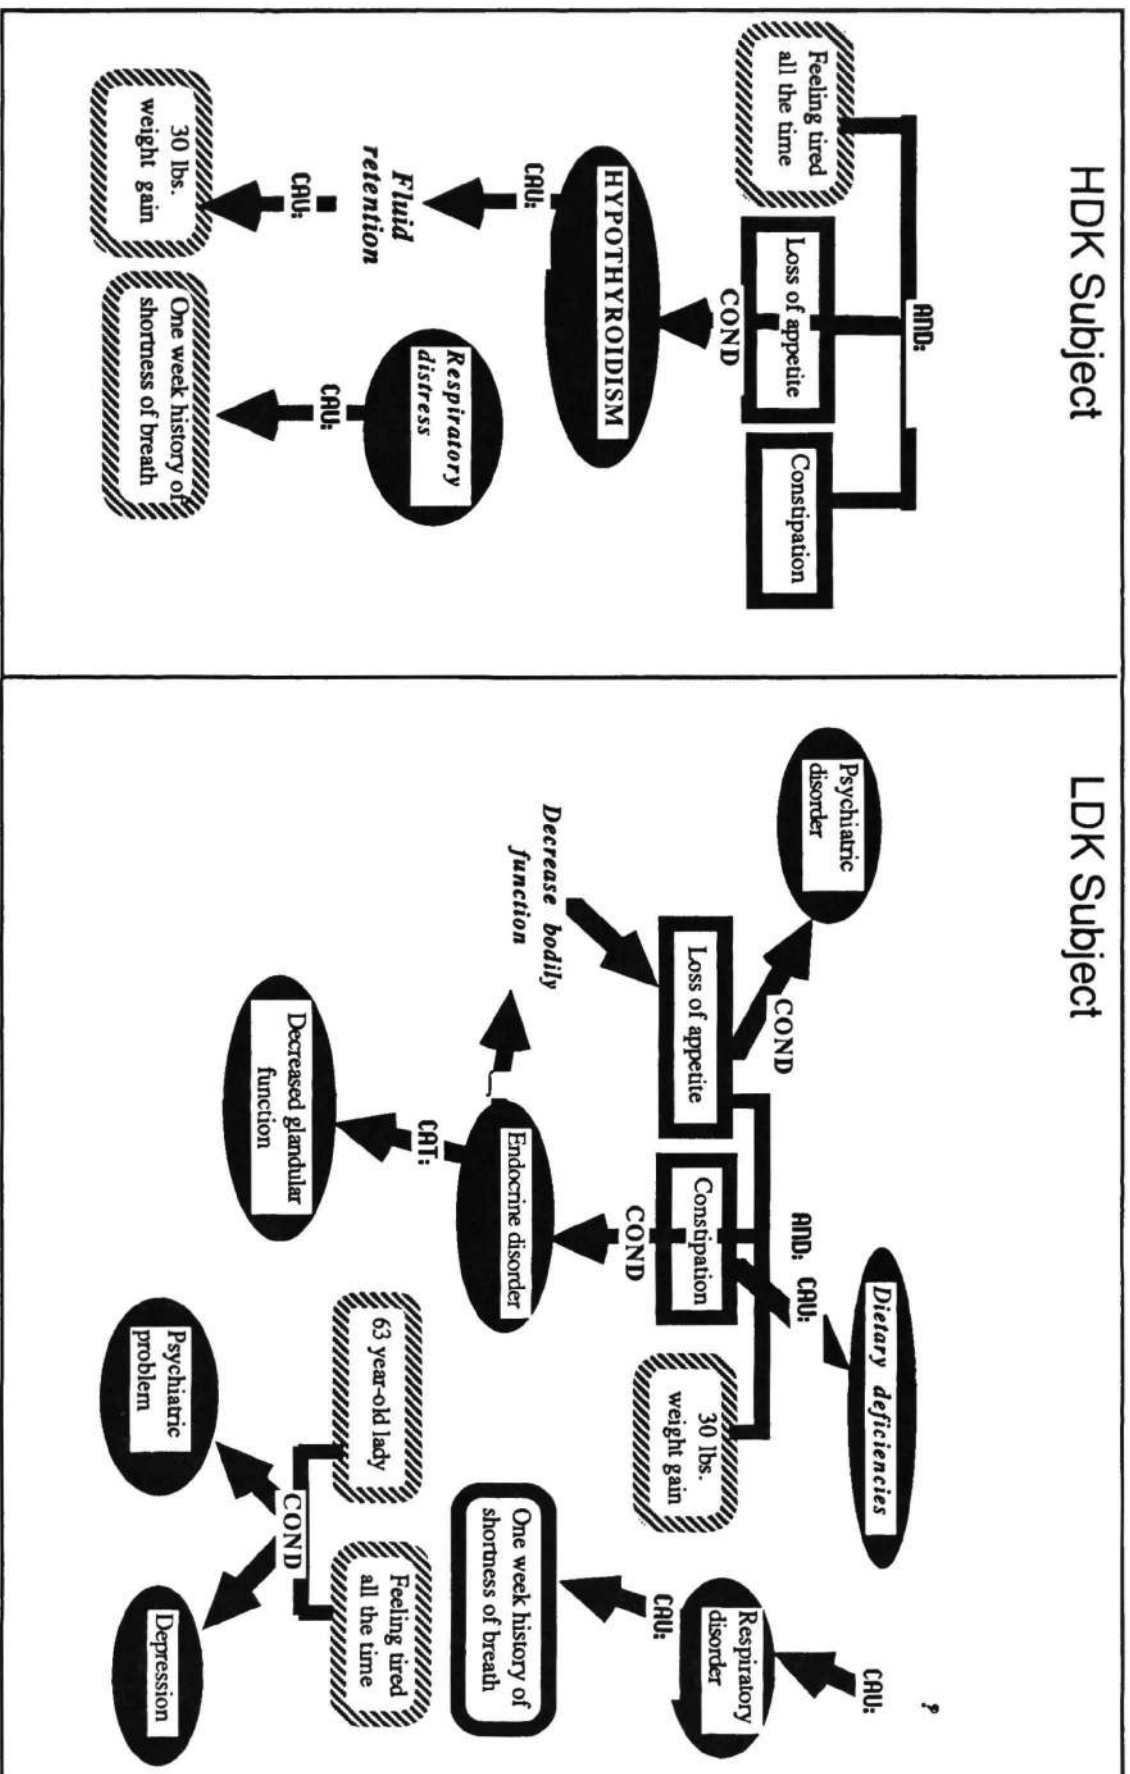


Figure 3. Parts of the problem representations generated by HDK and LDK subjects in response to sentence 3.

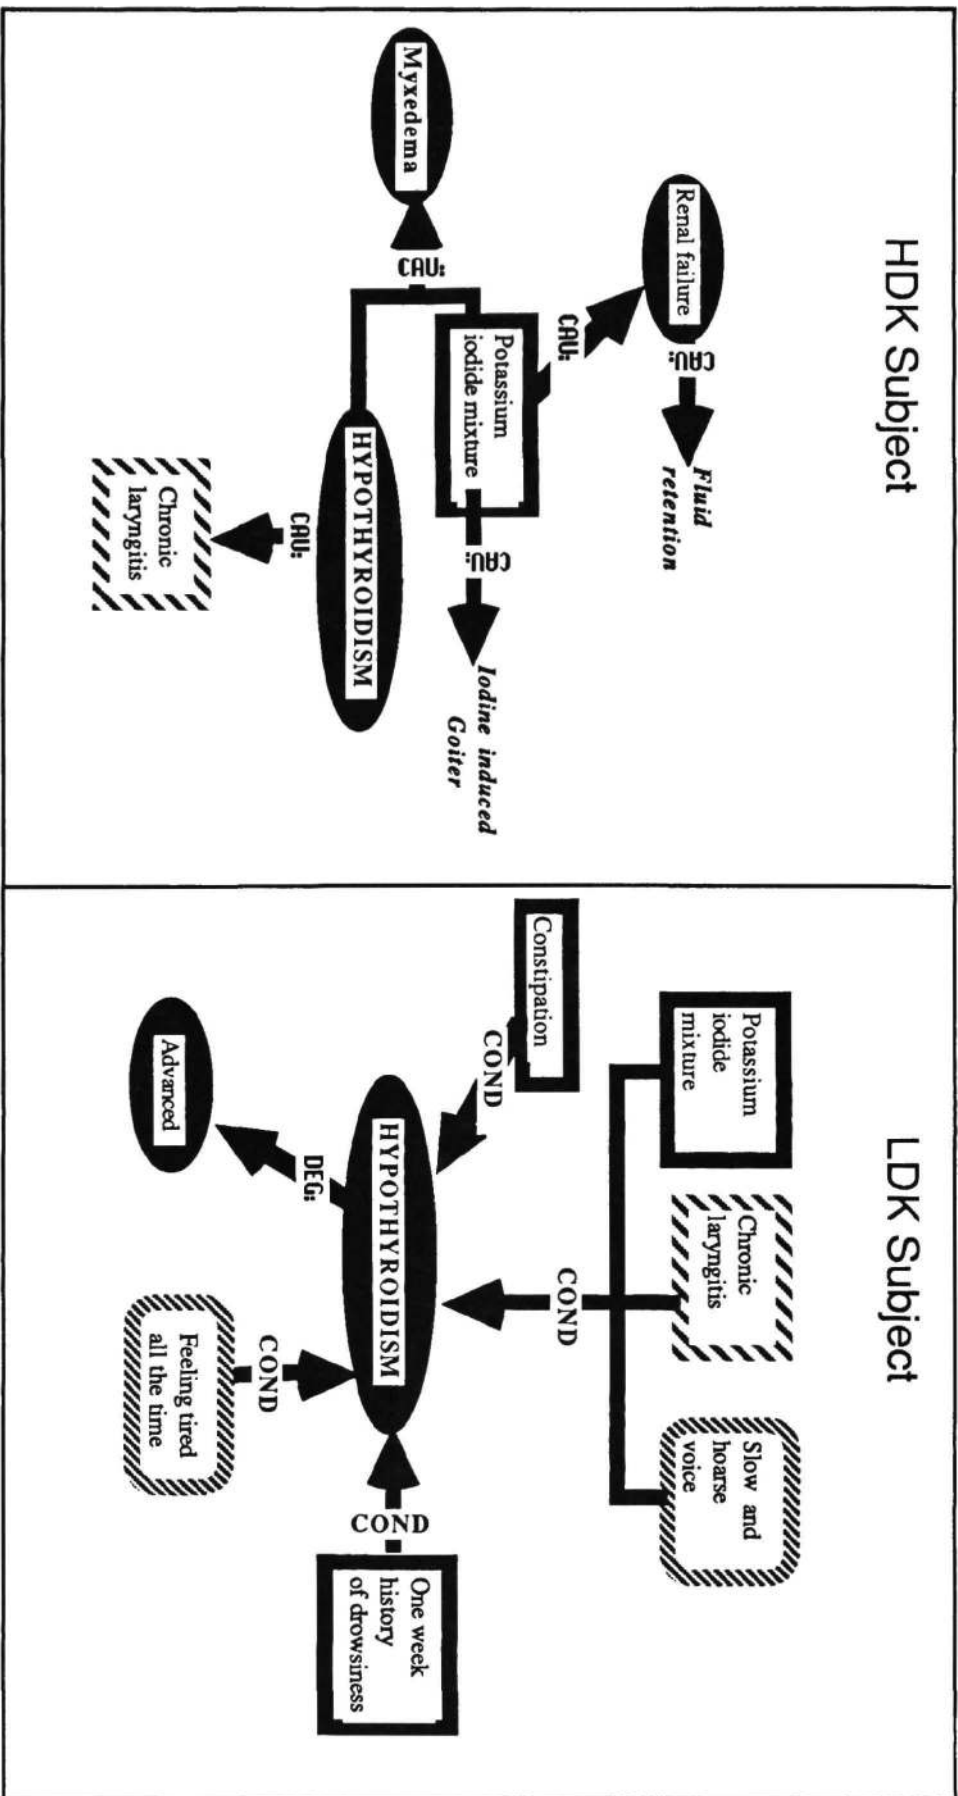


Figure 4. Parts of the problem representations generated by HDK and LDK subjects in response to sentence 4.

SEGMENT #7

The mapping of the HDK and LDK subjects' representation for this segment onto the reference model is illustrated in Figure 5. The HDK subject generated the specific subcomponent of the accurate diagnosis from the finding presented in that segment.

The LDK subject also recognized the autoimmune process. However, the LDK subjects had the wrong category of autoimmune process (pernicious anemia). This is interpreted as an indication of the difference in domain-specific knowledge between the HDK and LDK subject.

Overall, the results from the analysis of the problem representation generated by the HDK and the LDK subjects suggest that while there is overlap in the generation of the first two diagnostic subcomponents, on the one hand, the diagnoses of the HDK subject are more specific and precise, and on the other, the LDK subject provides a general description of the disease. Secondly, the pattern of generation of the two subcomponents is different for both subjects. The difference is mainly in the organization of the information that leads to the generation of the diagnostic subcomponents. For example, the HDK subject generates the second diagnostic subcomponent (myxedema) by associating the finding that the patient was hypothyroid and was given a potassium iodide mixture. On the other hand, the LDK subject generated a general description of the diagnosis by associating the finding of potassium iodide with the slow and hoarse voice. The HDK subject's elaboration seems to be more complete and focused than the one of the LDK subject. The representation of the HDK subject maps more directly onto the reference model than that of the LDK subject. Finally, the LDK subject re-evaluates the explanation of findings after the presentation of the last two segments, while the HDK subject does very little re-evaluation.

The pattern used by the two subjects in generating a representation seems consistent with the suggestion of Gick & Holyoak, (1983) that the HDK subject's construction of his representation is schema-driven. During the construction of the problem representation certain findings (text cues) seem to activate a disease schema from which the solver extracts the given and the goal information and connects it to existing knowledge so that an integrated representation can be formed. This explains the organizational difference (multitude of links generated from individual cues as well as the number of disease hypotheses generated from

the selected cues) between the HDK and the LDK subjects during their on-line interpretation. Hence, for the expert, solving a problem begins with the identification of the right solution schema, and then the exact solution procedure involves the instantiation of the relevant pieces of information as specified in the schema. While novices also solve problems in a schema-driven way, their schema of problem types is more incomplete and incoherent than those possessed by the experts.

The results of this study raise some questions regarding the assumption of the hypothetico-deductive model as a characteristic of expert reasoning in medicine as formulated by some researchers (cf. Elstein et al, 1978). The results presented above seem to make most sense in the context of a two-step model of clinical reasoning. The first step is that of generating a coherent problem representation, and the second is that of evaluating its goodness-of-fit. (confirmation and ruling out of diagnostic hypotheses)

The processes of organizing the incoming case information are clearly the most important to the study of diagnostic reasoning (see Patel, & Frederiksen, 1984). The physician receives a list of unrelated cues (at least none of the relations are made explicit in the case description) and has to impose on that list an organization that maximizes its coherence, normally by constructing links (or adding relations) between the cues and one or more diagnoses. The resulting structure of cue-nodes and hypothesis-links is what is usually referred to as the physician's problem representation. The analysis of the process of generation of the hypotheses produced in building these representations suggests that experts divide their time assessing a case description between a first phase of active generation of hypotheses and a second in which this generation activity is absent. It seems reasonable to assume that the first phase corresponds to the organizational process of representing the problem, and the second phase corresponds to a different process, presumably one of evaluating the representation. The fact that the time course of hypothesis generation (time in which two of the accurate diagnostic subcomponents) was very similar for the two groups in the first, organizational, phase suggests a limited role of domain-specific knowledge - instead it may be determined by general problem solving skill.

The second hypothesized step is that of evaluating the problem representation(s)

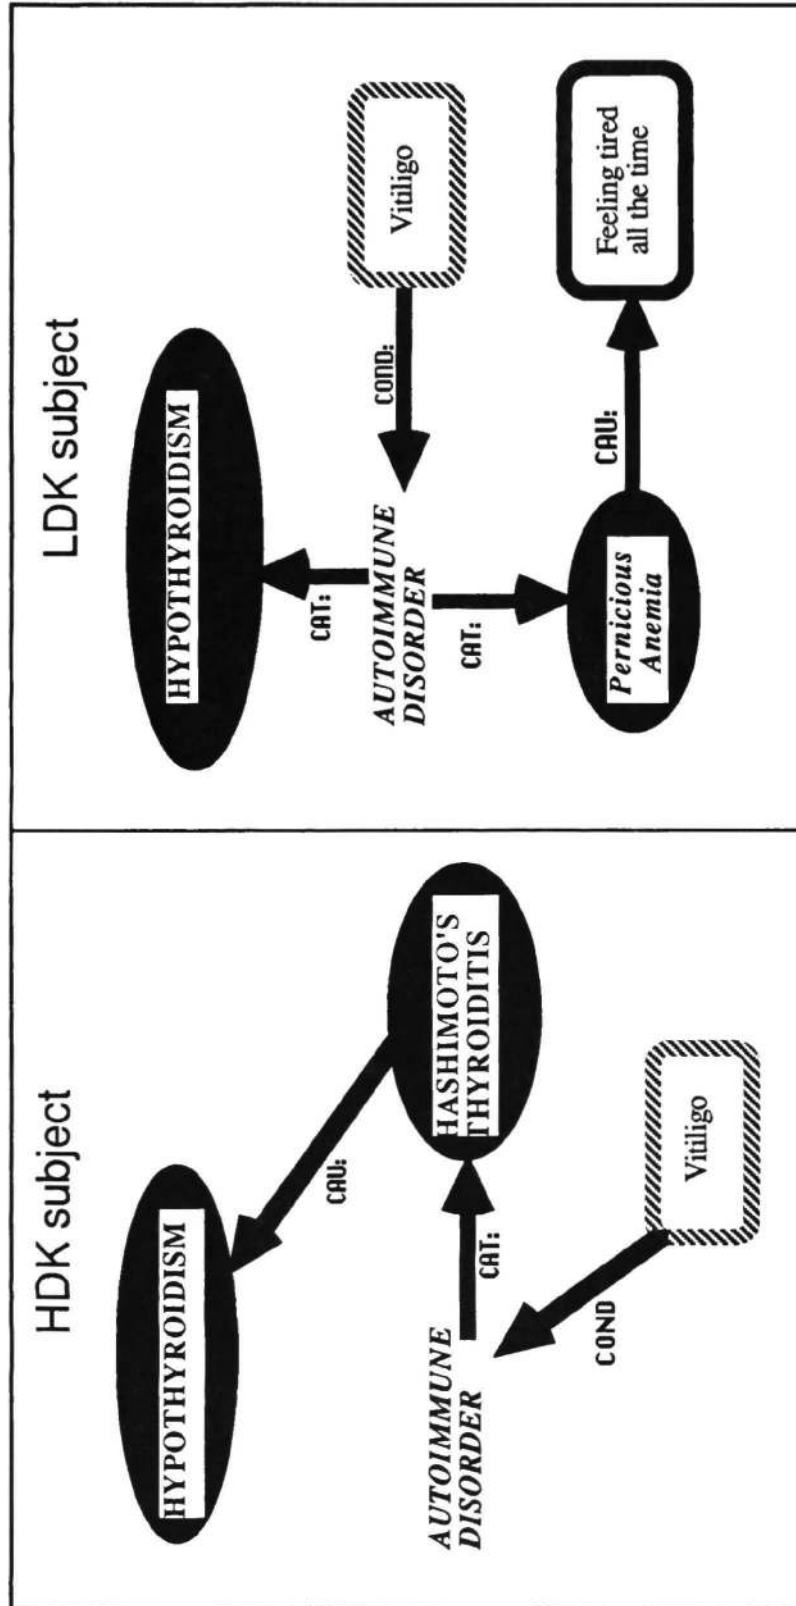


Figure 5. Parts of the problem representations generated by HDK and LDK subjects in response to sentence 7.

produced. Presumably, problem representations would have to be evaluated with respect to coverage of the cues presented, internal coherence, and relative certainty with which one hypothesis can be proposed over another. Evidence for a two-step model comes from the clear change in hypothesis-generating activity after sentence seven (for the HDK subjects), and from the fact that the LDK subjects continued to generate new hypotheses even after they had produced all of the components of the correct diagnosis. This suggests a more important role of domain-specific knowledge in evaluating hypotheses than in generating them.

This study also has implications for researchers in the field of artificial intelligence, especially for the development of expert systems (Buchanan & Shortliffe, 1984) and intelligent tutoring systems (ITS) (Anderson, Boyle & Yost, 1985; Clancey, & Letsinger, 1984; Shute & Bonnar, 1985). A central problem with the development of these systems is generating models of expert knowledge. The traditional method for developing expert models has been based on the use of protocol analysis methods to extract information from interviews with experts in a particular area. While the use of this method has not been unsuccessful, a number of researchers have become concerned about the validity and efficiency of the intuitive, non-formal methods used in protocol analysis and the limited information it provides about the expert's use of his knowledge (Kassirer et al., 1978). The use of an on-line task-environment with the combination of protocol and discourse analysis techniques seems to be a first step towards the solving this complex problem by complementing qualitative data to supplement

the quantitative data normally used. This study has demonstrated the richness of the protocols thus produced.

An obvious extension of this research involves the use of an interactive on-line task environment in which subjects can control the order and content of the information that is presented by requesting what they need at a particular point in solving a problem. The main difference between such a task environment and the one used in this study is that the information here was presented in a fixed order not under the subject's control. A more flexible environment will allow subjects to explore different hypotheses freely, ask questions at any time, and backtrack to review information presented earlier. The protocols and time data thus produced, when analyzed with more refined versions of the methods used here, promise to yield much more detailed information about the intricate processes of problem solving in technical-scientific domains.

ACKNOWLEDGMENTS

We would like to thank Drs. Y. Patel, S. Magder, and D. St. Vil for their assistance with the medical aspects of the study. We particularly want to acknowledge the support of the physicians who generously volunteered their time. We also wish to thank M. Dillinger, D. Blitz, L. Coughlin and A. Cruess for their help. This work was supported in part by grants from The Josiah Macy Jr. Foundation (No. B852002) and the Natural Sciences and Engineering Council of Canada (A2598) to Vimla Patel.

REFERENCES

- Anderson, J. R., Boyle, C. F., & Yost, G. (1985). The geometry tutor. In A. Joshi (Ed.), Proceedings of the Ninth International Joint Conference on Artificial Intelligence, vol. 1 (pp. 1-7). Los Altos, CA: Morgan Kaufman.
- Barrows, H. S., Feltovich, 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. Grants ODH-PR-273 and ODH-DM-226 Hamilton Ontario, Canada, MacMaster University.
- Barrows, H. S., & Tamblyn, R. M. (1980). Problem-based learning: An application to medical education. New York: Springer Verlag.
- Brachman, J. R., & Levesque, J. H. (Eds.) (1985). Readings in knowledge representation. Los Altos, CA: Morgan Kauffman Publishers.
- Buchanan, B. G. & Shortliffe, E. H. (1984). Rule-Based expert systems: The MYCIN experiments of the Stanford Heuristic Programming Project. Reading, MA: Addison-Wesley.
- Chase, W. G., & Simon, H. A. (1973). Perception in chess. Cognitive Psychology, 1, 55-81.

- Clancey, W. J. & Letsinger, R. (1984). NEOMYCIN: Reconfiguring a rule-based expert system for application to teaching. In W. J. Clancey & W. H. Shortliffe (Eds.), Readings in medical artificial intelligence: The first decade (pp. 361-381). Reading, MA: Addison Wesley.
- Dillinger, M. (1984). Segmentation and Clause Analysis. Technical Report No. 3, McGill University Discourse Processing Laboratory, Montreal, Quebec.
- Ericsson, K. A., & Simon, H. A. (1984). Protocol analysis: Verbal reports as data. Cambridge, MA: MIT press.
- Elstein, A. S., Shulman, L. S., & Sprafka, S. A. (1978). Medical problem solving: An analysis of clinical reasoning. Cambridge, MA: Harvard University Press.
- Feltovich, P. J. & Barrows, H. S. (1984). Issues in generality in medical problem solving. In H. G. DeVolder & H. Schmidt (Eds.). Tutorials in problem based learning. Assen, Holland: Van Gorcum.
- Frederiksen, C. H. (1975). Representing logical and semantic structure of knowledge acquired from discourse. Cognitive Psychology, 7, 371-458.
- Frederiksen, C. H. (1986). Cognitive models and discourse analysis. In C. R. Cooper & S. Greenbaum (Eds.), Written communication annual, vol. 1: Linguistic approaches to the study of written discourse. Beverly Hills, CA: Sage.
- Gick, M. L., & Holyoak, K. J. (1983). Schema induction and analogical transfer. Cognitive Psychology, 15 (1), 1-38.
- Greeno, J. G., & Simon, H. A. (1985). Problem solving and reasoning. In R. C. Atkinson, R. J. Herrnstein, G. Lindzey, and R. D. Luce (Eds.), Stevens' handbook of experimental psychology (Revised Edition), New York: Wiley.
- Groen, G. J., & Patel, V. L. (1985). Medical problem solving and cognitive psychology: Some questionable assumptions. Medical Education 19, 95-100.
- Halliday, M. (1967). Notes on transitivity and theme in English, part 2. Journal of Linguistics, 3, 177-274.
- Isselbacher, K. J., Adams, R. D., Brunwald, E., & Wilson, J. D. (Eds.). (1980). Harrison's Principles of Internal Medicine, Ninth Edition. New York: McGraw-Hill.
- Johnson, P. E. (1983). What kind of expert should a system be? The Journal of Medicine and Philosophy, 8, 77-97.
- Joseph, G.-M. (1987). The Time Course of diagnostic information processing: Levels of expertise and problem representation. Unpublished master's thesis, McGill University, Montreal Canada.
- Joseph, G.-M. & Patel, V. L. (1986a). Specificity of expertise in clinical reasoning. Proceedings of the Eight Annual Conference of the Cognitive Science Society, Hillsdale, NJ: Lawrence Earlbaum, 331-343.
- Joseph, G.-M., & Patel, V. L. (1986b). Online analysis of expertise in medical problem solving. Technical Report # 11, Center for Medical Education, Faculty of Medicine, McGill University.
- Kassirer, J. P., Kuipers, B. J., & Gorry, G. A. (1982). Toward a theory of clinical expertise. The American Journal of Medicine 73, 251-259.
- Larkin, J., McDermott, J., Simon, D. P., & Simon, H. A. (1980). Expert and novice performance in solving physics problems. Science, 208, 1335-1342.
- Newell, A., & Simon, H. A. (1972) Human problem solving. Englewood Cliffs: NJ, Prentice Hall.
- 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.
- Pople, H. E. (1982). Heuristic methods for imposing structure on ill-structured problems: The structuring of medical diagnostics. In P. Szolovits (Ed.), Artificial Intelligence in Medicine (pp. 119-190). Boulder, CO: Westview Press
- Schwartz, S., & Griffin, T. (1986). Medical thinking: The psychology of medical judgement and decision making. Springer-Verlag.

- Shortliffe, E. H., Buchanan, B. G., & Feigenbaum, E. A. (1979). Knowledge engineering for medical decision making: A review of computer-based clinical decision aids. Proceedings of IEEE, 67, 1207-1224.
- Shute, V. & Bonar, J. (1986). Intelligent tutoring systems for scientific inquiry skills. In Proceedings of the Eighth Annual Conference of the Cognitive Science Society (pp. 353-370). Hillsdale, NJ: Lawrence Erlbaum.
- Simon, D. P., & Simon, H. P. (1978). Individual differences in solving physics problems. In R. Siegler (Ed.), Children's thinking: What develops? Hillsdale, N.J.: Lawrence Erlbaum.
- Winograd, T. (1983). Language as a Cognitive Process, vol. 1: Syntax. Reading, MA: Addison Wesley.
- Wortman, P. M. (1972). Medical diagnosis: An information-processing approach. Computers and Biomedical research, 5, 315-328.