DEEP VERSUS COMPILED KNOWLEDGE APPROACHES TO DIAGNOSTIC PROBLEM-SOLVING

B. Chandrasekaran and Sanjay Mittal
Dept. of Computer and Information Science
The Ohio State University
Columbus, Ohio 43210

ABSTRACT

In this paper we argue that given a body of underlying knowledge that is relevant to diagnostic reasoning in a medical domain, it is possible to create a diagnostic structure which has all the relevant aspects of the underlying knowledge "compiled" into it in such a way that all the diagnostic problems in its scope can be solved efficiently, without, generally speaking, any need to access the underlying structures. We indicate what such a diagnostic structure might look like by reference to our medical diagnostic system MDX. We also analyze the role of these knowledge structures in providing explanations of diagnostic reasoning.

I INTRODUCTION

Recently Hart [1] and Michie [2] have written about the "depth" at which knowledge is represented and used in problem-solving by expert systems. Hart makes a distinction between "deep" and "surface" systems, while Michie characterizes a similar distinction by reference to "high road" vs "low road" approaches. The underlying idea is that surface systems are at best a data base of pattern-decision pairs, with perhaps a simple control structure to navigate through the data base. There is less agreement on exactly what characterizes deep systems, but it is suggested that deep systems will solve problems of significantly greater complexity than surface systems can. This distinction appears to capture a fairly widespread feeling about the inadequacy of a variety of first generation expert systems.

In the area of medical diagnosis -- which will be the exclusive concern of this paper, even though the spirit of what we say may be applicable to other tasks as well -- the straightforward approach of building a data base of patterns relating data and diagnostic states is not feasible given the large number of patterns that would be needed in any realistic medical domain. There is also the pragmatic problem of coming up with a complete set of such patterns in the first place. The next best approach, namely, devising some problem-solving mechanism which operates on a data base of partial patterns (i.e., patterns relating only a small set of data to diagnostic states; also called situation-action rules) has been tried with moderate success [3, 4]. Hart [1] and Patil [5] have raised many valid concerns about these "compiled-knowledge" approaches.

The major intuition behind the feeling that expert systems should have deep models is the observation that often even human experts resort to "first principles" when confronted with an especially knotty problem. Also there is the empirical observation that a human expert who cannot explain the basis of his reasoning by appropriate reference to the deeper principles of his field will have credibility problems, especially in life and death areas such as medicine. Added to this is the often unspoken assumption that the speed and efficiency with which an expert solves problems can be accounted for by hypothesizing that the physician uses a data base of commonly occurring patterns similar to that described earlier for quick problem solving in most cases. In this view this data base is no longer adequate in hard cases, and invocation of deeper structures is called for.

The above intuitions have resulted in calls for the representation and manipulation of deeper knowledge structures in expert system design. There is, however, no general agreement on the form and content of these deeper structures -- Hart suggests that they should model causality, while Michie, following Rouse [6], proposes that they should represent knowledge of the form "situation x action -> situation". The work of Patil [5] and Pople [7] is based on the idea that the appropriate form for the representation of this kind of knowledge is a causal net. There is an associated issue of the nature of problem solving mechanisms which can operate on these knowledge structures and produce solutions to problems. The investigations of Patil and Pople referenced earlier attempt to answer this question.

The thesis of this paper is as follows. Between the extremes of a data base of patterns on one hand and representations of deep knowledge (in whatever form) on the other, there exists a knowledge and problem solving structure, which (1) has all the relevant deep knowledge "compiled" into it in such a way that it can handle all the diagnostic problems that the deep knowledge is supposed to handle if it is explicitly represented and reasoned with ; and (2) will solve the diagnostic problems more efficiently; but (3) it cannot solve other types of problems-- i.e., problems which are not diagnostic in nature -- that the deep knowledge structure potentially could handle.

For the past several years, we have been developing a medical diagnosis system called MDX [8] [9]. This system embodies the above thesis. The knowledge structure of MDX has compiled in it all the relevant aspects of deep knowledge available to our human experts. Before we argue for the adequacy of such a compiled structure, it will be useful to give a brief characterization of diagnostic problems and how MDX solves them.

II MDX AND THE DIAGNOSTIC TASK

By the term "diagnostic task," we mean something very specific: the identification of a case description with a node in a pre-determined diagnostic hierarchy. For the purpose of current discussion let us assume that all the data that can be obtained are already there, i.e., the additional problem of launching exploratory procedures such as ordering new tests etc., does not exist. The following brief account is a summary of the more detailed account given in [9] of diagnostic problem solving.

Let us imagine that corresponding to each node of the classifiation hierarchy alluded to earlier we identify a "concept." More general classificatory concepts are higher in the structure, while more particular ones are lower in the hierarchy. The total diagnostic knowledge is then distributed through the conceptual nodes of the hierarchy in a specific manner to be discussed shortly. In the medical domain, a fragment of such a hierarchy might be:

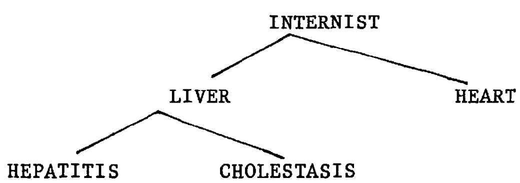

Fig. 1 Example Diagnostic Structure

The problem-solving for this task will be performed top down, i.e., the top-most concept will first get control of the case, then control will pass to an appropriate successor concept, and so on. In this case, INTERNIST first establishes that there is in fact a disease, then LIVER establishes that the case at hand involves some liver disease, while say HEART etc. reject the case as not being in their domain. After this, CHOLESTASIS may establish itself and so on.

Each of the concepts in the classification hierarchy has "how-to" knowledge in it in the form of a collection of diagnostic rules. These rules are of the form: <symptoms> ---> <concept in hierarchy>, e.g., "If high SGOT, add n units of evidence in favor of cholestasis." Because of the fact that when a concept rules itself out from relevance to a case, all its successors also get ruled out, large portions of the diagnostic knowledge structure never get exercised. On the other hand, when a concept is properly invoked, a small, highly relevant set of rules comes into play.

The problem-solving regime that is implicit in the structure can be characterized as an "establish-refine" type. That is, each concept first tries to establish or refine itself. If it succeeds in establishing itself, then the refinement process consists of checking which of its successors can establish themselves. Each concept has several clusters of rules: confirmatory rules, exclusionary rules, and perhaps some recommendation rules. The evidence for confirmation and exclusion can be suitably weighted and combined to arrive at a conclusion to establish, reject, or suspend it. The last mentioned situation may arise if there is insufficient data to make a decision. Recommendation rules are further optimization devices to reduce the work of the subconcepts. Further discussion of this latter type of rules is beyond the scope of this paper.

The concepts in the hierarchy are clearly not a static collection of knowledge. They are active in problem-solving. They also have knowledge only about establishing and rejecting the relevance of that conceptual entity. Thus, they may be termed "specialists", in particular, "diagnostic specialists."

The above account of diagnostic problem-solving is quite incomplete. We have not indicated how multiple diseases can be handled within the framework above, in particular when a patient has a disease secondary to another disease. A more powerful model is outlined in [9].

III EXAMPLE ANALYSIS

Reasoning with Deep Knowledge

Before we get into the example proper, the following background will be useful. (In this entire discussion, the interests of clarity have overridden the interests of completeness and medical accuracy. Much of the discussion is simplified to make the essential technical points.) The MDX system diagnoses in a syndrome called Cholestasis, which is a condition caused when the secretion of bile from the liver or its flow to the duodenum is blocked. Such a blockage can be caused by any of a number of causes. MDX attempts to pinpoint the cause in a given case. A subset of causes can be grouped under the category "Extrahepatic Obstruction," i.e., blockage due to an obstruction of bile flow outside the liver. In this example, we will assume that the physician has established that a cholestatic condition is present, and he is examining whether the cause is extra-hepatic, in particular which of a number of possible extra-hepatic causes may be the underlying reason. Bile flows from the liver into the duodenum via the bile duct. The following is a sequence of explicit reasoning steps by a (hypothetical) physician.

1. The bile duct is a flexible and somewhat elastic tube. If such a tube has a blockage at some point and if there is fluid pressure building up, then it will be dilated on the "upstream" side of the blockage. Thus if there is an extra-hepatic obstruction, the biliary tree inside the liver and a portion outside should be dilated. This should be visible in various imaging procedures as a specific visual pattern, call it <pattern 1>. Thus, EHO causes <pattern 1> in X-ray of the region.

2. Given a flexible duct, obstruction can be caused because there is a physical object in the duct, a contraction such as a stricture is present, some object outside the duct is pressing on the duct, or the internal diameter of the duct is reduced for some reason. Physical objects that I can think of in this context are biliary stone and a tumor in the duct. Looking at the anatomy of the region, nearby organs are gall bladder and pancreas. Cancers in these organs can press on the ducts and cause obstruction.

3. Biliary stones show up as a characteristic pattern (call it <pattern 2>) in a cholangiogram. Since the stones cause obstruction and an increase in the peristaltic action of the duct, they can cause acute colicky abdominal pain.

4. Cholangitis (inflammation of the bile duct) can cause swelling and reduce the duct diameter, and bile flow. A stricture can be caused by the surgical wound during a prior b.duct surgery not healing properly. ....

B. <u>Knowledge at Different Levels of Depth Used</u>

In para 1 above, very general knowledge about flexible pipes as well as knowledge about anatomy of the region and about imaging procedures is accessed, and from all these, a highly specific rule is compiled relating a diagnostic state to certain X-ray patterns. In para 2, again very general knowledge about flexible ducts is used in conjunction with an anatomical model of the region and other causal knowledge to generate a list of possible causes. In para 3 knowledge about imaging procedures, and about the physical properties of stones and ducts is used to <u>infer certain consequences of the diagnostic state being present</u>. Clearly, a variety of knowledge structures are used in the above reasoning fragment, many of them surely deserving the name "deep" or "underlying" knowledge. After all, what is deeper than very general knowledge about flexible ducts? Note also that not all these pieces of knowledge are at the same level : some had to do with flexible ducts, while others were highly domain-dependent pieces such as knowledge about cholangiograms. In this particular instance the physician reasoned about the relationship between colicky abdominal pain and stones from more basic pieces of knowledge about peristaltic ducts. A more experienced physician may simply use a piece of knowledge, "biliary stones —> colicky abdominal pain." Thus any such reasoning will be a mixture of such compiled pieces and more general items of knowledge.

C. <u>Diagnostic Structure</u>

Now we shall attempt to show that we can create a diagnostic structure as a result of the above reasoning. Once this structure is available, most of the steps in the above reasoning can be skipped.

Consider the hierarchy of specialists in Fig. 2 on the next page.

When a case is established to be cholestatic, control will pass to EHO in the above figure. If EHO is able to establish itself (i.e., if intrahepatic and a portion of the extrahepatic ducts are dilated — see establish rule in figure for EHO), it will call its successors (in parallel, for purposes of current discussion). Each successor will similarly establish or reject itself by using appropriate rules. Typically only one of them will be established at each level. The one(s) that are established refine the hypothesis by calling their succesors. The tip nodes that are established provide the most detailed diagnostic nodes for the case. (Again we caution that this portion of D is for illustration only; the reality is much more complicated.)

If one were to denote by U (for underlying knowledge) all the knowledge structures that were accessed during the physician's reasoning described above, and by D the knowledge in the above diagnostic structure (fig. 2), then the claim is that all the knowledge in U that plays a role in diagnosis is in D in a form that is directly usable. If D fails to solve a problem, a resort to U will not improve the situation. (An exception, which does not counter the basic argument, but may have some implementation consequences is discussed in point (5) of Sec. IV.)

D. <u>Relation Between D and U Structures</u>

In structure D of fig. 2, several diagnostic states are identified and organized hierarchically. (The justification for this particular hierarchy of diagnostic states as opposed to other alternatives -- e.g., one might have chosen to place biliary duct tumor and biliary stone as children of a node called physical obstruction, rather than group b. duct tumor with gall bladder cancer as we have done -- cannot be provided by refernce only to knowledge in U. Meta-criteria such as therapeutic implications or commonality of knowledge needed for establishing the states come into play. For example, many of the manifestations of cancer are the same whether it is b. duct cancer or GB cancer; hence the choice presented. While further discussion of this issue is beyond the scope of this paper, the point to be noted is that D is organized for efficiency in the diagnostic process, a consideration which may not be directly available from U.) The complex reasoning in para 1 based on anatomy and knowledge about ducts is simply compiled into the establishment rule in the EHO specialist. Para 2 resulted in candidates for further specializations of the hypothesis about the cause of cholestasis. Para 3 resulted in the procedure for establishing biliary stone as cause of cholestasis. The causal reasoning in para 4 resulted in a similar procedure for stricture.

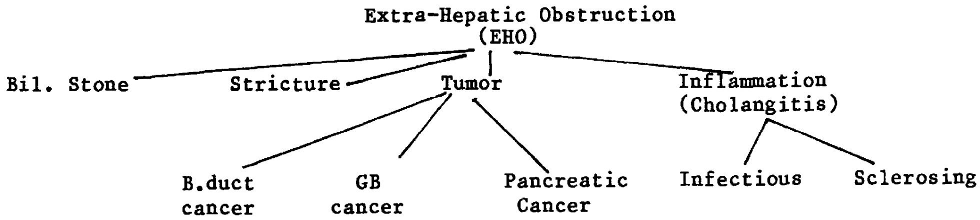

EH Obstruction
To establish: (from para 1)
Dilation of intrahepatic ducts and
portion of extrahepatic duct in X-ray

Bil. Stone
To establish : (from para 3)
<pattern 2> in cholangiogram;
colicky abdominal pain, etc

Stricture
To establish : (from para 4)
Evidence of prior biliary tree surgery;
<pattern 3> in X-ray

(Similar boxes for other nodes.)

Fig. 2 Diagnostic Structure for Extrahepatic Obstruction

IV <u>THE ADEQUACY OF D-STRUCTURES FOR DIAGNOSIS</u>

Let us briefly look at some of the ways in which D may fail to solve a case. (1) A concept is missing in D because a needed chunk of knowledge is missing in U. E.g., suppose in a particular case EHO is established, and evidence of inflammation is present. Suppose further that cholangitis is ruled out. D will be able to provide only an incomplete solution to the problem. It turns out that there is another inflammation possibility in that region that can also cause obstruction : pancreatitis. This piece of knowledge was missing in U. This resulted in D missing a concept, a sibling of Cholangitis. Thus the problem can be traced to an inadequacy of U. An automated diagnostic system endowed with the same U will gain no further ability to solve this case by referring to U, deeper though it is.

(2) Some establish/reject knowledge is missing in a specialist. Again such a deficiency can be traced to missing chunk of knowledge in U.

(3) D's problem solving strategy is not powerful enough. (As indicated in Sec. II the problem-solving strategy outlined there is inadequate, and a more powerful strategy is given in [9].) In this case a reference to U still will not be able to do the trick. The weakness of the PS strategy of D means that the system designer has not fully comprehended the <u>use</u> of knowledge for diagnostic problem solving. Thus access to U will not help, since the issue in question is precisely how to properly use the knowledge. A system which knows how to use the knowledge in U can equally be rewritten to embed that improved use of knowledge in the problem solving strategy of D. For this reason the arguments for compilation are not strongly dependent on the correctness of the MDX approach in detail.

(4) The knowledge is there in U, D's PS strategy is powerful enough, but the knowledge was improperly compiled in D. This in fact often happens. But again access to U during the problem solving by D will be useless, since the ability to use the knowledge in U effectively for diagnosis implies the ability to identify the proper use of that knowledge for compilation into D. In fact we feel the problem of automatic generation of D from U is a much harder research problem. One can view the research of Patil [5] and Pople [7] as attempts to understand the role of knowledge in U for diagnosis. If this type of research is successful, then we can provide an automatic mechanism for going from U to D.

(5) U has the relevant knowledge, but compiling all the relevant ways in which that knowledge can come into play in diagnosis will result in a combinatorial problem in representing them in D. A specific example will motivate this situation. In a CPC case invloving EHO that MDX was presented with, the system established "inflammation" (see fig. 2), but could not establish any of its successors. It turned our that in an unusual sequence of events a pice of fecal matter from the intestine had entered the bile duct through the Sphincter of Oddi, and was physically obstructing bile flow. (This possibility escaped not only MDX, but all of the physicians who analyzed the case in the CPC. It was discovered during surgery.) One could suppose that in theory all such possibilities could be compiled, but, even if it could, the number of possibilities in which some piece of matter from one part of the body could end up through some fistula say in the bile duct is quite large. One way out in this kind of situation is for D to simply have a procedure which calls the underlying knowledge structure U. But notice that all the possibilities that were not compiled in D because they would be too numerous will nevertheless have to be generated one by one at this point. Thus in principle it is simply a run-time compilation of the needed portion of D. In fact precisely because of the combinatorial nature of this situation, all the physicians, endowed with quite a powerful U-structure, failed to think of the real cause.

D may include a data base of patterns for further efficiency. Some of the major specialists such as EHO may contain a small set of patterns which can be used to dispose of some common possibilities without using the full PS capabilities of D. For example, assume for the sake of the argument that EHO due to b.stone accounts for 90% of all EHO cases, and that due to stricture for a further 5%. (This happens not to be true, thus we reemphasize that this is purely for making a point.) Now, the EHO specialist may, after establishing itself, access a small data base of patterns that may help to establish or dispose off these possibilities with a small effort. Thus the system can solve 95% of the cases with extreme efficiency. But the important point is that if this data base of patterns is exhausted before problem solution, the rest of D is still available for more thorough problem solving.

V PROVIDING EXPLANATIONS

The idea that deeper knowledge structures are needed for providing explanations has been proposed often. It is perhaps worth examining the concept of explanation itself. Referring to the structure D in fig. 2, suppose the diagnostic conclusion in a particular case was "EHO due to b.stone." Suppose we ask:

Q1. Why do you conclude that this patient has EHO due to b.stone?

The system at that point may refer to its procedure for establishing that hypothesis and say, "Dilation of intrahepatic and a portion of extrahepatic b.duct was observed in X-ray, thus I concluded EHO. Because further a characteristic pattern, <pattern 2> was observed in Cholangiogram, and because colicky abdominal pain was observed, I concluded b. stones." Note that the system is simply telling which of the rules was found to match for each specialist that was established. Some people may accept this explanation. At this level D can be quite satisfactory. Suppose a further question was asked,

Q2. "Why does a dilated IH b.duct indicate EHO?"

D as it stands cannot answer this question, since the rationale for the rule was left behind in U, at the time of compilation. But note that giving an explanation for this question does not need any additional problem solving. We can simply associate with each piece of compiled knowledge in D a text string that explains its rationale. This is possible because the answer to Q1 is patient-specific and thus requires problem solving with D, while the answer to Q2 is not patient-specific, and has already been generated in designing D.

Since in a human expert D is built by learning processes from U, inadequacies in D are often correlated with inadequacies in U. That is why we are suspicious of physicians who we feel cannot explain the basis of their conclusions by reference to deeper knowledge structures. Since U changes with time with new discoveries in medicine, most of us as patients would like to assume that our physicians are capable of translating the relevant changes in U to appropriate changes in D. An ability to explain by calling forth appropriate fragments of U is our assurance that the expertise that the physician is using to diagnose our illness is built on the basis of appropriate learning processes. Thus a system such as D when it answers Q1 is giving evidence of the correctness of its own problem solving, while its answer to Q2 is evidence of the correctness of those human experts who provided the underlying knowledge structures and the compilation process.

VI CONCLUDING REMARKS

Almost all of the discussion in the paper so far has dealt with the generic task of diagnosis. We have argued elsewhere [10] that there exist other generic tasks, each with a particular way of using knowledge for problem solving. Corresponding to each such task in a domain, one may construct a compiled problem solving structure. In [10] we have identified a few other generic tasks that we have encountered in the analysis of medical reasoning. While most of the arguments in this paper do not depend upon the correctness of the details of the MDX approach to diagnostic problem solving, they do implicitly depend upon the existence of diagnosis as a generic task. This is what makes it possible to compile the underlying knowledge in a structure such as D. In our view, identification of further such generic tasks and how these structures interact in complex problem solving situations is an important challenge for AI research.

The approach we are arguing for sharpens some of the questions surrounding how a novice learns to become an expert. We suggest that this issue is facilitated by having some idea of what are the target structures that are being learnt. In the case of diagnosis, we suggest that learning to become an expert diagnostician is the process of compiling a structure such as D. Thus research into what sorts of knowledge structures does U consist of and what sorts of processes are powerful enough to produce problem-solving structures such as D will significantly advance our ability to produce learning systems of considerable importance to expert system design.

ACKNOWLEDGMENT

We thank Jack Smith, M. D., for helpful comments on an earlier version of the paper. We acknowledge the contributions of Fernando Gomez to the approach outlined here.

The preparation of this paper was supported by NSF grant MCS-8103480. The computing resources of Rutgers University Laboratory for Computer Science Research were used in system development, and this was made possible by the Biotechnology Resources Division of NIH, under grant RR-00643.

REFERENCES

[1] Hart, P. E. "Direction for AI in the eighties" SIGART Newsletter No. 79, January 1982.

[2] Michie, D. "High-road and Low-road programs" AI Magazine 3:1 (1982) 21-22.

[3] Shortliffe, E. H. Computer based Medical Consultations:MYCIN New York: Elsevier/North-Holland Inc., 1976.

[4] Weiss, S. and C. A. Kulikowski "EXPERT: A system for developing consultation models" In Proc. IJCAI-79 Tokyo, Japan, August, 1979.

[5] Patil, R. S. "Causal representation of patient illness for electrolyte and acid-base diagnosis" Ph. D. Dissertation, TR-267, MIT Lab for Computer Science, Cambridge, Mass., October, 1981.

[6] Rouse, W. B. and R. M. Hunt "A fuzzy rule-based model of human problem solving in fault diagnosis tasks" Working Paper, Coordinated Science Lab, University of Illinois, Urbana, 1980.

[7] Pople, H. E. "Heuristic methods for imposing structure on ill-structured problems" In P. Szolovits (ed.) Artificial Intelligence in Medicine Westview Press, forthcoming (1982).

[8] Mittal, S., B. Chandrasekaran, and J. Smith "Overview of MDX - A system for medical diagnosis" In Proc. Third Symposium Computer Applications in Medical Care Washington D. C., October, 1979

[9] Gomez, F. and B. Chandrasekaran "Knowledge organization and distribution for medical diagnosis" IEEE Trans. Systems, Man, and Cybernetics SMC-11:1 (1981) 34-42.

[10] B. Chandrasekaran "Decomposition of domain knowledge into knowledge sources: The MDX approach" In Proc. Fourth National Conf. of Canadian Soc. for Computational Studies of the Intelligence (CSCSI/SCEIO) Saskatoon, Canada, May, 1982.