
Knowledge Organization and Distribution for Medical Diagnosis

Fernando Gomez and B. Chandrasekaran

During the mid-1970s, an AIM research group directed by Professor B. Chandrasekaran was initiated at Ohio State University. Fernando Gomez was a graduate student at the time and was involved in the group's work on MDX, a program for the diagnosis of liver disease. That system provided an experimental environment in which many of the ideas expressed in this chapter were developed.

In this paper, Gomez and Chandrasekaran adopt an analytical view for studying the nature of medical knowledge. Rather than saying "It's all a bunch of random heuristics," they try to formalize the rich structures that make efficient diagnosis possible. They center their representation around concepts, such as diseases and their causes, in the form of a hierarchical structure similar to a botanical or zoological classification. The key idea is that an expert diagnostician's knowledge is distributed through this hierarchy. Besides being of value for formalizing knowledge in an expert system, this perspective is of value for teaching. Specifically, a student needs to learn this refinement structure for focusing on and further specifying diagnostic hypotheses. The chapter also proposes a useful framework for viewing knowledge interaction in terms of communication via a blackboard model, a knowledge representation and control scheme that was first developed for speech understanding (Lesser et al., 1975). The actual system implemented by Chandrasekaran's group is much simpler, however.

It should be noted that Gomez and Chandrasekaran are trying to capture the compiled form of human knowledge and are not advocating that we

design expert systems in general by intermixing strategic and domain knowledge (cf. NEOMYCIN's separation of disease knowledge from domain-independent meta-rules, Chapter 15). Nor are they claiming that experts do not use general principles for ordering search and selecting alternatives (cf. Swartout's "domain principles," Chapter 16). Rather they are emphasizing that other constructs, in addition to rules, are needed to organize knowledge. Explicating the hierarchical structure of hypotheses and findings implicit in pure rule systems improves system organization for focused reasoning as well as ease of system building [cf. Aikins's "prototype hierarchy" (Aikins, 1980)].

The reader may be interested in pursuing a number of related AI topics, such as studies of epistemology and natural language understanding that are referenced in this chapter.

Concepts lead us to make investigations, are the expressions of our interests, and direct our interests.

Wittgenstein, *Philosophical Investigations*, prop. 560

13.1 Introduction

What are the criteria that should be used to organize the medical knowledge in an automated medical system? We start with the observation that diagnosticians, when they arrive at a diagnosis or diagnoses, have invoked some concepts. These can be diseases, causes of diseases, or other notions that are relevant to the diagnosis. We shall suggest that these concepts form a hierarchical structure similar to that of a botanical or zoological classification. The diagnostician's knowledge is distributed through this hierarchy. The concepts in the hierarchy provide the criteria to organize under them small pieces of knowledge represented in the form of production rules. Thus concepts may be viewed as clusters of production rules. They extend the capabilities of production rules to more complex problem-solving situations. The rules under each concept are further organized into three groups: exclusionary, confirmatory, and recommendation rules.

During the problem-solving process, the concepts can be considered to be *specialists*. They interact and communicate with each other by means of a blackboard, a notion borrowed from Erman and Lesser (1975). In that respect, the ideas presented in this paper can be considered as an extension of the notions of the HEARSAY-II speech understanding system (Carnegie-Mellon University, Computer Science Research Group, 1977) to the medical diagnosis task. Nevertheless, there is an important methodological difference. It is that concepts and not rules provide the principle of knowledge organization.

Section 13.2 contains critical notes on some aspects of knowledge representation. Section 13.3 describes the central features of our ideas on the organization of medical knowledge. Section 13.4 explains the different kinds of rules. Section 13.5 deals with the identity of the notions *concept* and *specialist*. Section 13.6 discusses distributive problem solving. Finally, the paper concludes by indicating some of the similarities and differences between this approach and other medical diagnosis systems.

13.2 On Knowledge Representation

In recent years, there has been much work in knowledge representation in artificial intelligence, but relatively little attention has been paid to how knowledge is used and organized. By *use of knowledge* we mean the invocation and instantiation of the right chunk of knowledge and the determination of the appropriate structure of the knowledge needed for the task being studied. Other authors, especially F. Hayes-Roth (1978), have expressed a similar view.

13.2.1 On the Representation and Use of Knowledge

The assumption that a separation can be established between knowledge representation and its use dates back to the distinction made by McCarthy and Hayes (1969) between epistemologically and heuristically adequate analyses. Underlying this distinction is the belief that the first does not involve the second, and *vice versa*. Most researchers in knowledge representation have, consciously or unconsciously, subscribed to this distinction, which is indirectly related to the Saussorian distinction of *la parole* and *la langue*, better known after Chomsky as the *performance-competence* distinction. The assumption underlying both distinctions is that it is appropriate to study the result of human thought, language, knowledge, etc., by “abstracting out” the homunculus that is using that thought. Both distinctions are influenced by the paradigm that modern logic brought to the study of linguistics and epistemological questions. While logic is no longer a dominant paradigm in AI, much research in knowledge representation nevertheless has concerned itself with the so-called epistemological adequacy of the representation, thus deepening the separation between knowledge representation and its use. In particular, while many of the current techniques of knowledge representation in AI arose as components of localized models

of human cognition, the emphasis has increasingly been on the *formalisms* in the models.

13.2.2 Content and Form in Knowledge Representation

The semantic network (Quillian, 1968) was proposed as a model of semantic memory. But since Quillian's original formulation, the formalistic aspect of it has gained a life of its own—so much so that much of the research in semantic networks scarcely differs from the logic formalism. Recently, some researchers have shown interest in the foundations of semantic networks (Woods, 1975; Brachman, 1979). Important distinctions have been made explicit, but no connection has been established between the proposed improvements to the representational formalisms and the use of the knowledge. It is unclear how the notational inventions will help in the understanding of the task being studied.

Since Minsky's (1975) important paper about frames, little progress has been achieved in extending his ideas, but formalisms (Goldstein and Roberts, 1979; Bobrow and Winograd, 1977) have been built on the outline proposed by Minsky. Minsky revived and began the task of giving computational meaning to an interesting theory of human cognition. The theory says that important aspects of vision, memory, problem solving, and comprehension can be explained as a process of *recognition*. In this process, the input is matched to an internal stereotyped structure called a *frame*, slots in the structure are filled, and others take default values. The notion of a default value was one of Minsky's most insightful ideas. Information not explicitly present in the input could be accounted for by reading the default values.

Frames have proved to be an excellent construct for dealing with extralinguistic knowledge in language. Other authors independently worked out a similar notion called a *script* (Schank and Abelson, 1977). In its representational aspect, frames are an extension of the property list notation. They look very much like a COBOL structure. But just as COBOL programs using structures are not exemplifications of the frame theory, neither are AI programs just because they happen to be written in a frame-type language. Otherwise, one would be confusing the form with the content of the theory. It is precisely the theory that needs to be extended. We know very little about the criteria that govern the recognition of frames (Charniak, 1978), the invocation of the appropriate frame, and the integration of frames in more inclusive structures. The available formalisms inspired by the frame theory, while they differ in the degree of their concern with such issues—FRL (Goldstein and Roberts, 1979) is meant as a programming language, whereas KRL's authors (Bobrow and Winograd, 1977) have shown a great concern with extending and giving depth to the

frame ideas—nevertheless do not provide answers to these questions. [See Lehnert and Wilks (1979) for a sympathetic critique of KRL.]

13.2.3 Production Rules and Organization of Knowledge

In expert knowledge domains, production rules have been extensively used. Despite their apparent simplicity, production rules grasp an important aspect of human cognition. But it seems to us that they have sometimes been used unilaterally to explain cognitive aspects for which other constructs are needed. They were used by Newell and Simon (1972) as techniques to model some aspects of human problem solving. Since this early and seminal work, production rules have been applied to almost every aspect of human cognition. [See Waterman (1978) for an excellent collection of papers about production systems.] Two of the most successful systems, DENDRAL (Buchanan et al., 1969) and MYCIN (Shortliffé, 1976) use production rules as the basic technique to represent knowledge.

Production rules are the right tool to represent some kind of *how-type* knowledge. Winograd (1975) refers to it as *secondary* knowledge. Other authors have used the term *judgmental* (Duda et al., 1978). Some aspects of the knowledge needed to repair a car, to diagnose a disease, etc., are of this type. In the medical diagnosis domain, there are many terms of which a doctor does not need to have a thorough understanding. The knowledge that a diagnostician has that tells him or her that certain lab findings are indications of a certain disease is of that kind. In domains where the required level of comprehension is deeper, for example, natural language understanding, the need for *what-type* knowledge has become apparent. Riesbeck's parser (1978) is a very good example of the integration of the two kinds of knowledge: production rules are used to build and to predict Schank's conceptual dependency structures. But the production rules are embedded in the conceptual structures.

The problem of organization of knowledge is of paramount importance in large knowledge base domains. The problem is not only one of efficiency, but one of focus and control. Things simply do not work if the knowledge is not properly organized [see Lenat and Harris (1978) for a discussion of these problems]. The solution generally offered by the advocates of production rules is a new rule called a *meta-rule* (Davis and Buchanan, 1977). Meta-rules organize the production rules according to some meta-knowledge criterion. While production rules are natural mechanisms to model an important aspect of human cognition, meta-rules seem to be an *ad hoc* solution to the problem. It is doubtful that they have any cognitive counterpart. We think that for an appropriate organization of knowledge, another construct, in addition to the rule, is needed. In the following pages we will show that even in domains in which the knowledge is basically *how-type*, concepts and not rules should provide the principles of organization.

13.3 The Role of Concepts in Medical Diagnosis

In a recent paper about AI work on natural language, Fodor (1978) has characterized it as suffering from “operationalism, empiricism. . . .” On the other hand, if an empiricist looked at the AI work, of course, beyond Winograd’s publications in 1971, he or she would consider it to be “irresponsible talk” about concepts. The existence of concepts and the need to take them seriously are almost granted, in particular by researchers in natural language understanding and in knowledge representation. The representation proposed for concepts has been a what-type of structure called a *frame* in FRL (Goldstein and Roberts, 1979), a *prototype* in KRL (Bobrow and Winograd, 1977), and a *concept* in KLONE (Brachman, 1979). In this paper, we will be speaking of concepts in a different sense, not as what-type structures but as labels that organize how-type knowledge represented in production rules. They can be considered as clusters of production rules.

13.3.1 Concepts and Organization of a Diagnostician’s Knowledge

Consider the following production rule: if bilirubin in urine and pruritus, then suggest cholestasis. A diagnostician has thousands of such rules. In our view they are associated with concepts such as “arteriosclerosis,” “hepatitis,” “cholestasis,” etc. These concepts themselves form a hierarchical structure similar to that of the botanical or zoological classification system. The most general concepts are placed at the top of the hierarchy and the most particular at the bottom (see Figure 13-1). Knowledge is distributed through this hierarchy. The structure serves the purpose of differentiating the knowledge, of assimilating new knowledge by inserting it in the appropriate place, or of retrieving the right piece of knowledge as a response to the appropriate query.

For diagnosticians, this hierarchy serves the function of organizing their troubleshooting knowledge. The concrete details for each disease are encoded in the production rules attached to the appropriate concepts. However, it is clear that medical doctors also have additional cognitive structures that organize their knowledge from other views: pathological, physiological, etc. The role of these additional structures during diagnosis then becomes a relevant issue. The cognitive structures corresponding to these other views do not need to be present for purposes of diagnosis, as long as knowledge from these structures relevant to diagnosis is compiled in the diagnostic structure. This can be done by appropriately structuring the relevant concepts and embedding the compiled production rules therein.

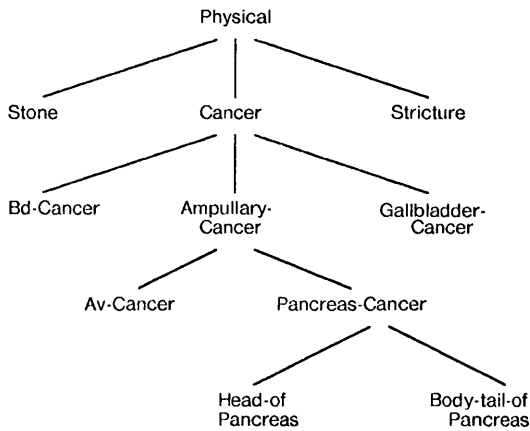
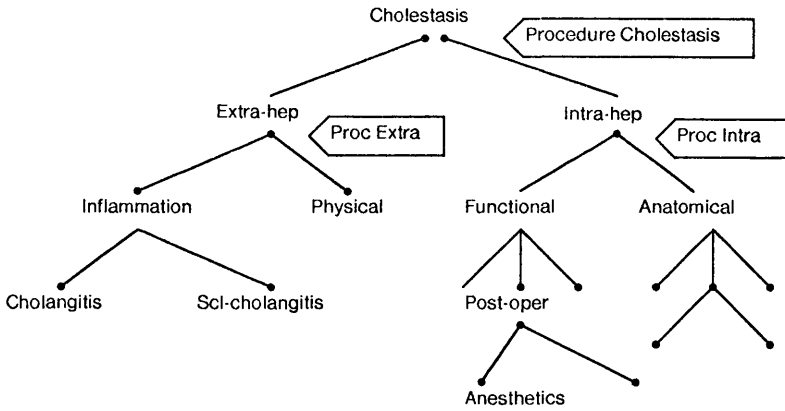


FIGURE 13-1 Conceptual structure of cholestasis.

13.3.2 Commonsense Knowledge

The role of commonsense knowledge structures is of equal interest. A distinction must be made between (a) the commonsense knowledge a physician needs in order to understand the data presented in a medical case and (b) that needed during the process of diagnosis. The patient data are entered in current AI programs for medical diagnosis in the form of a collection of atomistic facts, e.g., “high bilirubin, fever, jaundice.” In contrast, consider the following:

At the age of 19 years, one year prior to his appendectomy, he began to have occipital headaches, usually upon awakening in the morning and occurring once or twice weekly for a 10-year period. These headaches had not been severe enough to interfere with his activities . . . [taken from Harvey and Bordley (1972)].

Here we have a complex temporal interconnection of facts that cannot be decomposed into simple facts. It may be true that in most cases the atomistic collection of data may contain sufficient manifestations to make a correct diagnosis. However, for those cases in which comprehending the complex structure of data is essential to a solution, systems whose data input is atomistic will miss the solution. In order to uncover the needed structure for data input in these cases, it is necessary to make a semantic analysis of the commonsense notions of time and causality in this context. For simple instances of temporal and causal notions, temporal cases could be enough, for instance, structures like "< > while < >," "< > after < >," and "< > causing < >." But the kind of semantic notions and the mechanism needed to integrate these simple structures in more inclusive ones like those needed to understand *the course of the illness* need further study.

Let us consider (b), viz., the use of commonsense knowledge during diagnosis to verify or reject hypotheses. Suppose a doctor has established that a patient has hepatitis and is proceeding to find out the possible causes of the disease. Let a piece of data be "the patient is a farmer." The doctor can bring to mind the knowledge that farmers often drink water from wells and that the patient may have contracted a viral infection from drinking the water. Notice that in this case the piece of world knowledge "farmers usually drink water from wells" was only activated in the context of diagnosing the cause of hepatitis. The datum "the patient is a farmer" would not play any role and thus might have been unnoticed in the context of some other diseases. Medical knowledge has this type of knowledge embedded in it. The right medical context activates this knowledge, which can hence be easily compiled in the form of production rules. In particular, the production rules will be inserted under the concept "virus infection as a cause of hepatitis," explicitly checking whether the patient has been ingesting contaminated water.

The foregoing should not be interpreted as denying that, for a complete model of a physician's reasoning, physiological, anatomical, and commonsense knowledge structures need to be represented in addition to diagnostic knowledge. There is no doubt that a physician uses these other structures and that what-type knowledge must be captured in them. They are needed in order to acquire new pieces of judgmental knowledge, to reconfigure and extend the concepts in the diagnostic structure, and to do productive problem solving involving knowledge in these other domains (which may result in compilable production rules to be added to the diagnostic structure). However, for achieving expert diagnostic performance, we do not believe that these additional structures are needed.

13.3.3 Redundancy and Biasing of Knowledge

The above considerations point to the view that the resulting knowledge structure for the diagnostician must be biased by the function that it is meant to serve. This means that the concepts that make up the structure and their organization are determined by the fact that they are grasping the medical knowledge of a diagnostician and not that of, say, a pathologist. The organization of medical knowledge from a pathologist's point of view will call for a different set of concepts and a different organization in the structure. Similarly, the knowledge encoded under each concept must also be biased toward the diagnostic task. The knowledge diagnosticians have about stone, tumor, etc., is only that necessary to establish them or rule them out in the context of liver diseases. However, the structure does not have to be just a classification of diseases. Other concepts that are not names of diseases may appear in the structure, to the extent that these concepts are needed to properly diagnose a disease. For example, in the structure of Figure 13-1, the concepts "stone," "cancer," and "stricture" are causes of a disease and not themselves diseases.

An organization of knowledge following these principles results in a high level of redundancy. Small pieces of knowledge in the form of production rules will appear grouped under different concepts. Also, the same concept may appear inserted in different places in the cognitive structure of the diagnostician. But the production rules grouped under the concepts will have differences reflecting the differences in the roles of the concept. An example of this occurs in the cholestasis syndrome. Consider the concept "stone" in its role as a cause of cholestasis. There will be production rules to establish or reject stone here, and also to check if a stone is indeed causing obstruction. However, stones may not necessarily cause obstruction directly, but may result in cholangitis. "Stone" would also occur as a concept under "cholangitis." This concept, while sharing some of the same production rules with the other stone concept, nevertheless will also have some rules that are different, because of the different role of this stone concept.

13.4 Kinds of Rules

Three types of rules must be grouped under each concept: confirmatory, exclusionary, and recommendation rules.

13.4.1 Confirmatory Rules

Confirmatory rules look for those manifestations associated with the concept under which they are located. Those manifestations could be sufficient

to establish the concept completely or only enough to postulate it. These rules return a list of the findings on which they establish or postulate the concept.

13.4.2 Exclusionary Rules

The need for exclusionary rules has been recognized. For instance, Pauker et al. (see Chapter 6) use exclusionary rules to reject a disease categorically. In our approach, they are used in a more inclusive sense. They collect all of the negative evidence for a given disease. The evidence does not have to be sufficient to rule out the disease. An obvious reason to have such rules is that physicians need to rule out certain diseases before performing some invasive procedures such as biopsy. A more interesting reason is that doctors frequently use a ruling-out problem-solving strategy. This happens when the data suggest several diseases and there is no conclusive evidence for any of them. Then the strategy consists in ruling out those diseases with the lowest evidence and focusing on the remaining. The use of ruling-out rules is the key methodological principle of differential diagnosis as explained by Harvey and Bordley (1972). The strategy adopted throughout their book is the following: once hypotheses are postulated to explain a given disease, a procedure is invoked that systematically begins to rule out some of them. For instance, in the discussion of a case of splenomegaly (pp. 371–376), they establish up to seven possible major hypotheses: polyarteritis, systemic lupus enterocolitis (SLE), lymphoma, etc. Immediately they say:

Polyarteritis is rarely associated with very significant splenomegaly, and the arterial lesions should have been seen. Arteritis can occur at all levels and may simulate almost any bowel disease, but at some stage bleeding is usually noted. None of the other clinical manifestations which suggest polyarteritis were noted.

Notice that the reasoning is based on highly categorical production rules. These ruling-out rules are tried even for those hypotheses that will eventually be accepted as explanations of a disease; that is, an attempt may be made to refute a hypothesis even in the presence of positive evidence for it. Nevertheless, it would be incorrect to see this practice as an exemplification of Popper's principle of refutation, viz., hypotheses are not verified but refuted. This is because the clinician not only considers the negative evidence for a given disease, but also the positive evidence. In our opinion, he or she can be viewed as running two procedures. One collects the evidence in favor of a given disease, the other the negative evidence. Then both are weighed, and a decision is made.

13.4.3 Recommendation Rules

Although the knowledge under each concept is mostly the knowledge needed to establish it or rule it out, there are pieces of knowledge under each concept that anticipate its subconcepts. For instance, jaundice, intense and intermittent abdominal pain, and elevated alkaline phosphatase strongly suggest not only a liver disease but also extrahepatic obstruction (one of the liver subconcepts). These pieces of knowledge will be translated during the problem-solving process as “recommendations” of a superconcept to its subconcepts. This knowledge will be represented in production rules that will be applied to the list of positive manifestations found by the confirmatory rules.

13.5 Concepts as Specialists

Given these principles of organization of medical knowledge, the solution of a medical case becomes a problem of taxonomic classification. It is similar to the problem of placing, say, a specimen of maple in a hierarchy of botanical concepts. It consists of identifying its superordinate and subordinate concepts. This is a process of recognition that is intrinsically top-down. Let us consider the use of knowledge in the context of the representational framework we have proposed. The solution consists of taking each concept in the structure as a specialist in that body of knowledge. Each concept interacts with others in the solution of a case by activating simultaneously each subconcept under conditions we explain below.

The decomposition of a body of knowledge into small systems is an old idea in AI. It was one of the central notions in Simon's beautiful book *The Sciences of Artificial Intelligence* (Simon, 1969). Winograd (1972) used the term *specialist* to refer to his semantic program for the noun group. Lenat (1975) used the notion extensively in his notion of “Beings.” Rieger and Small (1981) are building a “word expert parser,” and Minsky has recently speculated about a “society of minds” (Minsky, 1979).

The idea of viewing the human mind as a society of experts is very attractive. It has its counterpart in the human body system with its multiplicity of functions going on in parallel. The notion of specialist or expert is another metaphor that AI has borrowed from computer science. Any program consisting of a collection of modules or routines can be viewed as a collection of experts. Then the following question arises: what contribution is made by calling them experts? It seems to us that, in order for the notion of the society of experts to be useful, (a) we need criteria for the decomposition of a body of knowledge into small independent units, (b) the decomposition should be such that it will be able to support par-

allelism, and (c) communication and control should resemble those found in human intelligence.

In the case of medical diagnosis, the identification of these components is facilitated by the fact that the medical community is organized as a society of experts. The solution of a medical case requires in many instances the interaction of several specialists. The concepts that make up a diagnostic specialty have already been identified by the medical community. Nevertheless, the right structure and the precise interdependence of concepts for a given disease are by no means clear. One can be easily convinced of this by the fact that often different books about diagnosis do not coincide in the decomposition of the relevant concepts that need to be considered to properly diagnose a disease. The mapping of the medical knowledge as it appears in books into a structure like the one we are proposing is by no means automatic. Epistemic work is needed in order to come up with the right structure and the concepts that form the structure. It is our deep conviction that if automated medical diagnosis succeeds some day, books on medical diagnosis will be written in a very different form from that of the current texts.

13.6 Distributive Problem Solving

A distribution of the diagnostician's knowledge in a hierarchy of concepts, which are considered as independent specialists in that body of knowledge, leads naturally to a distributive problem-solving situation. In order to illustrate this, we recapitulate our basic ideas by considering the medical knowledge of an internist.

Referring to Figure 13-2, we see that the top-level node has no rules in it since it is always established. Its immediate successors are formed by generic diseases such as those related to liver, heart, etc. Under the concept "liver," our internist will have that knowledge needed to determine if a given patient has a liver disease. That specialist will look for those key findings associated with liver diseases, for instance, abdominal pain, jaundice, alcoholism, etc. Also, the specialist will have knowledge to rule out a liver disease and to make some recommendations to its subconcepts. But it will not have the knowledge to discriminate between the different kinds of liver diseases. That knowledge will be located in the two concepts under it, the intrahepatic and extrahepatic specialists.

Pathognomonic knowledge is useful not merely in establishing the concept under which it is located. In medical diagnosis, pathognomonic manifestations are those that indicate the presence of a disease with near certainty. If a concept has pathognomonic knowledge about a successor concept and if the corresponding manifestations are present, control is

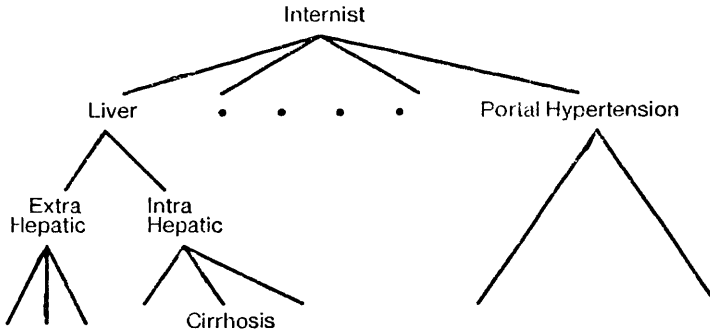


FIGURE 13-2 Top levels of the internist's conceptual structure.

transferred to the successor even if it is located several nodes down the tree.

If we consider the internist's knowledge organization, the nodes in the hierarchy are called *concepts*. Because these knowledge sources interact with others to solve a medical case, they are called *specialists*. Finally, because these concepts are names of diseases that must be verified or rejected, they may be called *hypotheses*. We use these three terms interchangeably in the remainder of the paper. A blackboard will be used to coordinate the work of the specialists in the solution of a case.

13.6.1 The Blackboard

The notion of a blackboard was used by Erman and Lesser (1975) as a way to provide an interface between different knowledge sources. The function of the blackboard in our design is similar to theirs: to provide a way of interaction among the specialists and to hold the current state of the system. The blackboard is divided into the following sections. ACTIVE-HYPOTHESES contains the names of all specialists that are active at a given moment. ESTABLISHED-HYPOTHESES contains the names of all hypotheses that have been established during the solution of a case. A hypothesis is established when the evidence exceeds some threshold. There could be cases in which the evidence in favor of a hypothesis is sufficient to categorically establish it, while in other cases the evidence could be only sufficient to postulate it. REJECTED-HYPOTHESES contains the hypotheses that have been rejected. SUSPENDED-HYPOTHESES contains all hypotheses for which a specialist has not found sufficient evidence to justify pursuing them. This section also includes those hypotheses that were initially postulated but later on abandoned because the evidence did

not exceed the threshold needed to pursue them. Finally, it should be noted that as hypotheses are entered in the various sections of the blackboard, the underlying hierarchical structure among them is preserved.

13.6.2 Activation

We can now consider the activation of the specialists. Once the top-level node is invoked, it activates simultaneously all its immediate successors and enters their names in the ACTIVE-HYPOTHESES section of the blackboard. These act in parallel on the patient's data base. These will look for those manifestations in the patient's data base that are associated with the generic concept they stand for. We can distinguish the following cases:

Case 1. A specialist, say S, can find data to consider that the disease it stands for must be pursued. If so, it will enter in the ESTABLISHED-HYPOTHESES section of the blackboard its name followed by a list of the manifestations on which it based its decision. Then it will activate its immediate successors (if some of the pathognomonic rules have been fired, it could activate some subspecialist down the tree). Upon their activation, S will pass to them the same information it entered in the blackboard plus a list of "recommendations." Finally S will deactivate itself by removing its name from the ACTIVE-HYPOTHESES section of the blackboard. The list of recommendations will contain pieces of advice about such aspects as what kind of rules (disconfirmatory or confirmatory) a given specialist should try first, indications to discourage the subspecialist to do an extensive search, etc. The type of recommendations depends on each disease. They provide another criterion to further organize the production rules under each specialist.

Each specialist, on establishing itself, will add to the list of manifestations, which then will be passed from parent node to child until it reaches a tip node. If the specialist in the tip node succeeds, it will print the list. At that point, the list will contain a classification of the medical case under study. The list could look like:

(Liver ($A_1 A_2 \dots A_n$) Extrahepatic ($B_1 B_2 \dots B_n$) Tumor ($C_1 C_2 \dots C_n$))

where A_i , B_i , and C_i are the manifestations on which each specialist based its decision.

Case 2. A specialist rejects itself. This happens when the exclusionary rules found the presence or absence of data sufficient to rule out the disease. In that case the specialist enters its name in the REJECTED-HY-

POTHESES section of the blackboard followed by a list of the negative evidence and deactivates itself.

Case 3. A specialist suspends itself. It then enters its name in the SUSPENDED-HYPOTHESES section of the blackboard followed by the list of manifestations before suspending itself. The suspension of a specialist can happen because the data it found did not exceed some threshold or because its immediate successors rejected or suspended themselves.

In cases 2 and 3, when the immediate successors of a node have rejected or suspended themselves, a mechanism has to be provided to remove that specialist from the ESTABLISHED-HYPOTHESES section of the blackboard. This can be accomplished by making the last active sibling (if it has suspended or rejected itself) check if any of its other siblings are in the ESTABLISHED-HYPOTHESES section of the blackboard. If none of them is there, it means that all of them have rejected or suspended themselves. In that case, the specialist will move its parent from the ESTABLISHED-HYPOTHESES section of the blackboard to the SUSPENDED-HYPOTHESES section. After that, it will check to see if none of its uncles is in the ESTABLISHED-HYPOTHESES section. If none is there, it will remove its grandparent, and so on.

13.6.3 The OVERVIEW Critic

It is generally accepted that a good practice in the diagnostic process is to explain again all the patient's manifestations from the point of view of the final diagnosis or diagnoses. Harvey and Bordley (1972) considered this to be the final step in the diagnostic process. In our approach, we need to organize a body of knowledge around that methodological idea. This is due to the fact that quite a few suspended hypotheses could result during the diagnostic process. They should be cleared, resulting in a more unified diagnosis. For this purpose we associate with each disease in the top level of the hierarchy an OVERVIEW critic.

OVERVIEW is activated only if the disease with which it is associated is advanced as one of the diagnoses. Basically, what OVERVIEW will do is to check if those manifestations that the specialists entered in the blackboard with each suspended hypothesis appear in the list of manifestations associated with any of the subspecialists of the disease that has been established. If all manifestations associated with a suspended hypothesis can be accounted for by this procedure, OVERVIEW will reject that hypothesis. Otherwise, it will advance that hypothesis as a second or third diagnosis. If the only function of OVERVIEW were this procedure, then it would not have to be associated with any particular disease. We feel, however, that other questions should be formulated by OVERVIEW, such as the relevance of the manifestation to the suspended hypothesis in particular

and to the diagnostic process in general and the chances of the appearance of both the suspended hypothesis and the established one. Further investigation will have to be conducted to determine the nature of these questions concretely. We conjecture that OVERVIEW would have knowledge global to the individual subspecialists into which a disease has been decomposed, as well as knowledge about other diseases in the top level of the conceptual hierarchy.

13.6.4 The Specific Role of the Blackboard

The blackboard can serve many functions in our approach to medical diagnosis. It will be a matter of further study to exploit all of its advantages. We can mention two instances in which its use is necessary. The first one deals with the problem of a disease being secondary to another. For instance, cirrhosis (a liver disease) can cause portal hypertension (which can have many other causes). In the medical jargon, it is said that the clinical manifestations of the latter are *secondary* to the former. However, the manifestations of each disease are different. Following our approach, let portal hypertension be a successor of the top node, internist (see Figure 13-2). Both nodes, viz., cirrhosis and portal hypertension, will be established in parallel in a patient with portal hypertension secondary to cirrhosis. At a given point, the portal hypertension specialist will pass control to subspecialists that will determine the possible causes of the disease. Then one of them is going to contemplate cirrhosis as being the cause. That subspecialist can verify this by looking at the blackboard for cirrhosis. Without this blackboard, the hierarchical call structure would be violated by a call to the cirrhosis specialist, or a redundant and *ad hoc* specialist would need to be created.

The second instance has to do with the fact that the specialists must communicate between each other to reduce the amount of search they must do. Consider specialists associated with different causes of the same syndrome. Although it is possible that a disease can have more than one cause, it is not frequent. Then if a given specialist has already found the cause of a disease, it makes very little sense for its sibling to pursue its search in the presence of very low evidence. As a specific example, consider the situation where extrahepatic cholestasis has been established, and each of its immediate successors, stone and cancer, is investigating itself as its cause (see Figure 13-1). As the stone and tumor experts are working in parallel, suppose the preliminary evidence for stone is low, while the tumor specialist establishes tumor. Now the stone specialist should suspend itself, but only if the information about tumor establishment is made available. This can be made possible by making the specialists (in this case the stone specialist) periodically inspect the ESTABLISHED-HYPOTHESES portion of the blackboard.

13.7 Implementation

In the preceding pages we have described a methodology for knowledge distribution and the associated distributed problem-solving strategies for medical diagnosis. There are two key aspects to the methodology: (1) knowledge is decomposed into a collection of specialists, and (2) these specialists perform problem solving in parallel in certain specified ways, using a blackboard as a record of the global state of problem solving.

A prototype diagnosis system called MDX (Chandrasekaran et al., 1979; Mittal et al., 1979) has been built by our group and has been operational for some time. This implementation has both points of contact with and differences from the methodology described in the preceding pages. The major points of contact are that the current domain of MDX, viz., cholestasis, is organized into a collection of specialists as indicated in Figure 13-1 and that diagnostic knowledge is distributed in this structure following the guidelines spelled out in Section 13.3. The problem-solving strategy is the area of most of the differences between the methodology described in this paper and MDX as implemented. The source of these differences is threefold: (1) the strategy in this paper is of more recent origin and goes beyond the current MDX strategy in power; (2) the methodology emphasizes the parallel invocation of specialists, which is of particular importance in a distributed implementation and of less operational significance in a serial implementation such as MDX; and (3) the domain of implementation is not large enough to need the global state record in the form of a blackboard. The power of a blackboard of the type we have envisaged will be needed as the domain is enlarged. In particular, it will be needed for decisions at the top (internist and one or two levels below) where proper coordination between subspecialists of vastly different scope would be needed.

These differences notwithstanding, MDX is a working implementation of a distributed approach to problem solving. As such, a brief outline of its performance is in order. A more complete discussion of the system and its performance is available in the papers cited earlier.

The top-level specialist in the system is GP (or internist), but all that it can do at this stage of implementation is either to hypothesize cholestasis and transfer control to it or to reject the case. Cholestasis may be hypothesized by a collection of production rules that respond to the relevant lab data and physical signs and symptoms. When cholestasis gets control, its charge is first to establish itself and then to further refine itself to account for all the manifestations. This *establish-refine* strategy is fairly general to the system as it currently exists. The rules used to establish cholestasis are of the confirmatory type mentioned in Section 13.4. The disconfirmatory evidence is not currently used in all the nodes, but where it is used it is in the form of negative weights for the disease for certain combinations of data in an evidence-weighting table.

Once cholestasis, say, is established, a priority scheme is needed to call subspecialists, since MDX is a serial implementation. This priority is provided by a collection of rules that suggest possible specialists on the basis of certain patient data or combination thereof. The criterion for the selection of the rules is that they represent common or easy possibilities. If this criterion is satisfied, the specialists that are called earlier by the priority scheme are more likely to solve most of the cases. Only in "hard" or uncommon cases will the rest of the specialists need to be called.

These specialists are typically called to establish and refine themselves and, when they succeed, to return those abnormal data that they can explain. The specialists that are established and the corresponding data are kept in an ACTIVE list. When the specialists in the top-level ACTIVE list together can explain all the abnormalities in a nonoverlapping way, the case is solved. Note that the specialists lower than cholestasis in the hierarchy may also have their own priority rules to select their subspecialists. The tip nodes, when called, match the data within their scope with confirmatory rules or equivalent tables to establish or reject themselves. This information is passed up to the calling specialist. Each specialist thus organizes, by means of production rules, the priority by which it uses its subspecialists to arrive at an explanation of abnormal data in its scope. When the subspecialists explicitly suggested by the rules fail to explain the case, then an exhaustive interrogation of all subspecialists one level below will be made. Thus the priority rules do not preclude the correct answer from being obtained eventually.

As stated earlier, the current implementation of MDX does not use a blackboard. Consider the case involving cirrhosis and portal hypertension that was discussed in Section 13.6.4. In our current implementation, portal hypertension will neither call up the cirrhosis expert nor look up the blackboard. Instead, it will have a cirrhosis-as-cause-of-hypertension subspecialist, most of whose knowledge would be a replication of the cirrhosis specialist. This is clearly an *ad hoc* solution, but as long as the domain is not very large, it does not produce serious problems.

Another constraint in the MDX implementation has to do with the atomistic nature of the patient manifestations (see Section 13.3). The cholestasis domain has so far not produced sufficiently complex cases for which this data representation presents a serious limitation. However, the future extensions of MDX will increasingly incorporate more sophisticated structured data representations as discussed in Section 13.3.

13.8 Concluding Remarks

The ideas presented in this paper contain points of coincidence with other research in automated medical diagnosis. They coincide with MYCIN (Shortliffe, 1976) in taking production rules as the formalism for repre-

sentation, but in our approach rules are organized under concepts. In INTERNIST (Pople, 1977) the hierarchy of diseases is essential to the problem-solving strategy. In our approach the hierarchy is not only of diseases, but also of causes of them and of any concept relevant to the diagnostic process. The concepts in our hierarchy are *specialists*, aggregates of knowledge about a significant step in the determination of the diagnosis. We coincide with CASNET (see Chapter 20) in the relevance of etiologic reasons in the diagnostic process, but in our approach that is one reason among others. The concepts in our hierarchy are highly compiled. Thus some specialists will have etiologic knowledge, while others will base their reasoning on other types of knowledge depending on the disease. Finally, our approach coincides with PIP (see Chapter 6) in taking each disease as a cluster of knowledge with distinct features. But the structure of diseases is a hierarchy in our approach; in PIP it is not.

An important aspect of our ideas is that medical (or for that matter any) knowledge can be viewed as a collection of essentially decoupled conceptual structures, each with an embedded problem-solving mechanism (reflecting its intended use). In the actual handling of a case, a physician is in the diagnostic mode only part of the time. The incompleteness of the diagnostic structure in a particular physician, as well as other considerations involving therapies, costs and other situational idiosyncrasies, and a perceived need for explanation at different levels will typically cause him or her to switch between different knowledge structures, but a satisfactory accounting of this overall process can be done, in our view, only after the underlying conceptual structures and the problem-solving mechanisms implicit in them are identified. We have advanced in this paper an analysis of one such structure, viz., the diagnostic one.

ACKNOWLEDGMENTS

We thank David C. Brown, James Reggia, Jorgen Hilden, Jack Smith, and Sanjay Mittal for their comments on an earlier draft. Our access to the computing facilities of the Rutgers University Laboratory for Computer Science Research, made possible by a grant (RR-643) from NIH Biotechnology Resources Program, Division of Research Resources, has been essential to our implementation activities. Finally, the National Library of Medicine Biomedical Computing Training Grant to The Ohio State University (LM 07023-02) has helped foster an active interest in these and other fundamental problems in medical knowledge representation.