

AN APPROACH TO MEDICAL DIAGNOSIS BASED ON CONCEPTUAL STRUCTURES

B. Chandrasekaran, F. Gomez, S. Mittal, and J. Smith, M.D.
Artificial Intelligence Group
Department of Computer and Information Science
The Ohio State University
Columbus, OH 43210

The power of an experienced medical diagnostician seems to us to arise from a knowledge organization whose very structure helps keep the combinatorial growth of processing under control. This structure mirrors the deep conceptual structure of the field. We propose a knowledge representation scheme called conceptual structures, which organizes knowledge in the form of production rules or more complex procedures in such a way that they are accessed and used as needed. This structure is dominantly hierarchical, the successors of a conceptual node standing for subconcepts which refine that concept. Associated with each concept is a set of procedures (experts) which attempt to apply the relevant knowledge to the case at hand; this might include turning over control to selected subconcepts for more detailed analysis. Thus, abstractly, the conceptual structure organizes the invocation of procedures available to the diagnostician. The calling of a specialist by a physician is akin, we think, to control transfer to a subconcept in the cognitive structure of a diagnostician during problem solving; the difference is in the level of detail of the cognitive map. In this paper, some principles governing the design of conceptual structures are presented. We apply the ideas to the design of MDX, a system to perform diagnoses in a syndrome called Cholestasis, and discuss the performance of the system.

1.0 INTRODUCTION

In this paper we shall describe an approach to the design of systems to solve problems in knowledge-rich domains. We shall consider in particular a system intended to perform medical diagnosis related to a liver syndrome named Cholestasis. The underlying metaphor is one of a society of experts among whom the knowledge of the domain and the procedures for acting out that knowledge are distributed. The organization of knowledge as experts raises many important questions, such as what knowledge should be organized into an expert and what other knowledge distributed between them; the nature of control transfer and communication between experts; the data-base organization and access; and the relation between redundancy of knowledge and data on one hand and the speed of processing on the other. While the idea of experts is not new, we believe that the epistemological principles needed as criteria for decomposing a body of knowledge into experts, in order that their interaction is controlled rather than chaotic, have been missing. Our discussion in the next few sections will attempt to highlight some of these principles.

The paper is organized as follows. In Section 2, we discuss how conceptual structures effectively organize a body of knowledge into experts. In Section 3 we discuss, for purposes of illustration, the knowledge representation of a fragment of Cholestasis. This leads, in Section 4, to a discussion of different types of experts and their structural representation. In Section 5, a prototype medical diagnosis system based on these principles is described. The description is accompanied by an illustration of the system diagnosing a real medical case. Other contemporary approaches to medical diagnosis using Artificial Intelligence techniques are discussed briefly in Section 6. In the concluding section, we outline some extensions that are contemplated for the near future.

2.0 CONCEPTUAL STRUCTURES AND ORGANIZATION OF MEDICAL KNOWLEDGE

We adopt the view that a given field of knowledge has an underlying conceptual structure which enables effective use of that body of knowledge. This is the deep structure to which the knowledge representation in the field should strive, and this is what effective teachers communicate and effective learners build. This conceptual structure can be viewed as a way of organizing knowledge in such a way that purposeful, focussed access to the relevant parts of the knowledge structure is achieved. Consider medical diagnosis as an example. It is our view that the power of an experienced medical diagnostician arises from a knowledge organization whose very structure helps keep the combinatorial growth of processing under control. Knowledge in the form of facts, heuristics, production rules or more complex procedures cannot be effectively used unless it is embedded in this underlying structure in such a way as to facilitate access as needed.

What are the principles by which domains of knowledge are to be organized so that they can be used effectively? Our thesis is that the major criterion has to do with the efficiency of transfers of control or access to different parts of the knowledge structure. There are correct and incorrect analyses of the conceptual structure of a field. Incorrect analyses lead to a proliferation of control transfers; correct ones should generally lead to focussed and controlled communication between different parts of the structure.

The organization of the medical community provides a concrete case study in the principles of knowledge structuring. A phenomenological account of medical decision making seems to us to be very revealing in this regard.

One of our guiding notions is that there is a correspondence between the broad conceptual organization in a diagnostician and the organization of specialties in a medical community. The calling of a specialist by a physician is not dissimilar, in our view, to control being handed over to a subconcept in the cognitive structure of a diagnostician during the course of his problem solving. Thus this is a concrete realization of the society of minds paradigm that has been put forward by Minsky and Papert[1]. In order for this paradigm to go beyond being a metaphor, we need to know how a society of minds, whether in an individual or as a collection of experts solving a common problem, should coordinate the activities of and communication among members to produce intelligent behavior. We feel that the organization of specialties in the medical community reflects the underlying conceptual structure, and, to a certain extent, the requirements of efficiency in control and communication structures as well*. The same principles that lead to effective communication among real experts in a medical community are to be used in organizing and structuring knowledge within a specialty. A simplified rendering of these principles are:

- P1) There is some sort of a hierarchical organization with respect to control transfer. That is, the GP who had original responsibility might, after deciding that it is a liver problem, transfer control to the liver specialist. The liver specialist might in some cases hand over control to a super specialist in a subarea, but it would be highly unusual for him to transfer the patient to a heart specialist. If such a thing is ever done, it'd be usually through return of control to the GP who would then accomplish such a transfer.
- P2) The expert who is contemplating control transfer knows enough about the domain of knowledge of the subexperts who are likely to be called by him to decide, generally correctly, which subexpert is relevant, but not enough to solve the problem himself.
- P3) Each expert knows enough about the domain of others in his level of the hierarchy to decide when he himself has been mistakenly called.
- P4) There are some experts who do not belong in the main hierarchy, but who are outside of it, so to speak. They are not generally given control for the purpose of diagnosing the original complaint. They are used as resource experts to accomplish specific subtasks requiring their expertise. They communicate back to the expert who called them with the answer. An example of this type of expert is the radiologist.
- P5) The communication between an expert and a subexpert is of several forms: answer specific questions

* It might be argued that medical specialties historically grew up around organ systems, i.e., along physiological and anatomical lines, and that control and communication efficiency is not a paramount criterion. In fact all this means is that the criterion of control transfer and communication efficiency is often met by organizing many specialties along these lines- after all one would expect that there would be a certain coherence within each such specialty. Further, not all medical specialties are organ-based; consider rheumatology, radiology, emergency room medicine, clinical pathology, toxicology and infectious diseases as counter examples. Our thesis is that the intuitive notion of internal coherence of a specialty is closely related to the control transfer and communication efficiency, and thus the latter becomes the concrete criterion of "natural selection" in the evolution of specialties.

dealing with the subspecialty, confirm or reject possibilities, or, in difficult cases, hand over total control. Further, in some instances a mere number standing for weight or probability is returned (see P7 below); in most cases, however, the communication is symbolic in nature, in a vocabulary which, in parts, is particular to the specialty subject matter.

- P6) An expert often has enough knowledge of the subject matter of the subexperts to be able to solve easier cases. For example, most GPs know enough about liver diseases to solve "easy" liver cases, enough about heart diseases to solve easy heart cases, etc. Among the criteria that determine which knowledge should be abstracted at this level are: how common is this kind of malfunction, how conceptually complex is this etc. This sort of organization keeps the combinatorics of control transfer to a minimum, since only in a relatively few cases will there be a large number of accesses to subexperts.

- P7) While many significant decisions involve manipulating symbol structures associated with concepts, there are very definite situations which are best viewed as weighting several pieces of evidence to produce a numerical weight for a hypothesis. This piece of information is eventually incorporated in symbolic structures which are used for problem solving.

Applying these principles to organization of knowledge within a specialty such as liver disease or Cholestasis, we have found that the conceptual structure can be usefully viewed as a tree, where a node stands for a particular concept and the successors of the node stand for subconcepts that help refine that concept; e.g., hepatitis, whose successors might be acute, fulminant and chronic types of hepatitis. Associated with each node is a set of procedures which decide on the applicability of the concept to the case at hand. Part of such decision-making in a node is often the decision to turn the control over to subconcepts and their associated procedures to check on their applicability. In this sense, the conceptual structure organizes the invocation of procedures available to the diagnostician.

Notice that the form of the structure as a tree already exercises a strong constraint over communication. An expert cannot call other experts arbitrarily, but can only communicate via their super and subconcepts. While it is hard at this stage to advance this principle as an absolute requirement, we have found that violations indicate either that the analysis was incorrect or that the concept actually belongs outside the structure in some sense, with specific restricted modes of communication with the experts in the main structure. This is consistent with principle 4 above.

We note in passing that the principles governing the organization of conceptual structures are meant not only to make effective use of knowledge, but also to aid in the acquisition of it. The concepts get refined as a function of further learning, and the node at which the refinement takes place gets appropriate abstractions about the subconcepts. Another point to be noted is that often heuristics are formed which enable the system to go directly to a concept several levels below, sidestepping the methodical, level by level communication. But concentration on this second order optimization device obscures the clarity of the underlying structure. In fact expertise often consists of an accumulation of such heuristics, and, for this

reason, experts tend to be poor elucidators of the underlying structure. These, however, are issues which are not the main concern of this paper, and will be discussed in detail elsewhere.

In the next section, we illustrate the usefulness of the above ideas by the concrete example of a fragment of the cholestasis system. While we have performed the analysis for cholestasis in general, we present here only a portion thereof, in order to concentrate on the principles.

3.0 CONCEPTUAL STRUCTURE OF CHOLESTASIS

Cholestasis is a condition caused by disruption of the secretion of bile or of its excretion into the duodenum. It happens that while the identification of the cause of cholestasis might be complicated, the determination that it is present or not present can be accomplished with relative ease. This fact makes the cholestasis knowledge base a useful object of study: it is large enough to be an adequate arena in which to test the ideas, and at the same time, its interaction with knowledge bases concerned with other diseases can be kept under control because of the relative certainty of its applicability. (We call nodes corresponding to concepts with the latter property "anchor nodes.") It has been found useful to group the causes of cholestasis as occurring within the liver itself (intra-hepatic) or in the biliary tree outside the liver (extra-hepatic). There are often broad indications of whether a case is extra- or intra-hepatic, which, while not 100% reliable, are generally helpful in deciding which is more likely* and thus should be investigated first.

We shall shortly give the details of Procedure Cholestasis, which will be representative of the procedures in most of these conceptual nodes. But first we need to refine the concept of Extra-hepatic Cholestasis in some detail. The movement of bile through the biliary tree can be obstructed by a physical object such as stone, stricture or cancer of various kinds or by swelling caused by inflammation due to a number of causes. Associated with each subconcept are procedures indicating how to establish it by symptoms, signs and lab data; information about deciding which of their subconcepts are applicable, etc. In accordance with the principles set forth earlier, our analysis of Cholestasis recommends the conceptual structure in Fig. 1. Because of the procedures attached to each of these nodes, the tree can be viewed as a structured organization of experts involved in diagnosing Cholestasis. The creation of such a hierarchy is the result of the analysis of the content of a field of knowledge. The difficulties involved in the uncovering of such a structure can be illustrated by considering the "stone" and "inflammation" (cholangitis) nodes. Textbooks on liver diseases inform us that inflammation may be caused by the presence of stones and, conversely, stones may be formed if there is persistent inflammation. A direct rendering of this information would result in the mutual calling of these two experts, violating the hierarchical expert call structure. The difficulty is cleared by noting that these two nodes do not correspond to the general concepts of "biliary stone" and "bile duct inflammation." Instead, they stand

* "Likely" here does not mean that probabilities are computed; rather, certain aspects of the problem description are looked at for information about which should be given priority in investigation.

for "stone as cause of Cholestasis" and "inflammation as cause of Cholestasis." The "stone" node then has the information about "inflammation" needed for its work localized there (see III in Procedure Stone-Consistency in sec.4.4) and similarly for the "inflammation" node. It seems to us that such a decoupling in knowledge organization is essential for operational efficiency.




Fig.1 Conceptual Structure of Cholestasis

4.0 PROBLEM SOLVING BY EXPERTS

4.1 Types of Experts

There are two types of experts - characterized by the control they exert over the problem solving process. Experts like Cholestasis, Extra-hep, Intra-hep, etc., play greater role in decision making. They decide, for example, which expert to call next; what sub-problems to solve; decide if the problem is solved; mediate between competing advice from subconcepts (in differential diagnosis situations), etc. The other class of experts, such as Stone, Stricture, etc., primarily evaluate the problem-data to determine how well the concepts they represent fit the data and to say what additional data will enable a better decision.

Concepts such as Stone are tip nodes in the concept tree (primitive concepts). Deciding if the presence of stones characterizes the case is done mainly by weighing different pieces of evidence, both pro and con, into a whole.* This does not entail any search through a space of hypotheses. On the other hand, nodes higher in the

* This is not to say that human diagnosticians use numbers in this situation. The important point is that

tree need to search below in order to fulfill their task. It often happens, however, that a node which was incorrectly analyzed as a tip node is indeed a concept which is sufficiently complex that further refining is necessary.

4.2 Problem Solving Strategy

The basic strategy in problem solving is one of establishing and refining the concepts. Experts associated with the concepts actually carry out these processes; i.e., there is no uniform mechanism operating on different concepts for this purpose. This is very important, as each concept may have different knowledge and thus associated strategies for establishment and refinement.

Some nonprimitive concepts are anchoring concepts, because they can be established with a high degree of certainty. In pathognomonic situations, they can be established by one datum or one set of data, but generally they are established by accumulating evidence from a small number of highly reliable symptoms/lab data. Once established, the concept is refined by the process of determining which of the subconcepts, alone or in combination, fit the data which this particular concept was originally asked to explain. These anchoring concepts provide a strong global focus to the problem-solving process. Control is systematically distributed to other experts so that each may solve a part of the refinement problem. The expert which provides the focus also retains the control over the problem-solving process, thus largely eliminating the "focus of attention" problem.*

How to handle additional data while the diagnosis is in progress is an issue in any diagnosis system. Our system is designed to begin with some initial data and ask for more, when needed. The new data do not prompt any costly computations or drastic change of focus; instead, the controlling expert uses them to re-compute only what was affected within its domain. As control is passed back and forth, each expert determines the relevance of the data and acts accordingly but only when it gets the control.

How the above problem solving strategy is embodied in the system can be best understood by considering some examples of experts. We shall consider Procedure Cholestasis and Procedure Stone-Consistency as examples of non-tip and tip node experts respectively. Cholestasis is also an anchor node and thus provides a good illustration of how control is carefully distributed with a clear focus.

4.3 Procedure Cholestasis

I. CHOLESTASIS has been handed control and asked to first ESTABLISH itself and then REFINES the concept it

symbolic template matching of some kind is needed. In Procedure Stone-Consistency that follows, the numbers really stand for discretized, symbolized degrees of evidence.

* Reference[2] discusses the focus of attention problems that arise in HEARSAY-II, which, even though it is not a medical diagnosis project, is of interest due to its experts-based organization. Our additional principles for constraining control and communication seem to make the difference.

embodies.

1. Establish cholestasis on historical data, signs, symptoms, and clinical data. This portion is a combination of heuristic rules and numerical weights. (This is straightforward and can be done quite reliably.)

2. Consider first the disease-hypotheses invoked by patient data in INVOKED-CONC list and made available to CHOLESTASIS by the calling expert. (As a doctor scans through the data, some possibilities come to mind. This is the typical heuristic knowledge of most doctors about cholestatic diseases. However these are merely suggestions to be evaluated.)

a. Generate calls to immediate sub-concepts of CHOLESTASIS, which are the superconcepts of those in the list. For example, if INVOKED-CONC contains SCL-CHOLANGITIS, BD-CANCER, PERICHOLANGITIS and PRIM-BIL-CIRRHOSIS, these calls would be ---

Call EXTRA-HEP (consider SCL-CHOL, BD-CANCER)

Call INTRA-HEP (consider PERICHOL, PRIM-BIL-CIR)

(The experts corresponding to these eventually get control and they evaluate data more carefully and return four lists: CONSISTENT-DATA, CONTRADICTING-DATA, RECONSIDERATION-DATA, and RECOMMENDATION, which together consist of weighted recommendation and reasons. Note that CHOLESTASIS knows which subconcepts to call, even though it doesn't know enough to make a decision about them. This is consistent with P2.)

b. Place members of INVOKED-CONC into three lists: ACTIVE, INACTIVE and UNLIKELY. INACTIVE holds those for which data are insufficient to decide one way or the other.

3. Check if concepts in ACTIVE list can explain all ABNORM-DATA (use CONSISTENT-, CONTRADICTING-DATA lists for this).

i. Yes, also no significant overlap, return ACTIVE list as answer (straightforward cases will be solved at this stage).

ii. No, some data in ABNORM-DATA cannot be explained. Go to DEFAULT-PROC.

iii. Yes, but overlap. Go to OVERLAP-HANDLE*, return modified ACTIVE list.

II. DEFAULT-PROC (Control transfers here when INVOKED-CONC fails to explain all facets of cholestasis.)

1. Use data to decide which of EXTRA or INTRA is more likely. (Imagery information is often very effective in this discrimination. Errors here will delay correct diagnosis, but not prevent it.)

2. Execute EXTRA-IN-CHOLESTASIS or INTRA-IN-CHOLESTASIS according to result from 1. If solved, return ACTIVE list. Else call the other. If that fails go to LOOK-FURTHER.

III. EXTRA- and INTRA-IN-CHOLESTASIS

(EXTRA- and INTRA-IN-CHOLESTASIS contain some knowledge about subconcepts that is of a more usual nature and thus helps deal with a large number of cases. Relatively easy cases will be generally handled by this. Again, putting this knowledge here

* OVERLAP-HANDLE, which deals with the notion of two diseases explaining substantially the same data, is not currently implemented and will not be further discussed here. Details will be given in [3].

reflects P2 and R6.) These consist of a collection of rule-based calls of the form:

```
If history of acute/chronic cholecystitis, Then
Call EXTRA (Consider Stone)
If air in x-ray of biliary tree, Then Call EXTRA
(Consider Stricture)
...
```

IV. LOOK-FURTHER

(Control transfers here when everything CHOLESTASIS knows directly has been tried, but there is still "significant" data left to be explained.) Order EXTRA and INTRA in terms of likelihood (as in II-1 above) and transfer control to one, and then the other if necessary, with request to diagnose unexplained abnormal data.

Principle 1 is implicit in the organization of the conceptual structure. We saw several explicit examples of P2 in I.2a. and II.2 above. This particular example does not involve P3. Interpretation of imaging data will involve calls to the radiology expert, which is not explicitly indicated in the above procedure, but which is an example of P4. The different forms of messages, as indicated by P5, are exemplified throughout the Procedure. DEFAULT-PROC employs P6 partially. This procedure has enough knowledge about subconcepts to call them for answers to specific questions, but keep the control for easier cases. The establishment of cholestasis in I.1 is done in our implementation partly by weights assigned to different symptoms. This is an example of P7.

4.4 Procedure Stone-Consistency*

(We show only that part of the expert which determines the fit; message and explanation details are not shown.)

```
I. CALCULATE X-RAY EVIDENCE
  If Cholangiograms available
    then If Stones seen in Biliary-tree
      then XEV ← 3; else XEV ← -1
    else If plain-film X-ray available
      then if stones seen in Biliary-tree
        then XEV ← 3; else XEV ← 0
      else XEV ← 0
  If XEV > 2 return positive consistency evaluation

II. CALCULATE HISTORY EVIDENCE
This decision is made by asking about stone-related
diseases: hemolytic- disease, cholecystitis etc.

III. CALCULATE CLINICAL EVIDENCE
This decision is made on the basis of the evidence
for Chlanguitis, colicky-pain, vomiting and nausea in
a weighted-sum logic.

IV. CALCULATE SUMMARY EVIDENCE
Evidence from 1, 2 and 3 is combined in a
weighted-sum logic to obtain the consistency
evaluation.
```

5.0 PERFORMANCE OF MDX - A PROTOTYPE SYSTEM

5.1 Nature of Interaction

In this section we discuss the performance of the

* Each of these four major decisions are made by a separate group of rules, which are evaluated only when needed. For details see [3].

prototype diagnosis system(MDX) on a real case. Currently the system can only diagnose extra-hepatic Cholestatic syndrome, which consists of about nine diseases represented as terminal nodes. The system is interactive in two senses. The user can enter as much data as he wants, and later on the system will ask for additional data, if needed. Interaction of the second kind arises when an expert which is not yet implemented is nevertheless invoked by some other expert. In this situation, if a message is sent to an expert which is known to the system but not implemented, the message is routed to the user who can provide answers on behalf of the unimplemented experts (see [3] for details). Part of the attraction of expert systems is that they permit such a modular development in a conceptually natural way. This approach has the additional advantage that the logic of the system's decision-making becomes relatively clear to the physician and he can view the system as a consultant and not as a competitor. We have space here to discuss only one case in some detail.

5.2 Example Case

This case is discussed in [4]. Our discussion will be in the form of computer printout fragments (in Gothic) interspersed with commentary. User response to queries during diagnosis is preceded by \$ sign. In most places, the explanation generated by the system should be sufficient. Patient data are first entered in episodal form. For this case the following data are entered for past medical history and the At admission episode.

PATIENT RECORD SYSTEM

```
PATIENTS NUMBER IS PAT01
AGE? 39
SEX? M
ENTER HISTORICAL EVIDENCE
? HALOTHANE
? (ULC-COLITIS F)
? (BIL-SURGERY F)
? (CHOLECYSTITIS F)
? E
TIME? (AT ADMISSION)
? JAUNDICE
? PRURITUS
? VOMITTING
? (WTLOSS (AMOUNT (LB 16.0)))
? (PAIN F ABDOMEN)
? (LIVER (SIZE NORMAL))
? (PAIN LIVER F)
? (PAIN T EPIGASTRIUM MILD)
? (SPLEEN (FELT F))
? (URINE (BILIRUBIN))
? (STOOL (BLOOD F))
? (BILIRUBIN (DIRECT 10.799998) 12.200000)
? (PROTEIN 8.1 (ALBUMIN 3.9) (GLOBULIN 4.2))
? (SGOT 32)
? (ALK-PHOS 19.200000)
? (PT 11.399999 (CONTROL 10.5))
? (ESR 16.0)
? (WBC 9200 (PMN(PER 69.0))(LYMPHOCYTE (PER 19.0))
  (MONOCYTE (PER 10.0))
  (BASOPHIL (PER 2.0)))
? (ULTRASONAGRAM (BIL-TREE (IHD DILATE)))
? (PLAIN-FILM (LIVER (NORMAL))(DUODENUM (NORMAL)))
? (ERCP (BIL-TREE (CHEP OBSTRUCTION CONICAL)
  (CBILE (NORMAL)))
  (GALL-BLADDER (NORMAL)))
```

```

(PANCREAS (PANC-DUCT (NORMAL))))
? (GI-XRAY (DUODENUM (COMPRESSION)))
? (LIVER-SCAN (LIVER LARGE) (SPLEEN LARGE))
? (PERC-TRANS-CHOL
  (BIL-TREE (IHD DILATE)(CHEP NARROW)
    (LHEP MULTIFOCAL)(RHEP MULTIFOCAL)))
? (SPLAN-AGRAM
  (LIVER (SURFACE (TUMOR-BLUSH F)
    (NEOVASCULATURE F))))
END OF DATA ENTRY FOR PATIENT PAT01

```

The entry of negative information, such as (PAIN F ABDOMEN) meaning "no abdominal pain", needs some explanation. Much of this information is obtained by routine physical examination, historical questionnaire and lab tests. However, information such as (ULC-COLITIS F) is not routine; it may be available from patient history charts, but more often it is extracted by questions or tests when a particular hypothesis is being considered. The system is capable of asking for more data as needed, though that aspect is not our main concern here. We shall consider the version of our implementation in which any data not given are assumed to be unknown, unless they can be inferred from known data (with the exception of queries answered by the User for unimplemented experts). In any case, the data that are entered here- both positive and negative- are those directly taken from the case description in the journal. The important point is that there exists a genuine intellectual problem of putting the pieces together into a coherent whole even when only relevant data are entered. The problem of extraneous data is a bit more subtle. Clearly, totally irrelevant data should not cause any problems, and in fact they don't. No data-directed component has any rules that would fire in this case, and no hypothesis-driven component would look for them. Potentially relevant extraneous data are more problematic, but that applies to our system as well as to human diagnosticians. In this case blind alleys are possible, but erroneous diagnoses usually are not. Space limitations prevent a fuller discussion of this issue in this paper.

5.2.1 The Diagnosis

The topmost expert in the system is GP. As the data are entered, some rules are fired, suggesting hypotheses. The hypotheses become the starting point for GP to diagnose the case. In this case, rules suggesting Cholestasis and Post-operative-cholestasis were fired,

```

BEGIN DIAGNOSIS FOR PATIENT PAT01
TRYING RULES TO INVOKE HYPOTHESIS
TRYING RULE:
((( PRESENT?(URINE BILIRUBIN))(NORM? ALK-PHOS AN))
  ((INVOKE CHOLESTASIS)))
SATISFIED
EXECUTING ACTIONS: ((INVOKE CHOLESTASIS))

TRYING RULE:(((HISTORY? ANESTHETICS))
  ((INVOKE (POST-OPER ANESTHETICS))))
SATISFIED
EXECUTING ACTIONS:
  ((INVOKE (POST-OPER ANESTHETICS)))

TRYING RULE:
  ((DEFORMITY? BIL-TREE ? STONE)) ((INVOKE STONE)))
UNSATISFIED:

```

Other rules involving Pruritus, Bilirubin and Jaundice and invoking cholestasis were also triggered.

```

INITIAL DATA INDICATES ---
CHOLESTASIS
(PUST-OPER ANESTHETICS)
CHOLESTASIS IS INDICATED BY
PRURITUS PRESENT
JAUNDICE PRESENT
URINE BILIRUBIN
NORM OF ALK-PHOS IS (AN)
NORM OF BILIRUBIN IS (AN)
RATIO OF DIRECT VALUE VS
  BILIRUBIN VALUE IS 0.885
(POST-OPER ANESTHETICS) IS INDICATED BY
HISTORY OF ANESTHETICS

```

CHOLESTASIS being the highest anchor node suggested, control is transferred to it with the message to Establish and Refine itself. The messages and some of the rules* used in establishing cholestasis are shown.

```

HANDING OVER CONTROL TO CHOLESTASIS EXPERT
EXPERT-CALLED: CHOLESTASIS CALLED-BY: GP
MESSAGE (QUERY (ESTABLISH CHOLESTASIS)
  (REFINE CHOLESTASIS))

```

Procedure Cholestasis described in an earlier section is now operational.

```

TRYING TO ESTABLISH CHOLESTASIS
ENTERING TABLE LOGIC
TRYING RULE: (((DUCT? IHD BIOPSY BILE-STASIS)
  (PRESENT? JAUNDICE)) ((T 3)))

```

```

UNSATISFIED:
EXITING TABLE LOGIC
PATHOGNOMONIC EVIDENCE IS 0

```

Similarly rules such as that involving large liver size result in accumulating physical evidence, and rules triggered by bilirubin, SGOT and other lab data are used to construct a more structured clinical evidence.

```

PHYSICAL EVIDENCE 4
CLINICAL EVIDENCE (8 4 0 2)

```

A symbol-matching table such as in Procedure Stone-Consistency is finally used to decide cholestasis on a scale of -3 to 3. Notice the difference between the use of rules for suggesting cholestasis in GP and for establishing cholestasis in this node. In the latter, evidence is methodically accumulated and weighted to reach a decision.

CHOLESTASIS FIT: (ESTABLISHED CHOLESTASIS 3)

GP had passed two related suggestions to Cholestasis, namely Post-Oper and Anesthetics as invoked suggestions (i.e. suggestions made in a data-directed fashion) and these experts are now called (see I.2 in Procedure Cholestasis). The message actually gets routed** to the

* In the rest of the discussion only a few of the rules from each conceptual group will be shown. Each decision within an expert is made by a rule-group or a procedure. The representational details will be described in [3]. Rule conditions are actually logical queries to the data-base. The details of the data-base organization and queries can be found in [5].

**The routing was done automatically by the communication system and CHOLESTASIS did not know that POST-OPER and its ancestor INTRA-HEP are not implemented. The answer does indicate, however, that the decision was made by the USER. Because POST-OPER is the immediate super-concept of ANESTHETICS, only one message was sent, though individual answers were returned.

MESSAGE (ANSWER(ESTABLISHED BD-CANCER 3))

reaches CHOLESTIASIS via EXTRA-HEP.
MD-CANCER made the decision after asking questions about where the obstruction and tumor in the biliary-tree may be (both questions are complicated queries posed to the Radiology expert, not shown here, which examines the relevant radiological information to make the inference).

SCL-CHOLANGITIS is called next, but returns an unlikely evaluation after detailed processing of relevant information.

```
TRYING TO ESTABLISH SCL-CHOLANGITIS
TRYING RULE:
  (((AND-TRI (DUCT? BIL-TREE BIOPSY FIBROSIS)
              (DUCT? BIL-TREE BIOPSY (CARCINOMA
                                      ATY-EPITH-CELL))
              (HISTORY? ULC-COLITIS))))
  ((T 3)))
UNSATISFIED:
```

PATHOGNOMONIC EVIDENCE IS 0

X-ray evidence is evaluated to be -2 by use of a logic similar to that shown in the decision between EXTRA and INTRA earlier.

XRAY EVIDENCE IS -2
SUMMARY EVIDENCE IS -2
UNLIKELY SCL-CHOLANGITIS -2

EXPERT-CALLED: EXTRA-HEP CALLED-BY: SCL-CHOLANGITIS
MESSAGE (ANSWER (UNLIKELY SCL-CHOLANGITIS -2))

At this point, there is one refinement of EXTRA-HEP in ACTIVE list and CHOLESTASIS has to decide if all important data have been covered by the active hypotheses. This decision is currently made by the user.

ACTIVE: BD-CANCER EXTRA-HEP
UNLIKELY: SCL-CHOLANGITIS POST-OPER ANESTHETICS
UNKNOWN:

The problem of when to stop is still an open problem in diagnostic knowledge-based systems, though some approaches based on thresholding confidence measures have been tried[6,7]. We are working on a symbolic matching criterion for this purpose and preliminary design has been completed for explaining the major findings in extra-hepatic cholestasis(see [3] for details). In this case, as the CPC discussion shows[4], all the data have been explained. However, in order to demonstrate the rest of the system, let us say NO to the question:

ARE ALL DATA EXPLAINED BY ACTIVE DIAGNOSIS? \$N

We now are in IV in Procedure Cholestasis. LOOK-FURTHER hands over control to EXTRA-HEP, asking it to see if it can refine the diagnosis any further (but limiting its search by asking it to consider only the typical diseases). EXTRA-HEP tries its own heuristic rules, but no suggestions are made. It then falls back on an exhaustive search and calls* each typical expert under it, one by one. As it turns out, all return an unlikely evaluation.

EXPERT-CALLED: EXTRA-HEP CALLED-BY: CHOLESTASIS
MESSAGE (QUERY (REFINE EXTRA-HEP)
(CONSIDER TYPICAL))
TRYING RULES TO INVOKE HYPOTHESIS

* Details of calls and processing by these experts have been skipped.

TRYING RULE:
(((DEFORMITY? BIL-TREE BIOPSY TUMOR)) ((T -3)))
UNKNOWN:

140

TRYING RULE:

```
((EQUAL? JAUNDICE INTENSITY SEVERE)
(EQUAL? JAUNDICE TREND UNRELENTING))
((SUGGEST AV-CANCER)))
```

UNKNOWN

SUGGESTIONS IN EXTRA-HEP NIL

EXTRA-HEP returns control back to CHOLESTASIS with the result of its processing namely failure to establish any more concepts, but able to pretty much rule out all the other possibilities. The point of continuing the diagnosis, even though the correct diagnosis had already been made, was to demonstrate the role of suggestion rules at all levels in the conceptual hierarchy as a means of achieving efficiency. If bd-cancer had not been suggested at the level of CHOLESTASIS, it would have been established when control was handed over to EXTRA-HEP. The same effect occurs at all levels (including within tip-node experts, which usually have pathognomonic rules) in the hierarchy.

EXPERT-CALLED: CHOLESTASIS CALLED-BY: EXTRA-HEP
MESSAGE (ANSWER (REFINED EXTRA-HEP))

ACTIVE: BU-CANCER EXTRA-HEP

UNLIKELY: AV-CANCER AMP-CANCER GB-CANCER STRICTURE

STONE SCL-CHOLANGITIS POST-OPER ANESTHETICS

UNKNOWN:

ARE ALL DATA EXPLAINED BY ACTIVE DIAGNOSIS? \$YES

The diagnosis is summarized on a scale of -3(Reject) to 3(Definite).

EXPERT-CALLED: GP CALLED-BY: CHOLESTASIS
MESSAGE (ANSWER (ESTABLISHED CHOLESTASIS 3))

DISEASES CONSIDERED AND FOUND UNLIKELY

AMP-CANCER -2

SCL-CHOLANGITIS -2

GB-CANCER -2

POST-OPER BY-USER

STRICTURE -3

ANESTHETICS BY-USER

STONE -2

AV-CANCER -3

CONSIDERED AND UNABLE TO DECIDE

THE FINAL DIAGNOSIS IS ---

CHOLESTASIS 3

EXTRA-HEPATIC INDICATED WITH OBSTRUCTION AT (CHEP)

BD-CANCER 3

There are three important points to note about this example. First, the diagnosis is correct (though systems based on other approaches may have done as well in this case). Second, the diagnostic knowledge was used at the correct place, once the proper context had been established. For example, it would be premature and often disastrous to try to establish tip-level concepts like bile-duct cancer, unless cholestasis had been established and extra-hepatic indicated as more likely than intra-hepatic. Finally, the problem of coping with a large volume of data at a single level was avoided by using the patient data at the appropriate level of specificity and importance. In other words, the system tries to explain different types of data at those levels in the conceptual hierarchy where they are most relevant.

6.0 COMPARISON WITH RELATED SYSTEMS

6.1 INTERNIST

INTERNIST is perhaps the most comprehensive AI medical program: it deals with 80% of internal medicine. The diseases are organized in an organ-based hierarchy. Disease nodes are linked by causal and associative links. In addition, each datum (type) has "manifestation of" and "invoked by" links with disease

nodes (with some weights attached to these links).

INTERNIST performs diagnoses by calculating a numerical score for diseases based on how many manifestations are present, fraction of prototypical findings present, support from causal diseases and a measure of how much unexplained data is explained by a hypothesis. The disease hierarchy and a partitioning of current hypotheses into complementary and supplementary diseases helps to create some focus and cut down the complexity of processing. The basic problem-solving strategy is one of pursuing the highest scoring hypothesis in a global way. A new datum introduced at any time leads to a recomputation of all estimates. This could also lead to sudden shifts of focus. Much of the intelligence of the program is embodied in the complex scoring function. Thus explanation of diagnosis is conceptually opaque.

6.2 Present Illness Program (PIP)

This program is designed to suggest possible diseases in the area of renal diseases. Knowledge about clinical and physiological states (not all being diseases) is organized into frames. Each frame contains information about relationships with other states; triggers to invoke states and differential diagnosis; and a table of probabilistic information to estimate the likelihood of the state, based on patient data. There are also rules which can determine the likelihood of states or rule them out in some cases.

The basic strategy is similar to that of INTERNIST, i.e., the invoked hypotheses are ordered by their likelihoods and the top one at any given time is pursued. A hypothesis is accepted or rejected if it crosses a threshold. The likelihood estimate includes the support from complementary states, match with the prototypical findings, and the degree to which patient data are explained by the hypothesis.

Again, as in INTERNIST, a new datum leads to potentially expensive recomputation of all estimates, and possible focus shifts. Further the use of Bayesian estimation in a global way would add to the computational and focus problems. Discussion of some problems associated with Bayesian estimates can be found in [11] and [12]. As a last comment, the frames are interpreted in a uniform way. It would be difficult to use specialized knowledge to bear on the problem at the appropriate place, unless it can be cast in the general mold.

6.3 MYCIN

MYCIN advises clinicians in the area of infectious diseases. The knowledge is stored in the form of deductive rules, with some associated certainty factor. The program also has available to it a context hierarchy, which allows it to select rules based on the context and ask for appropriate data.

The rules are used in a backward-chaining deductive fashion, to first find the cause(s) of infection and then suggest therapy for it. Much of the real power of the system seems to us to come from the ability of the context hierarchy to use the rules in a controlled way. More generally, while rules are a convenient way to express certain kinds of knowledge, many other kinds of knowledge are not easily encoded in rule form (see [13] for a case study in this regard).

6.4 CASNET

CASNET performs diagnosis and recommends therapy in

glaucoma diseases. The knowledge is represented in a causal network of dysfunctional states with well defined starting (etiologic) and ending states. Evidence about the likelihood of these states can be obtained from another network relating states to tests. A test may be a patient datum or combinations of those. Diseases are defined as paths in the causal net.

The system calculates the status of each state, based on the test values. This computation is repeated as new data are entered. It also computes the weight of each node, which is independent of its status, but depends on those nodes which occur in the same path. Diseases are hypothesized and verified on the basis of which paths are most likely or potentially able to account for all confirmed states. There is a computational burden in calculating and recalculating the status and weight of each node in the network. Given the exponential nature of such computation, it seems unlikely that such an approach would be easily extended to larger domains. Further, the use of patho-physiological states is, perhaps, not possible for most of medicine, as such mechanisms are known only spottily. A relatively minor problem is created by the fact that evidence from different tests can be combined by creating a "higher" test-entity. In many medical situations, where patient data combine in different ways to support or disprove disease hypotheses, many new test-entities would need to be created for such combinations.

7.0 CONCLUDING REMARKS

While the preceding pages can be viewed as a partial attempt at uncovering the structure of human cognition in the diagnosis task, the real test of the ideas proposed is in how well the system based on them works. In addition to the case presented in an earlier section, several other extra-hepatic cases (both from medical journals as well as from a practicing hepatologist) have been tried, and in all cases the diagnostic behavior was natural and closely corresponded to the physicians' reasoning. The system is constantly being modified, but not in the sense of "patching" but to increase the quality of the knowledge. For instance, it has been found useful to place certain rules higher in the hierarchy. In some cases, the table logic of some tin nodes has been found inaccurate, and, after consultation with medical texts and experts, modifications have been made. These seem to us to be natural ways of improving the power. Of course, the modifications to local procedures should converge eventually. We cannot say that the system has yet passed this test.

We envisage the next stage in our work to be enlarging the system to handle the whole of the Cholestasis syndrome. We hope to study the effects of upward scaling in this enlargement. For one thing, a more powerful data base organization for temporal phenomena will be needed. However, we shall not opt for a separate temporal expert, as this would violate our constraint on control transfers. Instead temporal knowledge will be distributed as needed.

Our current approach consists of giving the system all the data that are available in the journal description of the case or from the collaborating hepatologist. In either case, the case is in a sense "complete." In reality, however, as the diagnostic process unfolds, tests are ordered to confirm or reject hypotheses. Thus the ability to order tests would be a useful extension

to our system. The conceptual information needed is partially implicit in the knowledge structure (queries to the data base can be converted to orders for tests). Similarly, since the inter-expert messages contain purposes and responses, the basis for an explanation capability is already present. However, more work is needed to realize these facilities in an elegant manner, i.e., without too much detail and with appropriate abstractions. Explanation and test-ordering facilities would enhance the use of the system as a consultant.

In summary, we have attempted, in this paper, to lay the foundations for a knowledge representation scheme, which derives its power because of its ability to capture the structure of a body of knowledge at an appropriate level of depth. There are many interesting issues related to conceptual structures we have not had occasion to discuss in this paper. These and our experiences with problem-solving systems for larger domains of knowledge will be the subject of future reports.

ACKNOWLEDGMENT: We are indebted to Douglas Levin, M.D., for his interest and consultations.

REFERENCES

- [1] Minsky, M. "Plain Talk about Neurodevelopmental Epistemology." In Proc. IJCAI-77. MIT, Cambridge, Mass., August, 1977, pp.1083-1092
- [2] Hayes-Roth, F., Lesser, V.R. "Focus of Attention in the Hearsay- II Speech Understanding System." In Proc. IJCAI-77. MIT, Cambridge, Mass., August, 1977, pp.27-35
- [3] Mittal, S. Ph.D. Dissertation (forthcoming), Dept. of Comp. and Info. Sc., The Ohio State University, Dec., 1979
- [4] "Case 45-1977, Massachusetts General Hospital." New England Journal of Medicine. Nov.10, 1977, pp.1054-1059
- [5] Mittal, S., Chandrasekaran, B. "Conceptual Representation of Patient Data Bases." OSU-TR-79, Dept. of Comp. and Info. Sc., The Ohio State University, April, 1979
- [6] Pople, H.E., Myers, J.D., Miller, R.A. "DIALOG: A Model of Diagnostic Logic for Internal Medicine." In Proc. IJCAI-75. Tbilisi, USSR, Sept., 1975, pp.848-855
- [7] Shortliffe, E.H. Computer Based Medical Consultations: MYCIN Elsevier North-Holland Inc., 1976
- [8] Pople, H.E. "The Formation of Composite Hypothesis in Diagnostic Problem Solving: An Exercise in Synthetic Reasoning." In Proc. IJCAI-77, MIT, Cambridge, Mass., August, 1977, pp.1030-1037
- [9] Pauker, S., Gorry, G.A., Kassirer, J., Schwartz, W. "Towards a Simulation of Clinical Cognition." The American Journal of Medicine 60(1976)981-996
- [10] Weiss, S.M., Kulikowski, C.A., Amarel, S., Safir, A. "A Model-based Method for Computer-Aided Medical Decision-making." Artificial Intelligence. 11:2(1978) 145-172
- [11] Szolovits, P., Pauker, S.G. "Categorical and Probabilistic Reasoning in Medical Diagnosis." Artificial Intelligence. 11:2(1978) 115-144
- [12] Feinstein, A.R. "Clinical Biostatistics XXXIX. The haze of Bayes, the Aerial Palaces of Decision Analysis and the Computerized Ouija Board." Clinical Pharmacology and Therapeutics. 21:4(1977) 482-490
- [13] Reggia, J.A. "A Production Rule System for Neurological Localization." In Proc. Second Annual Symposium on Computer Applications in Medical Care. November, 1978, pp.254-260