Information Acquisition in Diagnosis

Ramesh S. Patil, Peter Szolovits
Massachusetts Institute of Technology
Laboratory for Computer Science
545 Technology Square
Cambridge, Massachusetts 02139

and

William B. Schwartz, M.D.,
Dept. of Medicine, Tufts University School of Medicine,
and the Medical Service, New England Medical Center,
Boston, Massachusetts

Abstract

We argue that the question selection processes used in the existing *AI in Medicine* programs are inadequate. We trace these inadequacies to their use of purely surface level models of disease and to the lack of planning in sequencing their inquiries. We then describe the information gathering component of ABEL (a program for Acid-Base and Electrolyte disturbances) with emphasis on its use of a causal model of a patient's illness in guiding its diagnostic inquiries.

1. Background

Diagnostic programs built using the *AI in Medicine* (AIM) methodology typically consist of two parts: one to form hypotheses that explain the data known to the program, and one to acquire new information from the program's users to clarify difficulties remaining with the hypotheses under consideration. In this paper we suggest that the questioning behavior of most existing diagnostic programs is inadequate because it is not very responsive to the totality of the program's hypotheses. Both rule-based systems such as Mycin [Shortliffe76] and frame-based systems such as Internist-1 [Pople77] and PIP [Pauker76] suffer from this defect, though for diverse reasons.

Mycin [Shortliffe76] is a rule-based expert program for the diagnosis and therapy of infectious diseases. It uses a backward chaining control structure to sequence through its rule-base in search of information relevant in the diagnosis of possible infections. Mycin always pursues the antecedent clauses of a rule in their order of appearance in the rule. Thus during any execution of the program the antecedent clauses of every rule are expanded in exactly the same order, resulting in a pre-determined order in which the questions may be asked.[2] This rigidity in Mycin results from its use of rules to encode both its medical knowledge and the control component of expert reasoning. The resulting limitations of the Mycin system were recognized by Davis [Davis76], who proposed meta-rules as a means of separating these two aspects of expert behavior.' Meta-rules are used in his system, Teiresias [Davis76], to encode strategic knowledge for the control of the backward chaining rule interpreter by dynamically re-ordering and selecting rules to be pursued. It is important to note that the meta-rules do not alter the program's domain knowledge; they provide additional knowledge that allows it to make more efficient use of this knowledge.

1. This research was supported (in part) by the National Institutes of Health Grant No. 1 P01 LM 03374-03 from the National Library of Medicine. Dr. Schwartz's research was also supported by the Robert Wood Johnson Foundation, Princeton, New Jersey.

2. Mycin's reasoning scheme chooses not to ask a question only if the rule whose antecedent part suggests it is already doomed to fail by answers to previous questions or if the question has been previously answered because it appeared in the condition of another rule.

Internist-I, a diagnostic program being developed for all of internal medicine [Pople77], and the Present Illness Programs (PIP), a program for the initial formulation of a diagnosis for patients with diseases of the kidneys [Pauker76], both focus on a problem to investigate next by considering all their active hypotheses. Therefore, control over which disease or disease set is next to be explored does not depend on the manner or order in which these are expressed in the program. However, once the controlling hypothesis set is determined, the information gathering strategies are limited to selecting one pre-packaged group of questions at a time. E.g., each time the program chooses to try to confirm some particular hypothesis, it will run through the same list of associated questions, in the order in which they appear in the program's static knowledge base without regard to what else the program knows. After a group of questions is asked, the program's hypothesis formation module takes control again, and whether this line of questioning is resumed or abandoned depends on the process of hypothesis evaluation and whether the same hypothesis set again rises to the top of the program's interest. Often too little control is thus exercised by the information acquisition part of the program, and overall diagnostic inefficiencies and incoherent question sequences result.

To study the value and use of multi-level causal descriptions in diagnostic and therapeutic reasoning, we have been developing a new medical expert consultation program for the domain of acid-base and electrolyte disorders, called ABEL. The information-gathering component of ABEL uses the program's causal model of what it believes might be wrong with the patient to help select and order possible questions to be asked (thus capturing the advantages of meta-rules used for this purpose in Mycin), and unlike Internist-I or PIP it constructs a rather detailed information gathering plan tailored to the specific set of hypotheses it must confirm or differentiate among. This approach gives ABEL an improved control over the sequencing of questions that are eventually asked of the user.

We have noted elsewhere [Patil81] that most of the medical knowledge contained in the first generation AIM programs can be characterized as phenomenological: it describes the associations among phenomena without the causal mechanisms underlying the observed associations. Therefore, unlike expert clinicians, these programs are unable to recognize and use causal and temporal relationships among diseases and to reason about the disease mechanisms at various levels of detail. They are unable to evaluate the interactions, commonalities and differences among diseases, except perhaps in terms of some simple measure of relative likelihood. ABEL is intended, by contrast, to provide multiple levels of description of its hypotheses, allowing it to use both the phenomenological associations of a shallow level and the detailed causal mechanisms of deeper levels in formulating hypotheses and in planning the gathering of new information.

In a previous paper [Patil81a] we have presented ABEL's mechanism for describing a patient's illness. Called the *patient-specific model* this description includes data about the patient as well as the program's hypothetical interpretations of these data in a multi-level causal network. In this paper we focus on the diagnostic information-gathering process of the ABEL program which complements the patient-specific model.

2. The information acquisiton problem

The overall diagnostic activity is concerned with developing a clear understanding of the disease process(es) responsible for the patient's illness, for use in prognosis and therapy. It involves interpreting available information about the patient to form diagnostic hypotheses which might explain the patient's illness and acquiring new information that will help discriminate among these hypotheses. The information acquisition activity must take into account factors such as the cost, efficiency and urgency of acquiring information. In addition, as it interacts with the user, it must also take into account factors such as topical coherence and the possibility of errors.

Previous programs e.g., Internist-I [Pople77] have employed strategies such as *discriminate, pursue* and *rule-out* to help select a question or group of questions to ask next. ABEL uses such strategies instead to decompose the overall diagnostic problem of resolving uncertainties in the hypotheses into sub-goals each of which is more specific and amenable to further decomposition. Thus, rather than generating questions directly from the top-level diagnostic problem, ABEL generates a tree of sub-problems the leaves of which can be solved by asking one or a small group of questions. Because the whole set of these questions is methodically selected before any question is actually asked, interdependencies among the questions can be taken into consideration. For example, ABEL currently organizes all the planned questions by the organ system, etiology, or time period on which they concentrate, so that groups of questions are asked in a coherent order. The number of places in the sub-goal tree where the answer to a question is thought to be useful is also used to efficiently order the sequence of questions.

Erroneous information is bound to be presented to the program during a diagnostic session. Therefore, an ability to identify questionable information and to challenge and correct it quickly is an important ingredient of clinical expertise. When presented with a finding which, if accepted, may require extensive reformulation of the currently held diagnoses, it may be unwise for the program to act on such information unless it can be substantially corroborated and its validity as a diagnostic sign checked out. For example, upon unexpectedly (and erroneously) finding "a substantial weight increase" in a patient over a short period of time, the program should check if the two weights were taken on the same scale. It is unreasonable, however, to ask the same question every time a weight change is reported for a patient. In ABEL we associate with each information gathering goal expectations about its possible outcomes. After each question is asked, if the answer is consistent with expectations, questioning continues with the planned sequence. However, if the answer contradicts the expectations, this information is considered questionable and an *excuse-finding* mechanism is activated. This allows ABEL to pursue questionable findings further before accepting them.

In the remainder of the paper we will describe the workings of the information gathering activity of ABEL with the help of an example.

3. An Example

Let us consider a patient who has been ill for 3 to 4 days and is known to have moderately severe metabolic-acidosis and slight hyponatremia (serum Na of 128 meq/L). Let us also assume that no additional history is available. Two possible formulations of the patient's problem are shown in Figs. 1 and 2. One hypothesis states that the underlying disorder is diarrhea, the other, that it is acute renal failure. The

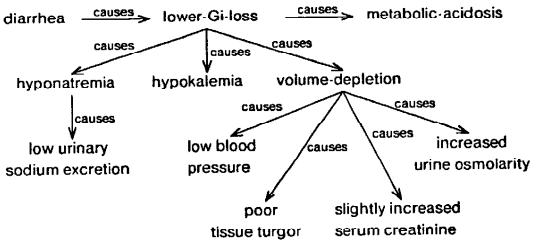

Fig. 1. In one explanation of the example case, diarrhea is lower gastro-intestinal losses, which result in metabolic-acidosis, volume depletion, and their consequences.

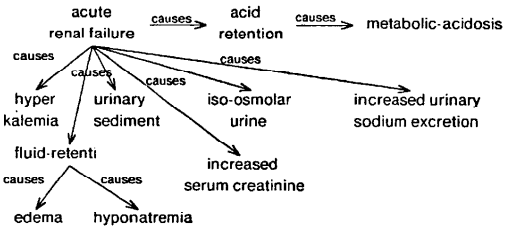

Fig. 2. An alternative explanation to that of Fig. 1 holds acute renal failure responsible for acid retention, which causes metabolic-acidosis and its consequences.

```
Goal 1: Differentiate diarrhea, acute renal failure
    Context: Cause-of metabolic acidosis
    Expectations:
        Possible: diarrhea
            Severity: moderate
        Possible: acute renal failure
            Severity: moderate
            Duration: a few days
    Subgoals: (AND 2 3 8 9 10)
```

Fig. 3. Top-level goal structure for discriminating between diarrhea and acute renal failure.

program has set as its top level goal, as shown in Fig. 3, the desire to discriminate between these two possible interpretations. To accomplish this objective the program compares the two interpretations, identifying the differences between the states predicted by the two interpretations, and formulates sub-goals to pursue each difference. For example, the program identifies urinary sodium concentration as a useful differentiator between diarrhea and acute renal failure, because diarrhea predicts that the urinary sodium concentration will be low whereas renal failure predicts a relatively high urinary sodium concentration. Similarly, the program can differentiate between diarrhea and acute renal failure by determining the state of hydration of the patient; this goal can be achieved by confirming either volume depletion or edema. Diarrhea predicts the loss of fluid and therefore volume depletion. Volume depletion, however, can not be directly

observed; therefore, the program further decomposes this goal into sub-goals for confirming poor tissue turgor, low blood pressure and a slightly elevated serum creatinine concentration. Acute renal failure, on the other hand, predicts the accumulation of body fluids if normal intake of fluid has continued during the period of oliguria. If a sufficient accumulation occurs, it will manifest itself as edema. A graphic representation of the complete goal structure is shown in Fig. 4, and the goals are listed below it. The program has now completed the generation of the sub-goal tree for differentiating between diarrhea and acute renal failure. It can group and order the questions by efficiency and stylistic measures. For example, it groups questions about the

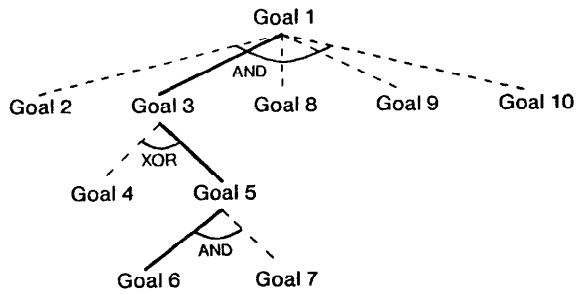

Goal 2: explore urinary sodium concentration
Context: differentiate diarrhea, acute renal failure
Expectations:
Possible: low (less than 10 meq/l)
Cause: diarrhea
Possible: high (greater than 40 meq/l)
Cause: acute renal failure

Goal 3: explore state of hydration
Context: differentiate diarrhea, acute renal failure
Expectations:
Possible: volume depletion
Cause: diarrhea
Severity: moderate
Possible: fluid retention
Cause: acute renal failure
Default: continued normal fluid intake
Severity: mild to moderate
Subgoals: (XOR 4 5)

Goal 4: confirm fluid retention
Context: caused by acute renal failure
Expectations:
Possible: edema
Severity: mild to moderate
Possible: no edema

Goal 5: confirm volume depletion
Context: caused by diarrhea
Expectations:
Possible: present
Severity: mild to moderate
Possible: absent
Subgoals: (AND 6 7)

Goal 6: confirm poor tissue turgor
Context: explore state of hydration
Expectations:
Possible: mild
Cause: volume depletion
Possible: absent
Cause: fluid retention

Goal 7: explore blood pressure
Context: explore state of hydration
Expectations:
Possible: slightly low
Cause: volume depletion
Possible: normal
Cause: fluid retention

Goal 8: confirm hemoglobin and tubular cell casts in urine
Context: differentiate diarrhea, acute renal failure
Expectations:
Possible: present
Cause: acute renal failure
Possible: absent
Cause: diarrhea

Goal 9: explore serum creatinine
Context: differentiate diarrhea, acute renal failure
Expectations:
Possible: slightly to moderately elevated
Cause: acute renal failure
Possible: slightly elevated
Cause: diarrhea

Goal 10: explore serum K
Context: differentiate diarrhea, acute renal failure
Expectations:
Possible: increased
Cause: acute renal failure
Possible: low
Cause: diarrhea

Fig. 4. The complete goal structure. AND indicates goals that must be simultaneously achieved and XOR indicates goals of which only one can be achieved. The solid arcs in the goal structure represent the path actually being taken by the information-gatherer.

results of urine analysis (namely urinary sodium concentration and urine sediment) to ask first, general physical questions relating to the state of hydration (i.e., edema, tissue turgor and blood pressure) next, and serum creatinine and serum K concentrations last. The following is a summary of the information gathered by the program in pursuing this goal structure.

Urinary sodium concentration: 50 meq/l
Urine sediment: negative
Edema: absent
Tissue turgor: slightly reduced
Blood pressure: normal
Serum Creatinine concentration: 2.5 mg per cent
Serum K concentration: 3.5 meq/l

After successfully achieving each of the sub-goals, the program reconsiders the top-level goal of differentiating between diarrhea and acute renal failure. The program realizes that the overall set of findings are not consistent with one another. In particular, the finding of high urinary sodium concentration suggests acute renal failure, whereas the low serum K concentration is inconsistent with acute renal failure. On the other hand, low serum K concentration is consistent with diarrhea, whereas high urinary sodium concentration is not. This conflict activates the program's excuse mechanism (see below). To resolve the contradiction, the program sets up two goals so that one of the two diagnoses can be confirmed. The two goals are shown in Fig. 5. Upon pursuing these goals the program finds that the patient has been vomiting, which explains the slightly low serum K concentration and a slight volume depletion. The high urinary sodium excretion argues strongly in favor of acute renal failure.

Goal 11: Find excuse for high urinary sodium concentration
context: diarrhea
expectation:
possible: diuretic use
possible: Addison's disease

Goal 12: Find excuse for low serum k
context: acute renal failure
expectation:
possible: diuretic use
possible: vomiting

Fig. 5. Additional goals created to resolve apparent contradiction in input data.

After completing a cycle of information gathering the program enters this information into the patient-specific models, revises its hypotheses for the patient's illness and continues the diagnosis.

4. Summary and Conclusion

In this paper we have briefly sketched the information acquisition method used in the ABEL program. We have proposed a novel way of viewing the traditional diagnostic strategies as decomposition operators. We have then used these strategies to decompose an overall diagnostic goal into a tree of sub-goals the leaves of which can be solved directly by asking questions. In addition we have proposed attaching to each diagnostic goal the expectations of the program about its possible outcomes. Because all the questions useful to the overall diagnostic goal are methodically generated before any question is asked, the program can order these questions for efficiency and coherence. Because the possible outcomes of pursuing each goal are determined before questions relating to it are asked, the program can easily determine the success or failure of its goals and can identify questionable information that contradicts all possible expectations. This allows the program to make efficient use of its case-specific knowledge and exert better control in its information acquisition activity.

5. References

[Davis76]
Davis, R., Applications of Meta Level Knowledge to the Construction, Maintenance and Use of Large Knowledge Bases. Ph.D. thesis, Stanford Artificial Intelligence Laboratory, Technical Memo 283 (1976).

[Patil81]
Patil, R. S., Szolovits, P., and Schwartz, W. B., Causal Understanding of Patient Illness in Medical Diagnosis, *Proceedings of the Seventh International Joint Conference on Artificial Intelligence*, available from American Association for Artificial Intelligence, 445 Burgess Drive, Menlo Park, Ca 94025 (1981).

[Patil81a]
Patil, R. S., Causal Representation of Patient Illness for Electrolyte and Acid-Base Diagnosis, Ph.D. Thesis, Laboratory for Computer Science, Massachusetts Institute of Technology, Technical Report MIT/LCS/TR-267 (1981).

[Pauker76]
Pauker, S. G., Gorry, G. A., and Kassirer, J. P., and Schwartz, W. B., Toward the Simulation of Clinical Cognition: Taking a Present Illness by Computer, *The American Journal of Medicine* 60:981-995 (1976).

[Pople77]
Pople, H. E., Jr., The Formation of Composite Hypotheses in Diagnostic Problem Solving: an Exercise in Synthetic Reasoning, *Proceedings of the Fifth International Joint Conference on Artificial Intelligence*, available from the Department of Computer Science, Carnegie-Mellon University, Pittsburgh, PA 15213 (1977).

[Shortliffe76]
Shortliffe, E. H., *Computer Based Medical Consultations: MYCIN*, Elsevier North Holland Inc. (1976).