

Opportunistic Use of Schemata for Medical Diagnosis*

Roy M. Turner
School of ICS
Georgia Institute of Technology

Medical diagnosis can be considered a planning task. This is not the traditional view, however. For example, Gomez and Chandrasekaran (1982) and others view diagnosis as a classification task: a problem, consisting of a set of signs and symptoms, is classified as being an instance of a disease or set of diseases. However, this viewpoint overlooks the fact that actions are performed in order to classify a disease: in other words, planning and plan execution must be done as part of the classification process. When viewed as planning, goals in diagnosis are such things as "diagnose the patient," "interpret a finding," and "evaluate a hypothesis." Operators, at the lowest level, are such things as asking questions, requesting tests be performed, and making inferences based on information known about the patient and the reasoner's general knowledge of the domain.

Medical diagnosis is unlike many traditional planning tasks in that an initial, complete statement of the problem is generally impossible. Instead, the diagnostician must gather information about the problem as part of the process of performing diagnosis. The result of this is that the diagnostician cannot formulate a plan for diagnosis, then carry it out: the problem statement would change as the plan for performing diagnosis is executed. The effect of executing one step (e.g., asking a question) would likely alter the assumptions upon which later steps are based (e.g., a new finding might radically alter the diseases considered as diagnoses, or might suggest specialized methods for interpreting the finding). The problem for a diagnostician, then, is to be able to interleave planning and execution (cf. McDermott, 1978) so as to make use of new information as it becomes available. In other words, a diagnostician should be *opportunistic*.

Our approach to this problem makes use of packets of procedural information called *schemata*, which are retrieved from memory in response to goals arising from changes in the problem solver's environment: e.g., new findings, new hypotheses, etc. Most schemata can achieve very specific goals, such as "interpret a finding" or "evaluate a hypothesis"; others control larger parts of the reasoner's processing, such as directing the reasoner in the overall consultation. Schemata are flexible enough to encode several variations of how to achieve their goal; in addition, specializations of schemata provide the reasoner with information about how to satisfy specific goals or goals arising in specific contexts.

When a goal arises that can be achieved by a schema, that schema is retrieved from memory and made active. As the reasoner may have many goals simultaneously, there may be many active schemata at any time. The reasoner must decide which goal to focus on, and hence, which schema to *apply*. In our approach, the reasoner uses information from two sources to help it focus its attention. One source is from memory structures representing generalized consultations similar to the current problem. Information from these generalized consultations, such as information about which findings are generally important in this context, can be used by the reasoner to help it select a goal to achieve. The second source is from packets of procedural knowledge, called *strategic schemata*, which contain generally useful strategies in the form of goal orderings: e.g., a medical reasoner would have strategies for performing hypothetico-deductive reasoning, reasoning under time pressure, etc.

In this paper, we discuss our approach to opportunism in medical diagnosis. Our approach is called *schema-based reasoning*. Our ideas are being tested in MEDIC (Turner, in press), a schema-based diagnostic reasoner whose domain is pulmonology.

*This research has been funded in part by NSF Grants IST-831771 and IST-8608362 and grant DTD 09-25-87 from the Lockheed AI Center.

OPPORTUNISM USING SCHEMATA

Opportunism involves responding to changes in the task environment as they arise during problem solving. There are at least three capabilities a reasoner must have in order for it to respond to changes; it must be able:

1. to interleave planning to achieve a goal and execution of that plan;
2. to respond immediately to new information appearing in the environment; and
3. to select the appropriate goal to pursue at each point in problem solving—that is, it must be able to focus its attention.

Traditional planners do not interleave planning and execution. Instead, the planner formulates a plan, then applies it. On the other hand, rule-based problem solvers and purely reactive planners such as PENGI (Agre & Chapman, 1987) do not really perform planning *per se*. The problem with these approaches is that there is very little coherence in their actions; consequently, their behavior may seem strange and unintuitive to a user. This presents a problem, especially in a medical domain, since a user is unlikely to accept a system if he or she cannot understand its reasoning.

A middle ground is needed between traditional planners and purely reactive planning. In our approach, problem solving is carried out by retrieving packets of procedural knowledge from memory, then applying them. These packets, or *schemata*, can be thought of as small plans or pieces of plans that achieve a goal; for instance, a reasoner may have a schema which can interpret a finding or one that can evaluate the likelihood of a hypothesis that a particular disease is present. Figure 1 shows a simplified view of a schema for interpreting a finding of dyspnea.¹ A schema contains *steps* to be performed by the reasoner in order to satisfy a particular goal.

Our approach is more flexible than traditional planning for three reasons. First, the order of the steps of a schema is not completely fixed ahead of time, but rather depends, to some extent, on the situation at the time of schema execution. For example, the step labeled "S1" contains information that allows the reasoner to select the next step based on the answer to the question asked in S1.

¹Shortness of breath.

```

Goal: interpret a finding of dyspnea
Patient: any patient
Findings: dyspnea
Preconditions: there is a finding of dyspnea
Steps:
  S1: action: ask how many stairs patient can climb
      goal: determine severity of dyspnea
      next:
        if pt. can climb flight of stairs => S3
        else => S2
  S2: action: ask how far patient can walk
      goal: determine severity of dyspnea
      next: S3
  S3: action: estimate the severity of the dyspnea
      goal: determine severity of dyspnea
      :
  S10: action: postulate hypotheses of pulmonary disease,
        cardiac disease
        goal: explain dyspnea
        next: done
Indices:
  patient/PATIENT1 -> SCENE2
  :

```

Figure 1: *sc-dyspnea*—a schema for interpreting a finding of dyspnea.

The second source of flexibility in our approach is also due to the nature of a schema's steps. In addition to specifying actions that should be taken—either primitive actions or other schemata—a step in a schema usually specifies the goal that the step is to satisfy. If the step fails, the reasoner can attempt to find another way of satisfying the step at run time. In addition, a step does not necessarily specify an action. Instead, it can specify *only* a goal, thus forcing the reasoner to attempt to satisfy the goal at run time.

The third reason our approach is flexible is that a single plan is not formulated for all of the goals in a problem, then executed. Instead, individual schemata are retrieved and applied to satisfy goals. As the situation changes, which schemata are active will also change. For example, as new goals arise during problem solving, new schemata can be found and activated to satisfy them.

In order to exhibit opportunism, a reasoner must be able to notice and respond to new information as it becomes available. In the context of a diagnosis program, new information comes from the user; information may be volunteered, or it may come from answers to questions asked by the system. In either case, when new information becomes available, the reasoner should interrupt what it is doing and incorporate the information into what it knows; the new information may also cause the reasoner to alter the course of its problem-solving behavior.

In our approach, all actions are governed by information contained in schemata; thus, to handle a piece of information, a schema must be found and activated. When a new item of information is encountered, the reasoner interrupts what it is doing and looks for a schema that can satisfy the goal created by the occurrence of the information: e.g., if the information is a finding, the goal will be to interpret it. When a schema is found, it is activated. The reasoner can then either return to what it was doing (i.e., to the schema it was applying), or it can choose, based on the altered state of the environment, to pursue a different goal (i.e., to apply a different schema—perhaps the one just activated).

There may be many goals present for a given diagnostic problem: e.g., goals to interpret findings, evaluate hypotheses, and to produce a diagnosis. Each of these will lead to one or more schemata being activated. In addition, information learned during diagnosis will result in more schemata being activated, as discussed above. Once we allow the reasoner to have more than one schema active at a time, we face the need for ability (3) above: the reasoner must be able to select the appropriate schema to apply at each point in problem solving;² i.e., the reasoner must decide on which goal to focus its attention.

This is not an easy task, since the importance of a goal varies with context. For example, in one situation, a goal to explain a severe rash may be quite important; however, in another situation involving both a severe rash and hemoptysis,³ the goal of explaining the hemoptysis should take precedence over the goal of explaining the rash.

One approach to this problem is to use knowledge about the type of consultation the reasoner is performing. This information, in our approach, is present in memory structures called *diagnostic memory organization packets*, or dxMOPs (cf. Schank, 1982; Kolodner, 1985; and the diagnostic categories of Kolodner and Kolodner, 1987). These structures participate in memory organization, and provide one way for the reasoner to retrieve schemata from memory (see (Turner, in press)

²We do not allow schemata to be applied in parallel. This is because we are trying, as far as possible, to model the behavior of human diagnosticians, who generally act as though they are thinking of one thing at a time. There are two reasons for modeling humans: (1) if the reasoner behaves similarly to a human, then the user is more likely to understand its reasoning, and hence, accept it; and (2) reasoning in a manner similar to that of a human should make explanations easier (though we do not currently address explanation).

³Blood in the sputum.

```

Goal: diagnose the patient
Patient: an alcoholic patient
Chief complaint: dyspnea
Findings:
  anemia: low importance, explained by
           alcoholism
  ataxia: low importance, explained by
           alcoholism
  :
Hypotheses: TB, sarcoid, generalized pulmonary or heart
             disease
Schemata:
  sc-consult: for goal of diagnosing patient
  sc-finding: for generic findings
  sc-TB: for evaluating TB
  :
Indices:
  finding/mass-on-X-ray → dxMOP3
  :
  patient/PATIENT4 → dxMOP2
  :

```

Figure 2: A dxMOP for consultations involving alcoholics with dyspnea.

for details). Each dxMOP represents a particular class of consultations, and provides information about goals, actions, etc., that can be expected in such a consultation. Information from a dxMOP can be used to decide which goal—and hence, which schema—to pursue. For example, suppose the dxMOP is the one shown in Figure 2; this dxMOP represents consultations involving alcoholics whose chief complaint is dyspnea. Among the expected findings are anemia and ataxia,⁴ which are both explained by alcoholism. If the reasoner discovers that the patient indeed has anemia, it would not need to follow it up, since the finding is anticipated by the dxMOP and marked with a low importance.

Information contained in a dxMOP can help the reasoner order goals; but what if there is no such information present in the dxMOP retrieved from memory? Or what if there are constraints associated with the current problem that are not anticipated in the dxMOP, such as time being limited?

In order to focus the reasoner's attention in this type of situation, we use the idea of meta-reasoning (e.g., Davis & Buchanan, 1984): reasoning that takes place to guide planning. In our approach, meta-reasoning information is present in the form of strategic knowledge structures called *strategic schemata*. A strategic schema is a packet of information which represents a strategy for the rea-

⁴Poor motor coordination.

Situation: any
 Goal ordering:
 select goals related to hypotheses
 select goals related to findings
 select goals for gathering information
 select goal for forming diagnosis

Figure 3: A simple strategy for hypothetico-deductive reasoning style.

Situation: time is short
 Goal ordering:
 select goal related to chief complaint
 select goals related to hypotheses
 select finding goal only if very important
 select goals of forming diagnosis

Figure 4: A simple strategy for reasoning under time pressure.

soner's behavior; since the reasoner's behavior is determined by which schemata it chooses to apply, a strategy is equivalent to an ordering of the schemata applied. Strategic schemata are useful for representing general strategies such as hypothetico-deductive reasoning or reasoning under time pressure. The reasoner can then use information, if available, from a dxMOP to modify its use of these strategies for a specific situation.

A simple example of a strategic schema is shown in Figure 3. This schema provides a goal ordering for the reasoner that induces a crude form of hypothetico-deductive reasoning: select any goals (i.e., select their corresponding schemata) that relate to hypotheses first; if there are none, then select goals related to findings in the hope of producing hypotheses; if none, then select goal of gathering information from the user; and finally, if that cannot be done, select the goal of forming a diagnosis. Figure 4 shows a simple strategic schema for reasoning under time pressure: follow up the chief complaint, if possible; evaluate hypotheses; only select goals related to findings if the findings are very important; and finally, when all else fails, form a diagnosis.

The overall algorithm for schema-based reasoning, including opportunism, is shown in Figure 5. Note that the reasoner can be interrupted; these interruptions can include interruptions both by the user and by schema application, as new information is added to STM.

```
begin
  loop forever:
    Wait until user requests a consultation;
    Add goal of diagnosing patient to short-term
      memory;
    Retrieve dxMOP using goal;
    Use strategy from dxMOP, if possible;
    Select a schema from the dxMOP to satisfy
      goal, add it to agenda;
  loop until done:
    Select a schema from agenda using strategies,
      local information in the dxMOP;
    Apply one action;
    if there was an interruption then:
      Handle interruption;
    fi;
    Specialize current dxMOP;
    if specialization succeeded then:
      Set current dxMOP to be the
        specialization;
    fi;
  end loop;
  Accept and process feedback;
  Update memory;
end loop;
end.
```

Figure 5: Basic schema-based reasoning algorithm.

MEDIC

Our approach to diagnostic reasoning is being implemented in the MEDIC program, a schema-based reasoner which performs diagnosis in the domain of pulmonology. MEDIC consists of three major modules: a long-term memory, which is organized as described in (Turner, in press); a short-term memory (STM); and a schema-based reasoner, which is directed at all times by schemata.

Conceptually, there are three types of schemata in MEDIC's memory: global, local, and strategic schemata. A *global schema* is one that directs a major portion of a consultation; an example is the schema which contains information that the reasoner can use to conduct the consultation: ask for a patient description, ask about the chief complaint, gather information, then form a diagnosis. Gathering information and forming a diagnosis are also directed by global schemata. A *local schema* is one which directs the reasoner in achieving very specific goals: e.g., interpret a finding of dyspnea or evaluate the hypothesis of lung cancer. Strategic schemata represent general reasoning strategies and are described above.

Currently, MEDIC can diagnose very simple cases of pulmonary disease. MEDIC follows an algorithm very similar to that in Figure 5. Let's look at an

```

Please describe the patient.
: (patient (sex female) (weight (value 204)) (height
  (value 64)) (race white))
Adding information about patient to STM.
What is the chief complaint?
: (finding (entity (dyspnea (duration (years 2))
  (character progressive))))
Adding chief complaint to STM...adding finding of
  <DYSPNEA0> to STM.
How many flights of stairs can the patient climb?
: (less-than 1)
How far can the patient walk on level ground?
: (yards 20)
I judge the qualitative value of SEVERITY of <DYSPNEA0>
  to be SEVERE.
  ... (same for cardiac disease)...
...explaining dyspnea...
Processing <HYPOTHESIS0> [pulmonary disease];
  relating to other hypotheses...
...generating expectations given <HYPOTHESIS0>...
...I'm scoring hypothesis <HYPOTHESIS0> (<PULM-DZ0>)
...hypothesis explains: (<FINDING0>) [dyspnea]...
...failed predictions for hypothesis: —
...hypothesis doesn't explain: —
...trying to specialize the hypothesis of
  <HYPOTHESIS0> (<PULM-DZ0>)...
...specialized <HYPOTHESIS0> to <HYPOTHESIS1>
  (<RPE>) [recurrent pulmonary embolism]
...generating expectations given <HYPOTHESIS1>...
  ...
Is there a finding of <SYNCOPE>?
: Yes
  ...
Enter information (<return> if no more).
:
  ...
My diagnosis is: Recurrent pulmonary embolism.

```

Figure 6: Part of a consultation with MEDIC.

example of a consultation with MEDIC, a portion of which is shown in Figure 6. Suppose a user requests a consultation. The reasoner looks in memory for a way of satisfying the goal of diagnosing a patient and finds a dxMOP, "dx-consult", representing how consultations are generally conducted. The reasoner then uses this dxMOP as a context for diagnosis: it is used as a source both of a strategy and of schemata to satisfy active goals. The strategy it contains is "st-HD-reasoning", the strategic schema mentioned above which provides a goal ordering to induce hypothetico-deductive reasoning. The only goal active is one to diagnose the patient; the schema to achieve this in dx-consult is "sc-consult". This is added to the reasoner's agenda of active schemata.

The reasoner now selects a schema from its agenda, using the goal ordering provided by the current strategy; the use of specific information from the dxMOP is not currently implemented. The only schema to select is sc-consult, so the reasoner selects that and begins to apply it. The user is asked for some initial information about the patient, including a description of the patient (a white female who

is overweight)⁵ and the chief complaint (progressive dyspnea). The information is added to STM. Adding the chief complaint causes the reasoner to be interrupted, and it searches memory for a schema to interpret the finding. Schema "sc-dyspnea" is found and activated.

Note the opportunistic nature of this: a piece of information appears in the reasoner's environment, which causes the reasoner to interrupt what it is doing in order to respond. The response involves finding a method of interpreting the finding and activating it, resulting in more than one active schema. The next step is to choose one to apply. Since the strategy in use dictates that goals related to findings have precedence over goals for gathering information or forming a diagnosis, the new schema, sc-dyspnea, is selected and used. The new information has effectively changed the course of the reasoner's problem-solving behavior.

This schema, sc-dyspnea, is a specialized version of a general schema to interpret findings; instead of asking general questions, the schema can ask very specific things related to dyspnea (e.g., asking how many stairs the patient can climb as a measure of the severity). The last step of this schema is to explain the finding by postulating diseases that could cause it; using this step, the reasoner postulates hypotheses of pulmonary disease and cardiac disease. Adding these hypotheses to STM again interrupts the reasoner, which finds and adds to the agenda schemata to evaluate the hypotheses: "sc-pulmDz" and "sc-cardiacDz".

The strategy orders goals related to hypotheses before any others; hence, one of the two schemata just added is selected, in this case, sc-pulmDz. The reasoner uses this schema to score the hypothesis of pulmonary disease,⁶ and then tries to specialize the hypothesis using information that is in STM. One possible specialization, based on the fact that the patient is overweight, is recurrent pulmonary embolism⁷ (RPE); this is hypothesized, resulting in a schema ("sc-RPE") being activated to evaluate it.

The reasoner then selects sc-RPE and begins to evaluate the hypothesis of pulmonary embolism. Eventually, it will have evaluated all the hypotheses that it can and will have exhausted the information the user can give it. The main schema, sc-consult,

⁵Input to MEDIC is in a version of Conceptual Dependency (Schank and Abelson, 1977); there is currently no natural language interface.

⁶Using a scoring scheme very similar to that of INTERNIST-1 (Miller *et al.*, 1982).

⁷Blood clots occurring in the lungs.

will then suggest the step of forming a diagnosis, which will be attempted.⁸ In this case, the best hypothesis is recurrent pulmonary embolism, and that will be proposed to the user.

There is still much work to be done on MEDIC. At the present, the program has very little domain knowledge, and relatively few schemata. Additional thought must also be given to the form and content of the strategic schemata, which are currently quite simple; eventually, we would like for them to specify actions for the reasoner to perform in order to select a goal to pursue, making them more like the reasoner's other schemata. MEDIC also does not make use of context-specific information in dxMOPs to focus its attention.

RELATED WORK

Opportunistic reasoning has been addressed in the blackboard approach to problem solving of HEARSAY-II (Erman *et al.*, 1980) and OPM (Hayes-Roth, 1985). These program's *Knowledge Sources* (KS's) are atomic, rule-like specialists that are invoked in response to some arbitrary condition occurring. They correspond only loosely to schemata. Schemata are larger-grained than KS's, and capable of being interrupted. Schemata also serve to cluster actions to be taken to achieve a goal; many KS's, on the other hand, may be needed to achieve a single goal. The use of schemata should allow the reasoner to behave in a manner that a reasoner can understand: e.g., question-asking should be more focused. Schemata should also facilitate explanation, since the actions taken to achieve a goal, though possibly temporally disjoint, can still be explained in relation to one another.

The VISIONS Schema System (Weymouth, 1986) uses an approach similar to ours for interpreting visual scenes. Their schemata are specialists in particular vision tasks and can work in parallel to interpret a scene. Unlike our schemata, theirs are largely represented using procedures written in a programming language; it is therefore not possible for their program to reason about or modify their schemata, as is potentially possible using our representation of schemata. In addition, parallel execution of schemata is not feasible in our domain, given the goals of focused question-asking and modeling a human diagnostician's behavior.

The NASL (McDermott, 1978) program concentrated on the interaction of planning and execu-

tion. In many respects, our schemata are similar to NASL's tasks: both are hierarchical, bottoming out at the primitive action or primitive task level. However, our schemata are somewhat more flexible than NASL's tasks, and we make explicit use of goals; the latter allows the potential of specifying the goal of a task without necessarily specifying the steps to achieve it, thus allowing the reasoner to make such decisions at run-time.

We face some of the same problems as did NASL, too, in the chore of selecting which schemata to pursue at each point in problem solving. NASL made use of choice rules, which contained the strategic knowledge of that system. Our strategic schemata can be viewed in the current implementation as packages of such choice rules. However, the ultimate goal is to make them less rule-like and more schema-like, specifying steps for the reasoner to perform in order to select schemata to apply.

Firby (1987) is also concerned with interleaving planning and execution in environments that change during planning. The behavior of his RAP planner is quite similar to that of our reasoner. However, RAPs would seem to be somewhat more simple than our schemata, and oriented towards real-time control rather than diagnosis. In addition, we extend the idea of using packets of control knowledge to the meta-level by using strategic schemata to direct the reasoner's attention.

CONCLUSION

Medical diagnosis can be fruitfully viewed as a planning task in which planning is interleaved with diagnosis. New information may be discovered during diagnosis which should impact the future problem-solving behavior of the diagnostician. The diagnostician must be *opportunistic* in order to take notice of and respond to this new information as it becomes available.

Schema-based reasoning provides one approach to opportunistic reasoning. By representing problem-solving knowledge as packets of procedural information designed to achieve a goal, the reasoner can activate schemata as goals arise due to changes in the environment. When several schemata are active, the reasoner selects the one to apply based on the reasoner's current focus of attention—i.e., the goal the reasoner is trying to achieve. Goals are selected by the reasoner based on information from two sources: general goal-ordering information, stored in strategic schemata; and specific goal-ordering informa-

⁸Again, using a method similar to that of INTERNIST-1.

tion, stored in the dxMOP representing consultations similar to the current one.

Schemata are flexible, and enhance the reasoner's ability to respond to changes in the environment in two ways: (1) the order of their steps need not be completely determined—this allows the reasoner to select the next step of a schema based on the state of the world resulting from the application of the previous step; and (2) steps may specify goals, which the reasoner can attempt to satisfy at run-time by retrieving schema specific to the current situation. This flexibility, plus its opportunistic character, allows schema-based reasoning to be viewed as a type of reactive planning (e.g., Firby, 1987).

Though this research addresses medical diagnosis, we believe that schema-based reasoning can be usefully applied to other tasks. Our approach should be useful for any task in which planning and execution must be interleaved, or in which all features of the problem cannot be known at the start of the problem.

ACKNOWLEDGMENTS

Many thanks to Janet Kolodner, Hong Shinn, and Elise Turner for their comments on earlier drafts of this paper, and to our domain expert, Eric Honig, of Emory University and Grady Memorial Hospital.

REFERENCES

- Agre, P.E., and Chapman, D. (1987). Pengi: An implementation of a theory of activity, in *Proceedings of the Sixth National Conference on Artificial Intelligence*, pp. 268–272.
- Cullingford, R.E., and Kolodner, J.L. (1986). Interactive advice giving. In *Proceedings of the 1986 IEEE International Conference on Systems, Man, and Cybernetics*.
- Davis, R., and Buchanan, B.G. (1984). Meta-level knowledge. In B.G. Buchanan and E.H. Shortliffe (eds.), *Rule-Based Expert Systems: The MYCIN Experiments of the Stanford Heuristic Programming Project*. Addison-Wesley Publishing Company, Reading, Massachusetts, pp. 507–530.
- Erman, L.D., Hayes-Roth, F., Lesser, V.R., & Reddy, D.R. (1980). Then HEARSAY-II speech-understanding system: Integrating knowledge to resolve uncertainty, *Computing Surveys*, Vol. 12, No. 2.
- Firby, R.J. (1987). An investigation into reactive planning in complex domains, in *Proceedings of the Sixth National Conference on Artificial Intelligence*, pp. 202–206.
- Gomez, F., and Chandrasekaran, B. (1982). Knowledge organization and distribution for medical diagnosis. In W.J. Clancey and E.H. Shortliffe (Eds.), *Readings in Medical Artificial Intelligence*, pp. 320–338. Reading, Massachusetts: Addison-Wesley Publishing Company, 1984. (Originally published in *IEEE Transactions on Systems, Man, and Cybernetics*, Vol. SMC-11, No. 1, pp. 34–42 (1981).)
- Hayes-Roth, B. (1985). A blackboard architecture for control, *Artificial Intelligence*, Vol. 26, No. 3, pp. 251–321.
- Kolodner, J.L. (1985). Experiential processes in natural problem solving. Technical Report #GIT-ICS-85/123, School of Information and Computer Science, Georgia Institute of Technology, Atlanta, Georgia.
- Kolodner, J.L., and Kolodner, R.M. (1987). Using experience in clinical problem solving: Introduction and framework. In *Proceedings of the 1987 IEEE International Conference on Systems, Man, and Cybernetics*.
- McDermott, D. (1978). Planning and acting, *Cognitive Science*, vol. 2, pp. 71–109.
- Miller, R.A., Pople, H.E., Jr., and Myers, J.D. (1982). INTERNIST-1, an experimental computer-based diagnostic consultant for general internal medicine, *New England Journal of Medicine*, vol. 307, pp. 468–476.
- Schank, R.C. (1982). *Dynamic Memory*, Cambridge University Press, New York.
- Schank, R.C., and Abelson, R. (1977). *Scripts, Plans, Goals and Understanding*, Lawrence Erlbaum Associates, Hillsdale, NJ.
- Turner, R.M. (in press). Organizing and using schematic knowledge for medical diagnosis. Submitted to the Seventh National Conference on Artificial Intelligence (and also appearing in the 1988 DARPA Workshop on Case-based Reasoning).
- Weymouth, T.E. (1986). *Using Object Descriptions in a Schema Network for Machine Vision*, Technical Report 86–24 (Ph.D. thesis) Dept. of Computer and Information Science, Univ. of Massachusetts.