In W. J. Clancey and E.H. Shortliffe (Eds.), Readings in Medical Artificial Intelligence: The First Decade (pp. 361–381). Reading, MA: Addison-Wesley, 1984.

15

NEOMYCIN: Reconfiguring a Rule-Based Expert System for Application to Teaching

William J. Clancey and Reed Letsinger

As described in the introduction to GUIDON (Chapter 11), Clancey's work on that system led to an appreciation of the severe limitations of MYCIN's knowledge base if the system were to be used for instructional purposes (Clancey, 1983b). The NEOMYCIN research described in this chapter has been an attempt to rethink the knowledge structure and diagnostic strategy of MYCIN in view of requirements for teaching. This effort has several important products:

- *a better understanding of medical diagnostic strategy and its relation to knowledge structures (such as Feltovich's "logical competitor set," Chapter 12);*
- *a design of a representation framework for separating strategy from domain facts, in which strategy is stated* abstractly *(Clancey, 1983c); and*
- *a body of meta-rules, constituting a* generic procedure *that eases construction of knowledge bases for related problems in other domains (e.g., another diagnostic consultation program).*

The work is also of interest because of its relation to psychological studies (Chapter 12) and explanation methodology (Chapter 16).

From *Proceedings of the Seventh International Joint Conference on Artificial Intelligence*, vol. 2, 829–836 (1981). Used by permission of International Joint Conferences on Artificial Intelligence, Inc.; copies of the *Proceedings* are available from William Kaufmann, Inc., 95 First Street, Los Altos, CA 94022.

NEOMYCIN is a medical consultation system in which MYCIN's knowledge base is reorganized and extended for use in the next version of GUIDON. The new system attempts to capture psychological characteristics of diagnostic reasoning, designed to provide a basis for interpreting student behavior and teaching diagnostic strategy. This psychological orientation provides a constraint for making choices about representation and the reasoning process. In particular, NEOMYCIN captures the forward-directed, "compiled association" mode of reasoning that characterizes expert behavior. Collection and interpretation of data are focused by the "differential" or "working" memory of hypotheses. Moreover, the knowledge base is broadened so that GUIDON can teach a student when to consider a specific infectious disease and what competing hypotheses to consider, essentially the knowledge a human would need in order to use the MYCIN consultation system properly.

In order to articulate this knowledge to a student, it was necessary to greatly revise MYCIN's representation. Kinds of knowledge that were procedurally embedded in MYCIN's rules are stated separately, to make them accessible to the teaching program. The key idea is to represent explicitly and separately a domain-independent diagnostic strategy in the form of meta-rules, knowledge about a disease taxonomy, causal and data/hypothesis rules, and world facts. In essence, the new representation explicitly structures and controls the use of the diagnostic rules, simplifying them by isolating the basic data/hypothesis relations from their application criteria.

A more detailed discussion of methodological issues in the development of NEOMYCIN can be found in Clancey (1984). More recent research, exploiting the features of NEOMYCIN, includes modeling student strategies (London and Clancey, 1982) and stating strategies in explanations (Hasling et al., 1984). With the combination of empirical and knowledge-engineering interests, this research also has implications for incorporating cognitive modeling in new tools for building knowledge bases.

15.1 Introduction

A knowledge base used in a teaching program must explicitly represent what a student might need to be told. Development of intelligent tutoring systems such as SOPHIE (Brown et al., 1975), WHY (Stevens and Collins, 1978), WUMPUS (Goldstein, 1978), and GUIDON (Clancey, 1979a; 1979b) can be viewed, in part, as a problem of knowledge representation. This research has shown the advantages of:

- multiple representations of knowledge (e.g., the simulation model and semantic network in SOPHIE);

- representations that can be both interpreted and used to generate teaching text [e.g., Brown's meteorological automata (Brown et al., 1973) and production rules used in WUMPUS and GUIDON];
- network representations of knowledge that capture "importance" [SCHOLAR (Carbonell, 1970)], "complexity" or "prerequisite" associations [WUMPUS, BIP (Barr et al., 1976)], "analogy" and "generalization" relations (WUMPUS); and
- representations that allow for variants on expert performance (for modeling the student) [WEST (Burton, 1979), BUGGY (Brown and Burton, 1978)].

In the GUIDON program we have been exploring the problem of using MYCIN's rule set as teaching material. MYCIN (Shortliffe, 1976) is a rule-based expert system that provides therapy advice for certain kinds of infectious diseases. It has spawned a class of systems, called EMYCIN systems, that all use the same production rule language and interpreter (van Melle, 1980). GUIDON can operate using the rule set of any EMYCIN system as subject material.

MYCIN's rules were thought to be potentially useful for teaching because formal evaluations indicate that MYCIN captures a high level of expertise (Yu et al., 1979b), and modular design and representational meta-knowledge enable the program to explain its reasoning (Davis, 1976). Ironically, we have found that it is in precisely these two areas—expertise and explanatory capability—so important for a successful teaching program, that MYCIN falls short. To solve these problems, we have implemented a new system we call NEOMYCIN.

15.1.1 The Limitations of MYCIN for Application to Teaching

MYCIN is designed to be used as a consultant; consequently, we encounter difficulties when using it for teaching a student how to be a primary diagnostician. MYCIN's knowledge base is designed to interpret culture results from the blood and the cerebral-spinal fluid (CSF). But the expertise that suggests that such a culture should be taken is not part of the system. It is the user of MYCIN, the person seeking advice, who will think about meningitis in the first place and order the CSF culture and who will consider competing hypotheses (and medical tests) that need to be considered before MYCIN is even brought into the case as a consultant. This knowledge is certainly a critical part of teaching infectious disease diagnosis, but it lies completely outside the scope of the MYCIN knowledge base.

Moreover, protocols of experts who solve the same cases as are presented to MYCIN indicate that the program does not organize or use its knowledge in the same way a human expert does. This result is not sur-

prising, for MYCIN was not designed to simulate the *process* of human reasoning. The rules make use of the same data a physician uses and some of the same intermediate concepts of disease, but MYCIN's weakly focused, exhaustive search is quite dissimilar from how people reason. For GUIDON, our tutorial program, to articulate and recognize the hierarchical organizations of knowledge and search strategies that humans find useful, we need to reorganize MYCIN's rule set and incorporate an explicit model of human diagnostic reasoning, the kind indicated by psychological research in medical problem solving (Miller, 1975; Rubin, 1975; Pauker and Szolovits, 1977; Swanson et al., 1977; Elstein et al., 1978; Kassirer and Gorry, 1978) (see also Chapter 6). In particular, the model must exhibit:

- focused, forward-directed use of data (including *trigger* associations that suggest diagnoses);
- follow-up questions that establish the disease process (part of what a physician calls "forming a picture of the patient"); and
- management of a changing "working" memory of hypotheses under consideration.

In this sense, the development of NEOMYCIN is an attempt to synthesize previous medical psychological research and to analyze its application to the infectious disease problem domain.

15.1.2 Developing a Psychological Model by Modifying EMYCIN

A psychological model of diagnostic reasoning cannot be represented using the EMYCIN representation alone, that is, by simply rewriting MYCIN's rules. For example, the idea of asking a follow-up question is not allowed by MYCIN's rule interpreter. Also, we need to apply rules selectively and nonexhaustively. In general, the rule representation and interpreter must be modified; rules need to be organized so they can be selectively applied in different ways.

Many of the changes to EMYCIN are straightforward. They illustrate how local changes to the "inference engine" of a program can dramatically change how the knowledge base is used in problem solving. For example, a simple change is to provide for data-directed reasoning so new data can cause new subgoals to be set up and pursued. In MYCIN, an *antecedent rule* is tried whenever some piece of information required by the rule's premise becomes known. A NEOMYCIN *trigger rule* is similar, but it allows for new data to be requested in order to apply the rule. For example, one trigger rule is "if the patient has a stiff neck and a headache, then consider meningitis."[1] When a physician hears that the patient has a stiff neck, the

[1]The medical examples in this paper are simplified; we make no claims about completeness or accuracy. They are for purposes of illustration only.

IF: 1) The infection is meningitis,
2) The subtype of meningitis is bacterial,
3) Only circumstantial evidence is available,
4) The patient is at least 17 years old, and
5) The patient is an alcoholic
THEN: There is suggestive evidence that diplococcus-pneumoniae is an organism causing the meningitis

FIGURE 15-1 Typical MYCIN rule.

association to meningitis might come to mind, prompting him or her to determine whether the patient has a headache as well. This behavior is brought about in NEOMYCIN by simply marking trigger rules to distinguish them from ordinary antecedent rules and "throwing a switch" in the rule interpreter so that pursuing new subgoals is enabled for trigger rules.

Besides interpreter changes, different kinds of knowledge had to be separated out of the rules and represented explicitly. Figure 15-1 shows a typical (paraphrased) MYCIN rule, an example of "compiled expertise." We can list some of the individual steps of reasoning and knowledge sources out of which it is composed, unknown to MYCIN, but explicitly represented in NEOMYCIN:

- Analysis of other rules shows that this rule (to determine the organism) is only invoked after it has been established that the patient has an infection. Thus four major subgoals are established in this order: Is there an infection? Is it meningitis? Is it bacterial? Is it *Diplococcus pneumoniae?* Each of these subgoals hypothesizes a more specific cause of disease. Thus, *the ordering of clauses constitutes a top-down refinement strategy.* However, MYCIN does not know about this specialization hierarchy. It does not even know that *Diplococcus pneumoniae* is a bacterium. Perhaps most serious of all for meeting our teaching goals, MYCIN omits intermediate categories such as acute/chronic meningitis and gram-negative meningitis that physicians find helpful. In NEOMYCIN these categories are represented explicitly in an *etiological taxonomy* by allowing parameters to be specializations of one another.
- The clause about the patient's age prevents MYCIN from asking if a child is an alcoholic. MYCIN does not know that the ordering of these clauses is important, or what the relationship is. In NEOMYCIN these world relations are captured by separate *screening rules.*
- When there is laboratory evidence (a culture with visible organisms), this rule does not apply (clause 3). However, a companion rule still allows the circumstantial evidence of alcoholism to be considered, but gives it less weight. This principle of considering circumstantial evidence even when there are hard, physical observations of the cause is not explicitly known to MYCIN. The principle is compiled identically into 40 pairs of rules, rather than being stated as a reasoning rule for combining hard

and soft evidence. NEOMYCIN has rules for reasoning about the evidence it has collected, so connections between data and hypotheses are separate from the contexts in which they will be used.

These forms of knowledge—a (top-down) strategy, an etiological taxonomy, world facts, evidence-weighing rules—form a basis for a psychological model about knowledge organization and access, but they are not sufficient. Consider the above rule again. How does a physician remember to ask about alcoholism? How does he or she remember the connection with *Diplococcus?* Experts use a rich set of organizational aids and mnemonics for accessing their knowledge.

For example, one can think of taking the patient's history as a process of determining the differential of possible causes. Under this *strategy*, the expert follows the principle (rule model) that "compromised host conditions broaden the differential by suggesting special causes." Alcoholism is one of these conditions. So the low-level behavior of asking "Is the patient an alcoholic?" occurs in the context of the general process of diagnosis. In explaining the question to a student, it is important to be able to step back from the immediate concern for supporting a particular disorder and to articulate the general goals and methods of diagnosis itself. At the lowest level, the association to *Diplococcus* might be remembered as a simple causal story: alcoholics breathe in their own secretions, so organisms found in the mouth find their way to the lungs, causing pneumonia.

In summary, NEOMYCIN incorporates these psychological aids for teaching diagnosis:

1. a *representation of diagnostic strategy* that provides a meaningful, useful orientation for collecting data ("attempt to broaden the differential");
2. *structural associations* for indexing evidence to consider (abstractions such as "compromised host conditions" and rule models that use them); and
3. *rule justifications* that relate data/hypothesis associations to underlying causal processes.

15.1.3 The Need for Focusing Strategies

As we mentioned above, we cannot use MYCIN for teaching about meningitis diagnosis because it does not know how patients with meningitis typically appear when the physician first sees them and what competing disorders need to be considered. But if we simply added knowledge about more diseases and when to order laboratory tests we would be in trouble: a top-down diagnostic strategy is inadequate for a broader range of problems. The combinatorics of the search problem for medical diagnosis make it impossible for an expert to consider every infection, to work top-down. Initial information most commonly brings the physician into the *middle* of his or her taxonomic hierarchy (via the "compiled associations" such as the

trigger rule given above). Working from the middle, the physician must first look upward to focus the possibilities ("Is it a traumatic process? cancer?") and then refine downward. The approach used by MYCIN's rules only works because the user of the program is the one who focuses on meningitis. MYCIN can verify that the historical and laboratory evidence is consistent with meningitis, but it does not have the knowledge for considering meningitis in the first place. The program has only two infections to consider and does not know about other causes of the findings reported by the user.

For the program itself to shoulder this focusing burden (so that GUIDON can teach it to a student), we should more properly think of its area of expertise as being related to the observations a user will bring to it rather than the problems it knows how to confirm and refine. Thus MYCIN's area of expertise is "meningitis"; in contrast, NEOMYCIN deals with "abnormal neurological signs" or "headache and fever." In order to give NEOMYCIN the capability to deal with a broader range of problems, to actually have it think of other causes of headache and fever, we did the following:

1. *expanded the etiological knowledge* to include broad categories of other, noninfectious problems, such as "toxic problem," and "neoplastic problem";
2. *incorporated the focusing strategy of "group and differentiate"* so the program could manage this broader range of possibilities; and
3. *added knowledge about disease processes,* knowledge that cuts orthogonally across the etiological taxonomy, so diseases can be compared according to location, extent of the disorder, duration, severity, etc., in order to enhance the program's ability to apply the focusing strategy.

15.2 An Overview of NEOMYCIN

A few words about the character of MYCIN's problem domain are in order. We assume that a diagnosis or problem solution consists of an ordered list of problem causes that have been selected from a fixed, hierarchical space of hypotheses (e.g., "cancer process," "chronic meningitis") or categories of disease and pathophysiological states (e.g., "mass lesion in the brain"). We assume that an *informant* presents a problem to the program, which acts as a *consultant,* the role played by a student using GUIDON. There are two types of data: soft data (circumstantial or historical) and hard data (laboratory or direct measurements). Some of the evidence may be missing, and conclusions will usually be uncertain.

A schematic of the NEOMYCIN system (Figure 15-2) illustrates the various knowledge sources and their relation to the strategic knowledge and differential (the set of diagnoses under consideration). These com-

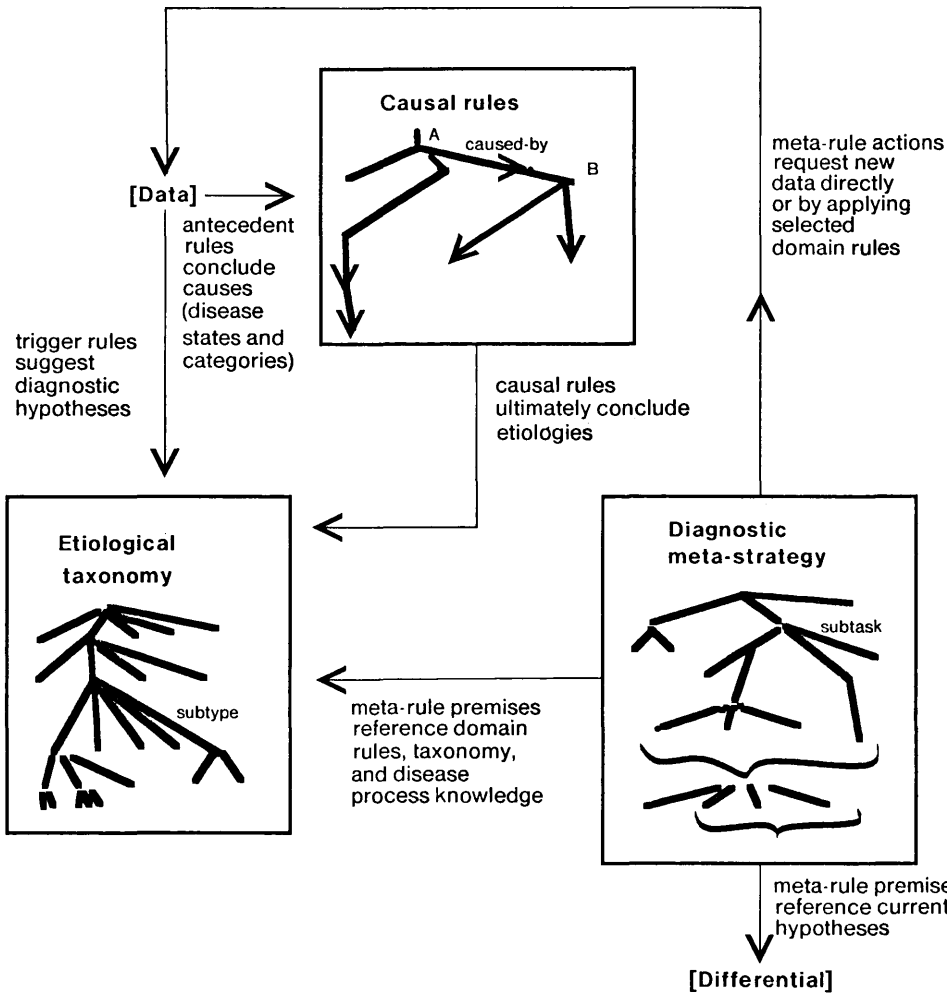

FIGURE 15-2 Components of the NEOMYCIN system.

ponents are shown as icons expanded in subsequent figures. The interpretation of Figure 15-2 follows.

- There are four kinds of domain rules:
  - *Causal rules* form a network of pathophysiological states and disease categories, ultimately linking raw observations (incoming data) to the etiological taxonomy.
  - *Trigger rules* associate data with etiologies, which are placed as hypotheses in the differential (maintained so that general causes are replaced by their more specific descendents).

METARULE397 (for the task group-and-differentiate)

IF: There are two items on the differential that differ in some disease process feature
THEN: Ask a question that differentiates between these two kinds of processes

FIGURE 15-3 A typical strategy rule.

- *Data/hypothesis rules* associate circumstantial and laboratory data with diseases, as do trigger rules, but only those rules focused by the differential are tried when the data are circumstantial (i.e., the associations that "come to mind" are those hypotheses already in the differential, as well as the nodes of the etiological taxonomy that hang below the hypotheses of the differential).
- *Screening rules* (not shown) form a hierarchy of abstractions and restrictions on data (e.g., "if the patient is not immunosuppressed, then he is not an alcoholic"); they are applied by backward chaining, in an attempt to determine a datum without asking the user.

- Other domain knowledge (not shown), orthogonal to the hierarchies of cause, considers diseases as processes having a location, extent, progression of symptoms, etc.
  - One form of disease process knowledge is represented as a framelike description associated with diseases in the etiological taxonomy and is used to differentiate among them.
  - A second form consists of a list of process-oriented, follow-up questions that should be immediately asked when some disease category or pathophysiological state is implicated (e.g., to establish when symptoms occurred and their ordering and change in severity).
- The meta-strategy for doing diagnosis consists of a hierarchy of domain-independent meta-rules. In general, these meta-rules examine the differential and make use of the etiological taxonomy, causal associations, and disease process knowledge to decide what datum to request next. A typical strategy rule is shown in Figure 15-3.

The annotated typescript in the next section shows how these forms of knowledge interact in practice. Subsequent sections provide a few more details about the representation.

15.3 A Sample Case

To illustrate the ideas presented above, a simple case was presented to a simple version of NEOMYCIN (there are generally only one or two rules to establish each hypothesis). Only the collection of circumstantial evidence

is shown in Figure 15-4. Although this trace of reasoning is very detailed, it is included here because it is precisely the *process* by which data, hypotheses, and rules interact that is at issue.

Observe the many different reasons why a question is asked; this is a reflection of the complexity of the diagnostic strategy. NEOMYCIN is hypothesis- and data-directed. In contrast, MYCIN only asks a question to evaluate a clause of a rule for the goal it is pursuing. Its rules are not sorted by conclusion, so its questions appear to skip back and forth randomly among hypotheses. It is not backward chaining *per se* that distinguishes the two systems, for NEOMYCIN essentially backward chains through its strategic rules. It is NEOMYCIN's forward, nonexhaustive reasoning and management of a space of hypotheses that make it reason more like a human does.

7-Dec-80 16:18:25

--------PATIENT-1--------

task MAKE-DIAGNOSIS METARULE384 succeeded.
task IDENTIFY-PROBLEM METARULE385 succeeded.

Initial tasks are unconditional. The first step is to collect identifying data, followed by the reasons for seeking advice.

Please enter information about the patient.

	Name	Age	Sex	Race
1)	** MIKE	4	MALE	CAUCASIAN

2) Please describe the chief complaint:

(Enter keyword or phrase.)

** DIPLOPIA

Diplopia is recognized to be a neurological sign which triggers meningitis, and leads to a follow-up question (associated with NEUROSIGN) . . . .

antecedent RULE376 succeeded.
Conclude: NEUROSIGN is YES (1.0)
antecedent RULE426 succeeded.
Conclude: INFECTIOUS-PROCESS is MENINGITIS (.7)

Differential: Adding MENINGITIS

3) How many days has Mike had abnormal neurologic signs?
** 13

Trying evidence rules: LENSIGN → CHRONIC-MENINGITIS
Trying RULE524; succeeded.
Conclude: MENINGITIS is ACUTE-MENINGITIS (-.3)
Conclude: MENINGITIS is CHRONIC-MENINGITIS (.3)

FIGURE 15-4 Trace of NEOMYCIN's reasoning for a sample case. [For brevity, the trace of rule conclusions is of the form "Conclude: PARAMETER is VALUE (CF)." Parameters are always associated with the context, here PATIENT-1. Strategic tasks are sometimes abbreviated. Figure 15-5 relates meta-rules to tasks.]

Differential:
replace MENINGITIS with CHRONIC-MENINGITIS

Another antecedent rule is associated with diplopia . . . .

antecedent RULE373 succeeded.
Conclude: BRAIN-PRESSURE is YES (.8)

Brain pressure suggests a mass lesion—these antecedent rules all model immediate, forward inferences that we observed in experts . . . .

antecedent RULE375 succeeded.
Conclude: MASS-LESION is YES (.72)

Differential: Adding MASS-LESION

Any other information?
** HEADACHE

Here a trigger antecedent rule requires more information before it can be applied . . . .

4) Does Mike have a stiff neck?
** YES
antecedent RULE424 succeeded.
Conclude: INFECTIOUS-PROCESS is MENINGITIS (.91)

Any other information?
** NO

The program now takes initiative, starting its main history-taking task of "establish the hypothesis space" (hereafter, EHS). The differential has new, unexamined elements, so the "group and differentiate" (G&D) task is invoked first. METARULE400 recognizes that the etiological ancestors of chronic meningitis have not been explored, so an attempt is made to confirm (task PURSUE-HYPOTHESIS = PH) that an infectious process is causing the problem.

Enter EHS loop: focus = NIL
task METARULE427 succeeded.

Review differential: MASS-LESION CHRONIC-MENINGITIS

Enter G&D loop: focus = NIL
task METARULE400 succeeded.

Enter PH loop: focus = INFECTIOUS-PROCESS
task METARULE410 succeeded.

5) Is Mike febrile?
** Y
antecedent RULE423 succeeded.
Conclude:
DISORDER-ETIOLOGY is INFECTIOUS-PROCESS (.7)

There are no further questions to ask; an interrupt returns control to the G&D task. No rules succeed, so control returns to the EHS task. The "explore and refine" (E&R) task silently refines mass lesion, expanding the differential by different etiological categories, and so triggering return to the EHS task.

observed STOP-PURSUING METARULE408

Repeating G&D loop: focus = INFECTIOUS-PROCESS

Repeating EHS loop: focus = INFECTIOUS-PROCESS
task METARULE428 succeeded.

Enter E&R loop: focus = INFECTIOUS-PROCESS
task METARULE429 succeeded.

FIGURE 15-4 continued

observed STOP-EXPLORING METARULE407

Repeating EHS loop: focus = INFECTIOUS-PROCESS
task METARULE427 succeeded.

The first step is again to review the differential, a process observed in experts. Process features of brain abscess and chronic meningitis are compared: they both occur in the central nervous system, are chronic problems, and are infectious, but brain abscess is a localized problem. NEOMYCIN asks a question to discriminate on this basis . . . .

Review differential: BRAIN-ABSCESS HEMATOMA
PUS-IN-BRAIN CHRONIC-MENINGITIS

Enter G&D loop: focus = INFECTIOUS-PROCESS
task METARULE397 succeeded.

6) Does Mike have focal neurological signs?
** NO
Trying evidence rules: FOCALSIGNS → BRAIN-ABSCESS
RULE179 failed due to clause 1

The program has not been supplied with knowledge for confirming other causes of mass lesion (e.g., traumatic hemorrhage, tumor), so it is unable to continue its grouping operation and begins an exploration cycle . . . .

Repeating G&D loop: focus = INFECTIOUS-PROCESS

Repeating EHS loop: focus = INFECTIOUS-PROCESS
task METARULE428 succeeded.

Enter E&R loop: focus = INFECTIOUS-PROCESS
task METARULE402 succeeded.

Enter PH loop: focus = BRAIN-ABSCESS
task METARULE409 succeeded.

Now directly focusing on brain abscess, the program "realizes" that data supplied earlier are relevant (RULE433). Chronic meningitis is then considered by refining it and pursuing specific causes. Pursuing TB, NEOMYCIN follows the strategy of confirming the first ("enabling") step in the disease process: contact with the organism . . . .

Trying evidence rules: MASS-LESION → BRAIN-ABSCESS
Trying RULE433; succeeded.
Conclude: INFECTIOUS-PROCESS is BRAIN-ABSCESS (.216)

Observed STOP-PURSUING METARULE408

Repeating E&R loop: focus = BRAIN-ABSCESS
task METARULE429 succeeded.

Enter PH loop: focus = TB-MENINGITIS
task METARULE411 succeeded.

7) Does Mike have a TB risk factor?
** YES
Trying evidence rules: TBRISK → TB-MENINGITIS
Trying RULE525; succeeded.
observed STOP-PURSUING METARULE408

Focusing strategies dictate that a sibling be considered next. Fungal meningitis is refined, and a child, cryptococcus, pursued . . . .

Repeating E&R loop: focus = TB-MENINGITIS
task METARULE401 succeeded.
Enter PH loop: focus = FUNGAL-MENINGITIS

FIGURE 15-4 continued

Repeating E&R loop: focus = FUNGAL-MENINGITIS
task METARULE399 succeeded.
Enter PH loop: focus = CRYPTOCOCCUS

A cancer patient is at some risk of getting cryptococcal meningitis. Rather than asking directly if the patient has cancer, the program models an expert's efficient casting of a wider net by asking a more general question. Specifically, there are "screening rules," that lead it to determine first if the patient is immunosuppressed (RULE395) and then compromised (RULE343). This is the only form of backward chaining that occurs in NEOMYCIN.[2]

task METARULE431 succeeded.

--[0] Findout: LEUKEMIA
--[1] Findout: IMMUNOSUPPRESSED

Trying RULE343;
8) Is Mike a compromised host (e.g. alcoholic, sickle-cell-disease, immunosuppressed)?
** YES
RULE343 failed due to clause 1

If the patient were not compromised, the program could have concluded that he is not immunosuppressed (RULE343). Now it is unsure and must ask directly. If the patient is not immunosuppressed, the program will know that he does not have leukemia (RULE395). The answer of LEUKEMIA below implies immunosuppressed, so RULE395 fails, and the original goal is determined.

--[1] Finished: IMMUNOSUPPRESSED
9) Is Mike immunosuppressed (e.g. corticosteroid therapy, cytotoxic drug therapy, radiation therapy, leukemia)?
** LEUKEMIA

I will assume that leukemia is one of the diagnoses of Mike
RULE395 failed due to clause 1

--[0] Finished: LEUKEMIA

Trying evidence rules: LEUKEMIA → CRYPTOCOCCUS
Trying RULE056; succeeded.
Conclude: FUNGAL-MENINGITIS is CRYPTOCOCCUS (.3)

Repeating E&R loop: focus = CRYPTOCOCCUS
task METARULE401 succeeded.

Attention turns to a sibling. Again, the "enabling step" is asked about first . . . .

Enter PH loop: focus = COCCIDIOIDES
task METARULE411 succeeded.

10) Has the patient ever been to a cocci-endemic area?
** NO

Trying evidence rules: COCCI-ENDEMIC → COCCIDIOIDES
RULE570 failed due to clause 1
RULE287 failed due to clause 1
observed STOP-PURSUING METARULE408

Repeating E&R loop: focus = COCCIDIOIDES

Repeating EHS loop: focus = COCCIDIOIDES
task METARULE430 succeeded.

[2]*Ed. note:* A later version of NEOMYCIN accomplishes this form of inference by meta-rules.

Having exhausted its limited knowledge, the program finds no other relevant, hypothesis-oriented questions to ask. Several general questions are asked . . . .

11) Is Mike receiving any medications?
** NO

Repeating EHS loop: focus = COCCIDIOIDES
task RULE430 succeeded.

12) Has Mike been recently hospitalized?
** NO

Repeating EHS loop: focus = COCCIDIOIDES

If additional data had been supplied, new hypotheses might have been placed on the differential and strategies for grouping or refining might have been called into play once again. This ends the history-taking process. Next the program would order laboratory tests, process them, and perhaps return to gathering circumstantial evidence.

FIGURE 15-4 continued

15.4 The Diagnostic Meta-Strategy

Formalizing the diagnostic strategy from protocol analysis was the most difficult part of designing NEOMYCIN. Figure 15-5 shows the general outline of the meta-strategy. Each nonterminal node in the tree stands for a task that is achieved by a set of rules. An important aspect of our model of diagnosis is that the process can be taught as a task-posing activity: the problem solver thinks in terms of what he or she is trying to do (e.g., to consider unusual causes and so broaden the differential) in order to bring knowledge sources to mind. Thus the meta-strategy is structured so the tasks make sense as things that experts try to do.

Figure 15-5 shows that the main object of the meta-strategy is to decide what data to collect next (invoke MYCIN's FINDOUT routine), generally by focusing on some hypothesis in the differential. Aside from collecting initial information, the basic idea is that collecting circumstantial evidence is a process of *establishing the hypothesis space.* This process takes the form of considering what could cause the reported data, grouping and refining the differential, and asking general questions.[3] A great deal of what we might call *heuristic confidence* is placed in the general questions, which constitute the outline of the "history-taking process" as it is generally taught to medical students. However, strategies for using causal and disease process knowledge enable the expert to be an efficient problem solver in a combinatorially large space, and these strategies are generally not taught.

[3]*Group and differentiate* is used here in the loose sense of establishing general focus on a process that is consistent with hypotheses suggested independently by the data. *Clustering* (in multiple ways) *and discriminating*, the usual meaning of the term, is one operation for achieving this focus.

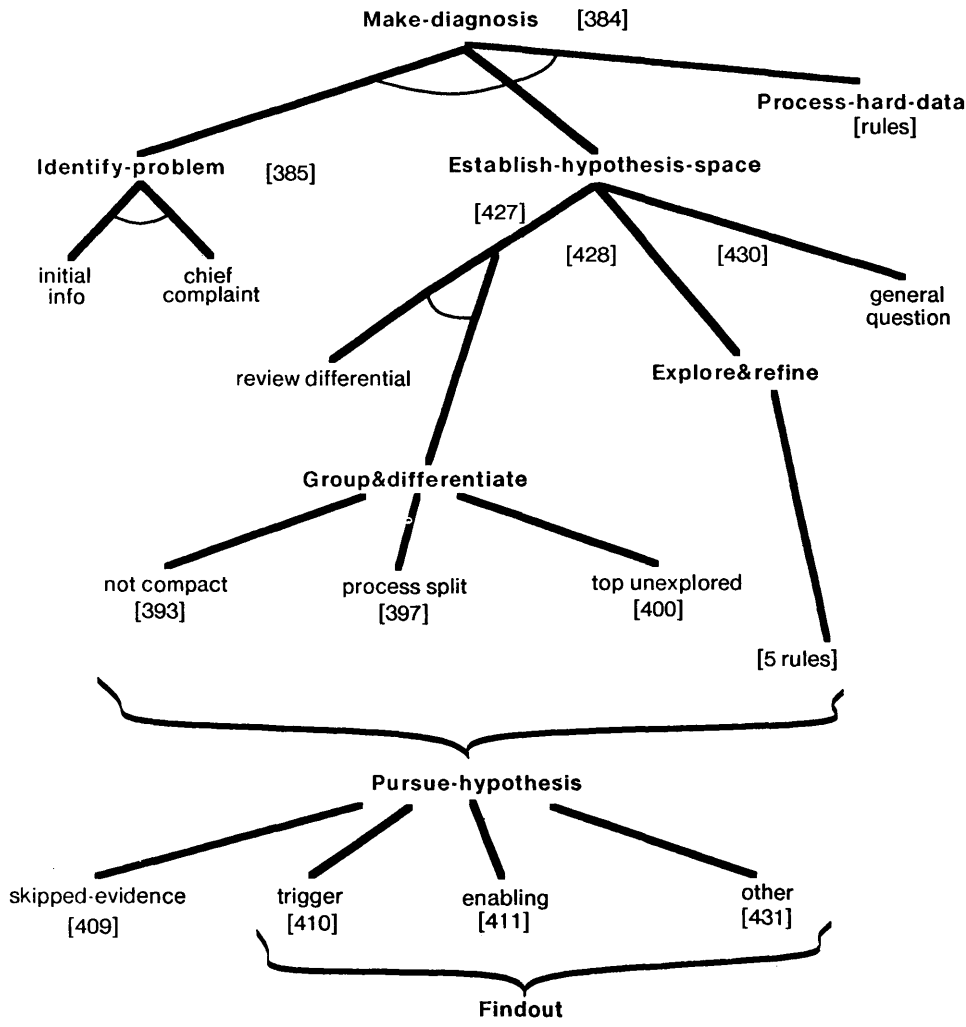

FIGURE 15-5 NEOMYCIN's diagnostic meta-strategy. (Rule numbers in brackets appear in the sample trace.)

The implementation is in terms of hierarchical meta-rules,[4] which as a whole constitute the meta-strategy. Figure 15-6 illustrates how the rules for a given task are treated as a pure production system—they are repeat-

[4]So called because they indirectly control the invocation of the domain-dependent object rules. Davis's conception of meta-rules was that they would directly order object-level rules. However, in our theory of diagnostic strategy, meta-rules reason about the state of the differential and knowledge sources (kinds of evidence) that could change it in desirable ways. Thus, our meta-rules choose *kinds of object rules* (hypothesis-confirming, process-oriented, causal).

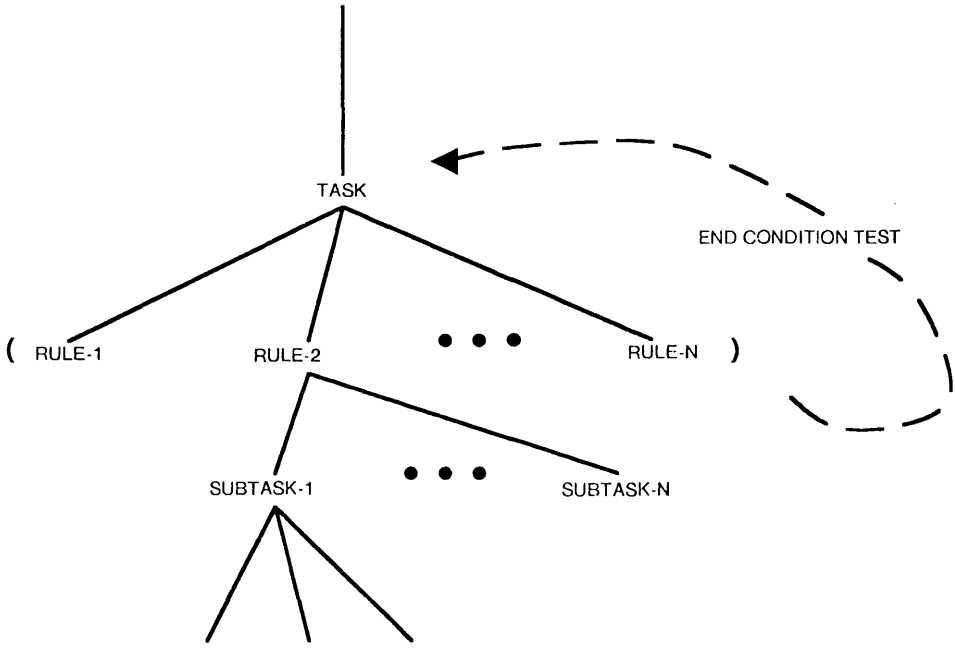

FIGURE 15-6 Rule-based invocation and interruption of strategic tasks.

edly tried in order, returning to the head of the list when one succeeds, stopping when no rule succeeds or an end condition is true.

The end condition is itself determined by rules, and is inherited as we descend into the hierarchy of tasks. The main use for this feature is to allow refocusing when new data change the state of the differential, as well as nonexhaustive consideration of hypotheses.

15.5 Etiological Taxonomy, Causal and Disease Process Knowledge

Some details of the implementation are given in this section. The etiological taxonomy (Figure 15-7) is implemented as EMYCIN parameters in which the values for one parameter (e.g., CHRONIC-MENINGITIS) are themselves parameters (e.g., TB-MENINGITIS and FUNGAL-MENINGITIS). We call these *taxonomic parameters.*

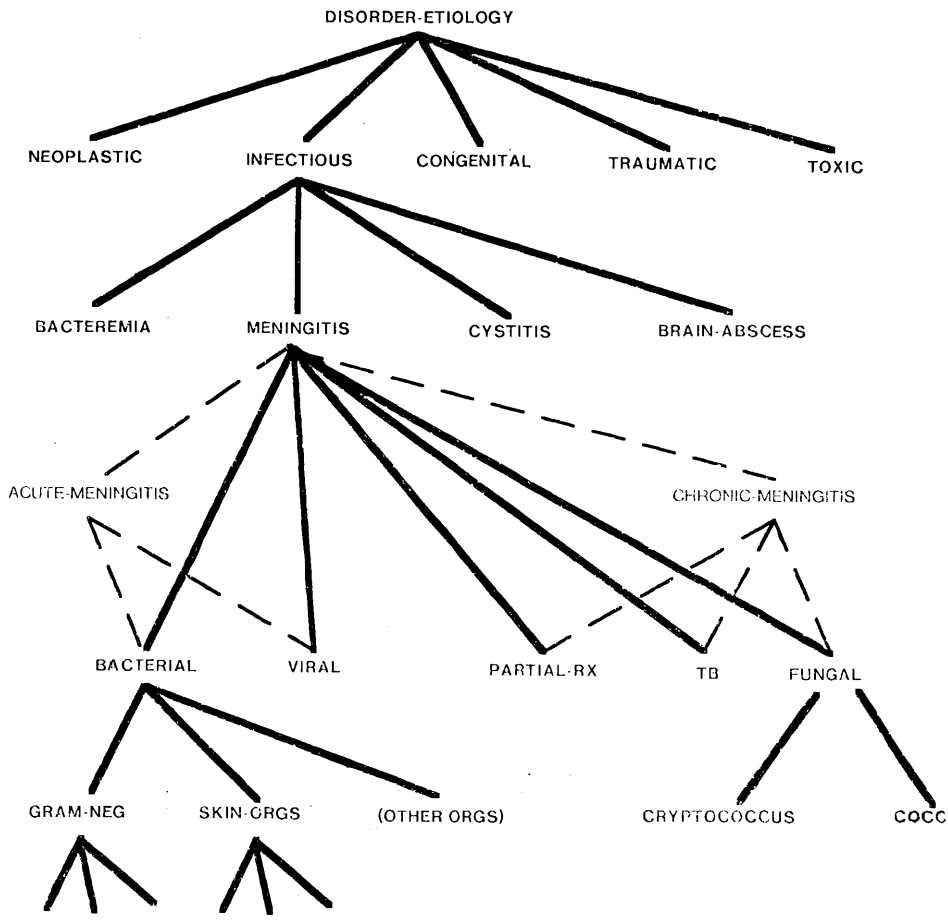

FIGURE 15-7 Portion of etiological taxonomy. (Links represent specialization of cause.)

Causal knowledge (Figure 15-8) is represented as rules modified by a certainty factor, as are all MYCIN rules. A causal rule of the form "if A then B" implies that A is caused by B, the direction of the association that is most generally useful for interpreting data and refining hypotheses. These rules mention *data parameters, taxonomic parameters,* or *state-category parameters.* State-category parameters stand for pathophysiological states or categories of disease (e.g., a mass lesion in the brain). In linking these concepts together, it is important to properly distinguish between causal and subtype links. (While we might say that an unknown mass lesion, a

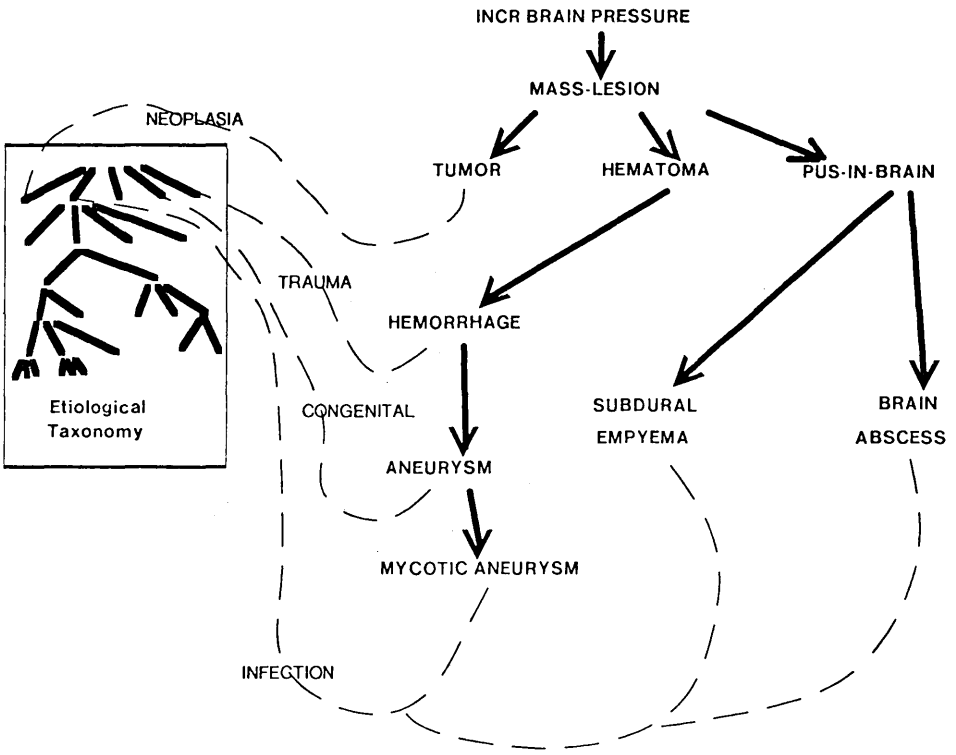

FIGURE 15-8 Portion of causal rule network, showing connection to etiological taxonomy.

space-occupying substance, is caused by a tumor, it is more proper to represent a tumor as a kind of mass lesion.) Causal rules are used by the "explore and refine" task to work backward from state-category hypotheses in the differential to prior causes, and ultimately to diagnostic hypotheses in the etiological taxonomy (as shown in Figure 15-8).

Disease process knowledge is represented as a frame associated with taxonomic parameters. Slots are process descriptors such as EXTENT, LOCATION, and COURSE associated with a literal value and a pointer to the parameter to establish it. For example, associated with BRAIN-ABSCESS is the triple (EXTENT FOCAL FOCALSIGNS), meaning that the extent of the disease is focal and this can be determined by asking about focal signs. Disease process knowledge is orthogonal to the etiological taxonomy, making it useful for grouping and discriminating hypotheses (see sample trace, before question 6).

15.6 Related Research

Besides the intelligent computer-assisted instruction (ICAI) projects cited in the introduction, our work has been motivated by previous research in teaching problem-solving strategies [e.g., Papert (1970); Brown et al. (1977); Wescourt and Hemphill (1978)]. We believe NEOMYCIN is the first attempt to formalize a runnable psychological model of diagnostic strategy that can be presented to a student. As should be obvious from our analysis, a considerable debt is owed to the medical problem-solving literature, cited above.

Both Reggia (1978) and Aikins (1980) modified the MYCIN system to make it more acceptable to physicians, particularly to improve knowledge acquisition. Aikins's use of an etiological taxonomy and trigger rules, derived from Rubin's work, is particularly close to our approach. However, we go a step further by representing strategic knowledge separately in domain-independent form. Our teaching application has also made clear the importance of disease process knowledge for broadening the diagnostic range of a consultation program.

Research in cognitive psychology has been helpful to us, particularly studies at the Learning Research Development Center (Anderson et al., 1981; Chi et al., 1981) (see also Chapter 12) in modeling the differences between experts and novices in geometry and physics problem solving. To some extent, our attempt to "decompile" MYCIN's knowledge is the inverse of Anderson's task of modeling how a novice composes and generalizes knowledge from experience.

15.7 Some Limitations

Pople's experience has been useful to point out limitations in our design. He shows that a simplistic causal network is not adequate when an attempt is made to represent all of general internal medicine (Pople, 1982). For example, when the causal connections between data and the taxonomy are long and complex, it may not be feasible to follow each path (possible cause). His "bridge concepts" [similar to Feltovich's "logical competitor sets" (see Chapter 12)] are attempts to model how an expert jumps over to distal, tentative hypotheses. They essentially provide a quick way to find the intersection of causes for a set of disease symptoms.

Similarly, Rubin's thesis illustrates a number of strategies for combining hypotheses (for example, relating complications and causes) that we have not yet found to be important in MYCIN's domain. To this extent,

our model is not the complete story of human diagnostic reasoning, but it can be built on as we expand our experience into other domains. We do not yet understand how an expert organizes his or her differential; how context is saved and restored from interrupts; how urgency, cost, and human values factor into the diagnostic process; and so on.

15.8 Summary of What We Learned

To teach diagnosis, it is useful to have a psychological model of problem solving. In particular, we need to incorporate into our model the medical knowledge and strategies an expert uses for initial problem formulation. An expert thinks in terms of a hierarchy of causes and the process characteristics of a disease so that he or she can order the data and the search. Moreover, an expert has learned "compiled associations" that allow him or her to efficiently associate hypotheses with data (e.g., trigger rules, Pople's "bridge concepts"), and cast a wide net of questions (e.g., general, screening, and follow-up questions).

Also, we need to represent the various kinds of knowledge explicitly so that they can be accessible for teaching. Our method is to represent strategic knowledge in domain-independent form, wholly separate from the medical knowledge described above. This requires that the medical knowledge be organized so that it can be indexed by the strategies (e.g., as the disease-process frame links abstract features of any disease, such as progression over time, to means for establishing this information in a particular case).

In a sense, we join cognitive psychologists [e.g., Anderson et al. (1981) and Rumelhart and Norman (1980)] in rediscovering the procedural/declarative problem in the context of how knowledge becomes transformed through experience. We recognize that the expert has composed associations, so he or she makes wide, tentative jumps between data and hypotheses. However, we represent these compiled associations declaratively for use in instruction (spelling out the diagnostic procedure in detail), and we record justifications of data-interpretation rules to allow for explanation of reasoning.

15.9 Future Research

Development of NEOMYCIN and GUIDON version 2 will proceed in parallel. Comparisons of NEOMYCIN's performance to MYCIN's will indicate if our more principled representation has changed the performance of the

system. This is a possibility because we have simplified some rules so they represent more closely the associations a human expert normally remembers. Preliminary runs give comparable results, though NEOMYCIN asks fewer questions because of its focused approach. We might also use NEOMYCIN's representation and meta-rules for diagnosis in a nonmedical domain, to test the domain-independence of the model.

GUIDON version 2 will use the NEOMYCIN representation, making it possible to articulate diagnostic strategy. A new phase of development will begin as we try to use the diagnostic strategies (and variants of them) for interpreting student behavior, leading to capabilities to evaluate partial solutions and provide assistance. The first version of GUIDON attempted these things, but was not able to recognize or suggest psychologically valid approaches.

ACKNOWLEDGMENTS

We are especially grateful to Timothy Beckett, M.D., for his patient explanations and willingness to be directed in our discussions. Bruce Buchanan and Bob London have also contributed to the GUIDON project. This research has been supported in part by an ARPA and ONR contract (NO0014-79C-0302). Computational resources were provided by the SUMEX-AIM facility (NIH grant RR00785).