

NEOMYCIN: RECONFIGURING A RULE-BASED EXPERT SYSTEM FOR APPLICATION TO TEACHING

William J. Clancey and Reed Letsinger

Computer Science Department
Stanford University; Stanford, CA 94305

ABSTRACT

NEOMYCIN¹ is a medical consultation system in which MYCIN's knowledge base is reorganized and extended for use in GUIDON, a teaching program. The new system constitutes a psychological model for doing diagnosis, designed to provide a basis for interpreting student behavior and teaching diagnostic strategy. The model separates out kinds of knowledge that are procedurally embedded in MYCIN's rules and so inaccessible 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 the structure of the problem space, causal and data/hypothesis rules, and world facts.

As a psychological model, 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.

INTRODUCTION

In order to use a knowledge base as subject material for teaching, it is important that the kinds of things a student needs to be told be represented flexibly, so that they can be singled out and articulated. Development of intelligent tutoring systems such as SOPHIE [5], WHY [27], WUMPUS [16], and GUIDON [11] [12] 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 [4], and production rules used in WUMPUS and GUIDON),

- network representations of knowledge that capture "importance" (SCHOLAR [9]), "complexity" or "pre-requisite" associations (WUMPUS, BIP [3]), "analogy" and "generalization" relations (WUMPUS),

- and representations that allow for variants on expert performance (for modelling the student) (WEST [8], BUGGY [7]).

In the GUIDON program we have been exploring the problem of using MYCIN's rule set as teaching material. MYCIN [26] 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," which all use the same production rule language and interpreter [29]. 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: 1) formal evaluations indicate that it captures a

high level of expertise [31], and 2) modular design and representational meta-knowledge enable the program to explain its reasoning [13]. 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.

A. The limitations of MYCIN for application to teaching

First, MYCIN is designed to be used as a consultant; its knowledge is too narrow to be used for teaching a student to be a primary diagnostician. The knowledge base is designed to interpret culture results from the blood and the cerebral-spinal fluid. But what expertise suggests that such a culture should be taken? What knowledge does a human draw upon for focusing on bacteremia or meningitis, and what competing hypotheses (and medical tests) need to be considered before MYCIN should even be used? This knowledge is certainly a critical part of teaching infectious disease diagnosis, but MYCIN knows nothing about it.

Second, protocols of experts solving the same cases presented to MYCIN indicate that the program does not organize or use its knowledge the way a human expert does. This result is not surprising, for it is consistent with a half-decade of psychological research into medical problem-solving [18], [24], [20], [21], [28] [14], [17]. If GUIDON, our tutorial program, is 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 diagnostic thinking. In particular, the model must exhibit: focused, forward-directed use of data; trigger associations that suggest new hypotheses; follow-up questions that establish the disease process ("picture of the patient"); and management of a changing "working memory" (hereafter, "differential") of hypotheses under consideration. To this extent, the development of NEOMYCIN is an attempt to synthesize previous research, and to analyze its application to our infectious disease problem domain.

B. Developing a psychological model by modifying EMYCIN

A psychological model of diagnostic thinking cannot be represented using the EMYCIN representation alone, that is, by simply rewriting MYCIN's rules. Instead, the representation and interpreter must be augmented and the rules organized by multiple, orthogonal structures.

For example, a simple interpreter change is to allow incoming data to cause new subgoals to be setup and pursued. Consider the trigger antecedent rule "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 association to meningitis might come to mind, prompting him to determine if the patient has a headache as well. To bring about this effect in NEOMYCIN, a new type of antecedent rule had to be allowed, and a local change made to the EMYCIN control structure.

Besides interpreter changes, different kinds of knowledge had to be separated out of the rules and represented explicitly. Fig. 1 shows a typical (paraphrased) MYCIN rule in which different kinds of knowledge are procedurally embedded.

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

¹ This research has been supported in part by ARPA and ONR contract N00014-79C-0302. Computational resources were provided by the SUMEX-AIM facility (NIH grant RR 00785-07).

- 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
 5) The patient is an alcoholic

Then: there is suggestive evidence that diplococcus-pneumoniae is an organism causing the meningitis.

Figure 1. Typical MYCIN Rule

This rule is 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 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. 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. NEOMYCIN incorporates these psychological aids: 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.

C. A realistic problem-solver needs focusing strategies

As we mentioned above, we can't use MYCIN for teaching about diagnosis because the range of problems it knows about is not realistically wide enough. 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 medical diagnosis search problem 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 taxonomic hierarchy (via the "compiled associations" such as the trigger rule given above). Working from the middle, the physician must first look upwards to focus the possibilities (is it a traumatic process? cancer?) and then refine downwards. 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 doesn't have the knowledge for considering it 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: 1) expanded the etiological knowledge to include broad categories of other, non-infectious problems, such as "toxic problem," "neoplastic problem"; 2) incorporated the focusing strategy of "group and differentiate" so the program could manage this broader range of possibilities; and 3) to enhance the program's ability to apply this strategy we added knowledge about disease processes, knowledge that cuts orthogonally across the etiological taxonomy.

II 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 state categories (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 (Fig. 2) illustrates the various knowledge sources and their relation to the strategic knowledge and differential. These components are shown as icons expanded in subsequent figures. The interpretation of Fig. 2 follows.




Figure 2. Components of the NEOMYCIN system

-- There are four kinds of domain rules: 1) *causal rules* form a net of physiological states and disease categories, ultimately linking raw observations (incoming data) to the etiological taxonomy; 2) *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 descendants); 3) *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 is circumstantial (that is, the associations that "come to mind" are those hypotheses already in the differential, as well as the nodes of the etiological taxonomy which hang below the hypotheses of the differential); 4) *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") which 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 frame-like 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 physiological state is implicated (for example, 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 Fig. 3.

META-RULE397 (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 3. A typical strategy rule.

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.

III A SAMPLE CASE

To illustrate the ideas presented above, a simple case was presented to a stripped-down version of NEOMYCIN (there are generally only one or two rules to establish each hypothesis). Only the collection of circumstantial evidence is shown. 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 strategical rules. It is NEOMYCIN's forward, non-exhaustive reasoning and management of a space of hypotheses that makes it reason more like a human.

(For brevity, the trace of rule conclusions is of the form "Conclude: PARAMETER is VALUE (CF)." The context is always PATIENT-2. Strategical tasks are sometimes abbreviated. Fig. 4 relates meta-rules to tasks.)

7-Dec-80 16:18:25

-----PATIENT-2-----

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) ** J.SMITH	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 J.smith had abnormal neurologic signs?
** 13

Trying evidence rules: LENSIGN -> CHRONIC-MENINGITIS
Trying RULE524/PATIENT-2; RULE524 succeeded.
Conclude: MENINGITIS is ACUTE-MENINGITIS (-.3)
Conclude: MENINGITIS is CHRONIC-MENINGITIS (.3)

Differential:
replace MENINGITIS with CHRONIC-MENINGITIS

Another antecedent rules 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 J.smith 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 J.smith 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-PURSUIING 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.

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 J.smith 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 is 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/PATIENT-2; RULE433 succeeded.
Conclude: INFECTIOUS-PROCESS is BRAIN-ABSCESS (.216)

Observed STOP-PURSUIING METARULE408

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

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

7) Does J.smith have a TB risk factor?
** YES
Trying evidence rules: TBRISK -> TB-MENINGITIS
Trying RULE525/PATIENT-2; RULE525 succeeded.
observed STOP-PURSUIING 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

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.

task METARULE431 succeeded.

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

Trying RULE343/PATIENT-2;
8) Is J.smith 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 J.smith immunosuppressed (e.g. corticosteroid therapy, cytotoxic drug therapy, radiation therapy, leukemia)?
** LEUKEMIA

I will assume that leukemia is one of the diagnoses of J.smith

RULE395 failed due to clause 1

--[0] Finished: LEUKEMIA

Trying evidence rules: LEUKEMIA -> CRYPTOCOCCUS
Trying RULE056/PATIENT-2; 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 = COCCIDIODES
task METARULE411 succeeded.

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

Trying evidence rules: COCCI-ENDEMIC -> COCCIDIODES
RULE570 failed due to clause 1
RULE287 failed due to clause 1
observed STOP-PURSuing METARULE408

Repeating E&R loop: focus = COCCIDIODES

Repeating EHS loop: focus = COCCIDIODES
task METARULE430 succeeded.

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

11) Is J.smith receiving any medications?
** NO

Repeating EHS loop: focus = COCCIDIODES
task RULE430 succeeded.

12) Has J.smith been recently hospitalized?
** NO

Repeating EHS loop: focus = COCCIDIODES

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.

IV THE DIAGNOSTIC META-STRATEGY

Formalizing the diagnostic strategy from protocol analysis was the most difficult part of designing NEOMYCIN. Fig. 4 shows the general outline of the meta-strategy. Each non-terminal 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 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 4. NEOMYCIN's diagnostic meta-strategy
[Rule numbers in brackets appear in the sample typescript]

Fig. 4 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 questions². A great deal of what we might call *heuristic confidence* is placed in the general questions, which constitute the

² "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.

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. (A later publication will discuss NEOMYCIN's meta-strategy in more detail.)

The implementation is in terms of hierarchical meta-rules,^{*} which as a whole constitute the meta-strategy. Fig. 5 illustrates how the rules for a given task are treated as a pure production system--they are repeatedly tried in order, returning to the head of the list when one succeeds, stopping when no rule succeeds or an end condition is true.




Figure 5. Rule-based invocation and interruption of strategic tasks

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 changes the state of the differential, as well as non-exhaustive consideration of hypotheses.

V ETIOLOGICAL TAXONOMY, CAUSAL AND DISEASE PROCESS KNOWLEDGE

Some details of the implementation are given in this section. The etiological taxonomy (Fig. 6) 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*.

Causal knowledge (Fig. 7) is represented as rules marked as being causal, and modified by a certainty factor, as 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 which 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). We are investigating the possibility of using People's "planning links" [22] to more precisely distinguish between causal and subtype links^{**}. Causal rules are used by the "explore and refine" task to work backwards from state-category hypotheses in the differential to prior causes, and ultimately to diagnostic hypotheses in the etiological taxonomy (as shown in Fig. 7).




Figure 6. Portion of etiological taxonomy (links represent specialization of cause)




Figure 7. Portion of causal rule network, showing connection to etiological taxonomy

^{*} So called because they indirectly control the invocation of the domain-dependent object rules. Davis' 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).

^{**} While we might say that an unknown mass lesion (a space-occupying substance) is caused by a tumor, it is more proper to represent a tumor as a kind of mass lesion.

Disease process knowledge is represented as a frame associated with certain 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 typescript before question 6).

VI RELATED RESEARCH

Besides the ICAI projects cited in the introduction, our work has been motivated by previous research in teaching problem-solving strategies (e.g., [19] [6] [30]). We believe NEOMYCIN is the first attempt to formalize a runnable psychological model of diagnostic strategy which can be presented to a student. As should be obvious from our representations, a considerable debt is owed to the medical problem-solving literature, cited above.

Both Reggia [23] and Aikins [1] modified the MYCIN system to make it more acceptable to physicians, particularly to improve knowledge acquisition. Aikins 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 strategical 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.

Other research in cognitive psychology has been helpful to us, particularly studies at the Learning Research Development Center [2], [10], [15] in modelling 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 modelling how a novice composes and generalizes knowledge from experience.

VII 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 [22]. 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" [15]) 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 differential; how context is saved and restored from interrupts; how urgency, cost, and human values factor into the diagnostic process; and so on.

VIII SUMMARY OF WHAT WE LEARNED

In order to teach diagnosis we need a psychological model of problem-solving. In particular, (lesson 1) 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 can order the data and his search. Moreover, an expert has learned "compiled associations" that allow him to efficiently 1) associate hypotheses with data (trigger rules, Pople's bridge concepts), and 2) cast a wide net of questions (general, screening, and follow-up (pinning-down) questions).

Also, (lesson 2) we need to represent the various kinds of knowledge explicitly so that they can be accessible for teaching. Our method is to represent strategical 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 to data points).

In a sense, we have re-discovered the procedural/declarative problem, as have other cognitive psychologists (e.g., Anderson, Rumelhart [25]). We allow for the fact that the expert has composed associations, so he does not think about the justifications of data/hypothesis links, and he makes wide tentative jumps to "bridge concepts." However, we represent these compiled associations declaratively for flexible use, and redundantly store intermediate steps (as text) to allow for explanation of reasoning.

IX FUTURE RESEARCH

Development of NEOMYCIN and GUIDON will proceed in parallel. We intend to compare NEOMYCIN's performance to MYCIN to determine 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 also intend to use our new representation for a computer failure diagnosis consultant, to test the domain-independence of our model.

GUIDON2, a new version of GUIDON, 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.

ACKNOWLEDGEMENTS

We are especially grateful to Dr. Timothy Beckett, MD, for his patient explanations, and willingness to be directed in our discussions. Bruce Buchanan and Bob London have also contributed to the GUIDON project.

REFERENCES

1. Aikins, J. S. *Prototypes and production rules: a knowledge representation for computer consultations*. PhD thesis, Computer Science Department, Stanford University, (STAN-CS-80-814), August 1980.
2. Anderson, J. R., Greeno, J. G., Kline, P. J., & Neves, D. M. Acquisition of problem-solving skill. Carnegie Mellon Technical Report 80-5. To appear in J. R. Anderson (Ed.), *Cognitive Skills and their Acquisition*.
3. Barr, A., Beard, M., Atkinson, R. C. The computer as a tutorial laboratory: the Stanford BIP project. *International Journal of Man-Machine Studies*, 1976, 8, 567-596.
4. Brown, J. S., Burton, R. R., & Zydel, F. *A model driven question-answering system for a CAI environment*. AFHRL-TR-72-39. Lowry AFB, CO: Technical Training Division, Air Force Human Resources Laboratory, March 1973.
5. Brown, J. S., & Burton, R. Multiple representations of knowledge for tutorial reasoning. In D. G. Bobrow & A. Collins (Eds.), *Representation and Understanding*. New York: Academic Press, 1975, 311-349.
6. Brown, J. S., Collins, A., & Harris, G. Artificial intelligence and learning strategies. In O'Neill (Ed.), *Learning Strategies*. New York: Academic Press, 1977.
7. Brown, J. S., & Burton, R. R. Diagnostic models for procedural bugs in basic mathematical skills. *Cognitive Science*, 1978, 2(2), 155-192.
8. Burton, R. R. An investigation of computer coaching for informal learning activities. *The International Journal of Man-Machine Studies*, 1979, 11, 5-24.
9. Carbonell, J. R. *Mixed-initiative man-computer instructional dialogues* (BBN Report No. 1971). 1970.
10. Chi, M. T. H., Feltovich, P. J., Glaser, R. *Representation of physics knowledge by experts and novices*. University of Pittsburgh, Learning Research and Development Center, Technical Report No. 2, March 1980.
11. Clancey, W. J. Tutoring rules for guiding a case method dialogue. *International Journal of Man-Machine Studies*, 1979, 11, 25-49.
12. Clancey, W. J. *Transfer of rule-based expertise through a tutorial dialogue*. PhD thesis, Computer Science Department, Stanford University, (STAN-CS-769), August, 1979.
13. Davis, R. *Applications of meta-level knowledge to the construction, maintenance and use of large knowledge bases*. PhD thesis, Computer Science Department, Stanford University, (STAN-CS-76-552), July 1976.
14. Elstein, A. S., Shulman, L. S., & Sprafka, S. A. *Medical problem-solving: An analysis of clinical reasoning*. Cambridge: Harvard University Press, 1978.
15. Feltovich, P. J., Johnson, P. E., Moller, J. H., & Swanson, D. B. *The role and development of medical knowledge in diagnostic expertise*. Presented at the 1980 Annual meeting of the American Educational Research Association.
16. Goldstein, I. *Developing a computational representation for problem solving skills*. Proceedings of the Carnegie-Mellon Conference on Problem Solving and Education: Issues in Teaching and Research, October 9-10, 1978.
17. Kassirer, J. P. & Gorry, G. A. Clinical problem solving: a behavioral analysis. *Annals of Internal Medicine*, 1978, 89, 245-255.
18. Miller, P. B. *Strategy selection in medical diagnosis*. Massachusetts Institute of Technology, MAC TR-153, September 1975.
19. Papert, S. Teaching children programming. *IFIP Conference on Computer Education*. Amsterdam: North Holland, 1970.
20. Pauker, S. G., Kassirer, J. P., & Schwartz, W. B. Toward the simulation of clinical cognition: taking the present illness by computer. *The American Journal of Medicine*, 1976, 60, 981-995.
21. Pauker, S. G. & Szolovits, P. Analyzing and simulating taking the history of the present illness: Context Formation. In Schneider/Segvall Hein (Eds.), *Computational Linguistics in Medicine*. North-Holland Publishing Company, 1977, 109-118.
22. Pople, H. E. *Heuristic methods for imposing structure on ill-structured problems: the structuring of medical diagnostics*. In Artificial Intelligence in Medicine, AAAS Symposium Series, Westview Press (forthcoming 1981).
23. Reggia, J. A. A production rule system for neurological localization. *Proceedings of the Second Annual Symposium on Computer Applications in Medical Care*, Silver Spring, Maryland, 1978, 254-260.
24. Rubin A. D. *Hypothesis formation and evaluation in medical diagnosis*. Artificial Intelligence Laboratory, MIT, Technical Report AI-TR-316, January 1975.
25. Rumelhart, D. E. & Norman, D. A. *Analogical processes in learning*. University of California, San Diego, Technical report 8005, September 1980.
26. Shortliffe, E. H. *Computer-based medical consultations: MYCIN*. New York: Elsevier, 1976.
27. Stevens, A. L., & Collins, A. Multiple conceptual models of a complex system (BBN Report No. 3923). In R. Snow, P. Federico and W. Mantague (Eds.), *Aptitude, learning and instruction: Cognitive process analysis*, 1978.
28. Swanson, D. B., Feltovich, P. J., & Johnson, P. E. Psychological analysis of physician expertise: implications for design of decision support systems. *MEDINFO77*, 1977, 161-164.
29. van Melle, W. *A domain-independent production-rule system for consultation programs*. PhD thesis, Computer Science Department, Stanford University, (STAN-CS-820), June 1980.
30. Wescourt, K. T., & Hemphill, L. (1978). *Representing and teaching knowledge for troubleshooting/debugging*. Institute for Mathematical Studies in the Social Sciences, (Tech. Rep. 292), Stanford University, 1978.
31. Yu, V. L., Fagan, L. M., Wraith, S. M., Clancey, W. J., Scott, A. C., Hannigan, J. F., Blum, R. L., Buchanan, B. G., Cohen, S. N. Antimicrobial selection by a computer -- a blinded evaluation by infectious disease experts. *Journal of the American Medical Association*, 1979, 242, 1279-1282.