Journal of Computer-Based Instruction
Summer 1983, Vol. 10, Nos. 1 & 2, 8-15.

GUIDON

William J. Clancey
Stanford University

GUIDON is an intelligent computer-aided instructional (ICAI) program for teaching diagnosis, such as medical diagnosis. The program is general. Without reprogramming, the program can discuss with a student any diagnostic problem that it can solve on its own. Moreover, by substituting problem-solving knowledge from other domains, the program can immediately discuss problems in those domains. This power derives from the use of Artificial Intelligence methods for representing both subject material and knowledge about how to teach. These are represented independently, so the teaching knowledge is general. There are teaching rules and procedures for: determining what the student knows, responding to his partial solution, providing hints, and opportunistically interrupting to test his understanding. Experience with GUIDON reveals the importance of separating out casual and strategic knowledge in order to explain diagnostic rules and to teach a reasoning approach. These lessons are now guiding the development of new representations for teaching.

GUIDON, a program for teaching diagnostic problem-solving, is being developed by William J. Clancey and his colleagues at Stanford University. Using the rules of the MYCIN consultation system (Shortliffe, 1976) as subject material, GUIDON engages a student in a dialogue about a patient suspected of having an infection. In this manner, it teaches the student about the relevant clinical and laboratory data and about how to use that information for diagnosing the causative organism. GUIDON's mixed-initiative dialogue differs from that of other ICAI programs in its use of prolonged, structured teaching interactions that go beyond responding to the student's last move [as in WEST (Burton and Brown, 1979) and WUMPUS (Goldstein, 1979)] and repetitive questioning and answering [as in SCHOLAR (Carbonell, 1970) and WHY (Stevens, Collins, and Goldin, 1982)].

MYCIN's infectious-disease diagnosis rules constitute the skills to be taught. As applied to a particular problem, the rules provide GUIDON with topics to be discussed and with a basis for evaluating the student's behavior. GUIDON's teaching knowledge is wholly separate from MYCIN. It is stated explicitly in the form of 200 tutorial rules, which include methods for guiding the dialogue economically, presenting diagnostic rules, constructing a student model, and responding to the student's initiative. Because of the separation of teaching and domain knowledge, MYCIN's infectious-disease knowledge base can be replaced by diagnostic rules for another problem domain.

The large and complex MYCIN knowledge base provides a unique opportunity to apply and extend ICAI technology for student modeling and mixed-initative dialogue. GUIDON is designed to explore two basic questions: First, how do the problem-solving rules, which perform so well in the MYCIN consultation system, measure up to the needs of a tutorial interaction with a student? Second, what knowledge about teaching might be added to MYCIN to make it into an effective tutorial program? MYCIN's rules have not been modified for the tutoring application, but they are used in new ways, for example, for making up quizzes, guiding the dialogue, summarizing evidence, and modeling the student's understanding.

Several design guidelines for the rules make it plausible that the rules would be a good vehicle for teaching. First, they are designed to capture a significant part of the knowledge necessary for good problem solving. Formal evaluation of MYCIN demonstrated that its competence in selecting antimicrobial therapy for meningitis and for bacteremia is comparable to that of the members of the infectious-disease faculty at the Stanford University School of Medicine (where MYCIN was developed; see Yu et al., 1979). Second, flexible use of the rule set is made possible by the provision of representational metaknowledge, which allows a program to take apart rules and to reason about the components (this knowledge describes the number and type of arguments of primitive functions in the rule language). Finally, MYCIN's rules, in contrast with Bayesian programs, are couched in terms familiar to human

Revised article reprinted with permission from *The Handbook of Artificial Intelligence, Vol. II*, A. Barr and E.A. Feigenbaum (Eds.), William Kaufmann, Inc., 1982.
Reprint requests should be addressed to William J. Clancey, Department of Computer Science, Stanford University, Stanford, CA 94305.

experts, so it seems likely that reading back MYCIN's line of reasoning to a student might be helpful to him (or her).

After a brief overview of MYCIN, this article discusses the following aspects of a GUIDON tutorial dialogue:

1. The nature of the interaction
2. The components of the student model
3. The organization of teaching knowledge into discourse procedures
4. The use of the student model
5. Opportunistic tutoring
6. Pedagogical principles behind the tutoring rules.

The capability of GUIDON to tutor from a library of cases and for domains outside of medicine is also discussed. The final section outlines the lessons learned about knowledge representation that are being applied to reconfigure the MYCIN rule base for its use in teaching.

Overview of MYCIN

MYCIN is a program that was developed by a team of physicians and AI specialists. The program was designed to advise nonexperts in the selection of antibiotic therapy for infectious diseases. The knowledge base consists of approximately 450 rules that deal with diagnosis of bacteremia, meningitis, and cystitis infections. The rules are applied by backward chaining, working from high-order goals, such as "Determine whether the patient requires treatment," down to more specific subgoals, such as "Determine whether the patient has high risk for tuberculosis." These goals and subgoals become the "topics" of a dialogue with GUIDON. A typical rule is, roughly stated, "If the patient has been receiving steroids, then his risk for tuberculosis meningitis is increased." The rules are modified by a *certainty factor* (CF), indicating the rule author's degree of belief, on a scale from − 1 to 1 that the conclusion holds when the premise is know to be true. (In the GUIDON excerpts shown below, the CFs are shown in parentheses, e.g., "(.95).") In a MYCIN consultation, the rules are chained together, working downward from the high-order goals. The program asks a queston when it needs more case data to apply a rule. Thus, a tree of goals and rules is constructed: The goals are *OR nodes* (any of a number of rules may help determine a goal) and the rules are *AND nodes* (all of the subgoals referenced in the premise must be known for the rule to apply). We call this *AND/OR tree* with rule evaluations and final conclusions about goals the ***solution*** of the diagnostic problem.

INTERACTION WITH GUIDON

An essential part of tutorial dialogue management is allowing the student to express himself. In a mixed-initiative program, provision must be made for every *kind* of initiative on the part of the student. This includes referring back to an earlier topic to provide more details, changing the topic, requesting case data, posing a hypothesis, asking for help, and so on. We might summarize this by saying that we must allow the student to specify what he knows, what he wants to know more about, and what he wants to ignore. The annotated protocol excerpted below illustrates GUIDON's flexibility in responding to the student's initiative. To this point in the dialogue, the student has been given basic data about the patient. He has already determined that the patient has an infection, has evidence that it might be meningitis, and is trying to determine what is causing the meningitis ("the type of the infection"). Student input (indicated here by **) is in the form of commands from a menu, discussed further below.

> *The student asks for the data and subgoals relevant to the topic being discussed. These are extracted from MYCIN's rules.*

**Factors

The following factors will be useful to determine the type of the infection:

3a. whether the meningitis is partially treated
3b. whether a lumbar puncture has been performed on Pt538
3c. the duration of the neurological signs
3d. the CSF protein
(*) 3e. whether Pt538 has a rash consistent with Herpes Zoster

> *GUIDON keeps track of what has been discussed, MYCIN's use of data, and conjectures about what the student knows.*

Factors marked with (*) have already been stated. Based on what we have already mentioned. you should be able to conclude about: 3b.

> *The student asks how a particular datum is useful. He is given the case-specific information and then a summary of its use in the rule or rules that apply in this case (negative conclusions are omitted here because the choices are mutually exclusive).*

** Use 3c

The duration of the neurological signs is 7.2 hours.

Using other information available to you now, the duration of the neurological signs enables you to conclude that the type of the infection is bacterial (.2) viral (.2) [RULE524].

> *The student indicates that he has determined a particular subgoal. As described in this article, a "discourse procedure" is invoked to respond to the student. If the student model were not consistent with his claim, he would be asked to state his conclusion and then possibly to support it.*

** IKNOW 3b

Good. MYCIN has also made a decision.

The site of CULTURE-1 is evidence that a lumbar puncture has been performed on Pt538 [RULE112].

Back to our discussion of the type of the infection...

At this point, the topic of discussion is completed and the student is receiving a summary of the relevant conclusions. The program is automatically extracting the "key factor" from each rule and displaying only the "interesting" (useful) conclusions.

Summary of evidence that the type of the infection is bacterial (.98):

29a. The fact that the meningitis is partially treated (.95) is evidence that the type of the infection is bacterial (.19) [RULE148]

29b. The WBC from the CSF and the percent of PMNs in the CSF is evidence that the type of the infection is bacterial (.9) [RULE501]

The student requests more information about one of the inferences. The program reasons that because several subgoals have not been discussed, these are the details to offer first.

** Details 29a

For which of the following factors would you like to see a review of the evidence?

(Select from the following):

1. the infection which requires therapy
2. the fact that the meningitis is partially treated (.95)

**2

Summary of evidence that the meningitis is partially treated (.95):

32a. The infection which requires therapy, the fact that organisms were not seen on the stain of the pending CSF culture, and the time since therapy with the cephalothin was started are evidence that the meningitis is partially treated (.95) [RULE145]

Given this information, perhaps the student would like to see the details of how it is used.

Do you want to see RULE148?

** No

GUIDON provides a menu of options for requesting case data, asking for MYCIN's evaluation of the problem (e.g., "What subgoals are PENDING?" "Give me DETAILS"), determining dialogue context (e.g., "What RULE are we discussing?"), changing the topic, requesting assistance (the options HELP, HINT, and TELLME), and conveying what is known (e.g., "I want to make a HYPOTHESIS"). The menu of over 30 options allows for input to be terse, while defining clearly for the student what the program can understand. As arguments to the options, the student can use phrases (e.g., "IKNOW about the lumbar puncture"), keywords (e.g., "IKNOW LP"), or indices of remarks made by the program (e.g., "IKNOW 3B"). All of the output text is generated from short phrases ("the following factors," "the CSF protein," "is evidence that") with verb tense and number adjusted according to context. GUIDON's initiatives involve probing the student's understanding (if a question or hypothesis is unexpected), offering overviews and summaries, introducing new topics when rules are being discussed and suggesting that a topic be terminated. These capabilities are discussed in the sections below on alternative dialogues, responding to partial solutions, and opportunistic tutoring.

THE STUDENT MODEL

Before a session with the student begins, GUIDON uses MYCIN to "solve" the case to be presented to the student. The results of this background consultation, consisting of MYCIN's rule conclusions and its records of why rules did not apply, are reconfigured into an explicit AND/OR tree of goals and rules so that the rules are indexed both by the goals they conclude about and the subgoals or data needed to apply them. During the tutorial session, as the student inquires about the patient and receives more case data, this same information is used to drive MYCIN's rules in a forward direction. Thus, at any time, some of the rules MYCIN uses for determining, say, the type of the infection, will have led to a conclusion, while others will require more information about the patient before they can be applied.

This record of what the expert (i.e., MYCIN) "knows" at any time during the student-run consultation forms the basis for evaluating a student's partial solutions and providing assistance. Such an overlay model (See Carr and Goldstein, 1977) assumes that the student's knowledge is a subset of MYCIN's knowledge and that there are unique reasoning steps for making any particular deduction. Neither assumption is always correct; the rule set nevertheless provides a first-order approximation to the student-modeling problem.

The three components of the student model are shown in Figure 1. The three components are stored as properties of each rule in the knowledge base. The first component, the cumulative record of whether a student knows a rule, is called the USE-HISTORY property of the rule. It is the program's belief that, if the student were given the premise of the rule, he would be able to correctly, in the abstract, draw the proper conclusion. USE-HISTORY is primed by the student's initial indication of his level of expertise, which is matched against "difficulty ratings" associated with each rule. Like the other two components, the USE-HISTORY property of a rule is represented as a certainty factor (the same belief measure used in MYCIN's rules) that combines the background evidence with the implicit evidence stemming from needs for assistance and verbalized partial solutions, as well as the explicit evidence stemming from a direct question that tests knowledge of the rule.

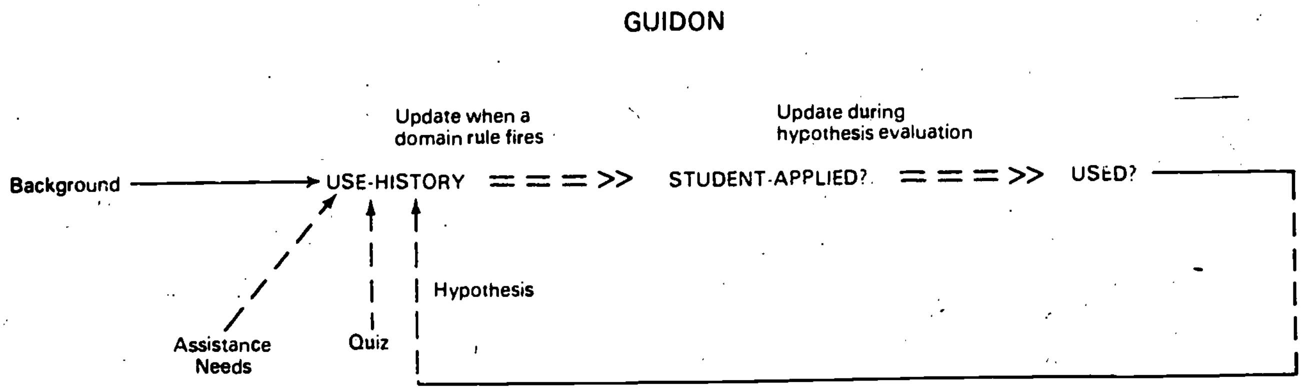

Figure 1. Maintenance relations for student-model components (Clancey, 1979b).

The second component, called STUDENT-APPLIED?, records the program's belief that the student is able to apply the rule to the given case, that is, that the student would refer to this rule to support a conclusion about the given goal. Thus, there is a distinction between knowing a rule (USE-HISTORY) and being able to apply it, since the student may know which subgoals appear in the rule but be unable to achieve them. STUDENT-APPLIED? is determined once for each rule during a case at the time MYCIN is able to apply the rule. (The evidence considered is: Is it believed that the student knows the rule [USE-HISTORY]? Was the rule mentioned during this sesson? Has it been discussed in previous tutorials? Is there a subgoal that the student is not believed to be able to determine?)

The third component of the student model, called USED?, is relevant whenever the student states a partial solution (a list of possible diagnoses, not intended to be complete). It records the program's belief that the student would mention a rule if asked to support his partial solution. This component combines indirect evidence by comparing conclusions made by rules with the student's conclusions, the record of what rules the student is believed to be able to use (STUDENT-APPLIED?), and evidence that the student may have remembered to apply the rule in this case (e.g., the rule mentioned earlier in the dialogue). This combined evidence affects how the program responds to the partial solution and feeds back into the USE-HISTORY component of the student model.

Discourse Procedures and Alternative Dialogues

The student is allowed to explore MYCIN's reasoning by using options like FACTORS, shown earlier in the protocol excerpt. However, the tutor is not a simple, passive, information-retrieval system. In addition to clearly laying out data and inferences, the tutor has to reason about what constitutes reasonable, expected elaboration on the basis of what has been previously discussed. For GUIDON's rule-based approach, this takes the form of selecting which rules and rule clauses to mention and deciding whether to introduce a goal for detailed discussion or just to offer a summary of evidence. In the excerpt, for example, GUIDON provided details for an inference (rule 148) by offering to support achieved preconditions that were not mentioned in the tutorial dialogue up to that point.

Similarly, when the student takes the initiative by saying he has determined some subgoal, the tutor needs to determine what response makes sense, based on what it knows about the student's knowledge and shared goals for the tutorial sesson (topics or rules to discuss). The tutor may want to hold a detailed response in abeyance, simply acknowledge the student's remark, or probe him for evidence that he does indeed know the fact in question. Selection among these alternative dialogues might require determining what the student could have inferred from previous interactions and the current situation. In the dialogue excerpt shown above, GUIDON decides that there is sufficient evidence that the student knows the solution to a relevant subproblem so that detailed discussion and probing are not necessary.

Decoupling domain expertise from the dialogue program, an approach used by all ICAI systems, is a powerful way to provide flexible dialogue interaction. In GUIDON, discourse procedures formalize how the program should behave in general terms, not in terms of the data or outcome of a particular case. A discourse procedure is a sequence of actions to be followed under conditions determined by the complexity of the material, the student's understanding of the material, and tutoring goals for the session. Each option available to the student generally has a discourse procedure associated with it.

For example, if the student indicates, via the IKNOW option, that he has a hypothesis about some subgoal but MYCIN has not yet been able to make a decision, the procedure for requesting and evaluating a student's hypothesis is invoked. Otherwise, if MYCIN has reached the same conclusion, the procedure for discussing a completed topic is followed. Whether or not the student will be probed for details depends on the model that the tutor is building of the student's understanding (considered below).

COMPLETEDGOAL.PROC005

Purpose: Discuss final conclusion for a goal.

Step 1: Decide whether to finish with a summary.

Step 2: Discuss final hypothesis for the goal.

Step 3: Wrap up discussion or record completion.

Figure 2. Discourse procedure for completing a goal discussion.

The procedure for ending discussion of a topic is paraphrased in Figure 2. Conditional actions in discourse procedures are expressed as tutoring rules (t-rules). T-rules decide whether an action should be taken, and when this involves invoking another discourse procedure, other t-rules will decide *what* should be said. For example, the second step of the procedure COMPLETEDGOAL decides whether to give the student the answer or to ask him to make a hypothesis. Figure 3 shows the t-rule that caused GUIDON to acknowledge the student's statement about what he knew in the dialogue illustrated above, rather than ask for details. To ask about and evaluate the student's hypothesis, another discourse procedure would have been invoked. Of course, the discourse procedure for discussing a completed topic is invoked from many other procedures besides the one corresponding to the IKNOW option: It may be invoked in the course of giving details about how a subgoal is determined, in responding to a student's hypothesis for a subgoal, when the program detects that the current subgoal (topic) is substantially completed (enough data have been given to make a strong conclusion), and so on.

T-RULE5.02 *Directly state single, known rule.*

IF
1) There are rules having a bearing on this goal that have succeeded and have not been discussed, and
2) The number of rules having a bearing on this goal that have succeeded is 1, and
3) There is strong evidence that the student has applied this rule

THEN Simply state the rule and its conclusion

Figure 3. T-rule for deciding how to complete discussion of a topic.

Responding to Partial Solutions

Shown below is an annotated excerpt demonstrating how the program responds to partial solutions stated by the student. Tutoring rules are used both to update the model and to select the appropriate discourse procedures for responding. In this excerpt, while the student's hypothesis for a subproblem is correct, nevertheless, after analyzing the previous history of interaction with the student, GUIDON is not sure that the student really knows why his hypothesis is correct. Therefore, instead of accepting this answer, GUIDON requests relevant factors to support the hypothesis and probes further to see if the student knows why an alternative hypothesis is discredited.

> *GUIDON asks the student to state a hypothesis for a subgoal. The student may have volunteered to do this, or the program may be testing the student's understanding.*

What is the type of the infection?

** Bacterial

> *The program asks the student to support this hypothesis either because it is incorrect or, as in this case, the student model indicates that the student may not know all of the relevant rules. (GUIDON changes the prompt for student input to ">>" to indicate that specific input is expected and student-initiated options are not.)*

What facts about this case tell you that the type of the infection is bacterial?
(Enter one factor per line.)

\>> PMNs

The percent of PMNs in the CBC is not relevant for determining the type of the infection.

\>> PMNs in the CSF
\>> Cerebral-fluid glucose value
\>> Protein of the CSF

> *These factors are related to the rules used by MYCIN to conclude that the infection is bacterial, and the student model is updated accordingly.*

You're right about the CSF protein, the percent of PMNs in the CSF, and the CSF glucose value; however, you should have mentioned the fact that the meningitis is partially treated, the WBC from the CSF, and the fact that the simultaneous blood glucose value is not known.

> *Other hypotheses are then discussed.*

MYCIN rejected some possibilities you did not mention:
What facts about this case tell you that the type of the infection is not viral?

Referring to the rules that MYCIN uses to solve subproblems (such as determining whether a meningitis infection is bacterial, fungal, viral, or tuberculous), GUIDON decides which of these rules, if any, might have been used by the student. That is, what inference chains are consistent with the student's behavior? This analysis is complicated by the fact that a particular hypothesis about the problem may be indicated by more than one rule, or negative evidence may outweigh positive evidence.

A potential weakness of the GUIDON program is that it attempts to explain the student's behavior solely in terms of MYCIN's rules. If the student is basing his questions and hypotheses on incorrect rules, GUIDON is not able to formulate these rules and address them directly. It is possible as well that the student's concepts are different from MYCIN's, so his conclusions might be correct, but he will want to support them with reasoning that is different from MYCIN's. This could involve something as simple as wanting to refer to the patient's age in general terms (infant, adolescent), while MYCIN recognizes only precise, numerical ages.

Modeling medical reasoning in terms of an alternative rule set (not just a subset of MYCIN's rules) is a theory-formation problem that goes beyond the current capabilities of AI. It is possible that the approach followed by Stevens, Collins, and Goldin (1982) of collecting data about student misconceptions and then incorporating these variations into the modeling process will prove tenable for the medical domain.

Opportunistic Tutoring and Pedagogical Style

It is sometimes advantageous for the tutor to take the initiative to present new material to the student. This requires that the tutor have presentation methods that opportunistically adapt material to the needs of the dialogue. In particular, the tutor has to be sensitive to how a tutorial dialogue fits together, including what kinds of interruptions and probing are reasonable and expected in this kind of discourse. GUIDON demonstrates its sensitivity to these concerns when it corrects the student before quizzing him about "missing hypotheses," asks him questions about recently mentioned data to see if he understands how to use them, quizzes him about rules that are related (by premise and action) to one that has just been discussed, follows up on previous hints, and comments on the status of a subproblem after an inference has been discussed ("Other factors remain to be considered...").

There are many subtle issues — when to interrupt the student, how much to say, and the like — that constitute a pedagogical style and are implicit in GUIDON's teaching rules. For example, several tutoring rules in different situations may present short orientation lectures, but nowhere does GUIDON reason that its interaction will be of the tutorial type, which provides orientation when appropriate, in contrast with the coaching type (e.g., Burton and Brown, 1979), which only makes interruptions. For this reason, it is useful to summarize the set of tutoring principles that appear implicitly in the tutoring rules:

1. *Be perspicuous:* Have an economical presentation strategy, provide lucid transitions, and adhere to conventional discourse patterns.
2. *Provide orientation to new tasks by top-down refinement:* Provide the student with an organized framework of considerations he should be making, without giving away the solution to the problem (important factors, subgoals, size of the task), thus challenging the student to examine his understanding constructively.
3. *Strictly guide the dialogue:* Say when topics are finished and inferences are completed, as opposed to letting the student discover transitions for himself.
4. *Account for incorrect behavior in terms of missing expertise* (as opposed to assuming alternative methods and strategies): Explain clearly what is improper from the tutor's point of view (e.g., improper requests for case data). This is, of course, more a statement of how GUIDON models the student than a principle of good teaching.
5. *Probe the student's understanding when you are not sure what he knows,* when you are responding to partial student solutions: Otherwise, directly confirm or correct the solution.
6. *Provide assistance by methodically introducing small steps* that will contribute to the problem's solution:
   a. Assistance should at first be general, to remind the student of solution methods and strategies he already knows;
   b. Assistance should encourage the student to advance the solution by using case data he has already been given.
7. *Examine the student's understanding and introduce new information* whenever there is an opportunity to do so.

Case and Domain Independence

Patient cases are entered into the MYCIN system for receiving a consultation or for testing the program, so the case library is available to GUIDON at no cost. This provides over 100 patients that GUIDON can discuss, clearly demonstrating the advantage that ICAI has over the traditional computer-based-instruction approach in which each lesson must be designed individually.

Besides being able to use the teaching procedures to tutor different cases, GUIDON can provide tutorials in any problem area for which a MYCIN-like knowledge base of decision rules and fact tables has been formalized (see van Melle, 1980). This affords an important perspective on the generality of the discourse and pedagogical rules.

Experimental tutorials using knowledge bases in two other domains — structural analysis (SACON) and pulmonary function diagnosis (PUFF) — have revealed that the effectiveness of discourse strategies for carrying on a dialogue economically is determined in part by the depth and breadth of the reasoning tree for solving problems, a characteristic of the rule set for each domain. When a solution involves many rules at a given level, for example, when there are many rules to determine the organism causing the infection, the tutor and student will not have time to discuss each rule in the same degree of detail. Similarly, when inference chains are long, an effective discourse strategy will entail summarizing evidence on a high level, rather than considering each subgoal in the chain.

RESULTS

GUIDON demonstrated that teaching knowledge could be treated analogously to the domain expertise of consultation systems: It can be codified in rules and built incrementally by testing it on different cases. The framework of tutoring rules organized into discourse procedures worked well, indicating that it is suitable to think of a tutorial dialogue as being separated into relatively independent sequences of interaction. Moreover, the judgmental knowledge for constructing a student model can also be captured in rules utilizing certainty factors, showing that the task of modeling a student bears some relation to MYCIN's task of diagnosing a disease.

In contrast to GUIDON's teaching knowledge, the evaluation of MYCIN's rule set for this application was not so positive. While MYCIN's representational meta-knowledge made possible a wide variety of tutorial activity, students find that the rules are difficult to understand, remember, and incorporate into a problem-solving approach. These difficulties prompted an extensive study of MYCIN's rules to determine why the teaching points were not as clear as had been expected. GUIDON researchers discovered that important structural knowledge (hierarchies of data and diagnostic hypotheses) and strategic knowledge (searching the problem space by top-down refinement) were implicit in the rules. That is, the choice and ordering of rule-premise clauses constitute procedural knowledge that brings about good problem-solving performance in a MYCIN consultation but is unavailable for teaching purposes. Rather than teaching a student problem-solving steps (rule clauses) by rote, it is advantageous to convey an approach or strategy for bringing those steps to mind — the plan that knowledge-base authors were following when they designed MYCIN's rule set. To make this implicit design knowledge explicit, a new system, NEOMYCIN[1](Clancey and Letsinger, 1981), is being developed that separates out diagnostic strategy from domain knowledge and makes good use of hierarchical organization of data and hypotheses.

Moreover, besides reconfiguring MYCIN's rules so that knowledge is separated out and represented more declaratively, it is necessary to add knowledge about the justification of rules. Justifications are important as mnemonics for the heuristic associations, as well as for providing an understanding that allows the problem solver to violate the rules in unusual situations.

Finally, NEOMYCIN has additional knowledge about disease processes that allows it to use the strategy of "group and differentiate" for initial problem formulation, in which the problem solver must think about broad categories of disorders and consider competing hypotheses that explain the problem data. Thus, we want to teach the student the knowledge a human would need to focus on infectious-disease problems in the first place, essentially the knowledge (previously unformalized) that a human needs to use MYCIN appropriately.

In conclusion, GUIDON research sets out to demonstrate the advantages of separate, explicit representations of both teaching knowledge and subject material. The problems of recognizing student misconceptions aside, this research demonstrated that simply representing in an ideal way what to teach the student is not a trivial, solved problem. An unstructured set of production rules is inadequate. GUIDON's teaching rules are organized into procedures; NEOMYCIN's diagnostic rules are hierarchically grouped by both premise and action and are controlled by meta-rules. GUIDON research demonstrated that the needs of tutoring can serve as a "forcing function" to direct research toward more psychologically valid representations of domain knowledge, which potentially will benefit those aspects of expert-systems research that require human interaction, particularly explanation and knowledge acquisition.

[1]GUIDON is described fully by Clancey (1979b); a shorter discussion is given in Clancey (1979a). Clancey and Letsinger (1981) describe the NEOMYCIN research. The study of MYCIN's rule base leading up to this new system and some methodological considerations are provided by Clancey (1983, in press a).

REFERENCES

Burton, R.R., & Brown, J.S. An investigation of computer coaching for informal learning activities. *International Journal of Man-Machine Studies*, 1979, 11, 5-24.

Carbonell, J.R. AI in CAI: An artificial intelligence approach to computer-aided instruction. *IEEE Transactions on Man-Machine Systems*, 1970, 4, 190-202.

Carr, B., & Goldstein, I. *Overlays: A theory of modeling for computer aided instruction*. AI Memo 406. AI Laboratory, Massachusetts Institute of Technology, 1977.

Clancey, W.J. Tutoring rules for guiding a case method dialogue. *International Journal of Man-Machine Studies*, 1979a, 11, 25-49.

Clancey, W.J. *Transfer of rule-based expertise through a tutorial dialogue* (Report No. STAN-CS-769). Stanford University: Computer Science Department, 1979b. (Doctoral dissertation)

Clancey, W.J. The epistemology of a rule-based expert system: A framework for explanation. *Artificial Intelligence*, 1983, 20 (3), 215-251.

Clancey, W.J. Methodology for building an intelligent tutoring system. To appear in W. Kintsch, J.R. Miller, & P.G. Polson (Eds.), *Method and tactics in cognitive science*. Lawrence Erlbaum Associates, in press-a.

Clancey, W.J., & Letsinger, R. NEOMYCIN: Reconfiguring a rule-based expert system for application to teaching. *Proceedings of the 7th International Joint Conference on Artificial Intelligence*, Vancouver, B.C., 1981, 829-836.

Goldstein, I. The genetic epistemology of rule systems. *International Journal of Man-Machine Studies*, 1979, 11, 51-77.

Shortliffe, E.H. *Computer-based Medical Consultations: MYCIN*. New York: American Elsevier, 1976.

Stevens, A., Collins, A., and Goldin, S.E. Misconceptions in students' understanding. In D. Sleeman and J.S. Brown (Eds.), *Intelligent Tutoring Systems*. London: Academic Press, 1982.

van Melle, W. *A domain independent system that aids in constructing consultation programs* (Report No. STAN-CS-80-820). Stanford University: Computer Science Department, 1980. (Doctoral dissertation)

Yu, V.L., Buchanan, B.B., Shortliffe, E.H., Wriath, S.M., Davis, R., Scott, A.C., & Cohen, S.N. Evaluating the performance of a computer-based consultant. *Computer Programs in Biomedicine*, 1979, 9, 95-102.