DIALOGUE MANAGEMENT FOR RULEBASED TUTORIALS

William J Clancey
Computer Science Department
Stanford University; Stanford, CA 94305

Mixed-initiative discussion requires that a tutorial program have the means to manage its share of the dialogue. Dialogue management includes operations ranging from record keeping and context focusing to heuristics for directing the dialogue economically according to the needs of the student. This paper illustrates how knowledge for generating remarks in a tutorial dialogue about rule-based domain knowledge can be represented conceptually as a network consisting of transitions between dialogue situations in which the links are various kinds of management heuristics. The chief finding is that a network of procedures is a useful representation for organiiing heuristics for carrying on a structured dialogue.

1 Introduction

This paper discusses a kind of mixed-Initiative tutorial dialogue that concerns a single, complex task to be solved by a student under a program's guidance Sequences of student/tutor remarks can be grouped into "discourse situations" or recurrent patterns In the discussion, making the mixed-initiative nature of this tutorial more complex than in previous work Typical discourse situations are: examining the student's understanding after he asks a question that shows unexpected expertise, relating an inference to one just discussed, giving advice to the student after he makes an hypothesis about a subproblem, and so on

The general problem of sharing initiative and making provisions to carry out one's discourse goals is here termed 'dialogue management" We illustrate principles for taking initiative that are desirable for tutoring rule-based diagnostic problems An important consideration is the possibility of 'alternative dialogues" We focus on the problem of making appropriate, prolonged presentations that go beyond interruption and repetitive question/answering, but use them as components in a larger scheme

Most research with Intelligent Computer Aided Instruction (1CA1) programs has emphasized the use of an underlying expert artificial Intelligence program for generating subject material to present to a student and for modelling his understanding. Indeed, researchers call this work "generative" CAI to underscore the constructive operation of creating questions or tasks for a student and creating a differential model that relates the student's behavior to the program's model of expertise [18] While one distinguishing feature of these programs is frequently cited to be their capability to engage the student in a mixed-initiative dialogue [I6J [17], we find that the "natural language understanding" In these programs Involves parsing student Input [5], not conversational interaction.

In the "reactive environment" of the early SOPHIE lab [2], the tutor never takes the Initiative at all; and in the games of WEST [6] and WUMPUS [13], the tutor's remarks are all Interruptions or reactions to the immediately preceding move taken by the student SCHOLAR [10] and WHY [22], the geography and meteorology tutors, follow a Socratic dialogue format of repetitive questioning using topic selection rules and strategies for testing a student's understanding; the session itself Involves no overarching problem or shared task to be discussed Thus, ICAI research has emphasized the use of domain expertise for modelling the student and tutoring principles for correcting his misconceptions, but it has dealt only tangentially with the problems of carrying on a prolonged, *purposeful* tutorial dialogue that as a whole considers a single task.

The discourse problems dealt with here include maintaining and sharing context as solution of the task proceeds, and providing means for the student to express initiative as he unfolds the complexity of the problem and encounters limitations in his understanding. Researchers who specialize In Natural Language studies have reported various theoretical aspects of task-oriented dialogues, such as recognizing and generating Intentions and plans [12] [7] and focusing on objects and subtasks [II] [15]. We distinguish our work from other Natural Language research in two ways. First, we view the discourse problem as requiring (a) Interpretation of intent and (b) reasoning about this understanding in order to act. We have minimized the Interpretation problem by restricting the student to a command-oriented input, allowing us to concentrate on the tutorial knowledge for generating remarks and guiding the dialogue Second, we are formalizing tutorial knowledge In terms of specific performance rules for various situations, rather than studying human dialogues and enumerating general dialogue Interaction principles.

In this paper, after setting the scene by a short overview of rule-based tutoring, as exemplified by a particular system, GUIDON, we present a transition diagram that we find useful for illustrating dialogue management issues Examples of the system in operation are given The second half of the paper deals with the problem of topic relationships, for example, how the dynamics of the Interaction complicate separation of discourse knowledge Into modular procedures. A final section discusses potential applicability of GUIDON to other representations and problems

2 Rule-based Tutoring: GUIDON

GUIDON [9] *is* a tutorial system that can be built on top of any MYCIN-like expert system [20], The knowledge base of a MYCIN-like system is in the form of situation-action rules, such as, "If the Infection that requires therapy Is meningitis and the patient has symptoms of mumps, then this Is weakly suggestive evidence that the type of the infection is viral' In this paper, all examples will be drawn from the MYCIN infectious disease consultation program, though the

* This research was sponsored in part by grants from ARPA (Contract Title MDA 9O3-77-C-0322) and NSF (MCS 77-02712)

GUIDON system Is capable of tutoring any knowledge base of this form.

The purpose of a MYCIN consultation is to determine the organisms that might be causing an infectious disease and to prescribe antibiotic therapy for them. The MYCIN interpreter invokes Its rules in a goal-directed way, chaining backwards to evaluate the predicate of a rule The program requests information (case data") from the user as it becomes appropriate in order to apply a rule The result of a consultation Is a record of rules that were applied, acquired case data, and conclusions inferred by the application of rules This record is configured into an explicit AND/OR tree for use In tutorials AND nodes are rules and OR nodes are the goals hanging below them.

In a GUIDON tutorial a student plays the role of a consultant In a diagnostic game The dialogue deals exclusively with a particular consultation that has previously been stored in the MYCIN system. The student is presented this same case That Is, he is given some information about a patient suspected to have an Infectious disease, and Is expected to request case data, as he sees necessary, to draw conclusions about the cause of the infection. The topics of this dialogue are precisely the goals that are determined by applying MYCIN rules, e.g., the type of the infection As the dialogue progresses. GUIDON is able to collect Information about the student's knowledge of the case and Its domain. The purpose of a GUIDON tutorial is to make the student aware of any gaps or inconsistencies in his knowledge with respect to the underlying rule base, and to correct these deficiencies Note that the course of the dialogue need not pursue the exact sequence of rule invocation and question-asking that MYCIN used in this case, and in general it does not.

The GUIDON system has already been described in some detail In *[9]* in which a framework was established for the development of the program. Briefly, the framework Incorporates: I) discourse knowledge in the form of sequences of domain-independent tutoring rules (hereafter, "t-rules") organiied Into stylized discourse procedures (patterns for tutorial Initiative); 2) domain knowledge In the form of the MYCIN-like rule base and records from the consultation to be discussed with the student; and 3) a communication record for recording dialogue goals and the dynamic state of the tutorial In this paper we will show how this framework has been Instantiated. The dialogue transition diagram presented in Section 3 represents the calling structure of the discourse procedures (comprising 200 t-rules) currently implemented in the program. Section 4 gives several examples of the topic relationship problems In dialogue management.

GUIDON is an example of an ICAI tutorial system. The central feature of such a system is the presence of an underlying "expert program" that can solve the problems that are posed to the student. A "differential student model" (6) is formed by comparing the student's problem-solving actions to those of the expert program. From this comparison of behavior the tutorial program infers whether the student knows about and uses the problem-solving knowledge incorporated in the expert program's actions. Thus, the model of the student's knowledge is an "overlay" [8] on the expert knowledge base. In a rule-based system like GUIDON, this consists of a marked subset of expert rules that the student is thought to know.

Tutoring consists of posing problems and/or making remarks that are Intended to *bring* the student's knowledge into alignment with the expert system; thus apparently missing rules are presented and misconceptions are corrected One of the major research problems Is the design of tutorial remarks: when to say something and what to say [4] GUIDON'S discourse procedures are designed to deal with this problem

Dialogue management Involves coordinating tutorial goals with the constraints Imposed by; I) time available for the session, 2) student initiative and conversational (social) postulates [14], 3) the communication channel, and 4) human memory and learning capability In early development of GUIDON, the limitation of time and verbosity have been the chief forcing functions: It is not possible to explicitly discuss every Inference that the expert program attempts By using heuristics developed for economical presentation and the student initiative options presented below, the crux of dialogue management in GUIDON, session time for relatively simple cases has been reduced from about five hours to less than one (for students who know how to use the program).

The main Issues of dialogue management discussed below are.

— On what basis does a tutor select alternate presentation techniques?

— How does a tutorial program maintain and share dialogue context?

— How can we provide for and cope with student Initiative?

— How can we ensure dialogue connectedness and comprehenslblllty?

3 The Dialogue Transition Diagram

In this section we present the concept of alternative dialogues, different ways something can be said and possible subdtalogues that can occur. The concept is illustrated by a state transition diagram that represents the invocation structure of the discourse procedures In GUIDON. The links represent control expressed by tutorial rules within the procedures. These links signify choice points that lead to alternative dialogues, dictated by domain logic, economy, or tutorial considerations Thus, these represent the management decisions In which the tutor takes the initiative to control dialogue situations

Fig. I illustrates a portion of a dialogue with GUIDON. After a case was chosen (the current system requires that the user select a case from the on-line library), the student was given Initial Information, such as the fact that there Is a pending culture in the laboratory and no organisms have been reported The discussion has reached the point of considering clinical Information (non-laboratory data) to determine the organisms that might be causing the Infection

The student asked whether the patient has a rash, and the tutor gave the reported case data (lines 9-11) (Student input is In the form of keywords or simple sentences) At this point several remarks were made by the tutor: I) lines 13-24, the datum requested by the student was related to the current topic by reciting the path that connects them in this case rash -> Herpes Zoster virus -> infection type -> organisms; 2) lines 26-28, the tutor focussed on the topic "Herpes Zoster virus." and stated a final conclusion; 3) lines 30-31, the tutor explicitly returned the student's attention back to the original topic of determining the organisms causing the infection.

You should now proceed to ask questions that will enable you to make an hypothesis about the organisms tnat might be causing the infection.

** RASHES

Pt538 does not have a rash.

Your question is indirectly relevant:

Whether Pt538 has a rash can be used to determine whether Pt538 has symptoms of Herpes Zoster virus... which can be used to determine the type of the Infection...

and this will enable us to determine the organisms that might be causing the infection.

The fact that Pt538 does not have a rash is evidence that Pt538 does not have symptoms of Herpes Zoster virus (RULE3691

Back to our discussion of the organisms that might be causing the infection...

Fig I Relating a factor to the current topic ("Pt 538" replaces the name of the patient.)

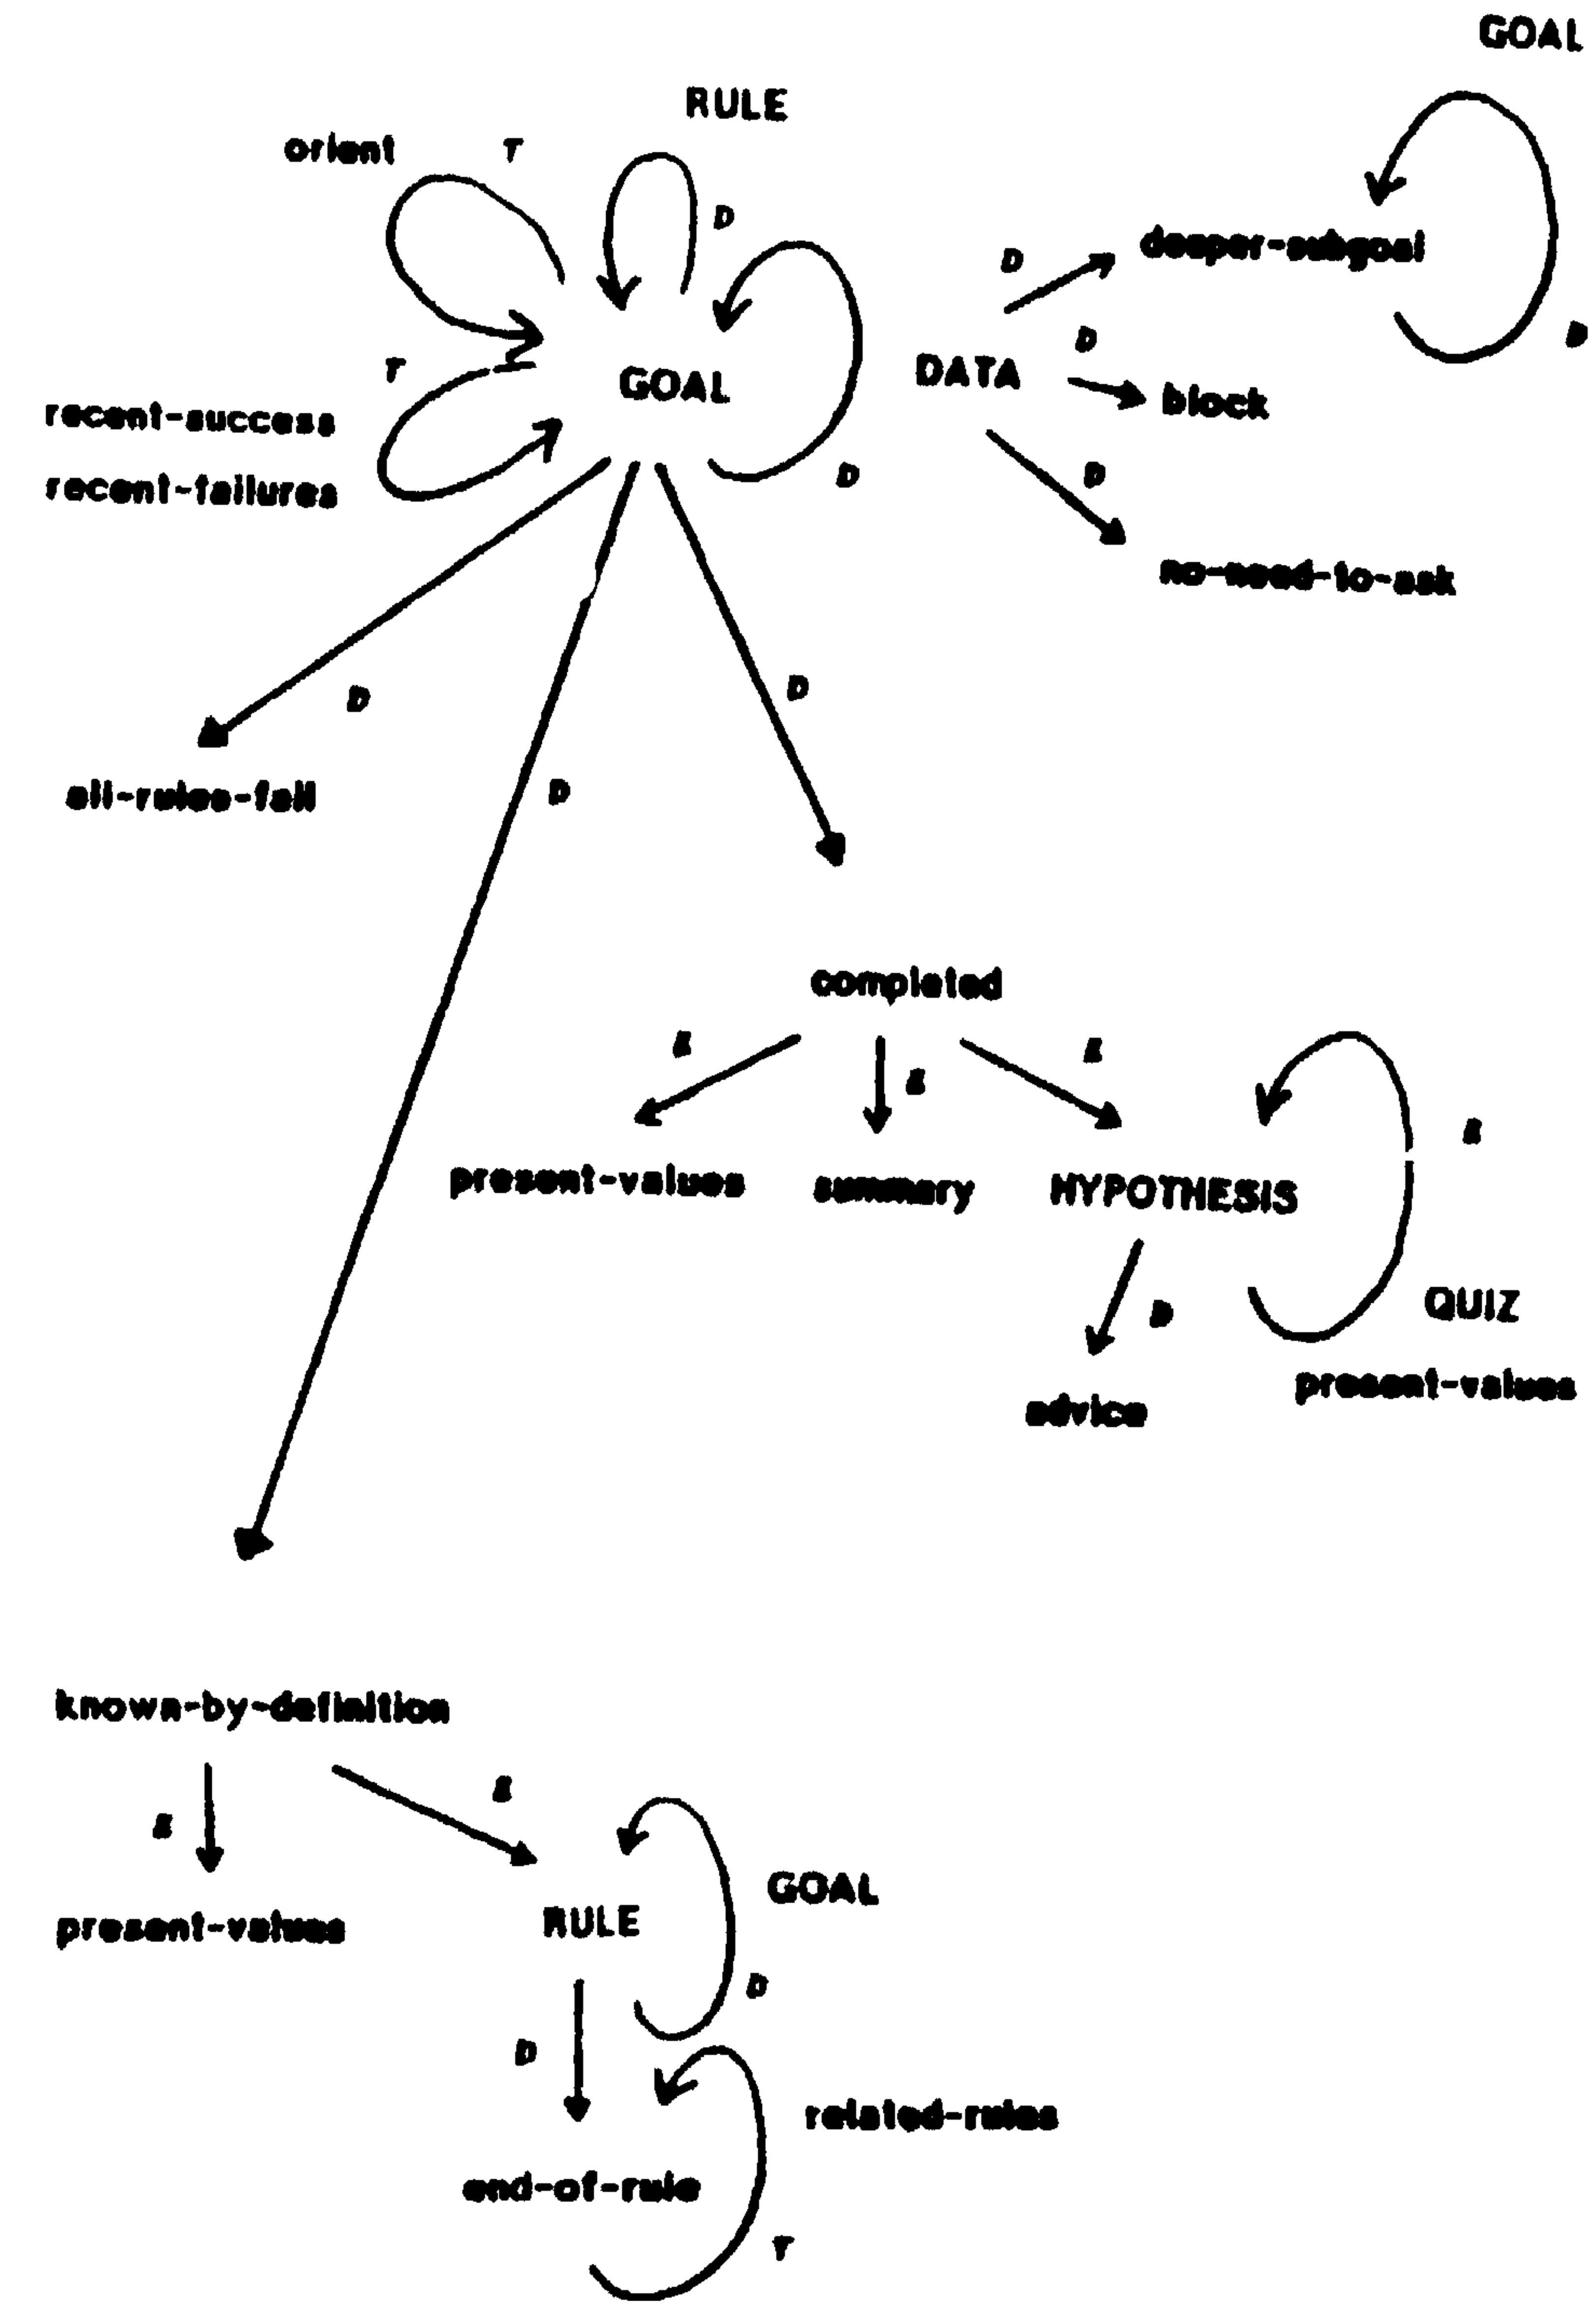

Fig 2 Dialogue Transition Diagram

It now may be helpful to look at the dialogue transition diagram (Fig 2). Each node In this diagram stands for a discourse procedure, or sequence of t-rules. For example, when lines 5-9 were printed In Fig. I, the procedure GOAL was being followed, the basic procedure for discussing any MYCIN goal which appears In a domain rule. To answer the question about rashes, the DATA procedure was applied. Whether or not the patient has a rash is a "deeper subgoal" In that it does not appear in any rule that can determine the current goal, but appears deeper in the AND/OR tree of goals and rules. The DATA procedure printed line 11 and set up the second-deepest subgoal as the new current topic (whether the patient has symptoms of Herpes Zoster virus). An arrow that loops back to a node signifies that the procedure that labels it is called one or more times, as one might expect. For example, we see that one or more "related rules are presented at the end of discussion of a given rule. The Italicized labels stand for the basis of the transitions—economy, domain logic, and tutoring goals—these distinctions are described below.

Returning to our example, the tutor observed that the new topic (Herpes Zoster—see lines 15-17 in Fig. I) was "completed" because all of the case data that the expert needed to make a final conclusion had already been given to the student Here a choice had to be made. Should the tutor present the final conclusion (as It did in lines 26-28)? Or should a summary of evidence be offered? Or should the tutor ask the student to make an hypothesis (as to whether or not the patient has symptoms of Herpes Zoster)? Fig 3 and Fig. 4 show the procedure COMPLETEDCOAL. We see that T-RULE5.02 was applicable: a single domain rule was used by the expert program to make a conclusion and the student model indicated that the student knew this rule, so it was simply stated The fact that the dialogue could have taken a different form at this point, at the tutor's Initiative, illustrates the possibility of alternative dialogues. There are many more alternative dialogues than those Illustrated by the different paths in Fig. 2. For example, the procedure for discussing a rule (shown here as RULE) incorporates 18 different methods, e.g., clause-by-clause discussion, supplying case data and then asking for an hypothesis, and discussing a failed subgoal and then mentioning the conclusion that can't be made.

COMPLETEDGOAL.PROC005

Purpose: Discuss final conclusion for a goal

Step I: <Decide whether to finish with a summary>

Step2: <Discuss the final hypothesis for the goal>

Step3: <Wrap up discussion or record completion>

Fig. 3 A Discourse Procedure

Apply the first rule that is appropriate:

T-RULE502 <Directly state single, Known rule>

If; I) 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* I, and
3) There Is strong evidence that the student has applied this rule
Then: Simply state the rule and its conclusion

T-RULE5.03 <Request hypothesis when rules may be unknown>

If: You have examined the rules having a bearing on this goal that have succeeded and have *not* been discussed, and have found a rule under consideration for which there is not strong evidence that the student has applied the rule under consideration
Then: Substep i.
If: I) An Introductory remark is to be made before requesting an hypothesis from the student, and
2) The student has not requested help for forming an hypothesis
Then: Say: hypothesis-ready
Substep Ii. Discuss the student's hypothesis for the goal currently being discussed [ProcO14]

Fig. 4 T-rules for Deciding How to Discuss • Final Hypothesis (Fig. S, Step 2)

The COMPLETEDCOAL procedure illustrates two Important constraints for choosing among alternative dialogues on the basis of the principle of economy. Specifically, "surface complexity" features of the knowledge being discussed (e.g., the number of rules involved, whether a rule is definitional, that all relevant rules fail, and how many subgoals remain to be completed before an Inference can be drawn) and the student model combine In these tutorial rules to select a presentation method that affects how extensive the discussion will be.

When a conservative programmed tutor Is not sensitive to the complexity of a situation or the student's knowledge, tedious, overly detailed discussion results; the program belabors topics that can be dealt with swiftly by a single remark or simply skipped over. For example, a striking Improvement to GUIDON tutorials was based on the recognition that many MYCIN topics are "definitional," so the usual process of discussing all of the evidence explicitly and asking the student to make an hypothesis can be bypassed when the model Indicates that the student knows the relevant rule.

Besides economical considerations, the dialogue transition diagram illustrates two other kinds of transitions: logical and tutorial The links leading from the DATA procedure are distinctions based on domain logic, as are the three straight links leading from GOAL. These links are based on domain facts and relations (e.g., that a topic Is the name of a block of case data) and decisions made by the expert program (e.g., why a question need not be asked). In following these transitions, the tutor is reasoning about the subject material.

Tutorial transitions are based on the tutor's goals for teaching particular material to the student and/or refining its model of his knowledge. The examples in Fig. 2 are: 1) quizzing the student about rules that "recently" succeeded or failed because of case data Just given to the student ("recent-success" and "-failure" in Fig. 2); 2) providing initial orientation to get the student started on a topic that is new to him ("orient"); and 3) quizzing about rules related to the one Just discussed ("related-rules). We call these instances of *opportunistic tutoring* because they signify initiative taken by the tutor at "appropriate" times with the purpose of fitting as much information into the session as is practical These quizzes constitute subdlalogues that can be quite similar in motivation and format to SCHOLAR [10], WHY (22), and A BLOCKS [3] tutorials.

The appropriateness of the interruption is determined by Its consideration within the sequence of a discourse procedure, a conventional dialogue pattern, and by heuristics that determine whether a remark should be made The basic idea is that the program has certain topics that it wants to bring up (eg., a couple of rules relevant to the current case that the student may not know), so it checks to see if the kind of remark that is appropriate at a given time provides the opportunity to raise a particular topic Thus, after the tutor returned attention to the previous topic In Fig. I (lines 26-28), it checked to see if there were any rules that it wanted to mention that had Just then been applied by the expert program (there were none).

4 Managing Topic Relationships

A number of complications have been glossed over in the previous discussion, primarily problems of relating topics to one another during the dialogue. The subsections below deal with I) maintaining and sharing dialogue status information; 2) providing for and coping with student initiative; and 3) making situation transitions coherent and arguments clear.

4.1 Dialogue Status Information

GUIDON's dialogue status Information is in two parts: I) records of Inferences associated with each topic and rule, and whether they have been discussed in the current session; and 2) a simple record of the context of the dialogue.

Recall that every topic In the dialogue is a goal for which a conclusion is drawn from MYCIN rules and/or case data has been reported. GUIDON has to keep track of how MYCIN gained information and correct the student if he asks questions that are not related to the current topic or whose answers can be Inferred from what has been given.

Associated with each goal pursued by MYCIN, and hence each potential topic in the tutorial dialogue, is a dynamic inference and discussion record of data given to MYCIN and the student, conclusions made by MYCIN and conclusions believed to have been made by the student, and whether a rule or topic has been discussed in the current session. This information is combined to fill in the AND/OR tree to maintain models of the expert's and student's current knowledge; for example a rule is marked as "fired" when

MYCIN has marked all of its subgoals as known. Some of the information is used in fairly complex tutorial strategies, such as providing assistance according to an hypothesis revision strategy.

Various options make it possible for the student to access this record to keep track of the evidence that has already been discussed and what remains to be considered. For example, the PENDING option causes GUIDON to state subgoals that remain to be discussed, case data that can still be requested, and subgoals for which the student has enough Information to form a final hypothesis.

Another student-controlled means of sharing status Information directs the program to indicate when a conclusion can be made, e.g., "MYCIN just made a conclusion about the type of the infection." In terms of teaching strategy, this kind of remark is Intended to encourage the student to examine his own understanding at this point, perhaps leading him to request more information However, note that some students may not care what MYCIN *is* doing, preferring to proceed on the basis of their own knowledge.

The focus record consists of variables that are set when the student uses various options and when the tutor makes hints about topics for the student to pursue. By using this record when providing subsequent assistance to the student, the tutor's choice of topics is connected to what has gone on before. This method for keeping the dialogue "focussed" is admittedly simple, but because of the other constraints on the dialogue (i.e., stacked topics and the option-oriented format), it seems to be sufficient.

4.2 Student Initiative

An essential part of tutorial dialogue management is making provision for the student to express himself. When one first sets out to construct a mixed-initiative program, it may not be clear that provision must be made for every potential kind of initiative that the student will be allowed to make. This includes being able to refer back to an early topic and provide more details (given that explicit presentation of evidence may be non-exhaustive), allowing the other participant to change the topic, 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.

GUIDON provides for student initiative by making available a variety of student options and by indexing tutorial remarks so that the student can easily refer to them later (use them as arguments to options). Options are grouped by the kind of initiative, some examples are shown in Fig. 5.

Option type	Examples
Get Case Data	BLOCK ALLDATA
Information Retrieval	PENDING DETAILS
Dialogue Context	RULE TOPIC
Say What You Know	IKNOW HYPOTHESIS
Request Assistance	HELP HINT TELLME
Change the Topic	DISCUSS STOP
Special	? PROFILE JUSTIFY

Fig. 5 Options Available During GUIDON Dialogues

Fig 6 illustrates the typical use of options Each option generally has a discourse procedure associated with it For example, the procedure for IKNOW invokes HYPOTHESIS if the expert program hasn't made a final decision yet; otherwise COMPLETEDGOAL is invoked. T-rule5.02 was applied in Fig. 6 when the student used the IKNOW option.

```
** 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

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

** USE 3C

The duration of the neurological signs is 7.2 hours.

Using other Information available to you now, the
duration of the neurological sigm enables you to
conclude that the type of the infection is bacterial
(.2) viral (.2) fungal (-.4) tb (-.4) [RULE524J.

** 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 [RULE1121

Back to our discussion of the type of the Infection...
```

Fig 6 Sample Use of Options

GUIDON allows a student to explicitly change the topic by the DISCUSS option. However, student requests for data can also (implicitly) change the topic if the datum requested is not relevant to the current topic (cannot be used directly in any Inference). Fig. I illustrated a dialogue that occurred when the requested Information was relevant to a deeper sub goal If requested Information is relevant to a previous, shallower subgoal, the tutor states this relation so that it is clear to the student what topic Is currently being pursued (Fig. 7).

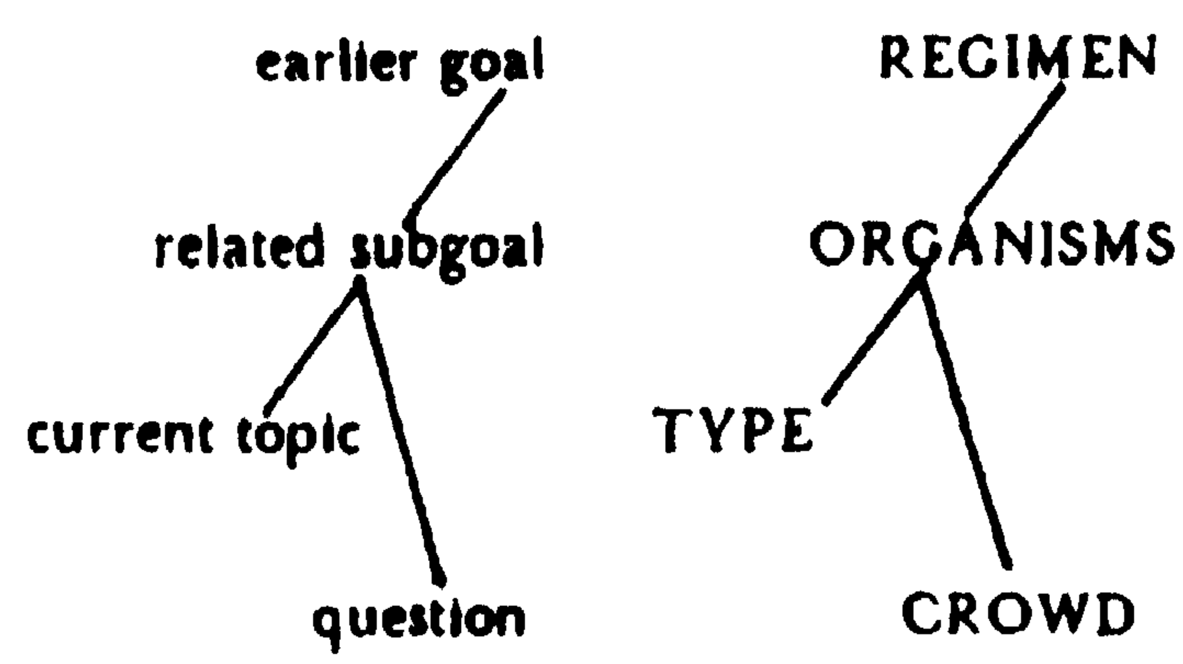

** DOES THE PATIENT LIVE IN A CROWDED ENVIRONMENT?

Pt538 does not live In a crowded environment.

Whether the patient does live in a crowded environment is not relevant to determining the type of the infection. It Is a consideration we can use later when we return to our discussion of the organisms that might be causing the infection.

Fig 7 Crossing Topic Paths in a Tree of Subgoals

4.3 Situation Transitions

Clarity of presentation is critical In productive teaching, as in any dialogue. In GUIDON tutorials this involves providing coherent transitions between dialogue situations and comprehensible summaries of inferences

Providing clear transitions requires that discourse procedures take Into account the context in which they are invoked, so that the tutor's remarks are not disconnected or redundant The chief design principle we have followed Is that a tutorial rule that Invokes a new discourse situation provides transition to the new situation and return to the former context. For example, T-RULE 5 03 prepares the student for the request for an hypothesis Note that the predicate "an Introductory remark Is to be made..." is a flag set by the calling discourse procedure, e.g., it is false if the student has explicitly requested the opportunity to make an hypothesis.

Fig 8 illustrates the kind of disconnected dialogue that arises from Inadequate communication between discourse procedures Following the procedure that Is Invoked after an hypothesis ("advice" in Fig 2), the tutor asks the student if he wants to continue discussion of this topic or if he would like to wrap It up at this point. The student wishes to continue, so the dialogue shifts to discussion of this topic (the procedure GOAL), and the first thing the tutor remarks about Is that the topic is essentially done (attempting to make an economical transition based on a comparison of the expert program's current hypothesis to the final conclusion it will reach after all evidence has been considered).

A smoother dialogue would result if the procedure for giving advice after an hypothesis looked to see if the topic was substantially completed and so the student would only be asked once If he wanted to go on. The interaction shown here is typical, and many of the changes made to the system after each trial tutorial session are designed to improve the dialogue In similar ways. While we are constantly allowing for more subtle variations in the dialogue, the framework of rules organized Into procedures is working well

** HYPOTHESIS 3A

Is the meningitis partially treated?
** YES

You're right

The time since therapy with thf cephalothin was started Is evidence that the meningitis Is partially treated (.95) [RULEI45]

It remains to collect further evidence for determining whether the meningitis Is partially treated.
Would you like to see a summary of all of the evidence that can eventually be collected (as opposed to working it out yourself)?
** NO

Would you like to see a summary of the evidence we have discussed before going on to ask more questions?
** NO

We have substantially completed our discussion of whether the meningitis Is partially treated
Would you like to make an hypothesis now (as opposed to gathering more Information)?
** NO

Fig. 8 A Disconnected Dialogue

Finally, when tutorial remarks are limited to lines of print on a teletype page (a constraint this research has adopted), there Is no opportunity to pause or to emphasize words (by changing voice volume or inflection) Therefore, it is worthwhile to incorporate visual keys that help to organize presentations, and provide clear transition between segments of the dialogue. For example, when summarizing evidence, GUIDON regroups goal/subgoal relations, separating "key factors" from contextual information The usual, MYCIN-style listing of these same rules could require several pages.

5 Limitations and Scope

As stated previously, GUIDON is not intended to be a complex natural language converser, such as GUS [I] In order to concentrate on formalizing specific discourse procedures, we have avoided the problem of recognizing new dialogue situations In free text by providing the student with an option-oriented vocabulary. Moreover, generation of remarks is on the level of choosing the course of the dialogue, rather than building up output from grammatical components. We devote just a small amount of computation towards determining the student's Intentions, chiefly by relating student requests for data to the current topic. However, we do Intend to build a more complex model of the student's strategies, and this might Involve parsing complex Input (as in Sleeman's program for understanding student explanations [21]) Furthermore, the AND/OR tree of topics and rules strongly constrains the tutor's choice of dialogue situations In less rigidly pedantic dialogues, generation of conversation Is sensitive to considerations such as the social context of the participants and emotional connotatioa

(Some of Burton and Brown's *kibitzing* strategies are like this [6])

Second, besides the limitations of GUIDON's mixed-Initiative conversation capability, It can only discuss rule-based knowledge It Is clear that not all domain knowledge can be expressed as topics that are determined by applying situation-action rules; but it Is plausible that this will be an adequate framework for discussing a wide variety of diagnostic problems For example, suppose that the representation of knowledge consisted of relations that were inference triggers (as In INT FRNIST [19]), we still might find It convenient to think of the dialogue as being structured into goals (something to find out), case data, possible outcomes (alternative hypotheses), and inference relations It seems probable that many of the Issues of dialogue management we have discussed will be important, though we may need to augment the situations we have formalized (Fig. 2) For example, for an hypothesis-oriented discussion we would refine the GOAL procedure to focus on one outcome value at a time, and Incorporate transition t-rules for detecting and discussing shifts to another outcome (hypothesis revision)

Third, some aspects of our work have been motivated and aided by the particular rule bases available to us The design of GUIDON's dialogue implicitly makes use of the fact that MYCIN'S AND/OR tree is never more than five topic levels deep and the current topic Is usually only two levels below the top goal of determining therapy for the patient. A significantly deeper tree might prohibit a stacked topic" design. Moreover, the success and usefulness of GUIDON's tutorial ability depends to some extent on the universality of the rule set, i.e., how the subject matter has been structured by the rule authors. The domain rules were designed to be comprehensible to potential users of the consultation system who are not experts, but it is Inevitable that some of the concepts represented in the rules are artifacts of the particular representation and the requirements of making everything explicit.

6 Closing Remarks

While most ICAI research has focussed on representation of domain expertise and construction of a student model, we have shpwn that there Is a group of issues that center about the problem of carrying on a coherent, task-oriented mixed-initiative dialogue with a student We have named this collection of issues the dialogue management problem and discussed *it in terms* of tutorials based on MYCIN-like rules. We presented a representation of dialogue knowledge in the form of a transition diagram in which the nodes are discourse situations and the links represent selection of alternative dialogues based on domain logic, economy of presentation, and tutorial objectives. Other issues were described in terms of managing topic relationships and sharing initiative.

We stated that the main forcing function behind improvement of the program was the impossibility of explicitly discussing each inference made by the expert program. Verbosity continues to be a problem for GUIDON. It is interesting to note that the most effective "pruning" of topics in the current version of the program is a result of the student's initiative. Perhaps it is not necessary or desirable for the program to attempt to manage the dialogue too severely; deciding what should be discussed is naturally a shared task. We hope that a proposed "case lesson plan" [9] will give the tutor additional leverage for non-exhaustive discussion by providing reasonable, time-sensitive goals for the session.

References

[I] Bobrow, Daniel G., Kaplan, Ronald M., Kay, Martin, et al (1977). GUS, A Frame-Driven Dialog System. Artificial Intelligence, Vol 8 No 2, 155-173

[2] Brown, J.S., Rubenstein, R., Burton R., (1976). Reactive Learning Environment for Computer-Aided Electronics Instruction. BBN 3314

[3] Brown, J.S., Burton R, Miller M, deKleer, J., Purcell, S, Hausmann, C, Bobrow, R. (1975). Steps toward a Theoretic Foundation for Complex, Knowledge-Based CAI. BBN 3135

[4] Brown, J.S (1977). Uses of AI and Advanced Computer Technology In Education In Computers and Communications: Implications for Education. Academic Press, Inc New York.

[5] Burton, R. (1976) Semantic Grammar: An Engineering Technique for Constructing Natural Language Understanding Systems. BBN 3453.

[6] Burton, R. (1979). An investigation of Computer Coaching for Informal Learning Activities The Int J or Man-Machine Studies II.

[7] Carbonell, Jaime R (1978). Intentionality and Human Conversations Theoretical Issues in Natural Language Processing-2, July 1978, 141-148.

[8] Carr, Brian and Goldstein, Ira (1977). Overlays: A Theory of Modelling for CA1. MIT AI Lab Memo 406.

[9]Clancey, William J. (1979). Tutoring Rules for Guiding a Case Method Dialogue The Int J of Man-Machine Studies 11,25-49.

[10] Collins, Alan. (1976). Processes in Acquiring Knowledge. In Schooling *and* Acquisition of Knowledge, (Anderson, Spiro and Montague, eds), Erlbaum Assoc., Hillsdale, NT.

(II)Deutsch, Barbara G. (1974). The Structure of Task-Oriented Dialogs. IEEE Symposium for Speech Recognition, 250-253.

[12] Faught, William S. (1977). Motivation and Intensionality in a Computer Simulation Model Ph. D. Dissertation, Stanford Unlversity,AIM-305.

[13] Goldstein, Ira. (1977). The Computer as Coach: An Athletic Paradigm for Intellectual Education. MIT AI Lab Memo 389.

[14] Gordon D and Lakoff, G (1971). Conversational Postulates. Papers from Seventh Regional Meeting, Chicago Linguistic Society, Chicago. *University* of Chicago Linguistics Department.

[15] Grosi, Barbara J (1977). The Representation and Use of Focus In a System for Using Dialogues. Proc. of 5th IJCAI, 67-76.

[16] Hart, Peter E. (1975) Progress on a Computer Based Consultant. Proc of 4th IJCAI, 831-841.

[17] Hart, R O. and Koffman, E B. (1975). A Student-Oriented Natural Language Environment for Learning LISP. Proc. of 4th IJCAI, 391-396.

[18] Koffman, Elliot B, and Blount Sumner E. (1973). Artificial Intelligence and Automatic Programming in CAI. Proc. of 3rd IJCAI, 86-94.

[19] Pople, Harry E. (1975). DIALOG A Model of Diagnostic Logic for Internal Medicine. Proc. of 4th IJCAI, 848-855

[20] Shortliffe, E H. (1974) Computer-Based Medical Consultations: MYCIN, American Elsevier, New York, 1976.

[21] Sleeman, Derek. (1977). A System which Allows Students to Explore Algorithms, Proc. of .5th, IJCAI, 780-786.

[22] Stevens, Albert L., and Collins, Allan. (1977). The Goal Structure of a Socratic Tutor. BBN 3518.