

Chapter 5

Explanation System

5.1 Introduction

As was emphasized in Chapter 2, one of the primary requirements for user acceptance of a consultation program is an ability to explain decisions. Rule-based knowledge has greatly simplified the implementation of such a capability in the MYCIN system. The portion of the system used for explanation is termed Subprogram 2 (shown in Figure 1-1). It is automatically invoked at the end of each consultation session, and may also be accessed optionally during the consultation itself (see the QA option, § 3.3.2-2).

Since MYCIN explains decisions only in response to queries from the user, the explanation system is also a question-answering (QA) system. Subprogram 2 is therefore often called the QA-module, a term that reflects MYCIN's debt to other AI programs for answering questions [Simmons, 1970; Fox, 1970].

The ability to answer questions obviously requires that the queries be understood. Since we have been anxious to minimize special training needed for use of the MYCIN system, we have been eager to let the physician ask questions using simple English. As discussed in § 1.3.1-7, however, writing programs to understand natural language is complex because of the myriad ways that individuals may choose to express themselves. Although several powerful techniques have been developed [Winograd, 1972; Woods, 1970; Schank, 1972], they all suffer from being either somewhat slow computationally or difficult to generalize in domains other than those for which they were designed. Since physicians will quickly reject a system that takes 2 or 3 minutes to answer a question, we sought an approach that would

MYCIN

emphasize speed of response rather than human-style discourse. Yet we did want to make the system powerful enough to answer most questions that a physician might want to ask. Since the goals of rapid response and powerful capabilities tend to work at cross purposes, we have been forced to try to strike a balance between the two. The approach described in this chapter is thus neither as fast as desirable (it requires 5 to 20 seconds to answer a question) nor as powerful (it has no sense of discourse, anaphora, or complex syntax). However, its performance is usually adequate, and an experienced user who becomes aware of its limitations is able to retrieve most of the information he desires. Furthermore, it should be emphasized that the consultation itself, which is after all the primary focus of the MYCIN system, requires no natural language processing. Use of the QA module is optional, and a physician who is in a hurry therefore need not take the time to seek explanations if he is satisfied with the advice the program has given.

As demonstrated in the sample consultation at the end of Chapter 1, the explanation system offers several options to the user:

QUESTION-ANSWERING (QA) OPTIONS

- | | |
|-----------|---|
| HELP | - prints this list |
| EQ | - requests an explanation of the specified question(s) from the consultation |
| IQ | - prefix to a question which asks about information acquired by the program during the consultation |
| NO PREFIX | - this question queries contents of decision rules in the system |
| PR | - requests that specified rule be printed |
| STOP | - escape from explanation system |
| RA | - entry to rule-acquisition module for recognized experts |

In this chapter I describe each of these options. Only the IQ and NO PREFIX options require natural language processing.

Section 5.2 describes how each option is used, giving examples of each. However, the implementation details are rapidly changing and have been described fully elsewhere [Shortliffe, 1974b]. They will therefore not be included in this report. The chapter concludes with a brief discussion, in § 5.3, of the Explanation System's limitations

Explanation System

and of how we intend to improve the program's capabilities in the future.

5.2 Using Question-Answering System

Unlike the Consultation System (Subprogram 1) in which MYCIN takes the initiative, asking questions and waiting for the physician to respond, the Explanation System expects the user to guide the interaction. This approach allows the system to instruct the physician or explain its advice only with regard to specific topics that may be puzzling to the user. Thus MYCIN prints its prompt characters (the double asterisk—"***"), waits for a question, performs the requested procedure, redisplay the prompt characters, and then waits for the next user input. This process continues until the user enters the word STOP.

In this section, I describe the capabilities of the Explanation System, i.e., the various QA-options listed in § 5.1. MYCIN checks every input sentence to see if it begins with one of the special prefixes (HELP, EQ, IQ, PR, STOP, or RA). If not, it assumes that the user has asked a rule-retrieval question. Examples of rule-retrieval questions are discussed in § 5.2.1. The EQ and IQ options are explained in § 5.2.2.

Option prefixes serve one of two purposes. Most allow MYCIN to perform certain repetitive tasks without invoking time-consuming natural language routines (e.g., HELP, EQ, PR, STOP, and RA). The IQ option, on the other hand, is required at present so that MYCIN can distinguish between the two principal kinds of questions that *do* involve English language understanding. Without the IQ prefix MYCIN would need to deduce the distinction on semantic or syntactic grounds, a complex problem that we have temporarily avoided by using the prefix mechanism. The distinction between IQ and rule-retrieval questions will be clarified as we proceed.

5.2.1 RULE-RETRIEVAL QUESTIONS

Since most of MYCIN's knowledge is contained in its corpus of 200 rules, many questions can be effectively answered by retrieving and displaying the relevant rule(s). Thus the challenge for MYCIN's

MYCIN

QA routines is to "understand" a question well enough to decide which rules should be retrieved.

5.2.1.-1 General Questions

MYCIN can answer rule-retrieval questions that are either specific to a given consultation (§ 5.2.1-2) or general in nature. General questions reference the corpus of rules without considering the status of the dynamic data base (Figure 1-1), i.e., they ask questions about MYCIN's knowledge rather than about how that knowledge has been applied to the patient under consideration. The following examples demonstrate how MYCIN uses both rule-retrieval and its LISP-to-English translation capability in order to answer general questions from the user:

****WHAT DO YOU PRESCRIBE FOR ACTINOMYCETES INFECTIONS?**

RULE002

IF: THE IDENTITY OF THE ORGANISM IS ACTINOMYCETES
THEN: I RECOMMEND THERAPY CHOSEN FROM AMONG THE
FOLLOWING DRUGS:
1 - PENICILLIN (.99)
2 - TETRACYCLINE (.99)
3 - ERYTHROMYCIN (.99)

****UNDER WHAT CONDITIONS IS AN ORGANISM APT TO BE A
CONTAMINANT?

Relevant rules: Several rules including:

RULE042 RULE106 RULE109 RULE123 RULE155

Which do you wish to see?

****155**

RULE155

IF: 1) THE SITE OF THE CULTURE IS ONE OF THOSE SITES
THAT ARE NORMALLY STERILE, AND
2) THE STAIN OF THE ORGANISM IS GRAMNEG
THEN: THERE IS STRONGLY SUGGESTIVE EVIDENCE (.8)
THAT THE ORGANISM IS NOT A CONTAMINANT

****WHY DO YOU ASK THE AGE OF THE PATIENT?**

Explanation System

Relevant rules: Several rules including
RULE118 RULE135 RULE137 RULE138 RULE150

Which do you wish to see?

**150

RULE150

IF: 1) THE THERAPY UNDER CONSIDERATION IS
SULFISOXAZOLE, AND
2) THE AGE (IN YEARS) OF THE PATIENT IS LESS
THAN .58, AND
3) THE PATIENT IS JAUNDICED
THEN: IT IS DEFINITE (1) THAT SULFISOXAZOLE IS NOT A
POTENTIAL THERAPY FOR USE AGAINST THE ORGANISM

**WHAT DOES IT MEAN IF PENICILLINASE IS ADDED TO THE
**CULTURE BOTTLE?

RULE039

IF: 1) THE SITE OF THE CULTURE IS BLOOD, AND
2) THE STAIN OF THE ORGANISM IS GRAMPOS, AND
3) PENICILLINASE WAS ADDED TO THE CULTURE
THEN: THERE IS WEAKLY SUGGESTIVE EVIDENCE (.4) THAT
THE ORGANISM IS A CONTAMINANT

**WHAT IS RULE030?

RULE030

IF: THE IDENTITY OF THE ORGANISM IS KNOWN WITH
CERTAINTY
THEN: IT IS DEFINITE (1) THAT THESE PROPERTIES - GRAM
MORPH AIR - SHOULD BE TRANSFERRED FROM THE
IDENTITY OF THE ORGANISM TO THIS ORGANISM

**IS THERE A RULE THAT DISCUSSES STREPTOCOCCAL
** MORPHOLOGY?

No. I have no such rules at present.

**DO YOU EVER CONCLUDE THAT GRAM POSITIVE COCCI ARE
**CONTAMINANTS?

THE FOLLOWING RULE MAY BE RELEVANT: RULE039

MYCIN

5.2.1-2 Questions Regarding Current Consultation

As discussed in § 3.3.3-3, MYCIN maintains a record of how it makes decisions during every consultation. This record is then used by the QA system in order to answer rule-retrieval questions that are specific to the current consultation. MYCIN answers these questions differently from the way it responds to general rule-retrieval questions. Instead of displaying all the relevant rules, the program mentions what rules were used, when they were invoked during the consultation (i.e., after which question number), and what certainty factor (CF) was associated with the specified conclusion. The following examples from the sample consultation in Chapter 1 demonstrate this capability:

****WHAT MADE YOU THINK ORGANISM-1 MIGHT BE A PROTEUS?**

I USED RULE163 TO CONCLUDE THAT THE IDENTITY OF ORGANISM-1 IS PROTEUS. THIS GAVE A CUMULATIVE C.F. OF .277. THE LAST QUESTION ASKED BEFORE THE CONCLUSION WAS MADE WAS 26.

I USED RULE084 TO CONCLUDE THAT THE IDENTITY OF ORGANISM-1 IS PROTEUS. THIS GAVE A CUMULATIVE C.F. OF .119. THE LAST QUESTION ASKED BEFORE THE CONCLUSION WAS MADE WAS 24.

****HOW DID YOU DECIDE THAT ORGANISM-2 WAS THE SAME AS ORGANISM-1?**

I USED RULE005 TO CONCLUDE THAT THE PSEUDOMONAS (ORGANISM-2) IS A PRIOR ORGANISM WITH THE SAME IDENTITY AS ORGANISM-1. THIS GAVE A CUMULATIVE C.F. OF .19. THE LAST QUESTION ASKED BEFORE THE CONCLUSION WAS MADE WAS 21.

I USED RULE004 TO CONCLUDE THAT THE PSEUDOMONAS (ORGANISM-2) IS A PRIOR ORGANISM WITH THE SAME IDENTITY AS ORGANISM-1. THIS GAVE A CUMULATIVE C.F. OF .133. THE LAST QUESTION ASKED BEFORE THE CONCLUSION WAS MADE WAS 18.

5.2.2 QUESTIONS REGARDING DYNAMIC DATA BASE

The questions described in § 5.2.1-2 involve both rule-retrieval and access to the consultation record which is part of the dynamic data

Explanation System

base (shown in Figure 1-1). There are two additional kinds of question, however, which reference dynamic data but do not access the program's knowledge base of rules. These are described in the following two subsections.

5.2.2-1 IQ Questions

IQ questions require natural language processing in much the same way that rule-retrieval questions do. The IQ prefix is simply a flag for MYCIN so that it knows that the following question does not require rule-retrieval. IQ stands for "Informational Question." Unfortunately, it is often unclear to the novice user when this prefix is needed, so we are currently looking for ways to dispense with it (§ 5.3).

An IQ question merely requests the current value of some clinical parameter from the dynamic data base (see § 3.3.3). It does not ask how the value was acquired (i.e., whether supplied by the physician or inferred using rules). Thus "How did you decide ORGANISM-1 is a pseudomonas?" is a rule-retrieval question, but "Is ORGANISM-1 a pseudomonas?" requires the IQ prefix. The following examples from the sample consultation in Chapter 1 demonstrate the IQ capability:

```
**IQ WHAT IS THE FINAL DECISION REGARDING THE IDENTITY OF  
**ORGANISM-1?
```

```
THE IDENTITY OF ORGANISM-1 IS PSEUDOMONAS (.357)  
KLEBSIELLA (.177) ENTEROBACTER (.166) E.COLI (.137)  
PROTEUS (.118) SERRATIA (.042)
```

```
**IQ DID YOU DECIDE ORGANISM-2 WAS THE SAME AS  
**ORGANISM-1?
```

```
ORGANISM-4 (.19) ORGANISM-3 (.19) ORGANISM-2 (.19) IS A  
PRIOR ORGANISM WITH THE SAME IDENTITY AS  
ORGANISM-1
```

5.2.2-2 EQ Command

During a consultation the user may request an explanation of any question that he is asked (see the RULE and WHY options, § 3.3.2-2). We also wanted to permit the physician to request such explanations after the consultation is complete. Therefore MYCIN maintains a record of every question asked. (This same record is used

MYCIN

for changing the answers to questions, as described in § 3.6.1.) The EQ option allows the physician to ask MYCIN for explanations of consultation questions (EQ = Explain Question). MYCIN answers by specifying the clinical parameter and context that were being considered when the question was asked. It also gives the rule that caused the question to be generated. Thus:

**EQ 48

48 QUESTION 48 WAS ASKED IN ORDER TO FIND OUT THE
PATIENT'S DEGREE OF SICKNESS (ON A SCALE OF 4) IN AN
EFFORT TO EXECUTE RULE068.

The EQ command accepts a list of question numbers as arguments and explains each in the manner demonstrated. The user may then display any rules with which he is not familiar by using the PR command (§ 5.2.3). Note that the EQ command requires no language processing. If anything following the command is not a legal question number, it is simply ignored.

5.2.3 ADDITIONAL OPTIONS

In addition to the options already described, the user of the explanation system may give the HELP, STOP, RA, and PR commands. The first three take no arguments. HELP simply displays the list of user options and thus parallels the HELP command available during the consultation itself (§ 3.3.2-2). STOP provides a mechanism for escaping from Subprogram 2 once the user is through asking questions. RA is available only to experts who are known to the system. It permits the user to enter the Rule-Acquisition System (Subprogram 3 shown in Figure 1-1) which is described in Chapter 6.

The PR command provides a quick way to ask the rule-retrieval question "What is RULE030?" (see § 5.2.1-1). It accepts one or more numbers as arguments and assumes that they correspond to the numbers of rules that the user wishes to see. Thus "PR 30" causes RULE030 to be printed. Several examples of the PR option are included in the sample consultation at the end of Chapter 1. Readers who are interested in the details of how the options described in this section have been implemented should consult Chapter 6 in [Shortliffe, 1974b].

Explanation System

5.3 Future Extensions

Improvements to MYCIN's language and explanation capabilities must necessarily bear in mind the important balance between comprehension and speed of execution. By customizing MYCIN's capabilities to the unique characteristics of its rule-based knowledge, we have managed to devise a surprisingly powerful although simplistic approach to question-answering. MYCIN does not "understand" questions in the sophisticated ways that characterize the most powerful and general of today's natural language systems. Yet it still manages to answer many questions adequately without a large expenditure of computer time during the analysis of each question. Since the language capabilities of MYCIN have been developed in response to a clear need for an explanation system (Chapter 2), rather than because of an inherent interest in the theory of language or computational linguistics, we are content at present to build upon the simple characteristics and limited power of MYCIN's current approach.

We are less than pleased, however, with those aspects of the current approach that will clearly interfere with the program's acceptability to physicians. Although doctors can learn to phrase their questions simply and to expect rules in response, limits on the kind of questions that can be asked or answered commonly lead to user frustration. We have therefore identified the following short-term goals for improvement of the Explanation System's language capabilities:

- (1) Development of a mechanism for permitting the physician to ignore the distinction between IQ and rule-retrieval questions; the IQ prefix should be unnecessary and MYCIN should itself deduce when a question is merely asking for the value of a parameter rather than for rule-retrieval.
- (2) Development of a mechanism for answering questions regarding those parts of MYCIN's knowledge that are not rule-based (see § 3.2.6-1); the current approach does not permit QA access to simple lists or knowledge tables.
- (3) Development of methods, as discussed in § 3.7, for moving algorithmic knowledge from functions to rules so that questions regarding therapy selection may be answered using standard rule-retrieval techniques.

Finally, work is currently underway to improve MYCIN's explanation capabilities during the consultation itself. The RULE command

MYCIN

we described in § 3.3.2-2 is less than satisfactory as an explanation or educational mechanism because it does not explain why the current rule has been invoked by MYCIN's goal-oriented control structure. A series of commands to allow the user to manipulate the entire reasoning chain is currently under development and should greatly enhance MYCIN's ability adequately to explain its questions and reasoning processes [Shortliffe, 1975b; Davis, 1976].