



01 Jul 1986

MICA: Prototyping an Expert System Consultant

Joseph Howard Marchal

Arlan R. DeKock

Missouri University of Science and Technology

Follow this and additional works at: https://scholarsmine.mst.edu/comsci_techreports

 Part of the [Computer Sciences Commons](#)

Recommended Citation

Marchal, Joseph Howard and DeKock, Arlan R., "MICA: Prototyping an Expert System Consultant" (1986).
Computer Science Technical Reports. 9.
https://scholarsmine.mst.edu/comsci_techreports/9

This Technical Report is brought to you for free and open access by Scholars' Mine. It has been accepted for inclusion in Computer Science Technical Reports by an authorized administrator of Scholars' Mine. This work is protected by U. S. Copyright Law. Unauthorized use including reproduction for redistribution requires the permission of the copyright holder. For more information, please contact scholarsmine@mst.edu.

MICA: PROTOTYPING AN EXPERT
SYSTEM CONSULTANT

J. H. Marchal* and A. R. DeKock

CSc-86-5

Department of Computer Science
University of Missouri-Rolla
Rolla, Missouri 65401 (314) 341-4491

*This report is substantially the M.S. thesis of the first author,
completed July 1986.

ABSTRACT

The United States, as the world's largest producer and consumer of scrap and flake mica, has an obvious economic interest in applications of artificial intelligence technology that would expedite beneficiation of mica. In the fall of 1985, personnel of the Tuscaloosa and Rolla research centers of the Bureau of Mines and of the University of Missouri-Rolla, Institute for Artificial Intelligence started the following long term research project: Develop an Expert System Consultant for the three basic stages of mica beneficiation: Stage 1: characterization of the material; Stage 2: treatment to obtain a concentrate; and Stage 3: evaluation of the resulting concentrate. Completion of the long term project will result in an Expert System Consultant (MICA) that will assist in providing the following advice for an ore body being considered for development: (1) usability of the non-mica present; (2) feasible end uses for the mica present; (3) concentration treatments for mica; (4) equipment configurations for the concentration treatments; (5) setting the parameters of the treatment circuits; and (6) fine tuning of the circuits. The consultant will identify the laboratory treatments to be investigated to determine best plant scale flowsheets that will result in commercial operations. To date, a prototype of Stage 1 and a detailed introduction to the consultant have been developed. Stage 1 provides the basis for recommendations concerning the usability of non-mica, the feasible end uses for mica, and the concentration treatments for mica. This paper describes and discusses the development of the project and the results achieved.

ACKNOWLEDGEMENTS

The author would like to thank his advisor, Dr. Arlan DeKock, for support and encouragement during the author's tenure as a graduate student in computer science and especially for help on the research project that resulted in this paper; Drs. David Barr and Ralph Wilkerson for serving on his thesis committee; and Dr. Larry Vonalt for suggestions that enhanced the paper's readability.

In addition, the author would like to acknowledge and give special thanks to Gerry Sullivan, Tuscaloosa Research Center, U.S. Bureau of Mines, the domain expert whose participation made the research project possible.

TABLE OF CONTENTS

	Page
PUBLICATION THESIS OPTION.....	ii
ABSTRACT.....	iii
ACKNOWLEDGEMENTS.....	iv
I. INTRODUCTION.....	1
II. BACKGROUND: AI AND EXPERT SYSTEMS.....	3
III. PROCESS OF CONSTRUCTING MICA.....	6
IV. DISCUSSION OF MICA FROM AN END USER POINT OF VIEW.....	11
V. DISCUSSION OF MICA FROM A PROGRAMMING POINT OF VIEW.....	20
VI. RESULTS AND FURTHER RESEARCH.....	28
BIBLIOGRAPHY.....	30
VITA.....	31

I. Introduction.

The United States is the world's largest producer and consumer of scrap and flake mica and is expected to continue to occupy this position in the mica industry for the foreseeable future (Davis, 1985). Efficient mica beneficiation is an obvious economic interest of the United States mica industry, as well as the mica industry worldwide. The United States Department of the Interior is the government agency charged with the responsibilities for assessing our energy and mineral resources and of working to assure that their development is in the best interest of all of the people of the United States. The Bureau of Mines is the subagency of the Department of the Interior which is delegated the front line responsibilities for research and development in the area of mineral resources. Under these circumstances, it should come as no surprise to find that the Bureau of Mines has long and productive history of research and development in mica beneficiation.

In keeping with these economic interests, responsibilities, and history of research and development, in the fall of 1985, personnel of the Tuscaloosa and Rolla research centers of the Bureau of Mine and personnel of the University of Missouri-Rolla, Institute for Artificial Intelligence started to explore ways in which recent developments in Artificial Intelligence technology could be brought to bear on research and development in mica beneficiation. The focus of attention was on the subarea of Artificial Intelligence known as Expert Systems. The following long term research project was identified:

Develop an Expert System Consultant for the three basic stages in the beneficiation of mica ore: Stage 1: characterization of

the material; Stage 2: treatment to obtain a concentrate; and Stage 3: evaluation of the resulting concentrate.

The expert system consultant would be designed to show which laboratory treatments should be investigated, to determine best plant scale flowsheets that will result in commercial operations.

The primary user group was identified as individuals, with various backgrounds and experience in mica beneficiation, who were considering the construction of new mica beneficiation plants. A secondary user group was identified as individuals who, for various reasons, might be able to use the consultant in a tutorial mode.

Completion of the long term project will result in an Expert System Consultant (MICA) that can assist in providing the following recommendations for an ore body being considered for development: (1) usability of the non-mica minerals present; (2) feasible end uses for the mica present; (3) concentration treatments for mica present; (4) equipment configurations for concentration treatments; (5) setting the parameters of the concentration treatment circuits; and (6) fine tuning of the concentration treatment circuits. This paper describes and discusses the (ongoing) development of MICA, an expert system consultant for mica beneficiation.

II. Background: AI and expert systems.

Artificial Intelligence is a relatively new area of study that deploys a set of concepts with a history of controversy. The way of speaking introduced below, while fairly standard, is not without exception. (Cf. Winston, 1984, Brownston et al, 1985, and Jackson, 1986.)

Artificial Intelligence (AI) is concerned with designing computer systems that exhibit characteristics associated with intelligent human behavior. (Barr and Feigenbaum, 1981). The particular set of characteristics associated with intelligent human behavior that expert systems are to exhibit includes representing, reasoning about, problem solving, and advice giving, with regard to some knowledge-rich domain (Jackson, 1986). No mean task. But such is the way of the human experts that the expert systems are to emulate.

An expert system contains three constituents: a working memory for transient information about a knowledge domain and particular domain tasks; a knowledge base of permanent information about the knowledge domain and the set of domain tasks; and an inference engine that accounts for the behavior of the system. Among the extant historically important examples of expert systems that illustrate this architecture are the following: MACSYMA solves a variety of symbolic mathematical problems (Engleman, Martin and Moses, 1968); DENDRAL infers molecular structure from spectrographic data (Lederber, 1964); MYCIN consults for medical diagnosis and therapy (Shortliffe, 1976); TEIRESIAS facilitates the (MYCIN's) acquisition of new knowledge (Davis, 1976); HEARSAY provides for continuous speech understanding (Reddy et al, 1976); SOPHIE instructs

students in the debugging of electronic equipment (Brown, Rubenstein, and Burton, 1976); and PROSPECTOR advises on locations of ore deposits, given geological data (Duda et al, 1978). These examples serve to illustrate the diversity of knowledge domains and expertise that have been encoded in expert systems.

As expert systems, they satisfy, more or less, various criteria that have been proposed for distinguishing expert systems from other denizens of AI. For example, Jackson (1986, p.1.) suggests the following. An expert system must

- . deal with subject matter of realistic complexity that normally requires a considerable amount of human expertise;
- . exhibit high performance in terms of speed and reliability in order to be a useful tool;
- . be capable of explaining and justifying solutions and recommendations in order to convince the user that its reasoning is correct.

The first criterion is primarily concerned with content, the second with performance, and the third with justification.

Two observations concerning these criteria are in order. First, while domains of human expertise may be complex, they tend to be relatively narrow, for example, internal medicine within the general practice of medicine or corporate law within the general practice of law. Typically, expert systems have exhibited this same narrowness of knowledge domain expertise. For example, MYCIN is a medical expert system with domain expertise in the relatively narrow domain of blood infections (Shortliffe,

1976). Second, whereas human experts may exhibit various forms of expertise with regard to a given knowledge domain, such as being able to interpret data, perform diagnosis, and give advice, being expert at any one behavior is normally sufficient grounds for deeming the human an expert in the knowledge domain. However, a human expert consultant must augment any knowledge domain expertise with the the ability to provide an appropriate explanation at pertinent points in a consultation. For example, the physician must be able to rationalize the medical consultation. Such rationalization is also an essential behavior for an expert system consultant.

Building an expert system is a multifaceted activity, involving at least three identifiable roles: domain expert; knowledge engineer; and programmer. The domain expert is the individual whose knowledge of a particular subject or domain is to be captured in an expert system. The knowledge engineer has the task of extracting this domain main specific knowledge from the domain expert and representing it in a form usable for computer implementation. The programmer has the task of programming the knowledge engineer's representation of the domain expert's knowledge and expertise. Ideally all three roles are played by individuals who have the knowledge and expertise associated with their respective tasks. In point of fact, the activities of the knowledge engineer are so new that expertise in this area is still a scarce resource.

III. Process of constructing MICA.

In theory, the process of constructing an expert system involves three conceptually distinct stages: knowledge acquisition, knowledge representation, and programming implementation. Knowledge acquisition is the activity of deciding what comprises the knowledge domain and expertise. It involves identifying the constituent facts, generalizations, practices, etc. of the knowledge domain. In knowledge representation, the knowledge engineer decides how best to represent the knowledge domain and expertise. The knowledge engineer chooses from the available representation schemes, a representation congenial to both the knowledge domain and to the potential programming languages. In programming implementation, the programmer selects the programming language and determines how to implement the knowledge representation in the chosen language. These three stages were used to organize the construction of MICA.

In practice, however these three stages tend to overlap and interact. In this section, each stage of the process is illustrated with a partial reconstruction of the experience corresponding to that stage of the process. The main researchers in the MICA project are G.V. Sullivan, U.S. Bureau of Mines, Tuscaloosa, AL, in the role of domain expert and J.H. Marchal, Institute for Artificial Intelligence, University of Missouri-Rolla, Rolla, MO, in the roles of knowledge engineer and programmer. What follows is presented from the perspective of the knowledge engineer-programmer. However, unless otherwise stipulated, all of the information on mica beneficiation has been provided by the domain expert.

Knowledge acquisition.

The knowledge acquisition stage of the process of constructing an expert system consists of the efforts of the knowledge engineer to extract the domain knowledge and expertise from the domain expert. This stage is typically a very time consuming experience, no doubt due to the relative lack of expertise in the area of knowledge engineering. The MICA project proves no exception to this rule. The knowledge acquisition stage was initiated through telephone conversations and correspondence that started the written record of the project. Thus began the education of the knowledge engineer-programmer in the domain of mica beneficiation and the education of the domain expert in the role of tutor, for a student who needs more detail than the average student. Specifically, every step of mica beneficiation must be laid out in complete and minute detail, along with an explanation of what is going on and why. To construct the expert system, at least this amount and type of detail on mica beneficiation must be available to the knowledge engineer. It is expected that the respective educational experiences of the researchers will be ongoing for the foreseeable future of the project.

A series of answers to questions on the initial written material resulted in (1) a detailed outline of the general process of mica beneficiation and (2) an example of how a particular test case of muscovite in consolidated pegmatite would be treated by the domain expert. These documents continue to be elaborated and refined, and have proved to be the foundation of the knowledge acquisition component of the project. They were the focal point of discussion during the first meeting of the project personnel. This meeting took place in Rolla and involved

a weeklong "interview" with the domain expert, by a knowledge engineering group. Not surprisingly, the face to face exchange of information proved more efficient and productive for the knowledge engineering role players than what had preceded it by telephone and by mail. New information about mica and mica beneficiation was gathered, while earlier confusions were cleared up. The domain expert was introduced to artificial intelligence programming. All members of the project seemed to gain a better appreciation of what the other members were trying to accomplish in their respective roles. Just as the initial outline and example of mica beneficiation facilitated the first meeting, in turn, the meeting and its outcome facilitated the ongoing activity of knowledge acquisition.

Knowledge representation.

Enough information had been accumulated so that the knowledge representation stage of the process could be started. The knowledge representation stage of the process consists of the efforts of the knowledge engineer to find a congenial representing of the knowledge arrived at in the acquisition stage. A congenial representation is one that that is both accurate to the original knowledge and lends itself to relative ease of programming.

Some time after the meeting described above had taken place, the knowledge engineer set about manually designing a typical interaction between an end user and the incipient consultant. The resulting design document contained a very detailed account of what a consultation would look like to both the end user and the programmer. In the case of an end user, the primary concern is the information to be displayed on the

screen: for example, the questions asked an end user and the types of answers the expert system would expect, the advice or recommendations given, and where in the consultation this would take place. In the case of the programmer, the primary concerns is how and where this information is going to be represented: for example, the facts and generalization about mica, the relevant task behaviors of the domain expert, along with an overall structure that leads an end user through a consultation. This detailed design of the consultation can be seen as a definite step in the knowledge representation stage of the process of constructing an expert system.

The next event in the knowledge acquisition stage involved a visit by the knowledge engineer to the laboratory site in Tuscaloosa in order to observe mica concentration treatments and discuss the design document. Wet and dry concentration treatments were demonstrated, including the uses of spirals, flotation, hydraulic classification, air classification, electrostatic separation, and selective crushing and grinding. While this visit produced abundant new factual information about mica beneficiation, it, more importantly, assisted the knowledge engineer in developing an intuitive feel for the concentration treatments. While knowledge acquisition is obviously augmented by this experience, what might not be so obvious is that it also facilitates both the knowledge representation and programming implementation stages of the project.

The design document was acceptable to the domain expert, as a rendering of the first stage in a consultation on mica beneficiation and an outline for the later stages. It was decided that the domain expert would

continue to gather and provide the information required to complete the first stage of the consultation, while the knowledge engineer would begin the programming implementation of the first stage.

Programming implementation.

The programming implementation stage of the process involves deciding how to implement the results of the earlier stages. What is wanted is a program that is as accurate and complete as possible, with regard to the knowledge domain and expertise represented in a consultation with MICA. The crucial decision, here, is the choice of a programming language. For the reasons listed below, M.1 (Version 1.3(1), Teknowledge, 1986) is the programming language of choice for the first part of MICA. (1) M.1's backward chaining orientation fits the dominate reasoning pattern of the representation; (2) M.1's preferred solution mechanism and space, namely, structure selection from a relatively small solution space, fits the solution space of the representation; (3) M.1's built-in explanation function facilitates providing the explanations that are essential to an expert system consultant; and (4) M.1's facility for rapid prototyping fits the project's initial emphasis on speed of development rather than speed of execution. How this all comes to pass will be illustrated in the next two sections.

IV. Discussion of MICA from an end user point of view.

This Section will look at a typical consultation with MICA from the perspective of an end user. In order to accomplish this in small compass, an attenuated version of the Introduction and Part 1 of a consultation are presented. To improve the readability and efficiency of the presentation, the examples of segments from the consultation have occasionally been modified, but only in ways true to their original intent.

MICA is presented to the user as composed of eight units: an introduction; three parts corresponding to the three basic steps in the treatment of a mica sample: characterization of material, treatment to obtain a concentrate, and evaluation of the concentrate; and four appendices: information on types of mica, end uses for mica, coproduct potential for non-mica minerals, and screen-mesh conversions.

Introduction to a consultation.

When initiating a consultation, the user first encounters a greeting and a brief description of MICA's purpose. (See screen example (1)). The user is then invited to select, from a menu, what is to be done next. Detailed information on MICA is provided for the user in the first option of the menu, Introduction to MICA.

- (1) Welcome to MICA, an expert system consultant for mica beneficiation.

MICA is designed to provide information that will help you decide which laboratory treatments to investigate to determine best plan scale flowsheets that will result in commercial operations. MICA will also make recommendations as to the usability of any non-mica minerals that might be present in an ore body.

The information provided depends on the selections that you make from the menu given below. On your initial visit with MICA we recommend that you start with option 1, Introduction to MICA. Otherwise, make the choice appropriate to your interest.

What would you like to do next?

1. Introduction to MICA: Detailed information on MICA.
2. Part 1: Characterization of the Material: usability of non-mica; end use for mica; concentration treatments for mica.
3. Part 2: Treatment to Obtain a Concentrate: equipment configuration; circuit parameter setting; circuit fine tuning.
4. Part 3: Evaluation of the Concentrate from Part 2.
5. Appendix A: Types and Characteristics of Mica.
6. Appendix B: End Uses for Types of Mica.
7. Appendix C: Non-Mica Co-Product Potential.
8. Appendix D: Tyler Sieve Series, U.S.A. Sieve Series, and Mesh designation comparisons.
9. Quit the Consultation.

The names of the options and the brief descriptions are intended to give a user a fair idea of the information available as a result of choosing an option. The point of having such a detailed menu is that users will come to a consultation with different information, interest, and experience. While it is anticipated that first time users will uniformly start with option 1, it is also anticipated that users with different backgrounds will expect different choices to be available. After using any option, the choice of returning to this (main) menu is always available.

Part 1: Characterization of the Material.

Part 1 provides a user with recommendations on usability of non-mica minerals, end uses for mica, and concentration treatments for mica. These recommendations are based, in part, on information that the user is asked to supply. This advice comes at three different points in the consultation. In each case, the user has the option of revising the user supplied information on which the advice is based, so that alternative recommendations may be considered. To illustrate from a user point of view how this is accomplished, the attenuated path through a consultation continues. It is assumed that the user has already selected Part 1 from the menu.

Every option that the user can select from the menu is introduced with an appropriate explanatory message. In this instance, the explanatory message is followed by three sets of questions requesting the user to supply information about the type of rock to be treated, the type of mica present, and the type of non-mica material present.

Screen examples (2.1) and (2.2) are taken from the first set of questions:

(2.1) What is the type of rock to be treated?

1. schist
2. pegmatite
3. carbonatite
4. unknown

2. [User response.]

(2.2) What is the state of the material?

1. consolidated pieces
 2. granular
 3. plastic (wet clayey)
2. [User response.]

These examples typify three features of the question-response situation encountered by the user; first, the user is presented with a well defined set of choices; second, the response can be accomplished by simply typing a number; finally, the response "unknown" is acceptable in some instances, in that the consultation can continue on less than complete information.

Screen examples (3.1), (3.2), and (3.3) are taken from the second set of questions:

(3.1) What types of mica are present? Multiple answers are permitted.

1. muscovite
 2. lepidolite
 3. phlogopite
 4. biotite
 - ...
- 1,4. [User response.]

**(3.2) What percent of the material is muscovite?
Use a real number between 0.0 and 100.0.**

27.0 [User response.]

These examples illustrate two additional useful features. First, the user has the option of choosing multiple responses. For example, (3.1) allows for more than one type of mica to be identified as present, for instance, both muscovite and biotite. Second, numerical responses are allowed for.

For example, (3.2) instructs the user to respond with a real number between 0.0 and 100.0.

Another important feature of the question-response situation is MICA's ability to explain why information requested from a user is relevant to a consultation. Whenever a user is asked to supply information, "why" is always a permissible response: that is, the user can always ask for an explanation of the request. MICA can supply two types of explanations: explanations which show that the information requested is relevant to a rule that it is trying to satisfy; explanations that show that the information is relevant to a goal that it is trying to satisfy.

Consider the second type of explanation:

(3.3) What is the liberation size of the muscovite present?

1. less than or equal to 3 mesh.
2. greater than 3 mesh.

why. [User response.]

In (3.3) the user responds with "why," instead of one of the numbered alternatives; that is, the user asks MICA for an explanation. In this case, MICA can respond as follows.

(4) Information about the liberation size of the mica present is used to decide which end uses and concentration treatments are feasible.

It seems reasonable that a user would see this as rationalizing MICA'S original request for information on the liberation size mica, given the goals of the consultation.

Alternatively, MICA can rationalize a request for information by citing the rule(s) that it is trying to satisfy. For example, if the user had responded with "why" to example (2.1), "What type of rock is to be treated?", MICA will rationalize the request for information as follows.

- (5) M.1 (MICA's inference engine) is trying to determine whether the following rule is applicable in this consultation:

```
rule 20:
  if the type of rock is known and
    the state of the material is known and
    the size of the largest pieces is known
  then the characteristics of the ore = yes
  with a certainty factor of 100.
```

The user is given to understand that MICA is requesting information on the type of rock present, because that information is relevant to determining the characteristics of the ore being considered for beneficiation.

After similar questions are asked about the non-mica present, a series of recommendations is made to the end user. The first recommendation concern the usability the non-mica. For example,

- (6) Beryl present at 7 percent has a high coproduct potential.

Because the question about the non-mica minerals present allowed multiple responses, more than one recommendation could be forthcoming.

At this point in the consultation the user is presented with summaries of the information collected on the type of rock, mica and non-mica present. The user is then given the opportunity to change any of the user supplied information. For example,

- (7) Summary of the information on the mica present.

1. type of mica present: muscovite.
2. percent of muscovite present: 20.0.
3. largest flake size of the muscovite: 3 mesh.
4. liberation size of the muscovite: 10 mesh.

Would you like to change any of these values?

1. yes
2. no

This option allows any mistakes that might have been made to be corrected. It also allows alternative combinations of characteristics of the ore, end uses, and concentration treatments to be considered.

If the user is satisfied with the information presented in the summaries, MICA proceeds to consider feasible end uses and concentration treatments for the mica. First, the user is queried about end uses with:

- (8) Choose end uses for the types of mica present.
Multiple answers are permitted.

1. oil well drilling mud
2. tape joint compounds
3. paint and caulking compounds
4. reconstituted mica paper
- ...

1,3. [User response.]

Then MICA responds with recommendations on these end uses, based on the information gathered earlier. Since the questions about the type of mica present and end uses were both multivalued, several recommendations might result. Given the choices suggested above, muscovite and biotite would be considered for both drilling fluid and paint additive. For example,

- (9.1) Oil well drilling mud is not a feasible end use for muscovite liberated at 100 mesh.
- (9.2) Paint and caulking compounds are a feasible end use for muscovite liberated at 100 mesh.
- (9.3) Oil well drilling mud is not a feasible end use for biotite liberated at 250 mesh.
- (9.4) Paint and caulking compounds are not a feasible end use for biotite liberated at 250 mesh.

As before, when MICA makes a recommendation the user has the opportunity to change the user supplied information on which the recommendation is based, in order to consider alternatives. For example, a complete series of size fractions end use combinations could be considered by systematically changing the response to the liberation size (fraction) question.

The final recommendations that are made in Part 1 have to do with concentration treatments for the mica. First the user is questioned as

to the user's water resources and resources to dry out an ore body. Depending on the user's responses to these questions, various wet and dry concentration treatments are recommended. For example,

- (10.1) Satisfactory concentration of muscovite at 8/10 mesh using air classification has a confidence factor of 95.
- (10.2) Satisfactory concentration of muscovite at 3/4 mesh using selective crushing has a confidence factor of 90.

Again, alternatives can be considered, contingent on the information that is varied. We would naturally expect the mica to be present at size fractions other than what was identified as the liberation size. It will be important to consider if and how these additional size fractions can be economically treated.

At this point the user is presented with a menu of options on what to do next and Part 1 of a consultation is completed. Relevant information generated by Part 1 is passed on to Part 2: Treatment to Obtain a Concentrate, as it is the basis for the recommendations made there on equipment configurations, and treatment circuit parameter setting and tuning. In turn, the relevant information resulting from Part 2 is passed on to Part 3: Evaluation of the Concentrate.

V. Discussion of MICA from a programming point of view.

This section will look at a typical consultation with MICA from the perspective of the program. That is, we will look at three aspects of the program that make a typical consultation with MICA possible: the way that MICA stores, represents, and activates the domain knowledge and expertise that result in its recommendations.

There are at least four types of domain knowledge to be represented in MICA: (1) knowledge of particular facts, such as the material is 20 percent muscovite or the user's preferred end use is for drilling fluid; (2) knowledge of substantive generalizations, such as all mica used for oil well drilling mud must be liberated at not greater than 5 mesh; (3) knowledge about how to do things with regard to the domain, such as how to set up treatment circuits and set treatment circuit parameters; and (4) knowledge of when to do things, such as when to change a treatment circuit parameter setting.

Working memory.

Most of the knowledge is partitioned between MICA's knowledge base and working memory. The working memory contains two types of knowledge of immediate interest: user supplied facts, such as the mica present is muscovite with a liberation size of -10 mesh; and the recommendations of a consultation, such as flotation is the recommended concentration treatment. The content of working memory is generated anew with every consultation and modified in various ways during the consultation. This content will vary from session to session, contingent the ore sample under discussion.

Knowledge base.

In contrast, the content of MICA's knowledge base is relatively stable. Changes in this knowledge would be understood as explicit programming changes. The knowledge base will contain rules representing the substantive generalizations about mica and the knowledge about how and when to act, mentioned above. It will also contain some facts about both mica and non-mica that are needed across consultations. All this information remains unchanged from one consultation to the next. Statements of these rules and facts are explicitly constructed when building the expert system.

The most important examples of facts, here, are those in the knowledge base which ground MICA's recommendations to an end user. For example, any of the following might be suggested to a user during a consultation:

- (11) The beryl present at 7 percent has a high coproduct potential.
- (12) The mica present can(not) be used in the manufacture of roll roofing.
- (13) Satisfactory concentration of muscovite at 3/4 mesh, using selective crushing has a confidence factor of 90.

The knowledge base entries for these three types of statements compose three different look-up tables. The tables corresponding to statement (11) and to statement (13) are fairly complete, while the table corresponding to statement (12) is quite sparse. The reason for the sparseness of the second table is that the knowledge on which it is built tends to

be proprietary and not readily available to the project. Consequently the second table contains only plausible examples. The knowledge base entries for the other two tables were supplied by the domain expert. These examples represent the expertise of the domain expert with regard to non-mica minerals potential and concentration treatments. The presence of this information in the knowledge base is one of the reasons for deeming MICA a prototype "expert" system consultant.

These tables are open to modification, as new information becomes available. Among other things, this built-in flexibility allows particular end users to change these component of the knowledge base to suit their needs. For example, the presence proprietary knowledge can be controlled in different replicas of MICA.

MICA implements the look-up table entries as ordered tuples. For example, tuples corresponding to advice about the potential of non-mica minerals present have the following form.

(14) non-mica(TYPE OF ORE, TYPE OF NON-MICA, % OF ORE)=ADVICE.

This expression can be understood as saying that the type of ADVICE given depends on the TYPE OF ORE present, the TYPE OF NON-MICA present, and the % percent OF the ORE that is non-mica. The expressions in capital letters are variables which take as their values examples of what they describe. For example, rutile, calcite and dolomite are all values of the variable TYPE OF NON-MICA. The entries in the look-up tables are all instances of the corresponding tuples. The entry

(15) non-mica(pegmatite, beryl, 13%) = high ...

is an instance of the tuple displayed in example (14), and if activated it would result in the message displayed in example (11).

The other types of knowledge found in the knowledge base are usually implemented with rules. Such knowledge would include the substantive generalizations about minerals along with the knowledge about how and when to act. For example, knowledge as to the usability of non-mica is represented with a variation of the following rule:

(16) rule 49:
 if the type of ore = TYPE OF ORE and
 the type of non-mica = TYPE OF NON-MICA and
 the percent of non-mica = % OF ORE and
 non-mica(TYPE OF ORE,TYPE OF NON-MICA, % OF ORE) =
 ADVICE
 then the usability of non-mica = ADVICE. with a
 certainty factor of 99.

Rule 49 can be understood as saying that if MICA has the information as to the type of ore, the type of non-mica, and the percent of non-mica, and these values collectively satisfy an entry in the look-up table, then MICA will return the corresponding table value of ADVICE as the advice component of a recommendation on the usability of the non-mica, with a confidence factor of 99. (MICA's ability to deal with uncertain knowledge, while noteworthy, will not be discussed at this time.)

This example of rule 49 illustrates, nicely, the interaction of MICA's working memory and knowledge base. The user is queried on the type of ore and the type and percent of non-mica present. This information then goes into the working memory. As indicated above, general information about the usability of non-mica is already found in the knowledge base, in the form of a look-up table and a rule accessing the table. The relevant knowledge in the working memory and the knowledge base is then combined to result in the advice provided to the user. To be specific, the advice in example (11) results from the user supplied information in the working memory, the table entry of example (15), and rule 49. The details of this interaction are the province of the inference engine.

Inference engine and other control structures.

The final repositories of domain knowledge and expertise are MICA's "control structures." The name is intended to suggest those components of the program that determine the sequence of steps in which a consultation is executed. They are part inference engine and part program design generated. For example, the main goal statement discussed below is provided by the underlying inference engine, while the configuration of conditional rules is a matter of program design.

One way that MICA controls a consultation is through the use of goals. Once a consultation is centered on a particular goal, MICA pursues it apace, until the goal is realized, replaced, or determined to be unattainable. MICA has several control structures that identify goals. We will examine only two: main goal statements and conditional rules. MICA uses the type of statement shown below to identify a main goal, "main"

in the sense that its satisfaction controls any additional goal activity.
Part 1 includes the following goal statement.

(17) initial data = part 1 stage 1 completed.

Once this goal is activated, MICA will relentlessly pursue the goal "part 1 stage 1 completed", until it realized.

The conditional rule is the second control structure that MICA uses to identify goals and thereby control the path of a consultation. To illustrate this, consider the following example

(18) rule 17:
 if the introductory message is printed and
 the mineral information is collected and
 the load of part 1 stage 2 is completed
 then part 1 stage 1 is completed = yes with a
 certainty factor of 100.

The rule has the form of a conditional sentence, with the consequent defined as goal of the rule. The rule requires for its (consequent) goal to be satisfied, that its antecedent (conditions) clauses must be satisfied. On this understanding of a conditional rule, it is reasonable to define the antecedent conditions as intermediate or subgoals of the rule. It is also important to understand that MICA will attempt to satisfy the subgoals in the order that they are listed in the rule. Consequently, the order in which antecedent conditions are placed in a rule can be used to control a consultation.

In example (18), the goal of the rule is also a main goal. If MICA sets about trying to satisfy this main goal by using rule 49, MICA's attention will be immediately redirected to the antecedent conditions as subgoals. In effect, these are new goals, which supersede the original goal in terms of the order in which they are to be satisfied. There are various ways that MICA can come by the information required to satisfy these new (sub)goals. For example, the information might already reside as facts in the working memory or the knowledge base. Alternatively, MICA might query the user for the information. As a matter of fact, all of the subgoals of rule 17 are conclusions of yet additional rules, resulting in additional subgoals to be tested.

For example, take the second clause of rule 17, "the mineral information is collected" It turns up as the (consequent) goal of rule 19.

```
(19) rule 19:
      if the characteristics of the ore is known and
         the characteristics of the mica is known and
         the characteristics of the non-mica is known and
         the mineral summary is known
      then the mineral information is collected = yes with
         a certainty factor of 100.
```

If rule 19 is used to realize the second clause of rule 17, a whole new set of (antecedent clause) subgoals must be considered. This will include the third antecedent clause of rule 19, which involves the information on non-mica discussed in examples (11 - 16) above. Of course, if this process is to stop, eventually the required information will be found either as user supplied knowledge in the working memory or as a factual component of the knowledge base.

This discussion of conditional rules has led us back to the example involving the usability of the non-mica present. In that example we had, in fact, assumed that the user had been queried for the information concerning the non-mica, etc., that would ground the recommendation concerning its usability. We are now in a position to summarize what is typically the case when MICA manages to give advice or make a recommendation. The consultation will have been activated by "firing" a goal statement. MICA will be in pursuit of a main goal, essentially the completion of the consultation. Along the way, the various recommendations that we are interested in will fall out as intermediate goals to achieving the main goal. These recommendations will have been properly ensconced as conclusions of conditional rules and their activation will be accounted for in the manner explained above.

VI. Results and further research.

The project is to develop an expert system consultant for the three basic stages of mica beneficiation that can provide recommendations on: (1) the usability of non-mica; (2) feasible uses for mica; (3) concentration treatments; (4) equipment configurations; (5) setting treatment circuit parameters; and (6) fine tuning treatment circuits. In addition, it is essential that the consultant be able to rationalize its behavior.

Results.

First, MICA prototypes the first stage of the projected consultant, Part 1 characterization of the material. Second, Part 1 provides three of the six types of recommendations set as goals for the project, namely, (1 - 3) above. Third, MICA can rationalize its behavior in contexts where it has asked an end user to supply information for a consultation by explaining the relevance of the information to the consultation. Fourth, MICA's knowledge base has built-in flexibility that allows different end users to control proprietary information in the particular replicas of MICA.

Further research.

First, Part 2 treatment of the material and Part 3 evaluation of the concentrate have yet to be completed. Their completion will require input from all three stages of the process of building an expert system: knowledge acquisition, representation, and programming. It is an open question whether their development will follow the development of Part 1. For example, M.1's backward chaining orientation that proved to fit the solution space of Part 1 might not be best suited to Parts 2 and 3.

Second, MICA's explanation facility needs to be expanded. MICA'S ability to justify requests for user supplied information needs to be complemented with an ability to justify MICA's various recommendations to a user. Third, and finally, continued development of MICA should open up artificial intelligence research in the larger area of mineral beneficiation, of which mica beneficiation is a but a particular case.

BIBLIOGRAPHY

Barr, A. and Feigenbaum, E.A. [eds.], 1981, The Handbook of Artificial Intelligence, Vol. 1, William Kaufman, Inc., Los Altos, Calif.

Brown, J.S., Rubinstein, R., and Burton, R., 1976, "Reactive learning environment for computer assisted electronics instruction," BBN Report No. 3314, Bolt Beranek and Newman, Inc., Cambridge, Mass.

Browston, Farrell, R., Kant, E., and Martin, N., 1986, Programming Expert Systems in OPS5, Addison-Wesley, Reading Mass.

Davis, L., 1985, "Mica", preprint from Bureau of Mines Bulletin 675, Minerals Facts and Problems 1985.

Davis, R., 1976, Applications of meta-level knowledge to the construction, maintenance and use of large knowledge bases. In Davis, R. and Lenat, D., 1980, Knowledge-based systems in artificial intelligence, McGraw-Hill, New York.

Duda R.O., Gaschnig, J., Hart, P.E., Konolige, K., Reboh, R., Barrett, P., and Slocum, J., 1978, "Development of the PROSPECTOR consultation system for mineral exploration," Final Report, SRI Projects 5821 and 6415, SRI International, Inc., Menlo Park, Calif.

Engleman, Martin and Moses 1968. See "D1. MACSYMA," in Barr, A., and Feigenbaum, E.A., [eds.], 1982, The Handbook of Artificial Intelligence Vol. 2., William Kaufmann, Inc., Los Altos, Calif.

Jackson, P., 1986, Introduction to Expert Systems, Addison-Wesley, Reading Mass.

Lederberg, J., 1964, "DENDRAL-64: A system for computer construction, enumeration and notation of organic molecules as tree structures and cyclic graphs. Part I Notational algorithm for tree structures," Report No. CR-57029, NASA.

Reddy, R., Erman, L., Fennell, R., and Neely, R., 1976, "The HEARSAY speech understanding system: An example of the recognition process," IEEE Transactions on Computers C-25, pp. 427-431.

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

Winston, P.H., 1984, Artificial Intelligence, 2nd ed., Addison-Wesley, Reading Mass.

VITA

Joe Marchal has been enrolled as a graduate student in computer science since Spring of 1984. He received his previous college education at Washington University, Stanford University, University of Oslo, Norway, and Michigan State University. He has received a Ph.D. (1971) in Philosophy from Washington University, and an M.A. (1962) and B.A. (1960) in Philosophy from Michigan State University. He received his secondary education in Traverse City, Michigan, and his primary education in Traverse City and Detroit, Michigan. He was born on January 11, 1938 in Traverse City, Michigan.