Incorporating User Models Into Expert Systems For Educational Diagnosis

Robin Cohen
Marlene Jones

Research Report CS-86-37
September 1986

INCORPORATING USER MODELS INTO EXPERT SYSTEMS FOR EDUCATIONAL DIAGNOSIS

Robin Cohen and Marlene Jones
Dept. of Computer Science
University of Waterloo
Waterloo, Ontario
Canada

ABSTRACT

In this paper we study a particular real-world domain, that of educational diagnosis. We argue that expert systems for educational diagnosis require user models, and that these user models should include several components, including the user's background knowledge of both the student and the domain as well as the user's goals. Our proposal is directed at enhancing the particular expert system of the CGD project. We then propose an architecture for this expert system that separates the knowledge base into relevant components and includes a user model. We further demonstrate that this divided model for the system facilitates providing the best response for a particular user, according to his background knowledge of the domain and of the student, and his goals. Finally we argue that the techniques outlined here will be useful in general in expert systems.

Key Words: user models, expert systems, educational diagnosis, generating natural language explanations

INCORPORATING USER MODELS INTO EXPERT SYSTEMS FOR EDUCATIONAL DIAGNOSIS

Robin Cohen and Marlene Jones
Dept. of Computer Science
University of Waterloo
Waterloo, Ontario
Canada

1. INTRODUCTION

There are many reasons for including a model of the system's user within an expert system. The foremost is to allow the system to tailor its explanations to meet the particular user's requirements. These explanations may concern the gathering and entering of data, the use of the system itself, or the system's findings, recommendations and reasoning process. One may argue that it is possible to alter the system's explanations without going to the effort of including an elaborate user model or any user model at all. For example, one can simply store predetermined explanations with varying amounts of detail and allow the user to request further (i.e. more detail) explanations, if desired. This is unsatisfactory in situations where one would like to vary several aspects of the explanations simultaneously, which is the case in educational diagnosis, the application domain with which we are concerned.

We provide examples to defend the need for user models in our expert system, and then address the issues of what information to incorporate into a user model, how the model can be derived and updated, and how the model can be used to provide explanations. In short, we argue for the feasibility and usefulness of modelling (not just the student but the expert system user) for educational diagnosis.

2. APPLICATION DOMAIN

This paper studies a particular real-world domain -- educational diagnosis. When a child is experiencing learning difficulties, it is important to ascertain the exact nature of the problems as quickly as possible. This may involve an in-depth educational diagnosis with the ultimate aim of developing a remedial program which is tailored to that particular child. An in-depth diagnosis/assessment generally includes examining the child's developmental history including the home environment and early school history; assessment of psychoeducational correlates such as intellect, visual, and auditory acuity; language skills; achievement within the fundamental areas such as reading, arithmetic, spelling; medical and behavioural information as well as a detailed assessment of any achievement areas in which the child's performance is deemed to be inappropriate. Usually such an assessment is carried-out by a psychologist and/or resource room teacher. The resultant diagnostic report is generally made available to the regular classroom teacher, any other participants in the remedial program, as well as the necessary administrators such as the principal. Of course, the diagnostic findings and educational decisions should be discussed with both the child and parents.

There are well-accepted guidelines and procedures regarding educational diagnosis (Committee on Teacher Education and Professional Standards 1971), as well as appropriate models; for example, see (Adelman 1982; McLeod 1982) and the references therein. It is rare, however, to find educational diagnosis carried out as thoroughly and rapidly as the guidelines suggest. This is largely due to insufficient resources, particularly personnel such as highly experienced resource room teachers.

One means of facilitating educational diagnosis within the regular school environment is via an expert system to guide a teacher (such as the resource room teacher) through the various stages of diagnosing learning difficulties. This is the long-term goal of the CGD (Computer-

Guided Diagnosis) Project based at the University of Saskatchewan (McLeod and Jones 1985). Preliminary systems have already been developed in the area of reading (Colbourn 1982; Colbourn and McLeod 1983, 1984) and arithmetic (Jones and Tubman 1986, Tubman 1986), with a more elaborate, comprehensive system under development (Jones et. al. 1986, McLeod et. al. 1986). It is this latter expert system that we consider in this paper.

The advantages of such a system are many. It will serve as an information source to experienced diagnosticians, an assistant to less experienced resource room teachers, and a tool for training teachers. Even during the handling of one particular case, there may be a variety of users at each of the several stages of the diagnostic process -- in-depth assessment, development of a remedial program, and monitoring the progress during remediation, which generally includes further assessment and may include placement and funding decisions. Furthermore, one may want to make some portions of the system's findings and suggestions available to either the parent or child. Hence, such an expert system may be used by or required to supply explanations to any of the following individuals: psychologist, resource room teacher, educational consultant, regular classroom teacher, special education teacher, principal, outside administrator such as the head of special education services or superintendent, researcher, parent, the student him/herself, or an education student or training teacher.

Obviously these users vary in a variety of factors, most notably in their *background regarding the field* of special education. However, even personnel with a strong background in special education will vary immensely in regard to their familiarity with testing procedures, knowledge of individual tests (some of which require special training in order to administer), knowledge of areas of assessment such as reading, arithmetic, language etc., administrative procedures and guidelines, laws and mandates, outside resources, funding criteria and availability etc.

Individuals will also vary in their *familiarity with the individual case*, such as knowledge about the medical history, the home environment or the current school behaviour. For example, the child, parent and teacher are familiar with the situation on a day-to-day basis whereas the psychologist may have a one-page referral form and the principal may have a short synopsis of the recent observable problems. Similarly, the "official" opinion on the case may be available to a select few. Some individuals, such as the psychologist and resource room teacher, may be familiar with a variety of similar cases which the teacher and principal are not.

Also the *goals* of the users may vary. Although, hopefully all users' main goal is the implementation of the most appropriate remedial program, the subgoals may vary. For example, the teacher's most immediate concern may be techniques for decreasing disruptive classroom behaviour. The principal may be concerned as to whether the child qualifies for special funding. The parent may be searching for assurance of the child's normalcy.

How does this myriad of variations affect the expert system's explanations? First of all, if the user has a particular goal in mind, this should be addressed. Even for a particular question or subgoal, the explanation itself should vary depending on the user's familiarity with the case and background in special education. Ideally, both the amount and detail of the explanation should be altered. Most parents will not benefit from a detailed description of the child's performance on a variety of test, nor is an in-depth description of each remedial activity appropriate. This information is, however, critical for the resource room teacher who is implementing the program in conjunction with the regular classroom teacher (often with the assistance of a consultant). The language in which the information is expressed is also critical. Most parents are overwhelmed by the use of technical terms. Straightforward, honest explanations are desired, which must include realistic, encouragement for both parent and child. It is important to remember that both parent and child may have been aware of problems for several years without understanding them. In other words, one may be contending with several years of frustration, fear and maybe even mistrust. Hence, the tone of the discussion is important.

We focus in this paper on the problem of how to set up the expert system so that different users can get different explanations. We then provide a framework for varying the content of the explanation. The topic of varying the surface form with appropriate natural language generation is left as future work.

3. DESIGNING THE SYSTEM

In this section, we propose an architecture for our expert system and a control structure for searching the knowledge bases of the system when producing explanations, which will provide different responses to particular users, thus making the current system available to a variety of individuals. Although this model for making the system user-specific has not yet been implemented, we provide an analysis of some sample queries to illustrate the usefulness of the approach.

One contribution of the research is thus specifically for educational diagnosis, a method for making a system more effective, and even more specifically a chronicle of the latest progress in the CGD project. Moreover, the proposal outlined here can be viewed as a general strategy for incorporating user models into expert systems to provide clearer explanations using less costly processing.

3.1 MOTIVATION FOR SYSTEM DESIGN

The educational diagnosis expert system of the CGD project is designed to interact with a school psychologist or resource room teacher, to discuss the case of a particular student. As the diagnosis proceeds, information regarding the student's background, behaviour, abilities, skills, *etc.* is collected from a variety of sources. At any point in the diagnostic process, the system analyzes the available information and employs its knowledge of educational diagnosis to draw conclusions which may include recommendations for further assessment or particular hypotheses regarding the student's difficulties. The main types of questions being posed by the user are thus: (i) what is the diagnosis for this student? (ii) why is this the diagnosis? and (iii) why is this diagnostic procedure (such as a particular test or type of test) being recommended? The system should be able to respond to similar questions posed by a classroom teacher, parent, or school administrator. Furthermore, there may be a variety of backgrounds among members of any of these "types" of users - e.g. two psychologists with different favoured approaches to diagnosis.

In Example I we illustrate the need for different responses to the same question for different types of users.

Example I:
User: Could the student's mispronunciation "errors" be due to dialect?
Response to Parent: Yes, non-standard pronunciations may be due to dialect rather than poor decoding skills.
Response to Psychologist: Yes, the student's background indicates the possibility of a dialectal difference.

Here, the parent is asking for more information regarding the domain of educational diagnosis (lacks the domain knowledge). The psychologist knows about the tests employed and has the necessary domain expertise, but is unaware of the background of the student.

In Example II, we emphasize the possible difference in goals of users employing the expert system. In particular, we indicate session-wide preferences a user may bring to bear.

Example II:
System: Administer an individual intelligence test. For this particular case, the Peabody Picture Vocabulary Test (PPVT) is recommended.
User: Why is this test being recommended?
For each user, the stated goal is the same ie. to find out why this test is being recommended. However, the inferred or intended goal may be different.
Psychologist Preference: administer intelligence tests with which he/she is familiar (e.g. the

psychologist may be more familiar with the Stanford Binet or the WISC-R (Weschler Intelligence Scale for Children - Revised)).

Response: The PPVT is being recommended because it does not require verbal responses, and the student appears to have language problems including difficulties with expressive language. Hence, the PPVT is more appropriate than more common tests such as the WISC-R or Stanford Binet.

If the user is a teacher, principal or parent, the system would respond by explaining why an intelligence test is being suggested, as well as why this particular test is appropriate. The user may have additional goals which also must be addressed, as illustrated below.

Principal Preference: employ tests which can be used in funding formulae

Response: There is no evidence to suggest that the student will require or will qualify for additional funding.

Parent Preference: minimize testing due to stress involved

Response: The student should not find this test stressful, as no verbal responses are required; answers are indicated through pointing.

The above discussion may lead the reader to believe that it is sufficient to stereotype our users and exploit such defaults when answering the user's questions. Defaults alone are, however, insufficient. For example, a parent may have substantial knowledge of educational diagnosis because of involvement in organizations such as ACLD (Association for Children with Learning Disabilities) or CEC (Council for Exceptional Children) or because of professional training within a related field. And yet, the defaults of a fully trained psychologist may be too much.

Furthermore, there are several possible reasons why a psychologist may not fit a given stereotype. For example, he/she may have additional information about the student which the system is not assuming. The psychologist may have been involved in some previous diagnostic situation with this student or perhaps knows the student from other circumstances. A second possibility is that the psychologist lacks experience within a particular area of exceptionality or assessment, such as the diagnosis of emotional difficulties. Having determined this, the system knows that additional explanations must be supplied for all aspects of assessment within this area, not just for the user's first query. Similarly, consider the case where the student is visually-impaired. Although the psychologist may be well aware as to why certain skills are being assessed, he/she may be unaware of the appropriate tests to do so, given this additional constraint. A third type of variation from the initial stereotype is based on differing philosophies. For example, the user may adhere to the belief that visual-motor training is effective in the remediation of reading difficulties; although previously a popular educational practice, current evidence indicates otherwise. Or the philosophical difference may stem from a disagreement regarding the usefulness of certain types of information during the diagnostic process or the conclusions that can be drawn from the student's performance on a particular test. The following example, illustrates this point.

Example III:

System: Administer an individual intelligence test such as the WISC-R or Stanford Binet.

User instead administers the ITPA (Illinois Test of Psycholinguistic Abilities) and the student's performance is recorded.

User: What are the conclusions based on the results of the ITPA?

System: The child's intellectual ability appears to be within the normal range; no further assessment of intellectual ability is required.

User: What other conclusions were drawn? For example, the student's poor scores on the Auditory Association and the Auditory Closure subtests indicate that the student's auditory channel is weak.

System: Although the test results raise some suspicions regarding the child's processing of auditory information, such conclusions cannot be drawn from the ITPA alone because the validity of the ITPA's subtests is questionable.

This type of question from the user would alert the system that the user does not adhere to certain current philosophies. The system does not consider the individual subtests of the ITPA to be valid. Moreover, the user does adhere to the ITPA model and related philosophies. Because the psychologist has drawn certain conclusions which the system has not, this will affect what he/she concludes in regard to reading assessment or even which tests he/she thinks are appropriate.

3.2 DETAILS OF THE SYSTEM DESIGN

In essence, we are requiring our expert system to accomplish two modelling tasks simultaneously. The system must model the student in order to produce an appropriate diagnosis. In order to provide appropriate responses to the user, it must also model the user. The system must, therefore, contain the following three knowledge bases:

(1) domain knowledge which includes knowledge regarding diagnostic procedures, practices and educational theories, information regarding standardized tests and remedial activities etc.,
(2) a model of the student,
(3) a model of the user.

A more detailed description of the required knowledge bases is provided below.

Domain knowledge

Rules
- * diagnostic principles and terminology
  - e.g. if possible, employ individually administered tests because individually administered tests tend to have higher reliability than group tests
- * test interpretation rules
  - * general
    - e.g. If the student's score is one standard deviation below the mean on an individually administered achievement test, then further assessment is required.
  - * specific
    - e.g. a reading quotient of less than 80 on the Schonell Reading test indicates that further assessment is required

Databases
- * test database
  - e.g. information regarding purpose, subtests, types of scores, reliability, *etc.*
  - (The system employs this information in conjunction with the above rules, particularly the general test interpretation rules, to draw appropriate conclusions.)
- * remedial activities and methodologies

Student knowledge

- * background
  - * historical background
    - e.g. medical history
  - * current background
    - e.g. school placement
  - * behavioural background
    - e.g. classroom behaviour

* test results (facts from testing)

* hypotheses
  theories of the system as to diagnosis of the student

User model

* knowledge of the domain (same division as above)

* knowledge of the student (same division as above)

* goals
  * for the session
    * domain goal (find "best" diagnosis for student)
    * preferences
      e.g. administer tests that can be used to get funding

  * for the particular question
    * stated goal
    * intended goal

As mentioned earlier, during the course of the diagnosis, the expert system is developing hypotheses (ie a diagnosis) based on the information at hand, as well as suggesting further testing in order to gather information to improve the accuracy of the diagnosis. The techniques employed in answering the user's questions regarding either the diagnosis or the recommended diagnostic procedure are essentially the same. Therefore, for the next section of this paper, we concentrate on one type of question "what is the diagnosis?" and illustrate our methodology for providing an appropriate answer to this query. Figure 1 highlights the system's knowledge bases and introduces a corresponding notation.

System (S)
  Knows About Domain (KAD)
  - diagnostic principles (d)
  - test rules (t)

  Knows About Student (KAN)
  - background (b)
  - test facts (f)
  - hypotheses (h)

  Believes that the User (B)(U..)

    Knows About Domain (KAD)
    - diagnostic procedures (d)
    - test rules (t)

    Knows About Student (KAN)
    - background (b)
    - test facts (f)
    - hypotheses (h)

Wants (W)
- domain goal
- domain preference
- query goal (stated)
- query goal (intended)

Figure 1: Labelling of System Knowledge Bases

In the next section we show a control structure of the system to generate a response for a user, mobilizing appropriate knowledge bases. The main features of this strategy are:
(i) default assumptions about an entire class of users (psychologists, teachers, parents, etc.) are used as a starting point for the model of the user's knowledge
(ii) the system is aware that its model of the user (as well as its knowledge of domain and student) is updatable - i.e. that new information revealed during a session with a user may cause the models to be revised. The defaults provide a head start for the modelling, and the dynamic maintenance of the knowledge bases allows for additional exceptions to the assumed model of the user. We focus on the psychologist as a user first, and then include discussion of the possible defaults and updates for the other types of users.

4. THE USER MODEL STRUCTURE

In this section we advance a proposal for the acquisition and updating of a model of the user to produce quality explanations. What we want ideally, is a system that can respond to all the examples offered in the first part of 3.1. We indicate here a proposal to isolate the parts of the system required for responses. But we do not deal with lower level generation issues, such as choosing the most appropriate terminology for a user. We will at times indicate what should ideally result; in these cases we also make clear to what extent the current framework also provides a solution to these harder questions.

4.1. THE PSYCHOLOGIST AS USER

We now show how to use the diagram of the system, outlined in Section 3.2, to generate appropriate responses for different users. We illustrate the procedure for the case of a user who is a psychologist. Consider, in fact, a user who is the psychologist who tested the student whose case is being evaluated by the system.

The general principle used to produce a desirable response in an efficient amount of time is basically: "Don't explain what's already known, and keep track of what's already known, to avoid providing extraneous detail". This extension of Grice's maxim of quality (Grice 1975) allows for some division of knowledge and tracking of shared knowledge bases.

At the start of the session, the system would evaluate the knowledge of the user, in terms of its knowledge bases, as follows:

1) SB (UKANf) = SKANf [read as: "system believes user knows about student.."]

The system believes that what the user knows about the test results of the student is equivalent to what it knows. This assumption is reasonable as the test results have been input by the user.

2) SB (UKANbt) = SKANbt

The system believes that the background knowledge it has acquired about the student, specifically through testing, is equivalent to what the user knows. This background information would be data like the student's age, public school, etc. that would be part of the information drawn during the testing stage. Note that the system may have gathered additional information about the student's background, from other sources, in the course of its search for a diagnosis, so there may be some discrepancy between the overall knowledge of the student's background between system and user (KASb).

3) SB (UKADt) = SKADt

The system starts off with the assumption that the user is familiar with the standard interpretation of tests, used in diagnosis.

4) SB (UKADd) = SKADd

The system also assumes that the user shares knowledge about basic diagnostic procedures.

The main features of the design are thus: (i) a coarse-grained division of the system into knowledge bases, to identify sources of shared knowledge and to thus focus on elaborating unknown information (ii) a tracking of specific information on the user's current beliefs, to produce more co-operative responses, within the areas where the knowledge is not shared. We illustrate these in a tracing of the system's algorithm below.

In order to describe how the system mobilizes certain components of the proposed architecture to generate better responses it is necessary to present some model for the system's reasoning. The computational framework employed in the CGD expert system is THEORIST, a logic programming system based on default reasoning, developed at the University of Waterloo (Poole et. al. 1986).

However, to clarify the discussion below, we refer to a more conventional model of reasoning in expert systems. Consider a production rule base, where the system encodes a series of "rules" of the form: "If X then Y", and may then record various "facts" of the form "X1", "X2", etc. By some deduction mechanism such as forward chaining, a diagnosis may be drawn, by proceeding from the facts and rules to the most likely conclusions. The main framework which we draw for discussion is: facts + rules lead to diagnosis (F + R --> D). (Note: which of the system's defined knowledge bases are invoked for the system's control structure should remain the same, regardless of the reasoning model employed.)

In our domain of educational diagnosis, the facts are items like scores on tests for students, rules are items like standard interpretations of test results for diagnosis. The model for reasoning is thus:

Facts + Rules ---> Diagnosis

this is expanded to:

Background + Test Results + Domain Rules --> *Hypotheses*

Further Domain Rules can be divided into Test Interpretation Rules and Diagnostic Procedures, yielding the final formula:

Background + Test Results + Test Rules + Diagnostic Procedures --> *Hypotheses*

There are in fact two ways to use Test Rules to produce a hypothesis: a) use a specific test interpretation rule or b) use a general test interpretation rule, together with information from the test database regarding the test at hand. We put these two cases together under the label "Test Rules". Note that extensive use of the test database in fact occurs in conjunction with responses that suggest administering new tests. For future work, we can examine more closely possible sources of difference between the system and psychologist users regarding testing procedures. One strategy worth studying is segmenting the test rules portion of the knowledge base as well in order to more appropriately respond to the user.

Then, a system should respond to a question posed by a user by activating those parts of its reasoning towards a hypothesis that the user may not already share. We illustrate this focusing mechanism below.

Dialogue I:
Question: What is the diagnosis for this student?
Response: *Hypothesis*
Reasoning: UKf (facts)
UKt (test rules)
UKd (diagnostic procedures)
- if all Background has been acquired through testing, then UKb (background)
- since *Hypothesis* is the only component not already shared by the user, this is included in the response

Note that after this exchange, SB (UKANh) is updated to include the hypothesis just presented. Once more, the separation into knowledge bases allows for a current view of the user's knowledge.

Dialogue II:
Question: What is the diagnosis for this student?
Response: *Hypothesis*, because background fact X is true for this student.
e.g. The system suspects that the student can be classified as "learning disabled"; there is a family history of reading problems.
Reasoning: as above, but in this case SKASb ~= SB (UKASb) because there is some background on the student not directly acquired through testing

Rule: The system is compelled to add in the response information that the user may not yet know.

This example indicates that for our domain the knowledge base referred to as background on the student may be further subdivided into background acquired through testing and background acquired from other sources. Again, the principle is to divide to facilitate labelling of shared knowledge with the user.

The above dialogue can also be employed to demonstrate how the system would handle Example I, as illustrated in Dialogue IIb below. Recall that the user is asking about a possible hypothesis:

Dialogue IIb:
User: Could the student's mispronunciation "errors" be due to dialect?
Response to Psychologist: Yes, the student's background indicates the possibility of a dialectal difference.
Reasoning: (as in Dialogue II above). In addition to responding with *Hypothesis*, provide the relevant background fact not acquired through testing.

A similar approach is employed when answering the parent's query:
Response to Parent: Yes, non-standard pronunciations "errors" as measured by the Neale Analysis of Reading Test may be due to dialect rather than poor decoding skills.
Reasoning: In addition to responding with *Hypothesis*, provide the relevant Domain Rules (in this case a Test Rule).
(A more detailed description of responding to users in the parent class is contained in Section 4.2.2.)

Note that when we paint with such a broad brush as for the case of the psychologist, it is not unreasonable to have the system simply mark those parts of its knowledge base used for the calculation of the response which are not shared by (known to) the user. This is because the majority of the knowledge is shared, so that the process is still computationally feasible. The suggestion for the general case (any user, any type of user) is to track the specific user more closely to focus on possible arenas for misunderstandings or lack of information.

For the user as psychologist, this tracking becomes an issue in questions of the type below.

Dialogue III:
Question: What is the diagnosis for the student?
Response: *Hypothesis*
Question: But why is *hypothesis* appropriate for this student?
Analysis: There is now a conflict between what the system has recorded as known by the user, and what the user should thus be able to conclude, and the lack of belief in the conclusion, as implied by the question asked.

One description of the problem is:

UK (A) and UK (A --> B) and yet ˜ UK (B)

This description is in fact the model for Dialogue I above - the *Hypothesis* is explained simply because the user may not explicitly believe the desired conclusion. (See (Levesque 1984) for further discussion of the notions of implicit and explicit belief). In Dialogue III, the problem is that, additionally, part of what the user is recorded as believing, to use in drawing the conclusion, has been incorrectly assessed as a current belief of the user. We will not include here a discussion of the possible reasons for the incorrect assessment, and which reason is most probable. (For example: Is it most likely that psychologists disagree about diagnostic procedure vs. test interpretation?). Instead, we emphasize the need to track the user to draw a more accurate picture of possible differences to the system's "default model" of shared knowledge.

One possibility is to draw from previous dialogue other misunderstandings that have been recorded, to realize a generalization of differences with this particular user. One example of this situation is a psychologist who still believes that certain tests now considered outdated are effective measures of student's problems. Once an indication of the user's preference for a certain test is noted, this information can be noted as a difference to the system's knowledge about diagnostic procedure - i.e. SKAD ˜= SB (UKAD). (See Example III).

Another effective strategy is to track the user's goals and preferences, and to postulate a possible plan for the user. This plan may then be compared to an ideal system plan, to note discrepancies and misconceptions. An outline of this work is provided in (van Beek and Cohen 1986). The general idea of performing some pragmatic processing on the input is not new; see (Cohen 1985) for a review of some of the key efforts in natural language pragmatics. But the particular emphasis of van Beek's work (van Beek 1986) is to define a mechanism for recording user specific information, regardless of domain, and to suggest some methods for judging the best alternative available to a user, given his goals. The model is thus able to indicate a plan which also satisfies a user's preferences (see Example II). Thus, in conjunction with ongoing natural language research on plan-based approaches to inferring users' goals, it is possible to repair possible defaults that need to be retracted.

So far we have presented a basic strategy for answering a user's query. The general idea is that the response must contain the relevant aspects of the knowledge base which are not known to the user. In order to determine what knowledge is or is not shared by both the system and the user, the system employs initial stereotypes for the various user types, as well as updating this information through tracking the system-user dialogue throughout the course of the diagnosis. In this manner, the system can determine what information the user is unaware of and hence, must be supplied as part of the response. However, this may not suffice. As a final step in the algorithm, the system must check that the user's goals and preferences have been addressed; if not, the answer must be modified to do so. Ultimately, one may wish to modify the interaction between the phase of the algorithm which detects which components of the knowledge base represent unshared knowledge and the portion of the algorithm which addresses the user's preferences and goals. The interplay between these two components as yet remains unaddressed.

The previously presented examples (Examples I-III and Dialogues I-III) illustrate various aspects of this algorithm. In particular, recall the initial three examples. In the first example, the stereotypic models of parent and psychologist alert the system to what portions of the knowledge base must be supplied in the response (e.g. rule vs background or testing facts). In Example III, we see the need to track the user-system dialogue in order to update the initial stereotypic models, and in Example II, the need to address the user's preferences is illustrated. Before leaving this discussion of the system's general algorithm for answering queries as illustrated by the case of handling the psychologist's questions, we present two more examples to illustrate the application of this algorithm as reflected in the system's reasoning.

Dialogue IV:

System: Administer an individual intelligence test such as the WISCR.

Psychologist: Why is the WISCR being recommended?

Analysis of Problem: This question can be interpreted in several different ways:

(1) the user may be questioning the need for intellectual assessment, or

(2) the user would prefer to administer a different test, perhaps one with which he/she is more familiar, or

(3) the user is unaware that only certain intelligence tests are legitimate when making a funding or placement decision within the given province or state, or

(4) the user is questioning the need to adhere to the placement/funding legal guidelines ie. the guidelines to not apply in this case.

In answering this query, the system could explain why it is necessary to get an estimate of intellectual ability (1), or the fact that only certain intelligence tests (in this case the WISCR or Stanford-Binet) can be legitimately employed within the provincial funding formulae (3), or why the adherence to the funding formulae is appropriate in this case (4). Because the user is a psychologist, the system initially assumes that the user is familiar with the rules regarding diagnostic procedure etc, and hence these need not be supplied as part of the response. The information which is not shared by the user and system is some of the background data regarding the case at hand which has lead the system to believe that the child may qualify for high-cost funding. Therefore this information will be supplied in the response. However, before answering this query, the system needs also to determine whether the response adequately addresses any relevant user preferences. If the user prefers to administer a more familiar test, this response should help to counter the preference as well.

System: There is some weak evidence (such as a family history of learning problems, left-right confusion, slow language development) to indicate that the child may qualify for "high-cost" funding; therefore, the WISCR is being recommended because it is appropriate for use within the provincial funding formulae.

User: I would prefer to administer the McCarthy scales. Is this test appropriate?

System: No, only the WISCR and the Standford-Binet are appropriate for use within the provincial funding formulae.

System Reasoning: The system now realizes that, contrary to its initial stereotype, the user is unfamiliar with the specifics of the provincial funding formulae and hence, marks within its knowledge base that this type of information is unknown to this user. This updated model can then be successfully employed later in the diagnosis. For example, the system now wants the user to administer an individual achievement test. Again the system would like the user to administer a test that can be employed within the funding formulae. Because the system knows that the user is unfamiliar with the these formulae, additional information will be supplied by the system.

System: Administer an individual achievement test. Either the WRAT or PIAT should be used because these can be employed within the provincial funding formulae.

For our second example, recall Example III from Section 3.1 in which the system and user have drawn different conclusions based upon the results of the ITPA.

Dialogue V

System: Administer an individual intelligence test such as the WISC-R or Stanford-Binet.

User instead administers the ITPA (Illinois Test of Psycholinguistic Abilities) and the student's performance is recorded.

User: What are the conclusions based on the results of the ITPA? (ie. What is the diagnosis?)

System: The child's intellectual ability appears to be within the normal range; no further assessment of intellectual ability is required.

User: What other conclusions were drawn? For example, the student's poor scores on the Auditory Association and the Auditory Closure subtests indicate that the student's auditory channel is weak.

System: Although the test results raise some suspicions regarding the child's processing of auditory information, such conclusions cannot be drawn from the ITPA alone because the validity of the ITPA's subtests is questionable.

System Reasoning: After this interaction, the system can conclude that the user does not adhere to certain current philosophies. In particular, the user *does* believe in the ITPA Model and consequently has drawn conclusions which differ from those of the system. The system's model of the user is updated to indicate both the difference in philosophy (adherence to the ITPA model) and candidate hypotheses (the child's auditory processing skills are weak). Then later in the diagnosis this information can be exploited when addressing the user's questions. For example, in the response below, the system supplies additional information regarding the child's auditory processing skills.

Scenario: The child has just completed an informal test of social maturity in which brief stories of moral dilemma are read to the child and he is asked to describe a solution.

User: What is the diagnosis?

System: The child's better than average score indicates the following:

- good social judgement
- good comprehension skills
- average ability to express himself verbally
- no difficulty with material presented via the auditory channel.

The child's performance on this test indicates that there is no reason to advocate further testing of the auditory channel. The child is *not* experiencing any difficulty with processing auditory information.

Recall the three examples of Section 3.1 which motivated the research discussed herein. The framework we have proposed is in fact appropriate for handling these examples. Dialogue V demonstrates that our algorithm can be used to handle Example III, and Dialogue IIb illustrates Example I. Recall that Example II concerns user preferences. We have addressed the handling of user preferences (as illustrated in Dialogue IV), although we have left for future research the full integration of user preferences within our algorithm.

4.2.2. OTHER USER TYPES

So far we have concentrated on answering the psychologist's questions. Queries posed by other users are handled in the same manner, although the system's initial stereotypes will be different. We now turn to a study of the default assumptions that can be drawn for the user classes of parents and teachers. For a standard question as in Dialogue I ("What is the diagnosis?"), the system must now provide a more explicit response, indicating facts and rules used to reach the hypothesis. There is in fact a spectrum of possible indications of domain knowledge, based on the user's background.

Teacher:

domain knowledge: knows the purpose of diagnosis;
knows about some very common tests;
conclusion: selective explaining of domain knowledge (D: d and t);
if labelled as known to most teachers, won't be explained,
otherwise it will

student knowledge: probably unaware of actual test results;
aware of some classroom behaviour and performance,
some of these may be precisely the "hypotheses" suggested by some of the tests
conclusion: KANf (knowledge about test results) must be provided,
KANbc (background on student acquired through classroom) assumed known,
KANbt (background on student acquired through testing) unknown.

Parent:

domain knowledge : knows that there are tests but is unaware of what they yield and how to interpret them -- ie. it is alright to mention tests, without explaining why tests are useful, but any specific rule, linking test to fact must be assumed unknown
conclusion: all KAD assumed unknown;

student knowledge : unaware of actual test scores;
aware of personal background information -- e.g. non-academic strengths, related family problems, etc.
conclusion: assume all KANf (knowledge about test results) are unknown;
KANb (background facts) are known and in fact certain background facts may be input to the system as new information (ie. known to user but not to system)

Therefore, for a question of the type from Dialogue I ("What is the diagnosis?") and a user who is a parent for whom no other information is known, the system will respond as follows:

Response: *Hypothesis*, because Test Results and Test Rule + Diagnosis Procedure
e.g. because student's score on Test X is above average and this level of performance on this type of test indicates *Hypothesis*, due to Y.

Note that if the user later suggests a contrary fact about the student, the system should be flexible enough to make corrections.

For questions of the type in Dialogue III ("why questions"), reasons why a user such as a parent, teacher or principal may differ from the system are more obvious; the user and the system are not operating with the same assumed base of expertise. In particular, once a parent has asked for the diagnosis and then proceeds to challenge (e.g. asks for more information), since both Facts and Rules are by default spelled out (see above) one is immediately signalled to interpret the goal of the user in order to best respond. Once again, our proposal allows for focusing on the differences in belief between user and system. For example, if a parent asks why a diagnosis is appropriate and specifically does not ask for (i) explanation of terminology or (ii) explanation of a test name or purpose, then this should be a signal to an unwillingness to accept a fact about the student. In other words, if neither (i) nor (ii) above is the case, the system should assume a conflict of belief in KANb or KANf.

Roughly, we have outlined how each of the various sub-parts of the system's knowledge base, incorporating user models, can be used to produce the appropriate content and form of response for each user.

5. RELATED WORK

There has been a wide range of previous work on the general topic of user models, but none of this work focuses on the particular needs of an expert system for educational diagnosis. One domain for which user models have been studied is intelligent tutoring systems (e.g. (Sleeman and Brown 1982)). The suggestion is to build up a model of the student interacting with the system, to better determine the learning problems of the student.

As well, there has been a study of producing effective responses from expert systems, but little accomplished to deal with a range of possible users. Mycin (Shortliffe 1976) at first had a very limited explanation capability. (Wallis and Shortliffe 1982) tried to accommodate different users, by having the user supply his level of expertise. A limited amount of control on the level of detail in the prepared response then resulted. (Swartout 1983) examined how a more elaborate storage of knowledge (descriptive facts and principles) could produce better responses from medical expert systems.

There has also been work on constructing stereotypes for classes of users, to be employed in fine tuning the actual model for a user (e.g. (Rich 1979)).

Finally, there are the kind of user models constructed for natural language understanding systems which perform pragmatic processing. The work of (Perrault, Allen and Cohen 1978) considers the plan underlying the production of a discourse, and the concept of mutual belief, to help a listener to comprehend the likely aims of a speaker. (Joshi, Webber and Weischedel 1984) emphasize the need for a system to predict the false inferences that a user may be drawing, by carefully studying the goals underlying utterances.

In a sense, the user models we are advocating for our educational expert system draw from each of the above approaches. We essentially allow for the more robust kind of user modelling practiced in natural language pragmatics (tracking the beliefs of system and user, and including goals in the model), but employ default assumptions on users as a handy starting point and also accommodate correcting misconceptions.

There are a growing number of researchers studying the topic of "user modelling" (e.g. the recent Invitational International workshop on user modelling, Maria Laach, Germany). Our work is quite distinct in its scope and approach, but relates to some of these other current efforts. For instance, the work of (McCoy 1986) addresses the different perspectives on objects in a knowledge base which a user may have. Our framework suggests different perspectives for whole classes of users (psychologists vs. teachers, for example), but not with respect to attributes of an object, but instead with respect to whole segments of the knowledge base which are more likely to be understood to a user.

The research of (Finin and Drager 1986) is also relevant, since this work includes the technique of assigning default values to the knowledge of users, and updating these models as discrepancies arise. We make the case for default models stronger for our particular domain where we have classes of users that can be identified, rather than the more general problem faced by Finin of determining the right classification for a user. As well, Finin offers more specific strategies on correcting the possible assumptions on a user.

There are also common concerns with the work of (Chin 1986) and (Paris 1985). Chin suggests a system of double stereotyping, assigning a user a level of expertise and then also labelling all the knowledge of the system according to a "difficulty level" (e.g. complex, simple, etc.). This calls for a segmentation of knowledge according to complexity. We instead advocate a segmentation into knowledge bases which are not labelled with one level for all users, but are determined as known or unknown to each user type. We thus avoid the problem of globally determining the classification of each individual rule.

Paris addresses a different dimension of accommodating different classes of users - varying the form of the response, rather than the content. So, expert users get a "process" description, while naive ones have "parts-oriented" explanations. We instead examine required content changes, and leave aside some of the lower level generation issues. Note that for both Chin and Paris, the user types are on an axis of expert to novice. For our domain, we work with a "user class" identification which is both appropriate for our given task and helpful in isolating the required user specific content in responses.

Some of the differences in our approach to user modelling, outlined above, allows us to advance a general framework for user modelling, summarized in the section below.

6. A GENERAL PROPOSAL FOR USER MODELLING IN EXPERT SYSTEMS

This paper makes some recommendations for the next version of our expert system, but the principles advocated are intended to be useful as general guidelines for allowing expert systems to vary their responses to users. We summarize the general procedure recommended:
1) There should be defaults for the user types, to initiate a user model
2) The system should make use of some partitioning of its knowledge bases, to focus on possible differences between user and system knowledge
3) Differences should be accounted for in responses; so differences may be more usefully recognized by tracking the user's goals and ideally the previous dialogue.

In connection with the third point above, it is possible to specify more clearly the potential avenues for disagreement between the system and a particular user class. For example, for the user as psychologist, we have already indicated in the Example III of section 3.1 that the user may well disagree on the applicability of certain tests in the test database. One area for future research is a deeper specification of these dimensions of difference.

7. CONCLUSIONS

Expert systems for educational diagnosis require user models, to provide variability in the form and content of explanations generated for the users. This paper describes a framework for incorporating user models into expert systems, which is specifically designed to

(1) identify and separate the knowledge bases which must be consulted in the preparation of an explanation
(2) allow for default assumptions of user's background knowledge in the absence of further information on the user, and
(3) continuously maintain and refine the user models, both as dialogue from the user reveals knowledge and as previous output from the system adds to the shared knowledge between system and user.

In short, there is a prescription for the design of expert systems to incorporate models in order to serve a broad base of users. Moreover, there is a detailed argument for why this framework is particularly effective in the area of diagnosing learning disabilities. In sum, we have the means for implementing a more powerful version of our current system and a general framework of use to designers of other systems, to focus on beliefs of users in producing quality responses.

REFERENCES

Adelman, H.S. (1982), "Identifying learning problems at an early age: A critical appraisal", *Journal of Clinical Child Psychology, 11,* 255-261.

Chin, D.N. (1986), "User modeling in UC, the UNIX consultant", *Proceedings of the CHI-86 Conference*, Boston, MA.

Cohen, R. (1985), "The Need for Pragmatics in Natural Language Understanding", *Proceedings of the CSCSI-sponsored Theoretical Issues in Natural Language Understanding Workshop,* Dalhousie University, Halifax, Nova Scotia.

Colbourn (Jones), M.J. (1982), *Computer-guided Diagnosis of Learning Disabilities: A Prototype.* M.Ed. Thesis, Dept. for the Education of Exceptional Children, University of Saskatchewan, 348p; also Report 82-8, Dept. of Computational Science, University of Saskatchewan, 1982; also ERIC Report No. ED 222-032.

Colbourn (Jones), M.J. and McLeod, J. (1983), "The Potential and Feasibility of an Expert System for Educational Diagnosis", *Proceedings of International Federation of Information Processing (IFIP)*, 891-896.

Colbourn (Jones), M.J. and McLeod, J. (1984), "Computer-guided Educational Diagnosis: A Prototype Expert System", *Journal of Special Educational Technology, VI, 1*, 30-39.

Committee on Teacher Education and Professional Standards (1971), *Standards for Educators of Exceptional Children in Canada*, The National Institute for Mental Retardation, Downsview.

Finin, T. and Drager, D. (1986), "GUMS: A General User Modeling System", *Proceedings of CSCSI Conference*, 24-30.

Grice, H.P. (1975), "Logic and Conversation", in Cole, P. and Morgan, J.L. (eds), *Syntax and Semantics: Speech Acts, 3*, Academic Press, New York, 41-58.

Jones,M., McLeod,J., Robertson, G., Tubman, J. and Wasson, B (1986), "The CGD Project: An Educational Diagnostic System Based on THEORIST", submitted for publication.

Jones, M. and Tubman, J. (1986), "Expert Systems for Educational Diagnosis: An Experiment with THEORIST", *Proceedings of the Sixth International Workshop on Expert Systems and Their Applications*, Avignon, France.

Joshi, A., Webber, B. and R. Weischedel (1984), "Living up to Expectations: Computing Expert Responses", *Proceedings of the National Conference on Artificial Intelligence (AAAI 84)*, 169-175.

Levesque, H. (1984), "A Logic of Implicit and Explicit Belief", *Proceedings of the National Conference on Artificial Intelligence (AAAI-84)*, American Association for Artificial Intelligence, Austin, Texas, 198-202.

McCoy, K. F. (1986), "The ROMPER System: Responding to Object-Related Misconceptions using Perspective", *Proceedings of 24th Annual Meeting of the Association for Computational Linguistics (ACL 86)*, 97-105.

McLeod, J. (1982), "Institute of Child Guidance and Development at the University of Saskatchewan", *Journal of Learning Disabilities, 15*, 290-293.

McLeod, J. and Jones, M. (1985), "The Development of an Expert System for Computer-Guided Diagnosis of Children's Learning Problems: Some Emerging Problems", *Proceedings of International Conference on the Computer as an Aid for Those with Special Needs*, Sheffield, England, 68-78.

McLeod, J., Jones,M., Robertson, G., and Wasson, B. (1986), "Computer-Guided Diagnosis: A Progress Report on a Canadian Project", *Fifth Canadian Symposium on Instructional Technology*, (to appear), Ottawa, Canada, May 1986

Paris, C. L. (1985), "Description Strategies for Naive and Expert Users", *Proceedings of 23rd Annual Meeting of the Association for Computational Linguistics (ACL 85)*, 238-245.

Perrault, C.R., Allen, J. and P. Cohen (1978), "Speech Acts as a Basis for Understanding Dialogue Coherence", *Proceedings of Theoretical Issue in Natural Language Processing-2 (TINLAP-2)*, 125-132.

Poole, D.L., Aleliunas, R., and Goebel, R. (1986), "Theorist: a Logical Reasoning System for Default and Diagnosis", Dept. of Computer Science Technical Report CS-86-06, University of Waterloo.

Rich, E. (1979), "Building and Exploiting User Models", Technical Report CMU-CS-79-119, Dept. of Computer Science, Carnegie-Mellon University.

Shortliffe, E.H. (1976), *Computer-Based Medical Consultations: MYCIN*, Elsevier, New York.

Sleeman, D., & Brown, J.S. (eds) (1982), *Intelligent Tutoring Systems*, Academic Press.

Swartout, W.R. (1983), "XPLAIN: A System for Creating and Explaining Expert Consulting Systems", *Artificial Intelligence Journal, 21(3),*, 285-325.

Tubman, J. (1986), *An Expert System for Educational Diagnosis Using THEORIST*, M. Math thesis, Dept. of Computer Science, University of Waterloo.

van Beek, P. (1986), *A Model for User Specific Explanations from Expert Systems*, M. Math thesis, Dept. of Computer Science, University of Waterloo.

van Beek, P. and Cohen, R. (1986), "Towards User Specific Explanations from Expert Systems", *Proceedings of the CSCSI Conference*, 194-198.

Wallis, J.W. and Shortliffe, E.H. (1982), "Explanatory Power for Medical Expert Systems: Studies in the Representation of Causal Relationships for Clinical Consultations", Technical Report STAN-CS-82-923, Dept. of Computer Science, Stanford University.