Chapter 3

RESEARCH DESIGN AND METHODOLOGY

3.1 INTRODUCTION

The central purpose of this research is to ascertain the feasibility of developing an effective computer decision support system to assist in the administration of schools. An interrelated purpose of this research is to ascertain whether such a system can be developed by the school personnel, rather than for the school by external personnel. The literature reviewed in Chapter Two indicated that successful expert systems had been developed for other domains, but no reference was found apropos expert systems for school administration. Potential benefits may be implied for schools, but have not been demonstrated. Thus, in determining a research methodology for this study, it was decided to develop an expert system in a school administration domain that may benefit from such a decision support system, and to use computer software which would enable this development to be undertaken by the school personnel.

Following preliminary readings in artificial intelligence, the researcher initially envisaged that a decision support system might be applicable in a domain commonly associated with educational administration, such as staff utilisation, financial management or curriculum planning. However, the task of recommending a student's subject selection was chosen for the research program because this domain was a real and current problem. To some people this domain may appear to fall outside the area of educational administration. But, if one accepts students as the prime reason for a school's existence, the provision of effective counselling is more significant than many bureaucratic and often mechanical tasks presumed of school administrators. In light of the apparent dearth of models and expert systems in this domain, it was decided to undertake the research in one school and develop an expert system appropriate for that school.

This study has drawn on the analysis of documentation, observation and interviews. As the researcher was principal of the trial school, an element of personal experience undoubtedly intruded in the research process. In particular, the researcher's position may well have positively influenced the level of co-operation and the perceived benefits of this research to the school. Recommendations from Chapter Two and Bucknall (1988), pertaining to the behavioural issues of introducing technological change, were applied in the current project to help ensure a positive mindset.

The interviewee's comments, their application of documented procedures, and their response to case-studies was used to prepare and test the reliability and validity of the data. Evaluation was also facilitated by the length of research period, combined staff input, plus verification and validation of the prototype expert system. The study recognises that no attempt was made to test the applicability of the findings from one domain in one school to other domains and/or other schools.

Although addressing a specific goal, the research was primarily what Schumacher (1979, 8) described as "one of discovery rather than verification". The researcher is, however, confident that the research findings extend the work undertaken in other domains and support the proposition that expert systems can be developed by schools as decision support systems.

3.2 SYNOPSIS OF THE SUBJECT SELECTION PROCESS

Dripstone High School opened in 1979 for an anticipated student population from year eight to year twelve. Until 1988 the students were divided into year levels. The Mathematics and English Faculties organised their classes by student achievement while the Science and Humanities Faculties endeavoured to cater for students in non-streamed classes. The other teaching faculties provided a mixture of compulsory subjects with non-streamed classes and optional subjects related to student achievement. In 1986 the Northern Territory Government introduced Secondary Colleges for students in years eleven and twelve, and changed Dripstone to a Junior High School catering for students in years eight to ten. As part of the school's planning to take full advantage of the changes that had been imposed on it, Dripstone High School staff and parents spent eighteen months planning for the introduction of a unitised curriculum and vertical timetable, which commenced at the start of 1988.

The school has four ten-week terms a year. The curriculum is divided into ten-week units of study, with each unit receiving equal time allocation during the term (five fifty-minute lessons per week and thus notionally forty-one hours per term). Students enrol in six units per term, and thus seventy-two units over the three years. To receive the Northern Territory Junior Secondary Studies Certificate (JSSC) it is necessary for students to satisfy minimum time requirements in a range of subject areas, as per Figure 3.2.1; this should account for sixty-eight of the seventy-two units normally studied. Some schools are unable to provide the minimum requirements for all subject areas; but their students remain eligible to receive a JSSC, notated accordingly. Dripstone High School does not offer all the subjects to the extent determined by the Northern Territory Board Of Studies, especially in Languages Other Than English (LOTE), though it is working to address this shortfall, and thus students at this school currently need to only complete fifty-six units to satisfy the certificate requirements; and they use the remaining sixteen units to extend their studies. Selection procedures need to accommodate the minimum requirements for JSSC certification, the student's achievement in pre-requisite units, individual interest in particular subject areas, and on the need to prepare for further study and career aspirations.

Figure 3.2.1
Northern Territory Board Of Studies Minimum Time Allocations Years 8-10
(Northern Territory Board Of Studies Handbook 1993, 6.29)

Subject	Hours over 3 years	Units over 3 years
English	400	10
Mathematics	400	10
Social Education	360	9
Science	360	9
LOTE	280	7
Health & Phys. Ed.	240	6
Visual & Performing Arts	240	6
Technical Studies	160	4
Home Economics	160	4
Business Education	80	2
Career Education	40	1
	2720	68

The school's vertical timetable does not group students by year level but does offer a range of units at different levels which students may select if they satisfy stated pre-requisites. If students are unsuccessful in completing a unit, they may study the unit again or complete one or more related units to accumulate the pre-requisites for a more advanced unit in that subject area. Students may concurrently study more than one unit in the same faculty, perhaps as part of the remediation process, perhaps to meet minimum requirements quicker, or perhaps to undertake optional units in the subject.

Informal feed-back from students, parents and staff at Dripstone High School has indicated the general success of the unitised curriculum offered through the vertical timetable. Although the majority of students have undertaken a relatively traditional mix of subjects, a significant number of students have used the opportunity to pursue unit selections that would not have been possible with a traditional timetable. The main criticisms of the vertical timetable are related to its administrative workload including the difficulties faced by students, parents and teachers in keeping track of subject pathways and their consequences. These informal and anecdotal criticisms at Dripstone High School were confirmed in a study by Fowler (1993a) which compared several Northern Territory secondary schools with vertical and horizontal timetables and noted that the administrative workload was a "major factor in those schools which have or are considering modifications to [the] vertical timetable" (146).

The current process of advising and making selections was designed to enable students, parents and teachers to be thoroughly involved. The process has evolved over several years from an initial manual system to the current system, which is still labour intensive but with some computer assistance. Each student is provided mid-term with a statement, from the school's computer records, listing by faculty their previous units and results plus their current units. The statement also lists all the units which will be available the following term for which the student has the pre-requisites based on previous results and presuming satisfactory results in current units. The list of available units does reduce the student's options by excluding units the student cannot enrol in the following term, but the list still includes units which may or would not be appropriate. For example, a year ten student achieving outstanding English results would still have year eight remedial English units listed. Subject teachers are expected to counsel their current students on appropriate units in that subject for the following term. Students may also seek assistance from other subject teachers, especially in subjects they are not currently studying. Each student is provided with a copy of the *Dripstone High School Course Outlines* that provides details of certification requirements, subject flow charts and a description of all the units that may be available. Students have two weeks to ascertain their preferred units for the following term and submit a form listing six preferred units plus six alternatives. These forms have to be signed by parents and the Homegroup teacher. Parents are encouraged and Homegroup teachers are expected to monitor the selections. The Assistant Principal Curriculum and three Curriculum Co-ordinators then examine all the forms to ensure they comply with the various faculty and certification guidelines. The verified requests are electronically scanned and a computerised matrix is generated to match student preferences and timetabling resources. Inevitably a number of students have a combination of preferred units which cannot be accommodated by the resources available and it is necessary to vary some resources and/or student preferences. The four staff involved in examining the forms and developing the matrix spend approximately a week, including the weekend, working exclusively on this task.

Significant further time is required by this team dealing with miscellaneous complications such as students absent during the process and students requesting changes to their allocated classes.

Despite the various handbooks, counselling steps and check-lists, it is generally recognised that the process of student subject selection needs to be improved. One particular problem is in the area of professional advice regarding appropriate unit selection. Conflicting advice may be provided within a faculty and this advice is further complicated by the need for teachers to ensure that students satisfy the requirements for their Junior Secondary Studies Certificate.

3.3 TIMING

The timing of the research project was a combination of two factors: (1) the school was experiencing difficulty in a specific domain, and (2) the school principal had commenced informal research into expert systems. A *prima facie* relationship between a real problem and a possible solution was apparent.

Two separate but closely related projects were undertaken, both with the aim of providing recommendations for student unit selection. A school based project was undertaken to develop paper models and a research project was undertaken to develop computer models. The school project was designed to satisfy the school's immediate needs and available resources. The research project, the subject of this thesis, was designed to enhance and extend the school project. As school principal, the researcher was actively involved in the school project. The researcher, as school principal, was able to bring pertinent knowledge to the research project. The relationship is summarised in figure 3.3.1. This thesis reports on the research project but it is appropriate that reference be made to the school project as it provided the initial models from which the researcher's models were developed. Reference to the school project is also critical in evaluating the research project, given the hypothesis that an expert system should be able at least to replicate the outcomes of a manual system.

Figure 3.3.1
Research development overview

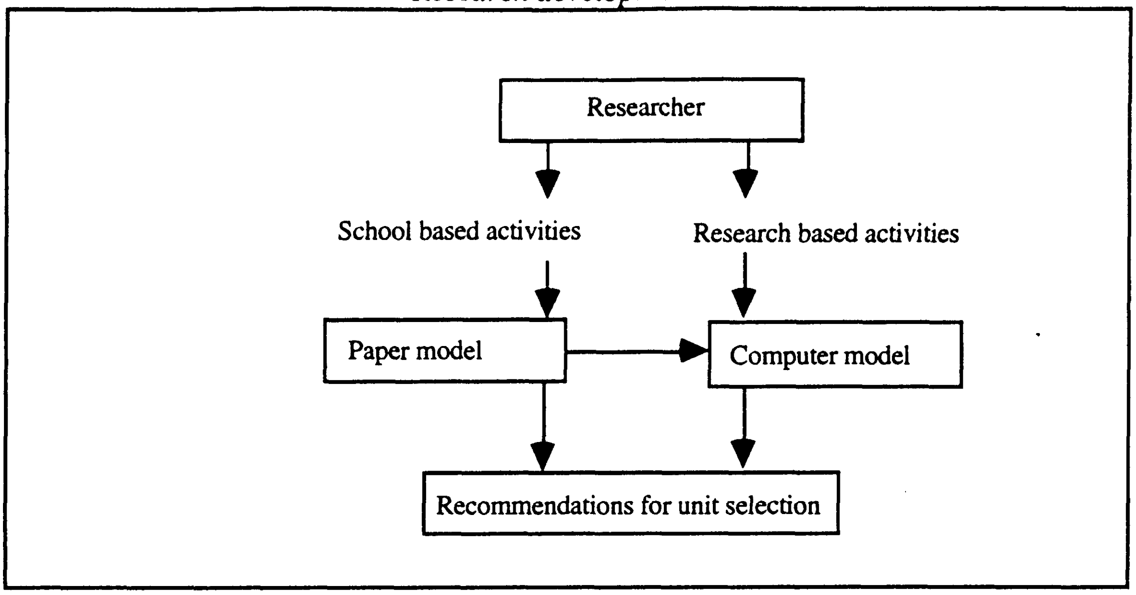

The school project was primarily undertaken between 1988 and 1991, with on-going fine tuning. The research project commenced in 1988 in conjunction with the school project but assumed a clearly separate identity in 1991. Following a pilot study in 1992, the research project was undertaken full-time during 1993 to enable preparation and evaluation of an expert system (RUS) to recommend unit selection.

3.4 RESEARCH DESIGN

The primary aim of this research project is to ascertain whether expert systems can be developed to assist in the administration of schools. Reports in the literature suggest they should and thus a central research problem was to demonstrate that they could. Resolution of this problem involved several elements: (1) modelling the research domain, (2) developing and implementing an expert system, and (3) evaluating the expert system to validate its recommendations and compare its performance with the current system and human experts.

3.4.1 MODELLING THE RESEARCH DOMAIN

Designing this element of the research involved the two main tasks identified by Gordon (1978, 6): (1) establishing a model structure and (2) ascertaining the data for these models. The school was using a variety of descriptive, prescriptive and normative models. The Department of Education and neighbouring schools made use of similar modelling devices including flow charts, reference charts and goal statements. Use of the school's models for the research project resulted in some minor changes to adopt a common style or where factual errors were detected. It was considered prudent to use current tools rather than introduce new modelling techniques because: (1) people were familiar with the modelling tools, (2) it would reduce distractions from the task of eliciting data and presenting it in a uniform manner, and (3) it would minimise the workload of the many people involved in the project. These school based activities and their association with the research project are illustrated in figure 3.4.1.

The overall domain of student subject recommendations was divided into sub domains on the basis of subject faculty. An expert system was developed for each faculty. The individual faculty expert systems were then concatenated in the RUS system for the overall domain. To assist the construction of the faculty expert systems, a variety of models were used to represent the different faculty requirements; these models are collectively described in this thesis as construction models.

Figure 3.4.1
Research development model

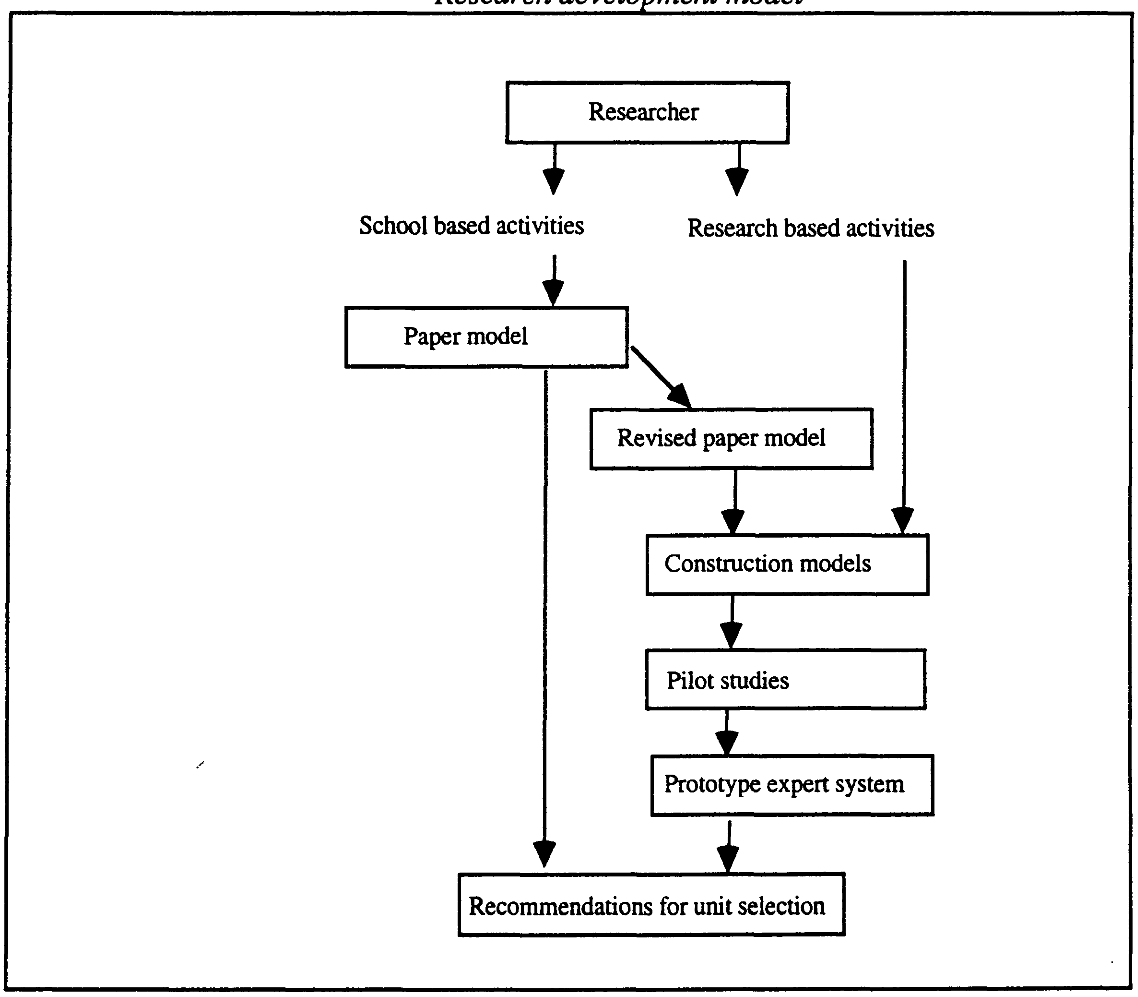

The construction models examined (1) the local implications of the official requirements by the Northern Territory Board of Studies, (2) the curriculum structure at the school, and (3) the application of the various formal and informal models by the teaching staff. They were used to prepare the data needed to construct the faculty domain expert systems. It was anticipated that the official requirements and curriculum structures would be relatively easy, given the level of documentation available, but that translating the mental models would be a more complex task requiring considerable individual and group discussions and negotiations. Figures 3.4.2 and 3.4.3 illustrate the models prepared for the Computing units. Similar models for the other subject areas are contained in Appendix One.

Figure 3.4.2
Computing flow chart model

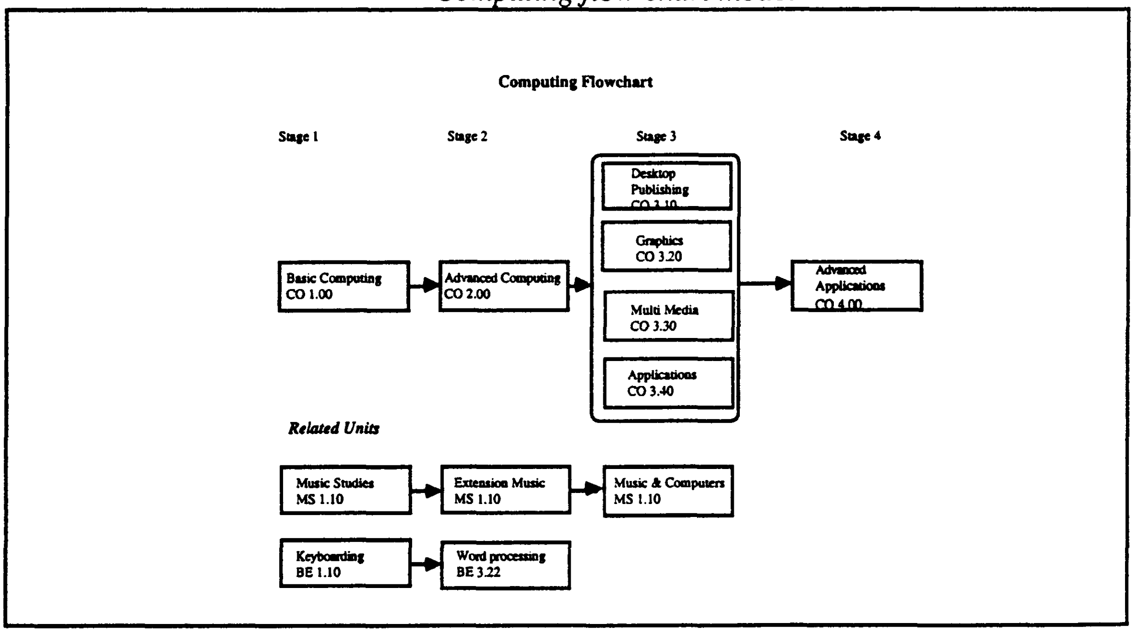

Figure 3.4.3
Computing summative model

Business Education - COMPUTING

Subject background

Seven units are available in this subject area, though both Typing and Music offer units which are specifically computer based. The relationship between these units is illustrated in figure 3.4.2. There are no minimum requirements or compulsory units in computing. Many of the other subjects presume keyboard skills for assignments. Some of the other subjects use computers for components of their units, and teach the relevant skills if necessary.

Recommendation antecedents

1. Students are to complete the introductory Typing unit (or demonstrate competency) before proceeding with computing units.
2. Extension units are not recommended unless students express the desire to undertake further art studies.
3. The strength <cf> of recommendations for extension units depends on:
   stated interest in Computing
   keen<1.00>, enjoy<0.75>, OK<0.30>,
   prefer not<0.00>, definitely not<-0.50>
   grades achieved in prerequisite units
   A<1.00>, B<0.85>, C<0.70>, D<0.55>, W<0.00>, E<-0.50>
   unit requested <1.00>, unit alternative <0.75>.
4. The following timing recommendations are applied to these units:
   first to third "Extension"
   fourth to seventh "Specialisation".

A concurrent task was to develop other construction models which had no purpose other than to assist the construction and testing of further models. For example, simulation models were developed for some faculties to list every feasible subject combination to facilitate testing for possible errors in the school's formal models and/or the expert system. Figure 3.4.4 is an extract of one of the simulation models (the complete model over four pages is contained in Appendix One) used to construct a table of correct Drama option combinations. The extract tracks one of the potential unit patterns after a student has completed DR110. This student then has the choice of one of the stage two units DR210, DR220, DR230 or the alternative stage one unit DR140. If the student then achieves an A or B grade in DR210 the choice is to progress to the stage three unit DR310, enrol in one of the remaining stage two units or the other stage one unit. There are 756 possible combinations of the ten Drama units of which 176 unit combinations comply with the various guidelines.

Figure 3.4.4
Extract from Drama options simulation model

110	C - 210 C - 220 C - 230 C -140				
110	210 (A-B)	C - 310 C - 220 C - 230 C - 140			
110	210	310 (A-B)	C - 320 C - 330 C - 220 C - 230 C - 140		
110	210	310	320 (A-B)	C - 330 C - 220 C - 230 C - 140	
110	210	310	320	330	**140**

3.4.2 DEVELOPING AND IMPLEMENTING AN EXPERT SYSTEM.

The literature reviewed in Chapter Two contained conflicting views in regard to who should develop expert systems, that is, who should be the knowledge engineers. The researcher accepted the contention that an important reason why schools should get directly involved in knowledge engineering is that, when done properly, the exercise examines the entire work flow, habits and procedures of the organisation in general and for a given issue in particular. Further, the researcher supported the view that, where possible, in-house developments are generally preferable to using external consultants.

The literature reviewed also highlighted conflicting views with regard to the source of the expert system; that is, the use of general systems to which local data is appended through to the development of individually prepared software.

A pilot study, which is described in Chapter Four, was used to select computing hardware and software for the expert system. This preliminary study also provided an opportunity to ascertain whether the researcher, as knowledge engineer, would be able to elicit, document and use the known facts and various expert's heuristics.

In determining a research technique for this study, it was decided to use the expert system development cycle outlined in figure 3.4.5, which is described in more detail in Chapter Five.

Figure 3.4.5
Research paradigm

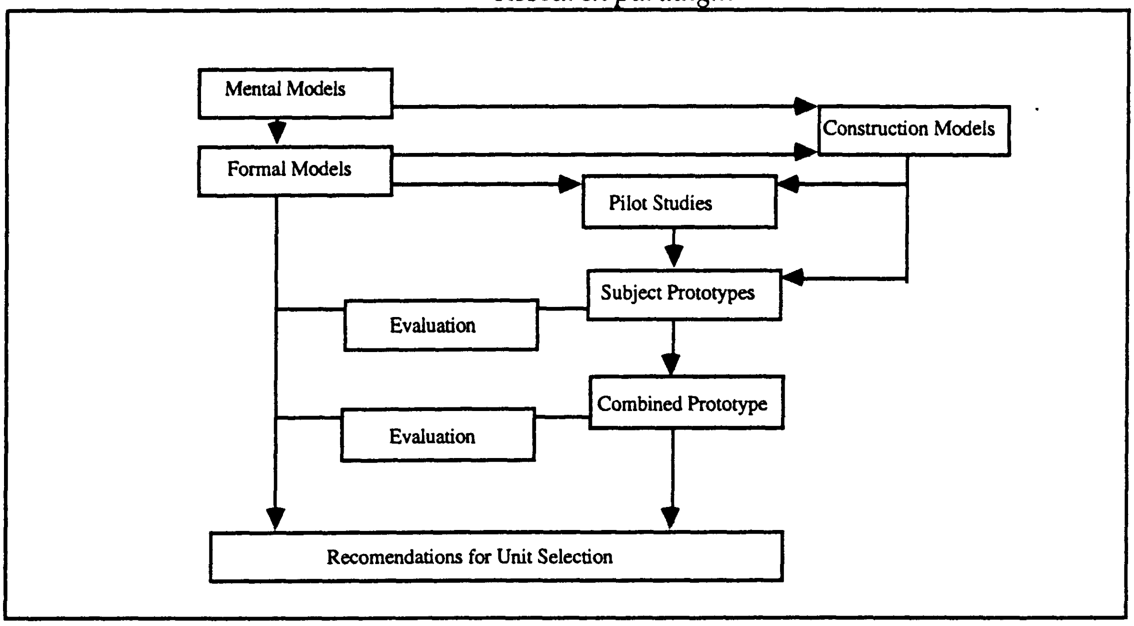

3.4.3 EVALUATING THE EXPERT SYSTEM.

Since expert systems are designed to provide recommendations, it is critical that the advice provided can be trusted. Thus the research design needed not only to verify that the RUS system operated correctly and produced valid output, but also to evaluate the system's performance against that of the current system utilising human experts. For these reasons the research design included provision for evaluating the output of both the manual and expert systems from the same input data as well as surveying users on their reactions after using the prototype.

Teaching staff had already participated in the school project for the preparation and publication of the faculty flow charts and descriptors. These publications were used as the base for the research project construction models, which were then examined by the subject co-ordinators for accuracy. During the development of each faculty expert system, archetype student records were constructed and used to test all perceived situations. Where necessary, the co-ordinators also provided clarification and/or additional information when testing indicated insufficient or conflicting data in the publications. Some subject teachers reviewed the system procedures and provided further input of personal heuristics. In situations where conflicting opinions were expressed, a consensus was established rather than including multiple heuristics in the knowledge base.

Because the RUS system was developed on a stand-alone computer, rather than within the school's computer network, a copy of all the current students' academic records was made to the stand-alone computer. The RUS system was formally validated using these student records. The following outcomes were examined: (1) the minimum number of units required in each subject area and the number undertaken by the student, (2) the subject areas in which the students were enrolled the following term and which subject enrolments would be recommended as priorities by the RUS system, and (3) the units for which each student was enrolled for the following term and whether these units were included in the recommendations made by the RUS system.

Throughout the development of the RUS system, a variety of potential student scenarios were discussed with the school's Assistant Principal (Curriculum) and subject co-ordinators to clarify situations and ascertain reliability of recommendations by the human experts. The RUS system was also formally tested by teachers, students and parents for acceptable performance. The evaluation was non-binary; that is, it did not seek to simply determine if the outcome was true or false. The RUS system was not designed to be absolutely correct; there were too many subjective variables (for teachers, students and parents) that were not included in the system for it to aspire to binary status. The RUS system was evaluated to ensure that it did narrow the options to all the recommendations that (1) satisfied the statutory requirements and (2) the personal preferences, and (3) that it did not provide incorrect recommendations. In addition to the evaluations of the expert system's performance, a questionnaire using Likert type scales completed by end users of the RUS system contained both validation and verification components. The questionnaire included an 'open ended' section to provide these end-users the opportunity to extend their responses, to comment about RUS or to provide other reactions to the project.

3.5 CONCLUSION

A research methodology was adopted to enable school personnel to construct an expert system for use at their school to assist the selection of students' curriculum units. The main elements of this methodology were (1) the preparation of paper and computer formal models from the various mental models used by the personnel concerned, (2) a preliminary study to assist the selection of appropriate computer hardware and software followed by the development and implementation an expert system, and (3) evaluation of this expert system including comparison of its performance against the human experts.

Chapter 4

PILOT STUDY

4.1 INTRODUCTION

Despite the extensive range of literature reviewed in Chapter Two and since no reports of a similar project had been found, it was unclear what would be the most appropriate method to construct the proposed expert system. It was decided to undertake a preliminary study in one faculty aimed at (1) selecting appropriate computer hardware and software and (2) interpreting the school's paper models to ascertain their suitability for subsequent knowledge engineering. If successful, the preliminary study would then be extended to full implementation in a working system; if unsuccessful, alternatives would be examined.

It was decided initially to endeavour using a shell for the construction of the expert system. Most proponents of in-house development argue in favour of commercially produced shells around which are established the local knowledge and rules; and that shells developed from similar domains are better than general shells. For this research, no similar domains were apparent and therefore it was necessary to examine whether a general shell would be suitable for the project.

One strand of the school's mathematics units was selected for the pilot study because that faculty provided a good range of simple and complex problems. The pilot study ultimately involved the following: (1) preparation and evaluation of a shell-based expert system, (2) preparation and evaluation of a toolkit-based expert system, and (3) an examination of the paper models developed by the school.

4.2 TRIAL IMPLEMENTATION WITH A SHELL

The selection of a specific shell for the pilot study followed an examination, in 1989, of product reviews in two Australian and American computer journals. No expert systems appeared to have been developed to support school administration or subject selection and thus there was no simple basis for evaluating which shell to use. The journal reviews examined many simple shells priced to $250, a limited range of more complex shells priced to $1500, and several shells priced up to $9000. MacSMARTS™ was selected, at a cost of $1200, because it was favourably rated in the journals against the other software. The pilot study was undertaken using a Macintosh SE/30.

The MacSMARTS Professional manual describes this product as "a powerful tool for creating expert systems and knowledge systems and using them to diagnose, advise, design, plan or troubleshoot ... [with] its easy to create powerful knowledge systems that will revolutionise the way you and/or your company do business" (Thomas 1989, 1). This shell could use both rule-based and example-based knowledge representation. In addition, 'hyper text' provided a facility to link rules and advice to graphics, text, spreadsheets, databases, other knowledge bases and Hypercard stacks.

The rule-based knowledge base is founded on a Logic Worksheet containing columns for Facts, Rules and Advice (figure 4.2.1). The manual advises use "is best when you know the exact logic to be followed and can express it in rules based either on true or false conditions ... or limits on named variables or functions" (Thomas 1989, 3-1). This is done by first preparing a flow chart with branches to direct the end user with minimal effort and avoiding unnecessary user input. The questions/links/advice in the flow chart are then entered on the Logic Worksheet by completing the Facts, Rules and Advice columns (figures 4.2.2, 4.2.3, and 4.2.4) to prepare a Logic Worksheet as per figure 4.2.5.

Figure 4.2.1
Example of a MacSMARTS Logic Worksheet

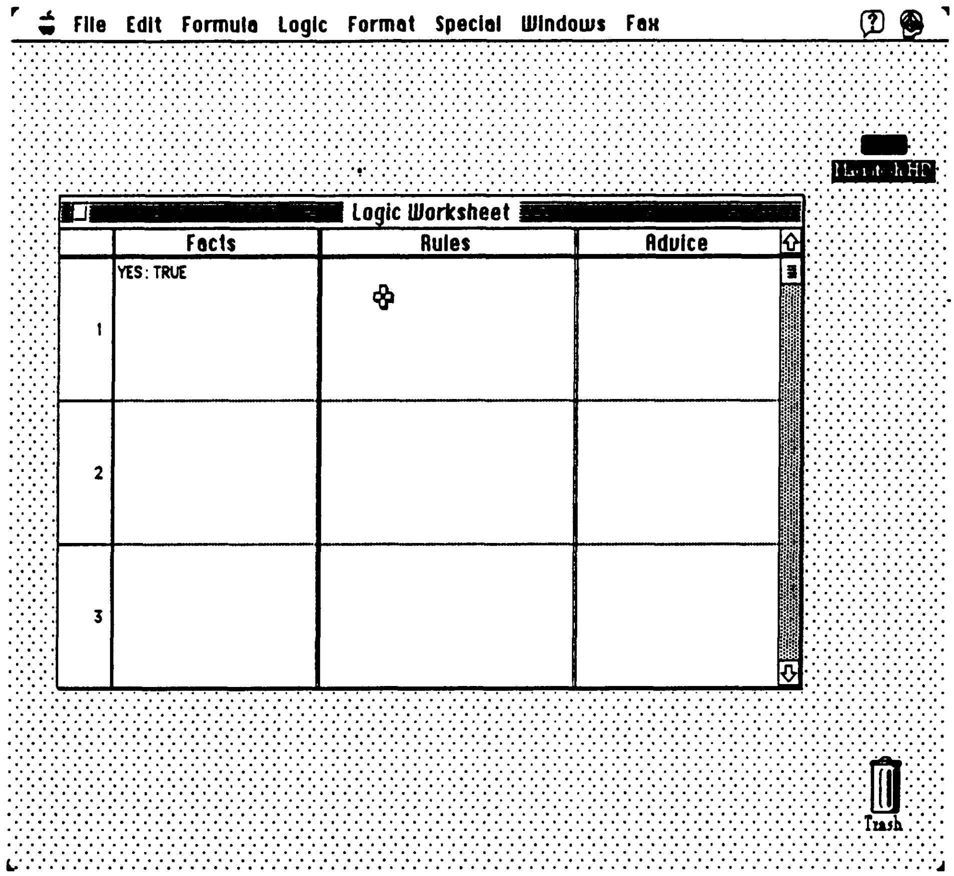

Figure 4.2.2
MacSMARTS Logic Worksheet Facts entry

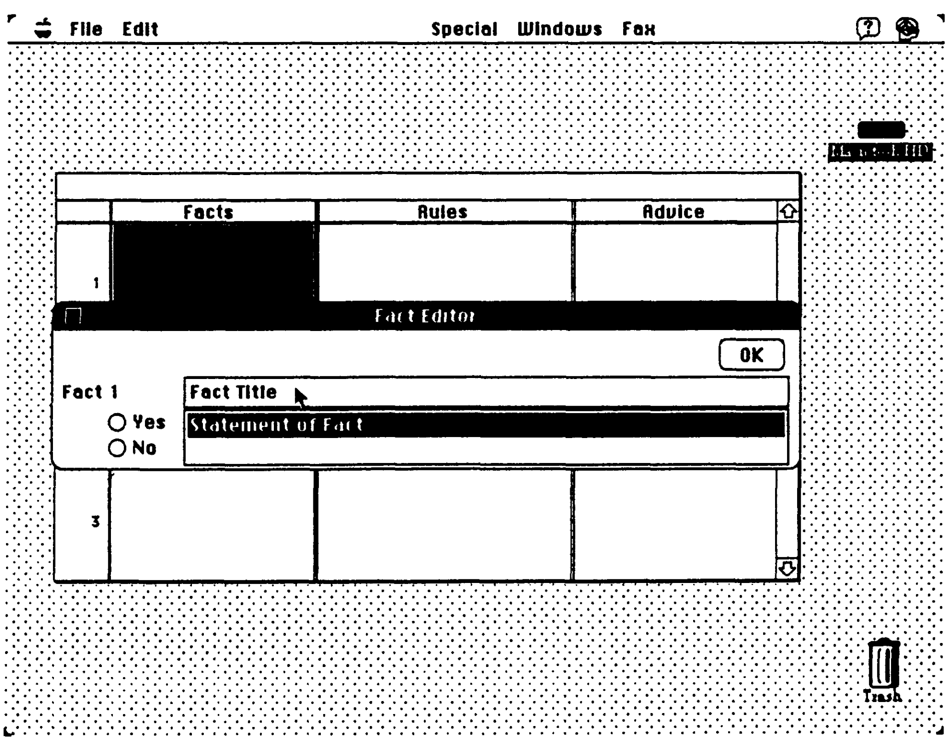

Figure 4.2.3
MacSMARTS Logic Worksheet Rules entry

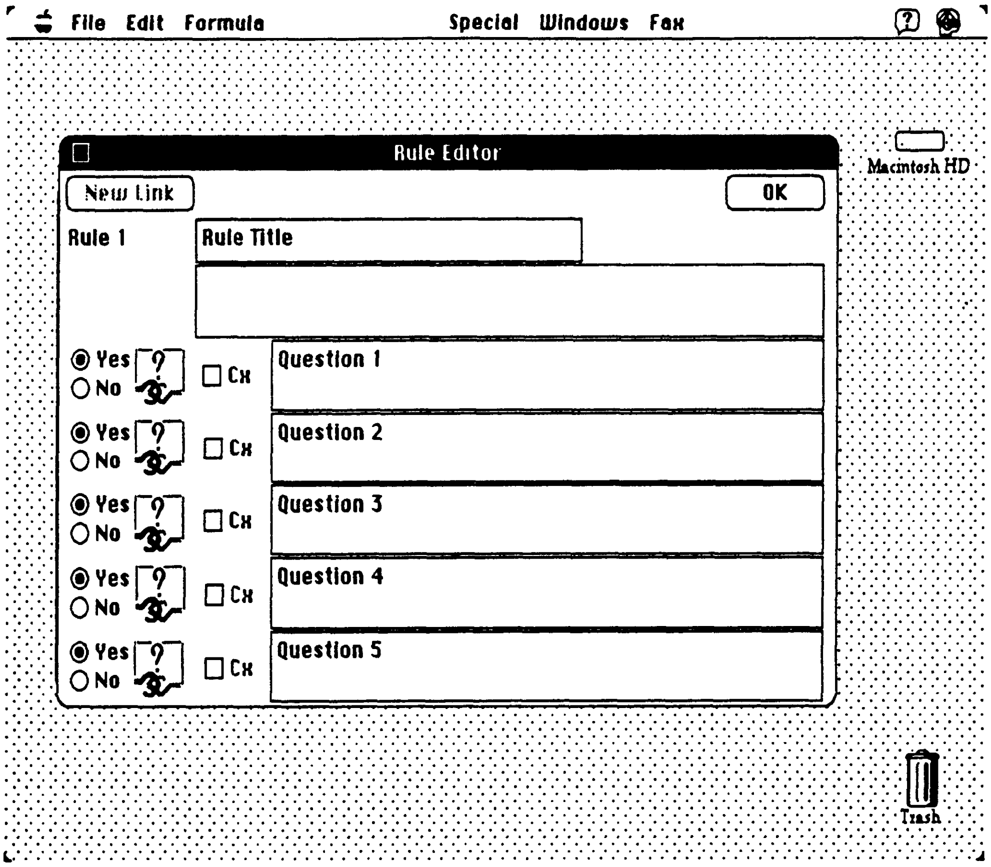

Figure 4.2.4
MacSMARTS Logic Worksheet Advice entry

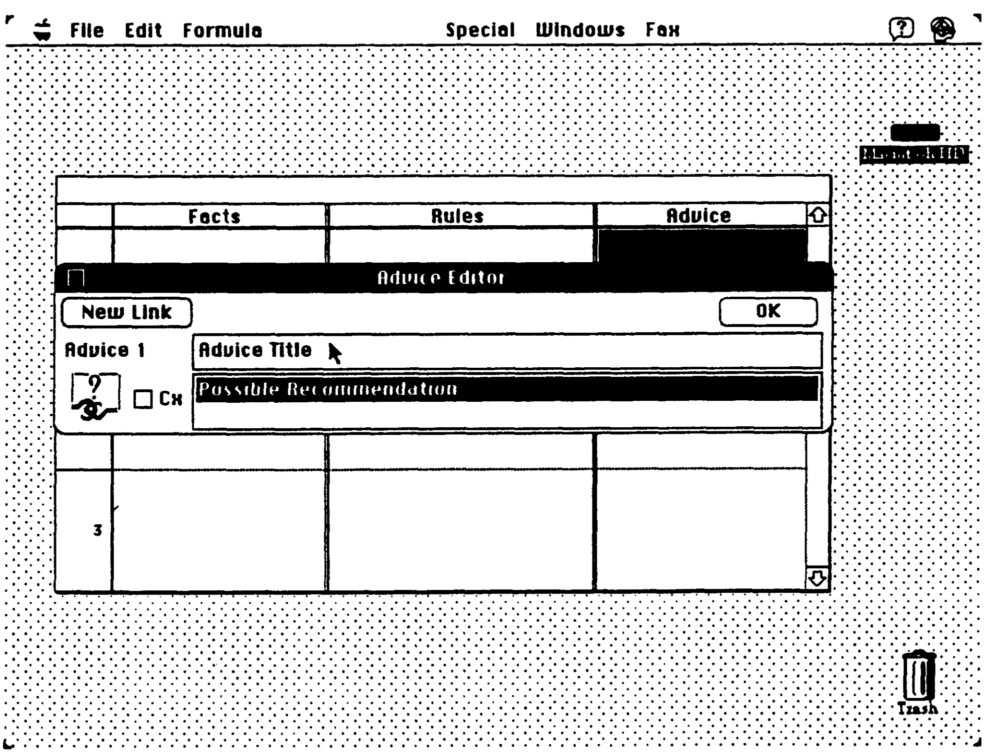

Figure 4.2.5 is an extract from the pilot study which examined mathematics options.

Figure 4.2.5
Extract from pilot study MacSMARTS Logic Worksheet

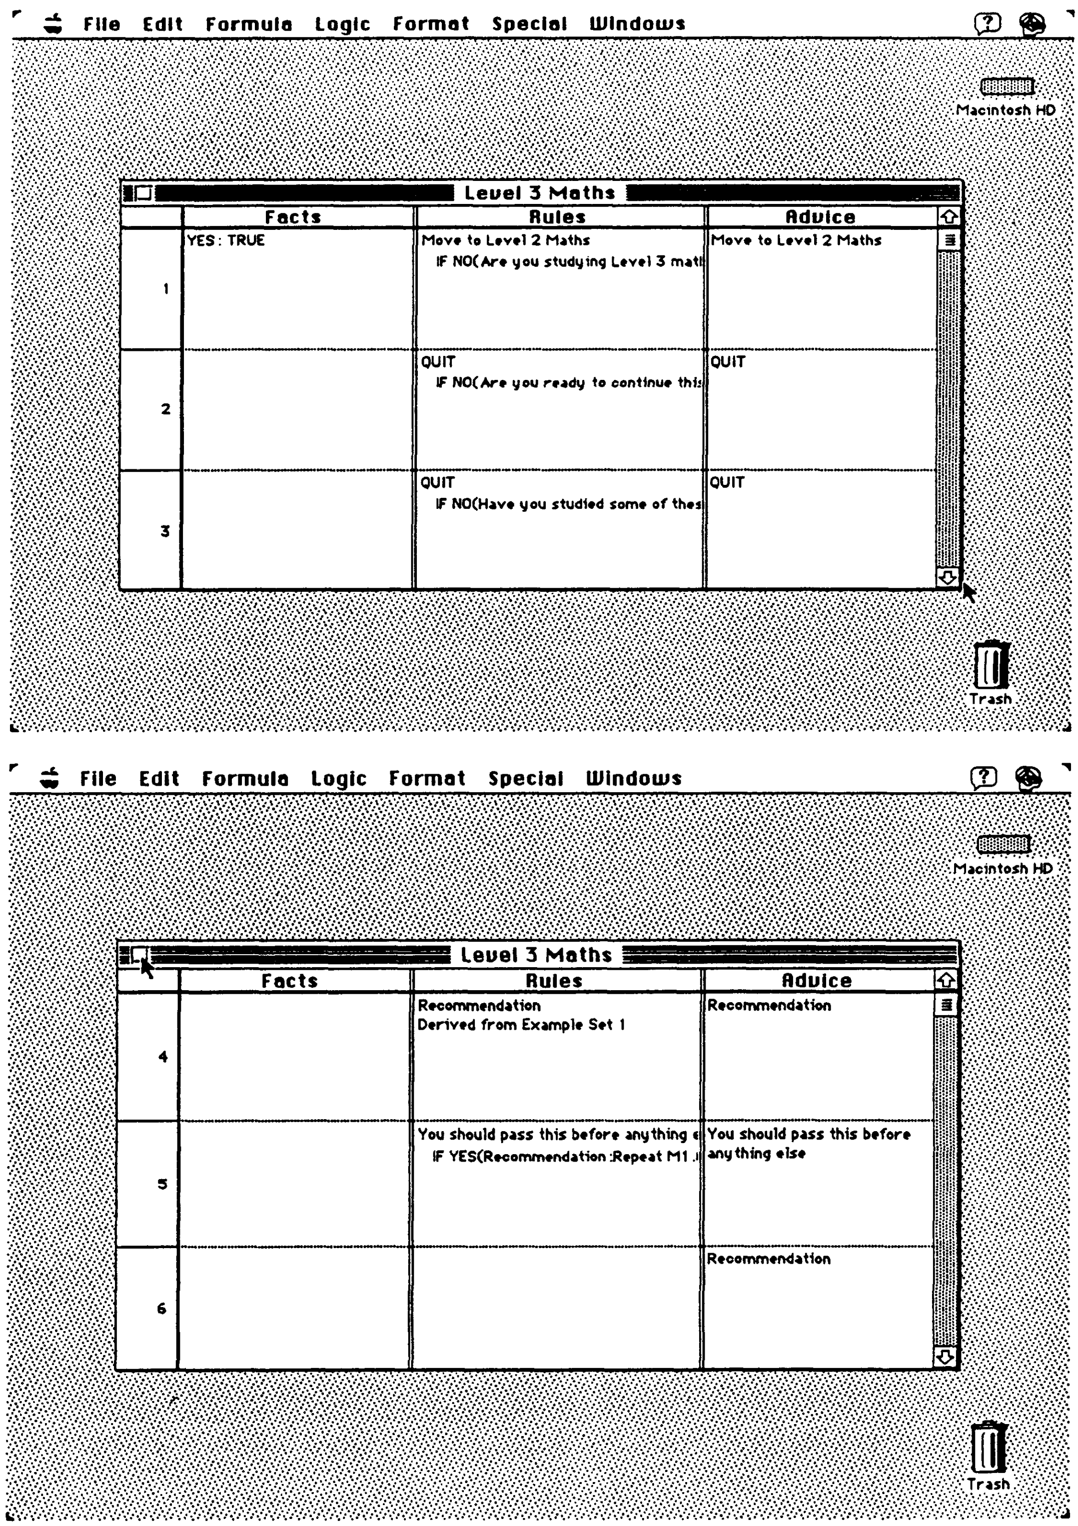

One of the features illustrated in figure 4.2.5 is the opportunity to access other data bases (for example, rules "Derived from Example Set 1"). Figure 4.2.6 illustrates part of the "Example Set 1" which was prepared using a series of three Editors illustrated in figures 4.2.7, 4.2.8 and 4.2.9.

Figure 4.2.6
Example of a MacSMARTS example-based knowledge base

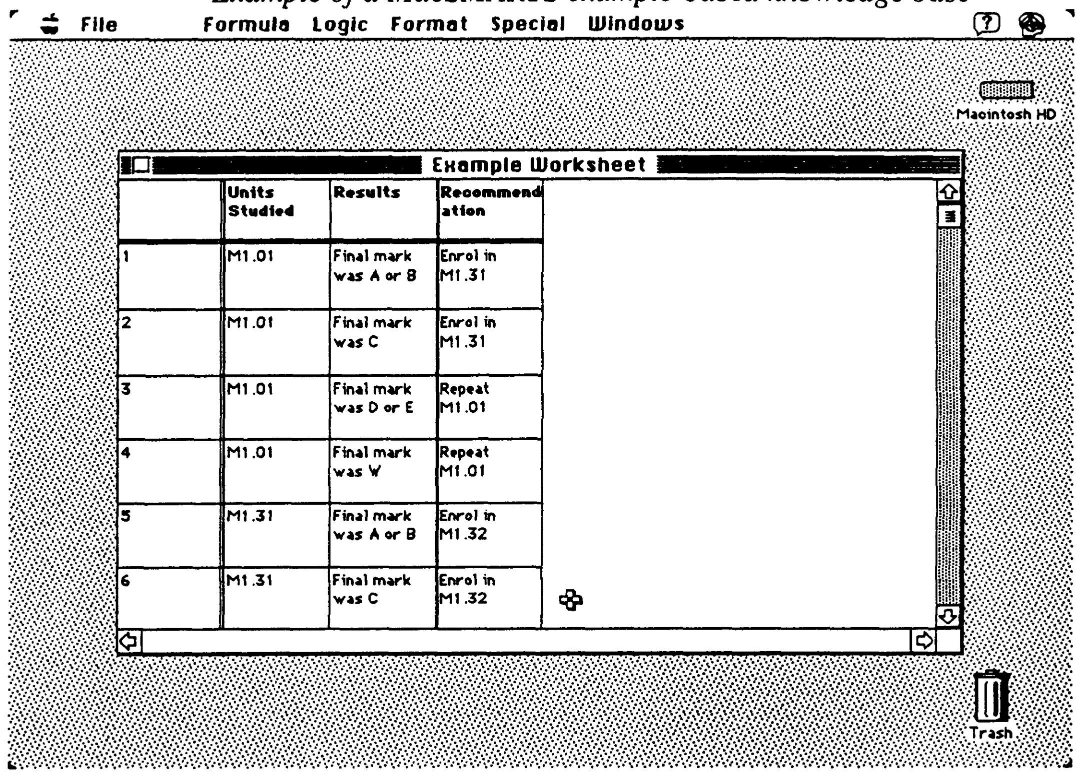

Figure 4.2.7
Example of a MacSMARTS knowledge base Factor Editor

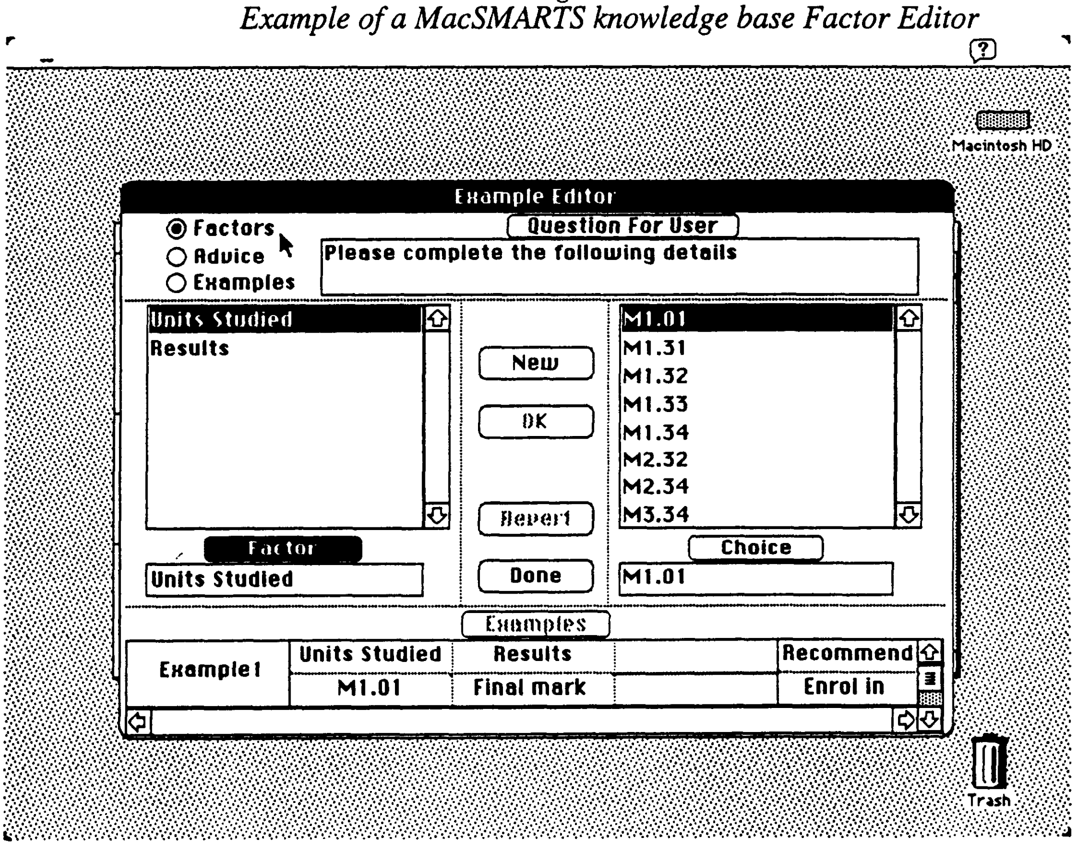

Figure 4.2.8
Example of a MacSMARTS knowledge base Advice Editor

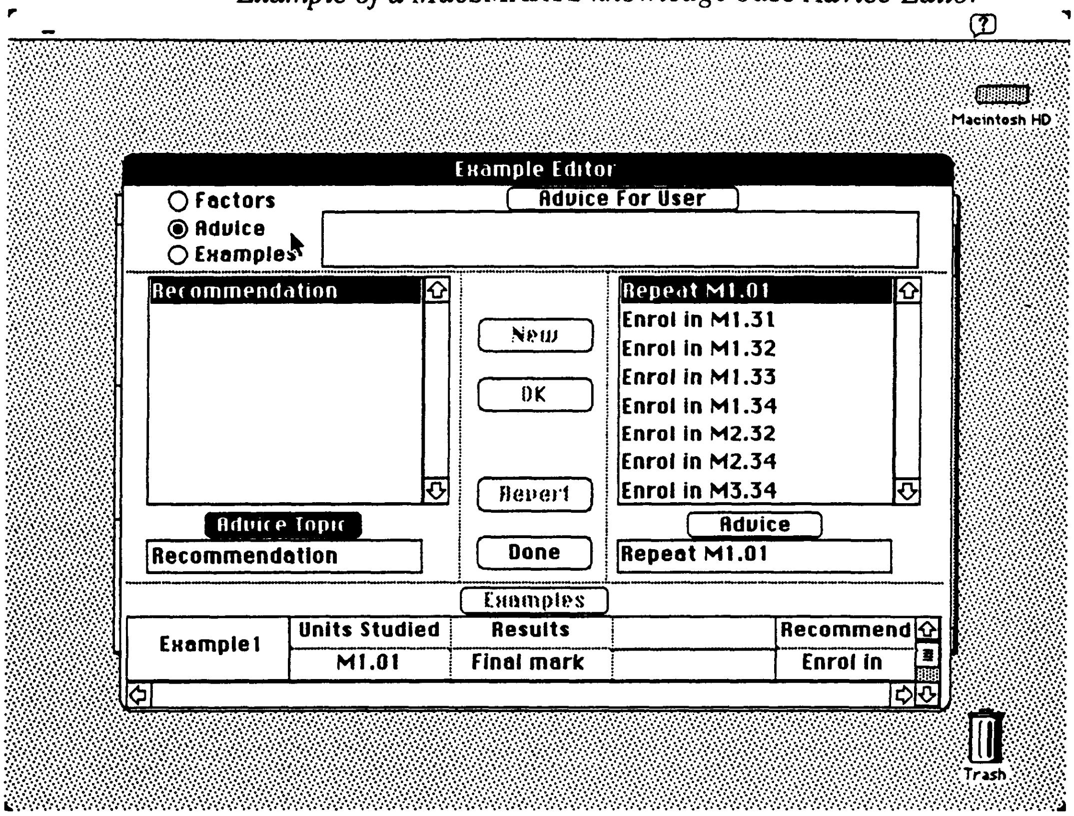

Figure 4.2.9
Example of a MacSMARTS knowledge base Examples Editor

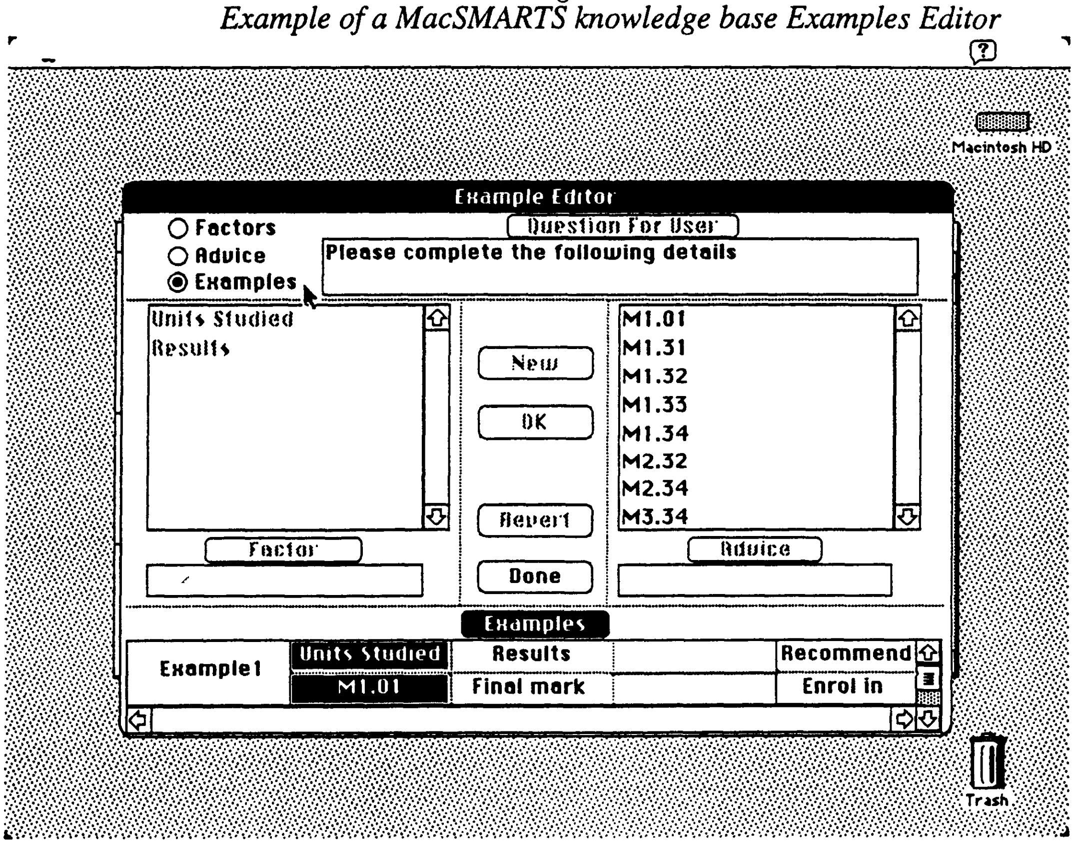

The following screens illustrate a user interaction with MacSMARTS. User response to the screen illustrated by figure 4.2.10 initiates one of three mathematics programs to run or for the session to end.

Figure 4.2.10
MacSMARTS module selection

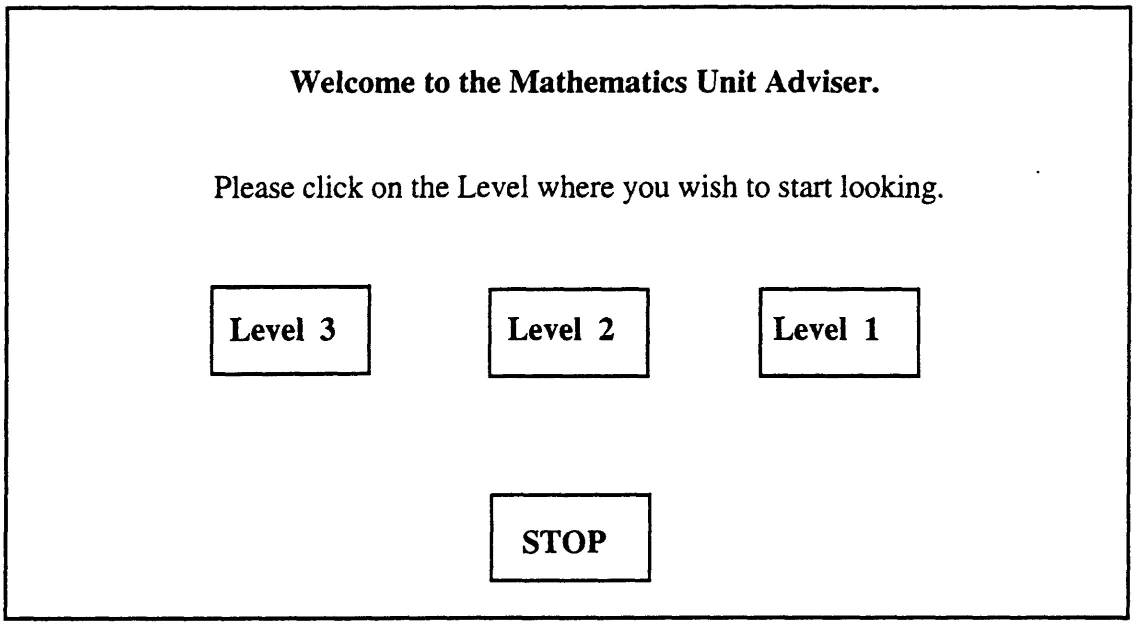

A user selecting Level 3 would then be presented with the screen in figure 4.2.11.

Figure 4.2.11
MacSMARTS in-use sample screen 2

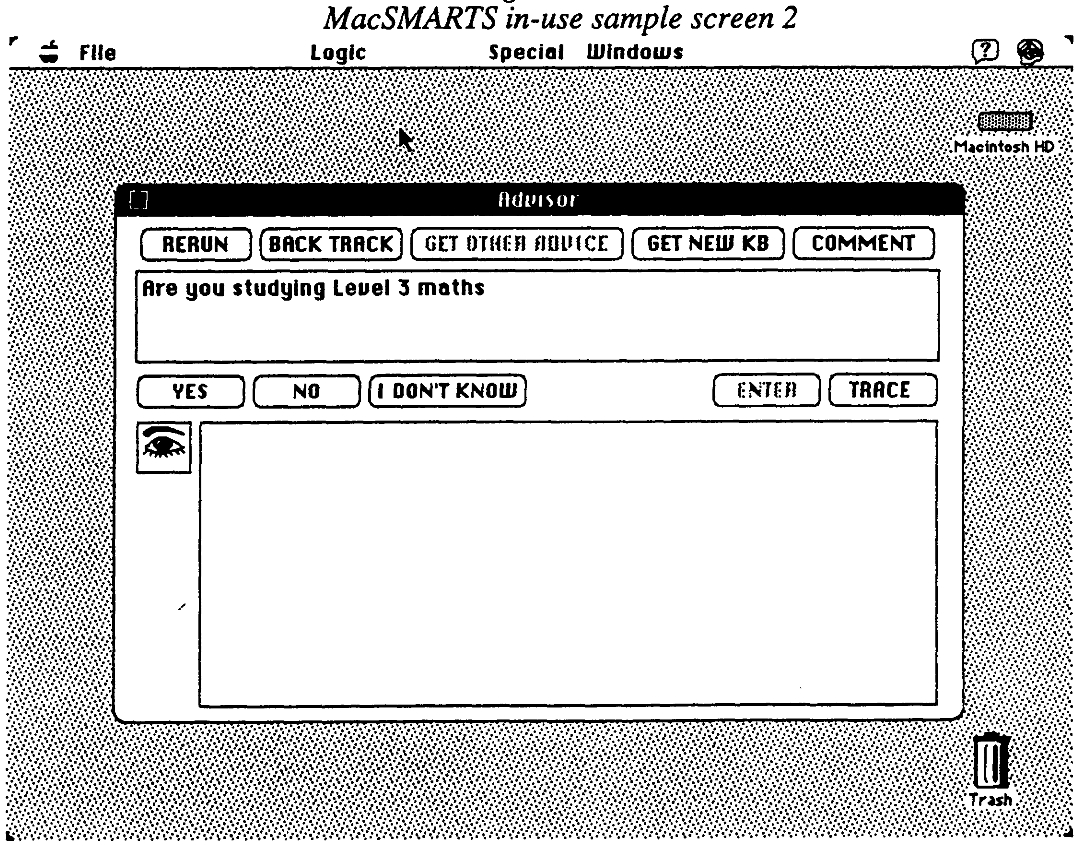

In this example, YES has been selected and the user is then presented with the screen in figure 4.2.12. Had the user responded NO, the next screen would have asked the user if they were studying Level 2 maths in which case the Level 2 program would be run. If the user's response to Level 2 was negative then the Level 1 program would run.

Figure 4.2.12
MacSMARTS in-use sample screen 3

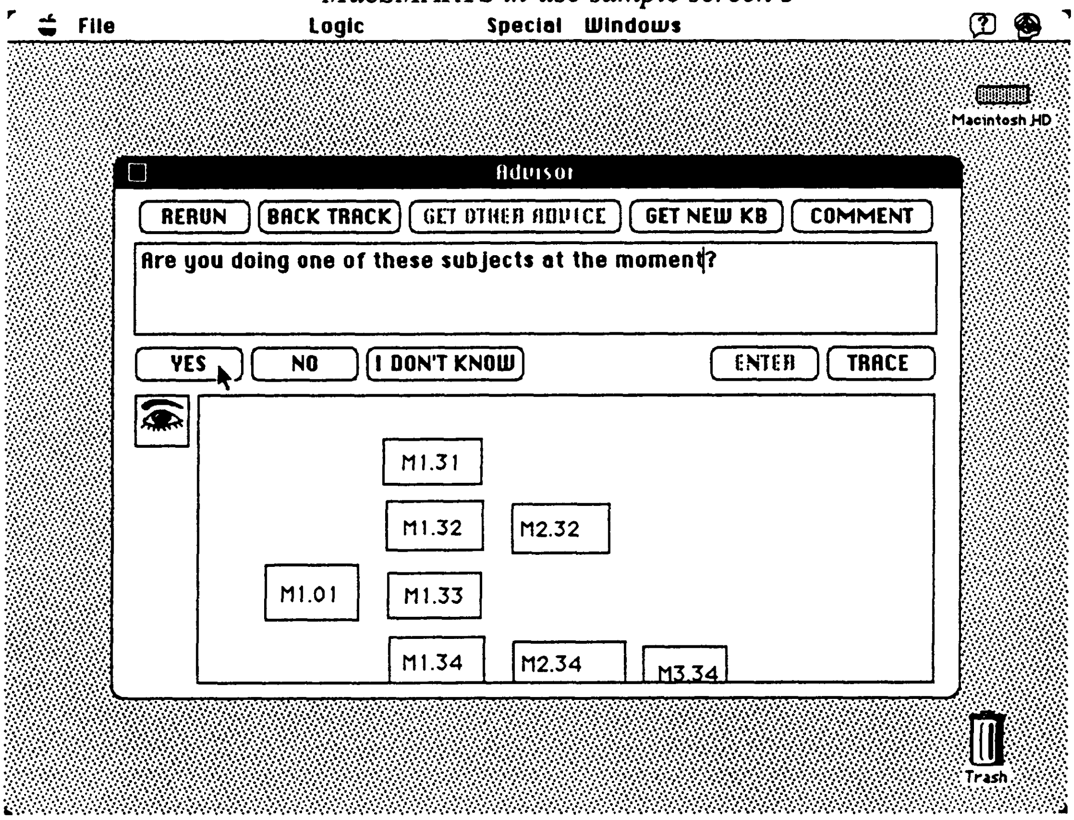

A negative or 'Don't Know' response to figure 4.2.12 would result in the Level 2 program being run; the presumption being that the student had incorrectly selected Level 3.

In this example the user is doing one of the subjects listed and is then asked (figure 4.2.13) to select his or her current mathematics unit and enter the grade received for this unit (figure 4.2.14).

Figure 4.2.13
MacSMARTS in-use sample screen 4

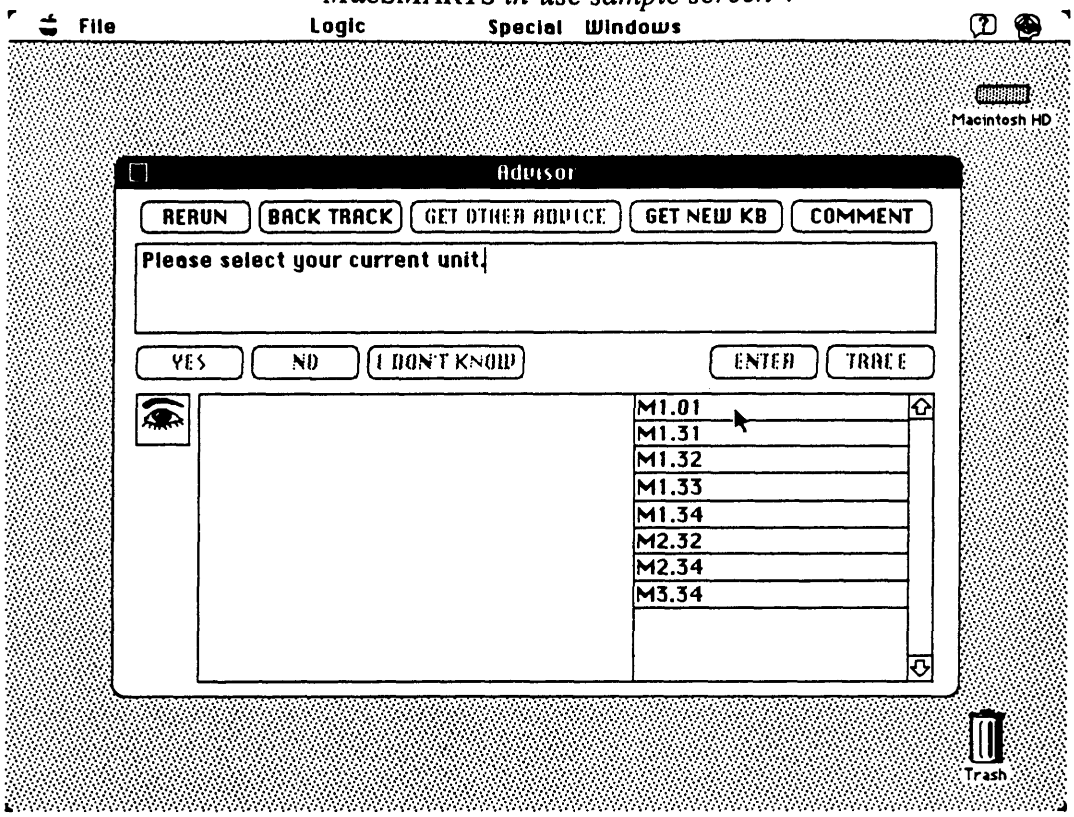

Figure 4.2.14
MacSMARTS in-use sample screen 5

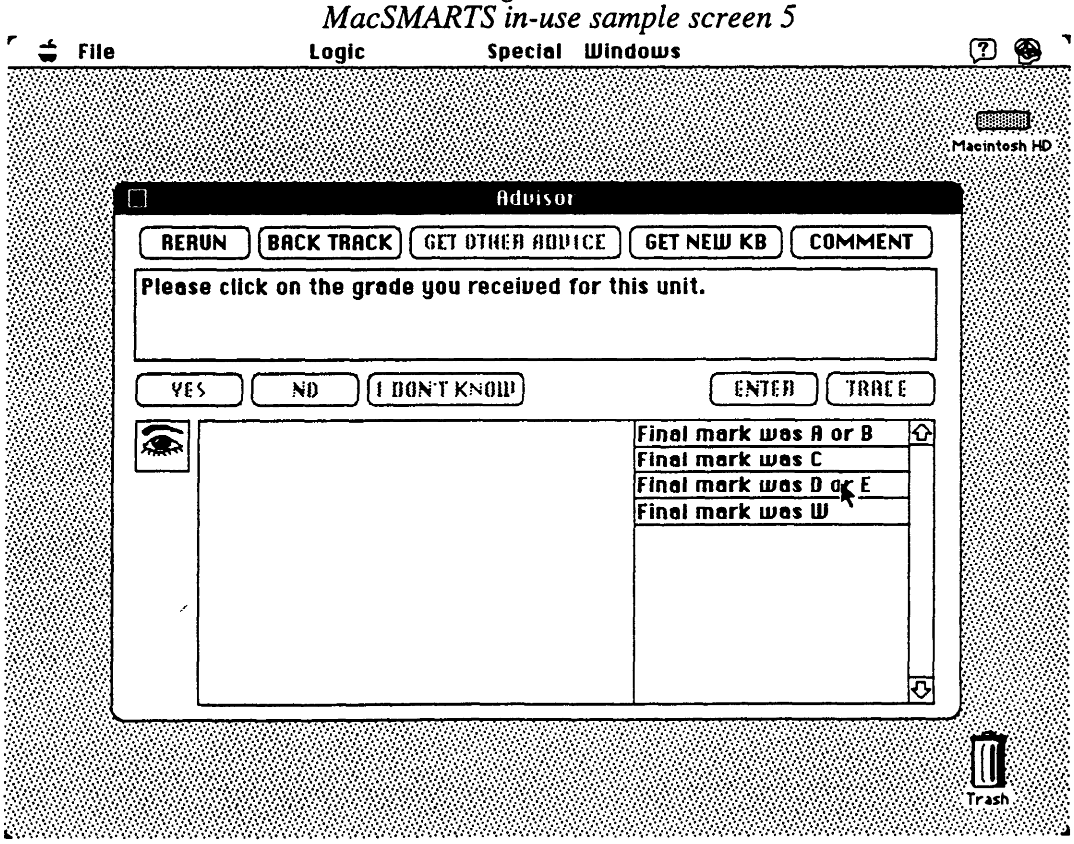

The recommendation is then given — in this case the Final mark was D or E and the student is advised to repeat the unit. On seeking further advice the user is advised as per figure 4.2.16.

Figure 4.2.15
MacSMARTS in-use sample screen 6

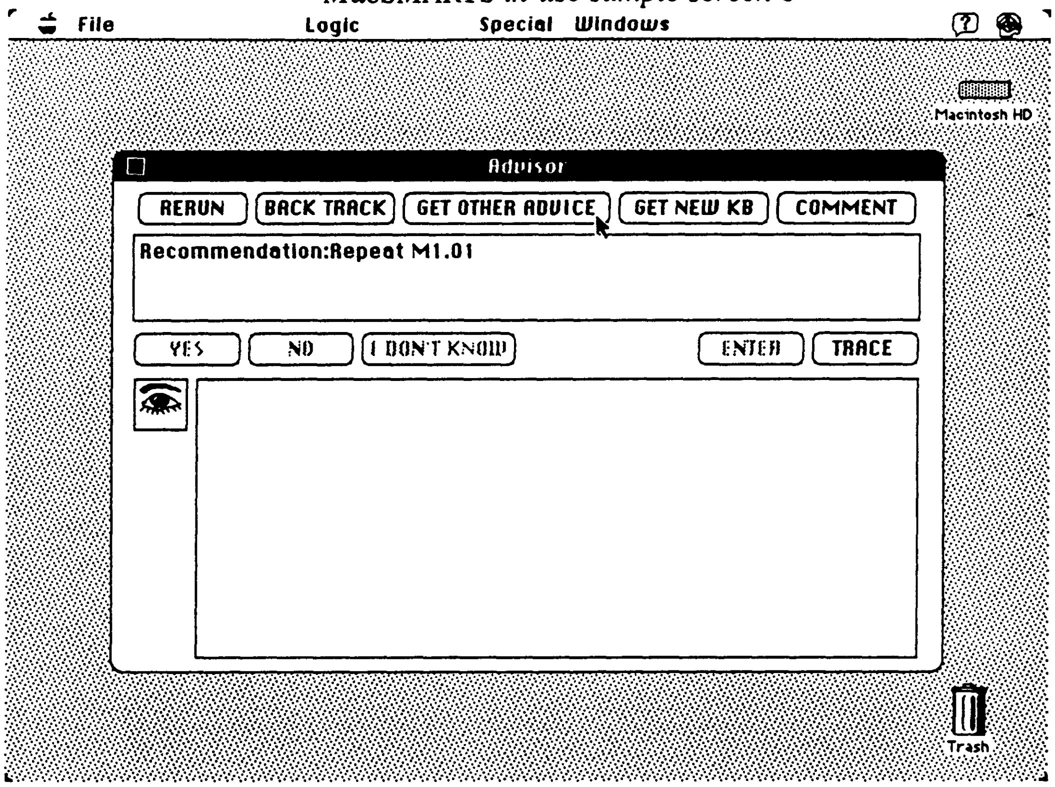

Figure 4.2.16
MacSMARTS in-use sample screen 7

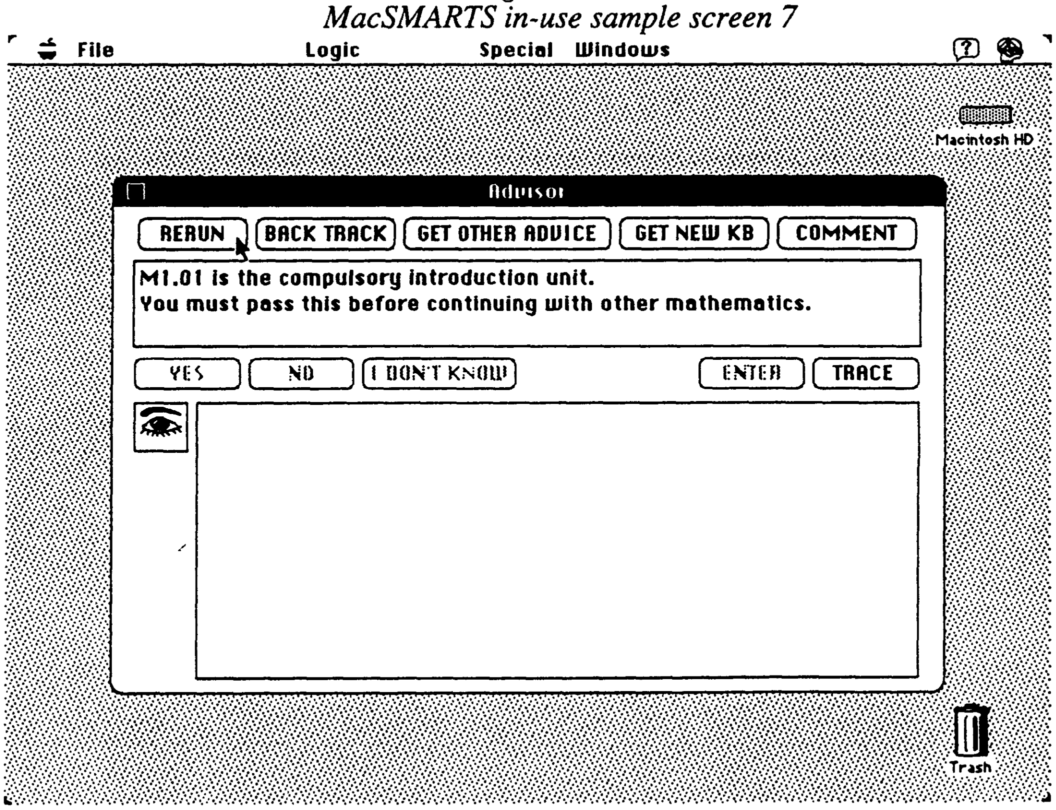

4.3 SHELL TRIAL FINDINGS

The MacSMARTS shell used in the pilot study did provide some of the expected benefits. The shell was easy to install on the Macintosh. It did not require significant computer programming skills and was relatively easy to customise. Although a suggested link with Hypercard could not be achieved, thus limiting the user interface, the program included graphic as well as text displays and the opportunity for user input. The shell was able to link with other MacSMARTS shells and thus would only be limited by the computer's memory capacity. The MacSMARTS program only required 250 Kilobytes and the sample data 89 Kilobytes of computer memory. However, the pilot study domain was limited to a small component of one subject area and a complete school system would require considerable more memory. It is unlikely that computer memory capacity in current Macintosh computers would present any problems for a complete expert system, but may in older computers.

The shell also displayed some of the weakness highlighted in Chapter Two. The support manual was encouraging but inadequate, and the Australian distributor appeared unfamiliar with the product. The support material included several attractive demonstration programs, but access to the program code was denied and thus it was not possible to see how these demonstration programs were prepared. Despite the options to include graphics and text, rules and reference tables, the user interface was restricted to standard screen displays and response options. Although these were intended to be user friendly, they limited the use of inference strategies and knowledge representation. The program interface also required user training which was significantly more than keyboard and keystroke familiarisation. Input of more than one item of data to the same screen by users was not possible, and thus it was cumbersome to obtain additional information. The shell did not include provision for certainty factors and was limited by the manner and order in which it examined rules. The explanation facilities were limited to tracing rules which had been fired.

The trial with MacSMARTS did not proceed beyond attempting a small component of the Mathematics area. The difficulty of translating a relatively simple flow chart into an expert system helped to illustrate the complexity of trying to undertake such a task for the rest of the faculty let alone the other faculties. The demonstration programs which accompanied the MacSMARTS shell highlighted the program's strategy using a branching pattern with the aim of providing a single diagnosis or recommendation. The limitations and difficulties arising during the trial clearly indicated that it was undesirable to continue with that, or probably another, shell. This is not to imply that MacSMARTS™ was an unsatisfactory product rather that it was not the product for the task.

Despite the disappointing results from trialing the shell MacSMARTS, and a better appreciation of the difficulties in developing an expert system, it was appropriate to continue the research project. The researcher's experience thus far with the preliminary study, supplemented by further investigation, indicated that utilising a toolkit for a second pilot study would be merited.

4.4 TRIAL IMPLEMENTATION USING A TOOLKIT

A second pilot study, using the expert system development toolkit KES II™ (Knowledge Engineering System), was developed for the same domain as the first pilot study so that differences in outcomes would reflect differences between the shell and the toolkit. Using the material prepared for the first pilot study also avoided the need to involve other people further in what might have been an unsuccessful exploration, saving their goodwill for later stages.

KES II was selected on very pragmatic grounds. No specific reports or comparisons of toolkits had been found in the literature reviewed. General reports in the literature had suggested KES II, like other toolkits, as a worthwhile but expensive program. Following the cost of purchasing MacSMARTS, the researcher was reluctant to spend an even higher amount on a program that might or might not be useful. Because most secondary schools in the Northern Territory had Macintosh computers, and all had Unisys computers, it was considered desirable to develop the expert system in one of these environments. No toolkits appeared to be available for the Macintosh but Unisys were prepared to lend a copy of their then current toolkit KES II.

KES II is an expert system building tool which "allows non-programmers to build powerful systems that can help solve problems" (Unisys 1987a, 2). KES provides three separate inference engines to cater for rule based Production System (PS engine), Hypothesise and Test (HT engine), or statistical reasoning (BAYES engine) expert systems. These engines can be run independently or linked together and may share data. The programs are written in C but may be linked with other languages or embedded in other C or Pascal programs. KES was designed to run on Unisys B2x or B3x Series work stations. The combined program had a 1989 retail value of approximately US $55 000.

All three KES inference engines use a goal directed (backward chaining) approach; the difference between them is the manner in which the knowledge is represented and the information processed. Two of the KES sub systems (PS and HT) can also perform event-driven (forward chaining) inferencing in which it responds to an occurrence rather than pursuing a goal; that is, a value is assigned to an attribute which causes other events to occur. Thus, depending on end-user requirements, a KES expert system may be prepared to perform backward and/or forward chaining.

KES was supplied by Unisys with an extensive set of manuals for each of the three inference engines. The first manual in each set provided instructions for installing and verifying the installation of KES onto the computer's hard-disk from floppy disks. The second manual contained general information about expert systems and KES and described how to use the KES subsystem. The third manual contained reference material related to the particular inference engine. The fourth manual was a reference for the knowledge base author and contained a glossary, format summaries, definitions and other aids such as error message explanations. It also included several sample knowledge bases and a list of suggested readings.

4.5 TOOLKIT TRIAL FINDINGS

The second pilot study exhibited the major disadvantage suggested by several authors in Chapter Two; that is, the product was not in general use and there were virtually no local (or Australian) training or user-help services available, and thus the researcher was required to teach himself from the program manuals. These extensive manuals appear to have been prepared to support rather than teach users of the toolkit.

Problems were exacerbated by hardware and software incompatibility. Although not technically difficult to solve, they required considerable time and external advice. To some extent the problems were compounded by the research location (Darwin), but mainly were the outcome of using a software product which did not have an effective support base in Australia.

The second pilot study, despite these difficulties, was judged to be considerably more successful than the first pilot study. The toolkit provided a richer data input and knowledge structure, the inference engines were more versatile, it had better explanation facilities and the user interface had more potential. KES is described more fully in Chapter Five but the following extract (figure 4.5.1) indicates the relative ease, compared against the shell, of providing recommendations. (The school changed the structure and unit titles in mathematics during the pilot study and thus the following does not exactly match the data shown in the shell.)

Figure 4.5.1
KES Pilot Extract

```
M12default:
\ If a student has not completed any mathematics
\ then we recommend Introduction Mathematics for next term.
        if
        DoneUnits = No Mathematics
        then
        RECOMMEND = MA12T.
        endif.

M12T diagnosis:
\ If a student has completed Transition Mathematics
\ but has not done MA12T
\ then we recommend Introduction Mathematics for next term.
        if
        DoneUnits = MA11T
        and DoneUnits # MA12T
        then
        RECOMMEND = MA12T.
        endif.

M12_R21M diagnosis:
\ If a student has failed MA12T
\ and has not completed Stage One remedial Mathematics
\ then we recommend RR21M for next term's mathematics.
        if
        GRADES:MA12T>RESULTS = E
        and DoneUnits # RR21M
        then
        RECOMMEND = RR21M.
        endif.

M1EL diagnosis:
\ If a student has completed either MA12T or RR22M
\ but has not done MA1EL
\ then we recommend MA1EL for next term.
        if
        DoneUnits = MA12T | RR22M
        and DoneUnits # MA1EL
        and STAGE = Stage one
        then
        RECOMMEND = MA1EL.
        endif.
```

The researcher's experience with the second pilot study indicated that it may be possible to develop an effective expert system for the selected domain by an amateur using an expert system development toolkit.

4.6 KNOWLEDGE ENGINEERING

The pilot study was commenced in 1991. By that time the school had prepared formal models to assist students in their subject selection. These models continue to be modified annually in reply to user responses and changes at the school, but their cardinal style and relationships are constant. These formal models had been intended as stand-alone documents to which individual data could be applied to ascertain appropriate subject options. If they satisfied this criterion then they should have been capable of providing the basic information from which the expert system could be developed. One aim of the pilot study was to attempt such development for part of a faculty to determine an implementation procedure for the full implementation.

Undertaking the preliminary studies demonstrated that simply translating the formal models to the expert system was less than satisfactory. The preliminary studies demonstrated four problems within the formal models: (1) content errors which gave inconsistent information, (2) errors of omission which failed to cater for unexpected situations, (3) errors in logic which provided conflicting information, and (4) presumptions of common sense or knowledge of procedures. To resolve these problems, the faculty co-ordinator and other staff were consulted. These consultations further revealed that staff were not consistent in their interpretation of the formal models; this inconsistency was demonstrated between teachers and even by individual teacher's advice to different students. These differences mainly arose from subjective opinions of a student's personal characteristics such as ability, potential, and diligence. Some differences arose from different weighting applied to grades in pre-requisite units.

Undertaking the pilot study revealed that the formal models required corrections for use by the school and required further interpretation as a source for the expert system. The pilot study revealed that technically correct models were still open to subjective interpretation and application. Thus in determining an implementation program for the expert system it was decided to develop the knowledge base in five stages: (1) translate the formal models into computer-based faculty knowledge bases, (2) prepare a means to mechanically test these knowledge bases, (3) liaise with the faculty co-ordinators and other senior staff to clarify discrepancies, (4) determine a technique for incorporating subjective weighting in the computer model information, and (5) liaise with staff to evaluate the subjective weighting. The working system is described in Chapter Five.

4.7 CONCLUSION

The preliminary study using a shell and toolkit to develop limited expert systems was successful in helping to ascertain appropriate techniques for, and the possible success of, an expert system dealing with student subject selection.

Although the shell used in the pilot study did not meet expectations, the time and financial cost were an appropriate component in the trial and selection process and a valuable contribution to the researcher's learning curve. The results of this preliminary study supported the view that there is no simple or single answer in determining the best form of knowledge engineering and construction of an expert system.

The preliminary study demonstrated the likelihood that the toolkit KES II would be satisfactory for developing the expert system. It was considered appropriate to continue use of the toolkit KES II, even though such use beyond a prototype would probably be more expensive than using either a shell or language. The selection of appropriate computer hardware was also a relevant factor in this decision. The trial school (and most other secondary schools in the Northern Territory) had Macintosh computers. On the other hand, all Northern Territory Government schools had Unisys computer networks on which, *inter alia,* student records were maintained. It was considered more practical to undertake the expert system on the same computer hardware where the student records were held and thus reduce the need for user input or the requirement to communicate between computer systems. Unisys made available a copy of the toolkit KES II, an advanced (for 1989) expert system development program, and the Northern Territory Department of Education provided hardware for the research. Thus pragmatic factors of cost and availability were significant elements in the selection of hardware and software for the expert system.

The preliminary study revealed some difficulties in the information provided by the school about courses. However, the process of knowledge elicitation which was used during the preliminary study could be refined to accommodate these difficulties and continued in the working system. The form of knowledge representation was not a primary factor in the preliminary study of MacSMART and KES but was significant. The rule base system for knowledge representation used by KES had been examined in the literature reviewed and was considered appropriate for the working system.

The preliminary study also confirmed that developing an effective expert system to recommend unit selection would be a complex and time consuming task.