Problem Based Learning with Technological Support in an AI Subject: Description and Evaluation

Lawrence Cavedon James Harland
Lin Padgham
Department of Computer Science
RMIT University, Melbourne, Australia
email: {cavedon, jah, linpa}@cs.rmit.edu.au,

Abstract

The diversity of students' backgrounds has caused a number of problems with an introductory Artificial Intelligence subject. In this paper we describe our experiences using a problem-based learning approach to teaching this material. This involved the students forming themselves into groups in order to tackle a series of related problems, each of which dealt with a specific area of the course. An integral part of this project is the technological support, which was used not only to monitor the progress of students through the course, but also to provide the necessary *contextualisation* of course material. We describe our methodology, as well as the two WWW-based support tools, the CourseWeb, which functions (roughly speaking) as an automated textbook, and the ProblemWeb, which specifies the problems to be solved and keeps track of the group's progress in solving them. We report on our experiences in teaching a summer version of this subject and on our evaluations of the process.

1 Introduction

At the Masters by Coursework level we have found a number of problems with the Artificial Intelligence subject which we have attempted to address using an approach involving *contextualisation.* There are also some issues in teaching the undergraduate level AI subject which have been partially addressed using this teaching approach.

Our Masters by Coursework classes comprise a student group with extremely diverse backgrounds, needs and capabilities. There are students with significant experience of industry, with a variety of educational backgrounds (including a substantial number of foreign students) and with a variety of motivations for undertaking the course. There are severe shortcomings with the traditional lecturing/tutoring mode of teaching in this situation, as some students are inadequately prepared while others find that too much familiar ground is covered. A number of students also find that the material seems irrelevant and not related to their needs and interests. Many of the teaching examples used are regarded as trivial and uninteresting. These shortcomings are reflected in both higher than usual attrition rates and student dissatisfaction. Several attempts have been made to modify the level at which material is covered in lectures, labs and tutorials, but the problem has remained acute.

We have addressed this problem by a dual approach to what we call contextualisation. Firstly we have adopted a problem based learning approach [3, 1] to teaching the course, where all learning is focussed around real problems, so that the relevance of techniques studied is evident. Also the students themselves identify and decide on the approach to a problem solution (with teacher guidance), so there is less likelihood that students will see the material as irrelevant.

Secondly we have developed an innovative approach to the presentation of course material that allows each student to adopt a learning strategy that suits their individual requirements. We have developed a WWW-based software tool called the *CourseWeb*, where course material is organised in a way that allows a student to access components of the course material that are relevant and match the students background and requirements. Modules of the CourseWeb differ in the amount of detail that they contain, the mode of presentation and how they are linked to other modules. This permits a student to navigate (via Internet) through the modules at the level of detail that suits their background and requirements.

We have also developed a software tool called the *ProblemWeb* which organises problems and sub-problems to support the problem based learning aspect of the course. We describe each of these software tools briefly in the following section, and then describe the running of the initial trial version of this course. We report on the various forms of evaluation which we have carried out, and conclude with a discussion of the factors to be investigated in our continuing work with this course.

2 Software Support

The CourseWeb and ProblemWeb are orthogonal WWW-based tools, with no direct links between them. However it is envisaged that when designing a particular course/subject, the teacher will ensure that there is a match between the problems addressed in the CourseWeb, and the resources needed to solve these problems, which should be referenced via the CourseWeb.

2.1 CourseWeb

The basic unit of the CourseWeb is a *module.* A module contains a unit of course material, along with a list of *conditions* which comprise the prerequisites of that module (i.e. information assumed by the material contained in the module) and a list of conditions which comprise the outcomes of the module (i.e. information provided by the module). The CourseWeb system automatically links modules together by matching prerequisite and outcome conditions: for example, in Figure 1, following the hypertext link under the first Prerequisite takes the user to the module which has the associated condition as an outcome. In the current prototype, a condition is simply a text string; however, we plan to extend this so that a series of problems can be associated with a condition, allowing students to test themselves when exiting a module, or to check their proficiency against a given prerequisite on entering a module. This clear statement of prerequisites, coupled with the links to modules providing the prerequisites, is designed to cater for students with different background knowledge. Further, the modular design allows the CourseWeb to support random access into the course material — a student can work backwards through the prerequisite topics if the material has been accessed at a point for which he or she feels unprepared.

Another unit of information in the CourseWeb is a *thread*, which is a sequential list of modules. The most common use of a thread is to combine a number of modules together into a set of materials for a specific topic, or to combine a set of topics (each with its own thread) into a course. A thread is created simply by specifying the modules which are to be part of it. The ease with which threads can be defined makes it realistic for the course designer or lecturer to define different versions of a course (using different threads) to cater for different needs.

As well as the sequencing of modules, other relationships between modules can be defined. Most interestingly, a module may be specified to be a *specialisation* of another module. In this case, the specialisation module includes the information contained in the more general version, but to a greater level of detail. For example, in the AI course, the module which describes search has a specialisation module which contains the algorithms for implementing search techniques. This module itself has specialisations which describe the implementations in particular programming languages. The various modules at different levels of detail are contained in different threads; a student who is simply after an overview of search can access an overview thread that contains the more general modules, whereas a student who is interested in implementing search algorithms or needs the more detailed understanding can access the thread that contains the specialisation modules. A student following the less detailed thread may decide that they would like further details; specialisation modules are linked to the more general versions so such students are easily able to move "downwards" for such details.

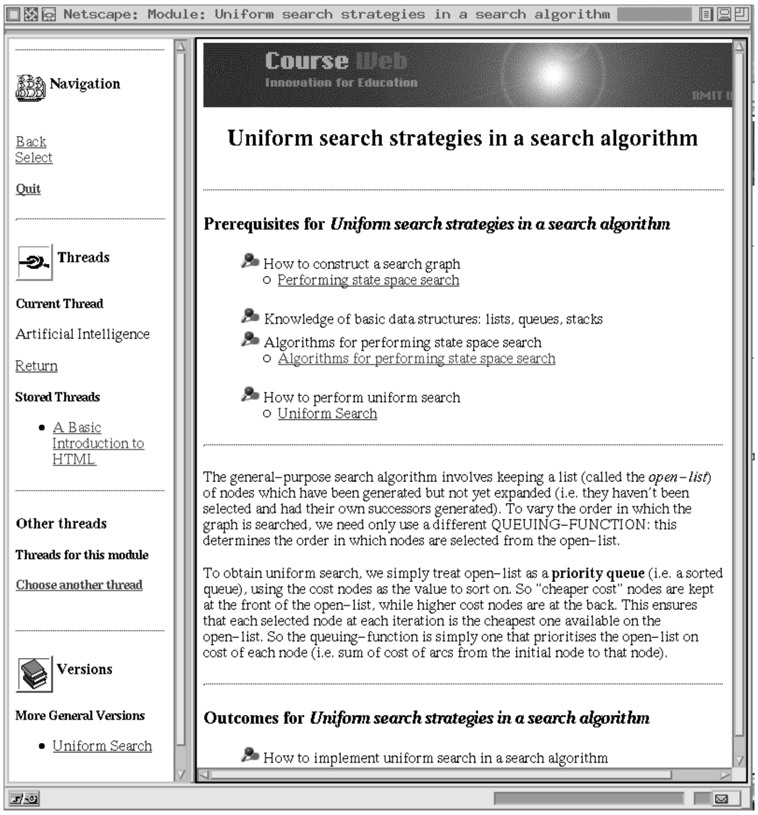

Figure 1: Figure 1: Example of a CourseWeb page

A clear advantage of using WWW-based technology is that a variety of information sources can be used — traditional notes, references to books and papers, WWW pages (such as the CMU repository), animations, movies, language manuals, exercises, example programs, etc.

2.2 ProblemWeb

The ProblemWeb consists of a set of problems related to a particular subject. In the AI course these problems are all related to a single "umbrella" problem which is to automate the provision of a taxi service, including driving the taxi, communicating with passengers and co-ordinating the taxi fleet. This umbrella problem serves to introduce the students to various areas of AI, relating the relevance of these areas to a concrete problem.

Other problems in the ProblemWeb are necessarily more limited in nature, but are clearly able to be related to the initial problem. For example the students are asked to investigate general purpose algorithms for finding a route between a start point and a destination, leading them into an exploration of various search algorithms. They then represent a simplified map, and implement the domain-specific functions (goal check and successor functions) required to use a general purpose search algorithm. Knowledge representation is motivated by the large amounts of knowledge and reasoning obviously required, and subsets of this are used for motivating exploration of logic, reasoning with uncertain or default information and hierarchical representations. Planning, natural language processing and a variety of other AI areas are treated in a similar manner.

The ProblemWeb allows for some problems to be specified as core problems (i.e. they must be addressed) while others are elective problems, allowing some choice. It is possible to specify that certain problems must be addressed before others are attempted, if the tutor wishes to build this guidance into the problem set.

A progress monitoring tool is built into the ProblemWeb and is used to give feedback to the students about their overall progress in the subject. A student can click on a button from the menu and see an overall summary of progress in the subject, and can also view a log of the submissions made. This also shows the number of problem nodes required to complete the subject. This means that groups can use the ProblemWeb to keep track of the overall progress of the group, and to plan future work, as well as allowing a tutor the ability to monitor the students' progress more efficiently than is possible on paper.

3 Teaching Experiences

We have run one class using this methodology and software, and another will be complete by the end of the first semester of 1997. The initial class consisted of 15 undergraduate students and 1 postgraduate. The students were selected to have a normal distribution based on their previous marks. Two students had previously failed the undergraduate AI course. We used undergraduate students for the initial version of the course, as we ran it as an intensive summer subject in the undergraduate program, to allow for some evaluation prior to the initial use in the Masters program for which it was primarily designed. Two students withdrew after the initial session, leaving 14 students who worked in two groups of 7, with 2 staff members.

Groups were scheduled for fifteen three-hour sessions with the tutor, over a period of five weeks. One of the groups met with the tutor at all fifteen sessions, though not always for the full three hours. The other group met with the tutor on eleven of the fifteen occasions.

In the first week, students were introduced to the methodology and software support, had some introductory discussions in various small groups as well as the full class, selected their groups, discussed and researched the umbrella problem, produced a written report in groups, and had a trial run of the individual and peer assessment component of the final assessment. In the following weeks students worked through the four problem areas, which were search, knowledge representation, planning and natural language. There were two deadlines for completion of a "near final" version of some number of submissions. This was in order to assist students in meeting the overall deadline.

During the five weeks of the course the students met most weekdays, and spent significant amounts of time discussing what they had read, possible solutions to the problems being investigated, and clarification of details. They visited the library regularly, and often brought additional reference books, or notes taken from reference books, to their discussion meetings. It was evident that all students were actively engaging with the subject matter of the course, and discussion was stimulating and active. The role of the tutor was primarily to direct the students to various resources, such as a code repository and particular areas covered by the CourseWeb, to clarify areas of confusion, and to ask questions to draw attention to weaknesses in proposed solutions.

In the summer version of the course, there was no choice of which problems to cover, although students had some choice regarding the order in which problems were addressed. The constraints imposed were that search and knowledge representation must be worked through before planning and natural language were attempted. Both groups in fact worked through the problems fairly much in the order in which they were presented in the ProblemWeb, with the exception of an overview knowledge representation question which was addressed by both groups after the detailed questions had been tackled.

All submissions could be made multiple times, up until the final deadline for the course. Feedback was given on each submission to assist students to improve their work. Feedback was in the form of an indication of areas that could be improved, and questions that were intended to prompt further discussion which could discover problems, or facilitate more in depth exploration. There was also some automated feedback incorporated into the ProblemWeb, which took the form of some explicit questions about the problem solution to ensure that the students had considered critical areas.

According to educational research it is important that assessment be congruent with the teaching and learning approach used [2], and that multiple forms of assessment are desirable [5]. Consequently we had four aspects to the assessment — a mark for group process, worth 25% of the total, a mark for individual contribution to the group, worth 20%, a mark worth 30% for the group submissions regarding problems, and 25% for an exam.

Assessment is an important motivator to students, making it important to give some marks to aspects such as group process, as well as to products such as group or individual reports. In order to be congruent with the self-directed nature of the teaching approach used, the group process mark and the individual contribution mark were determined by a process of peer and self-assessment. The exam was incorporated partly to ensure that individuals had achieved the requisite standard of knowledge, and partly because it provided a means of comparison with students taught using the more traditional approach.

We had some concerns about the congruence of an exam with the problem-based learning approach, but attempted to design the exam questions around a single problem in order to provide some similarity to the learning model used in the course.

The group process mark was decided in group discussion, together with the tutor, with reference to how well the group had managed the various issues that had arisen within the group. The individual contribution marks were also decided in group discussion, together with the tutor. In one group the students chose to have individuals leave the room while they were discussed, and the results of the discussion, including the mark, were then summarised when the student returned. In the other group discussions of each individual took place while the individual was present. Both groups ended up with a spread of marks. There was some awkwardness surrounding the process, and the bottom student in one group, and the two bottom students in the other group were not entirely satisfied with their resulting marks.

4 Evaluation

Prior to running this course we undertook formative evaluation, where we examined the teaching model explicitly with respect to our aims. We also ran a trial half day course on learning basic HTML, which we ran with a reference group made up of academics and student representatives.

We have three types of information which we have used in our current evaluation of this course. One is tutor observations of student progress and engagement; another is the comments of students, made in either structured interviews held towards the end of the course, in a final questionnaire, or unsolicited comments in discussion; and the group solutions to the problems set and the individual exams.

The lack of lectures and the opportunities for self-paced learning were factors which attracted several students to the course. An additional factor was that it was offered in summer, which gave the students the opportunity to distribute their academic load over a larger portion of the year, as well as to be able to concentrate on one subject. Nevertheless a high proportion of the students were extremely positive about the mode of teaching, claiming that they enjoyed it more and learnt more than in traditional courses. A few students disliked the mode of teaching and indicated that they prefer to have well-structured lectures and course notes, and not to be required to interact in a group. All students felt however that this was the best group they had participated in during their education (they have had a number of group projects in other courses).

Tutor observations indicate that the students were significantly more engaged with the subject matter than is traditionally the case. One indication of this is that almost all students visited the library, (without being required to do so) in order to read more widely around the topics being discussed. Traditionally many students barely even read the set text, and almost never visit the library. Many students commented that having a large problem to solve made the motivations for particular topics much clearer than in a traditional course. One student, who had failed the equivalent traditional course previously, commented that “I never understood what all this search and knowledge representation was about, but now I understand what it is, what its used for and why you need it.” Another (extremely bright) student commented that “I really like not being told too much by the lecturer, because when you do it wrong first, then you understand why its wrong, and you understand why the right way is right. If you just get shown how to do it, you don't really understand why.”

The quality of the group submissions was high, indicating a wide and relatively thorough consideration of the area(s) and problem(s). The exams in contrast were mostly very poorly done, with a couple of exceptions. Part of the reason for this was certainly that the attempt to cast the exam questions around a single larger problem, introduced an unforeseen dependence between questions. When one question was poorly done, others became ill-defined and generally difficult to do well. It was evident from the exams that a number of students who had been active in discussions, who were contributing and appeared to be understanding, nevertheless did not understand the technical details associated with some aspects of the subject matter. We hypothesize that these students, while involved and active, perhaps did not actually themselves work out the low level details of some of the problem solutions.

There was considerable discussion about the peer and self-assessment process. Some students in one of the groups felt that the necessity to assess fellow members of the group spoilt the positive trust and sense of co-operation and interdependence that had been built up. There was an initial desire to just give all members a uniform high mark, but this was disallowed by the tutor. The two weakest students in this group found it distressing to be graded lowest by their peers, as they felt this betrayed the acceptance and inclusion they had felt during their time working together. However on reflection both decided they preferred the system of peer and self-assessment to being assigned a mark by the tutor, as they learnt more from it, and felt they had some input into the process.

One of the groups were unanimous in their feeling that they should relieve a near perfect mark for group process. This group had previously identified some problems in group process — mainly that there had been some dissatisfaction between students regarding standard of work, or the way that their individual work was received. When this was noted by the tutor, questioning the high assessment, the students defended the high mark, indicating how they had all worked at dealing with a number of problems which had arisen, such as initial poor attendance by one member and very poor written English of two other members. This group claimed that the relatively high mark (25%) given for group process motivated them to actively deal with this aspect, rather than just ensuring that the project work was performed to a high standard by, for example, having the best students do most of the work.

The other group was somewhat more negative about the group process, initially claiming that it hadn't worked at all, with one member saying they should fail this aspect. On further discussion they decided that despite some significant shortcomings, and problems they had failed to deal with in their group, the group had worked better than any other project groups they had been involved in. They ended up giving themselves a credit grade for this aspect of the course. On analysis it seems that the major problem in this group was that there was a subgroup of three friends, one of whom was particularly negative to the teaching approach and contributed comparatively little to the overall group process. His friends were not particularly happy about this, but overlooked it because he was a friend. The other group members felt unable to address it because he was shielded by the tight friendship group. This resulted in all but this one person feeling satisfied with their marks for the contribution to the group.

The tutors noted that the problems chosen, and their exact description, were very important. Discussion of how to solve the problems, and what to investigate in order to understand the problem was very useful. However discussion as to what was meant, and what was expected to be done was distracting, and could be avoided by careful problem specification. The tutors also noted that the students took seriously the relationship of subproblems to the motivating umbrella problem, and were disturbed by inconsistencies. For example one of the knowledge representation subproblems talked about the taxi driver, although the initial problem specified a completely automated taxi. Students reacted quite strongly to this, despite the fact that it clearly did not affect the relevance of the material. This emphasised how important it is to ensure that the problems are **exactly** right.

Students commented that they did not like or use the CourseWeb. Some reasons for this were the difficulty in seeing what was available, and the lack of any search facility to help them find relevant material. In addition, the CourseWeb did not have as much detailed information as we would eventually hope, so that sometimes after difficulty in finding the relevant module, students found that the information available was not worth the effort. However, they generally found the ProblemWeb useful and appropriate.

5 Discussion and Continuing Work

The development of this course is part of an ongoing project called SLiCE (Student Learning in a Contextualized Environment), and the evaluation from the initial running of this course indicates areas for continued work.

One of the most urgent areas for further work is the support software.

One of the intended functions of the CourseWeb is to assist the lecturer in building a coherent map of the concepts appropriate to the discipline, the transfer of which is one important aspect of teaching at the tertiary level [4]. However unless students are regularly using this software, this opportunity will be missed. We plan to develop better indexing, searching and organisational mechanisms, to facilitate student use. We also plan to improve the existing tool for tutor input of material.

The assessment mechanisms also require further work, as indicated by the above discussion of the peer assessment process. In addition, the provision of an appropriate exam remains a problem. From our experiences some form of individual assessment of this kind appears necessary, but finding the appropriate such form remains to be done. The relative weighting of each component of the course also requires some review.

We also noted that during the running of the course there was little opportunity for feedback on individual detailed work, which was borne out in the exam results. This suggests that some revision of the group structure would be appropriate, in order to ensure that each individual student has a grasp of the detailed subject matter, rather than the conceptual overview alone.

We would also like to extend this methodology to larger classes, in the manner of Judy Kay and her colleagues at the University of Sydney [6]. We also plan to undertake further comparative evaluation, particularly in regard to retention of material after the conclusion of the course.

Acknowledgements

The authors acknowledge the significant contributions of fellow team members of the SLiCE project Wayne Jenkins, Sheila Howell, Daryl D'Souza, Vic Ciesielski and Peter Ling.

References

[1] H. S. Barrows. A taxonomy of problem-based learning methods. *Medical Education*, Volume 20, 1986.

[2] J. Biggs. Approaches to the enhancment of tertiary education. *Higher Education Research and Development*, Volume 8, 1989.

[3] Boud and Felletti. *The Challenge of Problem-Based Learning.* Kogan Page, London, 1992.

[4] D. Laurillard. *Rethinking university teaching: A Framework for the effective use of Educational Technology.* Routledge, 1993.

[5] Rawnsley, Spaziani and Rangachari. Evaluation in a problem-based course: what teachers want, what students think. In *Proceedings of the International Conference on Problem-Based Learning in Higher Education*, Linkoping, 1995.

[6] Jeff Kingston Tony Greening, Judy Kay. Bassernet report on teaching initiatives for 1997 blue and problem-based learning. Technical report, University of Sydney, 1997.