Intelligent Case Selection and Presentation

Robert Farrell

Yale University
Department of Computer Science
New Haven, CT 06520-2158

Abstract

Designers of AI educational systems have largely ignored the problem of choosing what information to present and how to present it because systems to date have not been faced with a large space of possible presentations. This paper explains how a student model can provide the knowledge needed to make presentation decisions. Systems that use a student model for this purpose can make problems more interesting, challenging, and hand-tailored to students. We describe DECIDER: a program that improves student's decision-making abilities by constructing an environment where they can confront problems and decide on a plan of action. DECIDER chooses from a database of decision-making problems and varies how it presents those problems based on a model of the student's knowledge of goals, plans, and priorities.

A good teacher can take the same material and present it in entirely different ways to different audiences. The better the teacher's model of the audience, the more appropriate and instructive their presentation will be. Because most AI educational systems do *no* reasoning about how to present information, they have not had the need for a student model for this purpose. We outline a set of techniques that can select cases and make presentation choices based on a student model. The resulting system focuses the student on new problems while providing an environment where they can attempt those problems, see the outcome, and explain their failures [SCHA86].

We have chosen decision-making as our task area, but the system design and pedagogical techniques described in this paper are applicable to any tutorial system that must select and present multi-faceted cases from a large library. Our system, DECIDER, teaches decision-making by presenting background information, a problem, and several options. The student decides on a plan to remedy the problem and observes the outcome. We have chosen U.S. national-security decision-making as our application domain, but the knowledge base and tutorial component are separate, so any decision-making domain could be equally substituted.

This research was supported in part by the Advanced Research Projects Agency of the Department of Defense and monitored by the Office of Naval Research under contract N00014-85-K-0108.

I. Constraining The Teaching Domain

With any AI educational application modelling a real-world domain, the choice of what to teach must be constrained by both an abstract domain model and pedagogical goals defined within that model.

A. The Decision-Making Model

Decision-making is a complex process involving the interaction between advisors and decision-makers, opponents and proponents, planners and executers [NEUS86]. We have made a number of approximations to actual time-pressured, multi-agent decision-making, aimed at discretizing the actual continuous process of assessing information and arriving at decisions:

- Only one decision-maker *(the DM)*
- Decisions occur at a single *Decision Point*
- The DM's goals and priorities are encoded in the *DM goal tree* and are known before the *Decision Point*
- Decisions are made independently
- The *Decision Variables* are known at the *Decision Point*
- Several *Plans* are known at the *Decision Point,* only one of which (the *Choice)* is executed
- *Outcomes* can be predicted from the *Decision Variables* and *Plan*
- Goals are *Achieved* or *Failed* when a plan is executed

B. Pedagogical Goals

With these simplifications, we can define learning decision-making as learning 5 things:

- What goals should be considered
- What plans are available
- What plans are appropriate for which goals
- Which goals are less important and should be sacrificed
- Which goals are more important and should be pursued

C. Pedagogical Techniques

Our Pedagogical Techniques can be easily derived from our Pedagogical Goals:

- Introduce new goals (e.g. 'establish good relations with a country')
- Introduce that a goal is a subgoal of another goal or goals (e.g. 'give aid to a country' is a subgoal of 'acquire good relations')
- Introduce new plans (e.g. 'covert action')
- Introduce that a plan is instrumental to accomplishing a goal or goals (e.g 'covert action' is instrumental to 'establish control over another government')
- Introduce a relative importance relationship between 2 goals (e.g. 'establish good relations' is more important than 'communicate military superiority')

These concepts must be acquired by the student *through experience* because our aim is *not* to establish a given ideology for the student, but to have them discover their own ideology. We only want to make that ideology more consistent and more detailed as the student experiences more and more cases.

II. The Student Model

For our Student Model representation we use Carbonell's goal-trees [CARB79]. We view goal trees as compiled hindsight: they summarize what decision variables were examined and what priorities were attached to achieving, avoiding, or preserving various states of the world in past cases. More formally, a goal tree (GT) is an abstract data structure that is a set of goals and two associated graphs: a goal/subgoal graph (GS-graph) and a relative-importance graph (RI-graph). The GS-graph is a directed acyclic graph with subgoal links (S-links) as arcs while the RI-graph is a directed acyclic graph with transitive relative-importance links (RI-links) as arcs. Both graphs share the same goals as nodes.

There are only 3 operations that can be performed on a GT, thereby limiting how the Student Goal Tree (SGT) can grow:

- Add a new goal
- Add an RMink between 2 goals
- Add a subgoal link between 2 goals

DECIDER assumes students become aware of goals when they pick a plan where those goals succeed or fail. It assumes they prioritize two goals (with an RMink) when they pick a plan that sacrifices one goal and pursues the other. Finally, it assumes they know a subgoal relationship when they pick a plan designed to achieve a subgoal of some goal in the *DM goal tree*.

III. Case Selection and Presentation

There are a large number of choices of which concepts to introduce at any one time. The system has 2 methods for making the process more effective and efficient: building a pedagogical memory organization and using heuristics for determining which Student Model expansion opportunities are most important.

A. Memory Organization

The main factor that differentiates historians, students of history, and teachers is their memory organization. An expert historian organizes cases by political issues, goals of participants, and other functionally-useful categories. Novices organize cases by extreme outcomes, emotional import, and personal relevance. Good teachers organize cases around their pedagogical goals so they can quickly index to a group of cases to illustrate a point.

We tried to capture aspects of both good teachers and good historians by organizing cases around pedagogically *and* functionally useful categories. We index cases on types of plans and the failures and successes of using those plans in different time eras (e.g. 'military attack' failures in the Reagan administration due to violation of 'limit civilian casualties').

B. Relative Importance Heuristics

Good teachers know how to introduce material appropriately so that students are not overwhelmed or bored. A system mirroring this expertise must know a priori what kinds of deficiencies in the Student Model are most important. For DECIDER, this knowledge is captured in 4 Relative Importance Heuristics:

- Introduce goals you think the student should find more important first
- Introduce a goal before introducing how that goal can be achieved
- Introduce goals before introducing their relative importance
- Introduce the relative importance of two goals before introducing how those goals can be achieved

History teachers often introduce points that they feel strongly about before points they feel undecided upon. Heuristic #1 biases the student to learning the ideology of the teacher. Heuristic #2 gives the student a chance to learn a whole range of issues before specializing. This is a special case of iower order factors before higher order factors' [COLL80]. Heuristic #3 gives the student a chance to learn the issues separately before learning how they interact. This is a form of the 'separate goals before goal interactions' heuristic found in some machine learning systems [SUSS75, HAMM86]. Heuristic #4 gives the student a chance to relate important issues before learning minor variants of them.

DECIDER picks goals and relationships to add to the SGT from the Possible Goal Tree (PGT), guided by its relative importance heuristics. The PGT is a set of goals to be introduced, with any number of suggested S-links and RI-links between these goals. The PGT is meant to model the teacher's ideological biases. If the PGT has its links completely specified, DECIDER will try to make the SGT look more and more like the PGT by giving a one-sided presentation of the cases. If the PGT has no links, DECIDER will try to show both sides of all issues. The system finds opportunities to expand the student model by comparing the SGT and the PGT and then uses those opportunities to index to cases.

C. Pedagogical Opportunities

There are 4 information sources that are consulted when assessing the pedagogically useful information in a case:

- The DM goal tree
- The Choice taken
- The list of Sacrificed Goals and Pursued Goals for the Choice taken
- The list of Achieved Goals and Failed Goals associated with each Actual Outcome of Choice

Based on these knowledge sources, the system can derive a set of Pedagogical Opportunities for a case and index it in the database. The system knows it can introduce a goal when it succeeds or fails, a plan when it is the Choice, and priorities between goals when two goals are active and the Choice sacrifices one goal and pursues another. Later, when the system needs to teach these concepts, it will retrieve the case according these indices.

D. Pedagogical Utility

The final 'teacher-initiative algorithm' tries to introduce the cases in the database that are most pedagogically useful by consulting the Student Goal Tree and Possible Goal Tree according to the Relative Importance Heuristics.

The algorithm consists of 4 rules run serially; the first 2 rules embellish relationships between existing goals, the next 2 rules introduce new goals. Each rule does 4 operations: (A) COLLECT goals in the SGT or PGT, (B) INDEX to cases based on those goals. If there are cases found by indexing, (C) CHOOSE the best case and (D) DISPLAY parts of the case relevant to the goals. If one rule doesn't INDEX to any cases, the next rule is tried. If all rules don't INDEX to any cases, there are no pedagogically-useful cases left.

1. Introduce New RI-links
A. COLLECT highest pairs of goals in GS-graph of the SGT with no RI-links between them.
B. INDEX to cases based on relative importance of those goals
C. CHOOSE case with greatest number of RI-Hnks used in deciding on Choice.
D. DISPLAY decision variables and outcomes relevant to both goals of the RI-links

2. Introduce New Sub goal Links
A. COLLECT highest goals in RI-graph of SGT with missing subgoal-of links to higher goals
B. INDEX based on Failure and Success of these goals
C. CHOOSE the case with the greatest number of subgoal-of links used in Choice.
D. DISPLAY only decision variables and outcomes relevant to the subgoal relationship.

3. Introduce Important New Goals
A. COLLECT highest unintroduced goals in the RI-graph of the PGT.
B. INDEX based on Failure and Success of these goals
C. CHOOSE the case with the greatest number of unintroduced goals
D. DISPLAY only decision variables and outcomes relevant to unintroduced goals

4. Introduce Next Most Specific Goals
A. COLLECT highest unintroduced goals in the GS-graph of the PGT.
B. INDEX based on Failure and Success of these goals
C. CHOOSE the case with the greatest number of unintroduced goals
D. DISPLAY only decision variables and outcomes relevant to those goals

In each round, the teacher-initiative algorithm finds and displays a case. The student then picks from the list of plan options and observes the outcome. Based on the student's choice, the SGT is updated. With each round, the SGT expands, the algorithm indexes on more specific goals, and the student must confront more difficult decision problems.

IV. Conclusion

The future of AI in education depends on making programs that understand the students they are teaching and tailor their presentation to the student's background. The ideal educational system will carefully monitor interactions with students and change the environment to make it more challenging without the student ever knowing they are being taught. When this happens, students will learn more than they would from teachers or textbooks, and AI systems will become an integral part of modern education.

Acknowledgements

Thanks to Drs. Moon, Etheridge, and Westerfield for their detailed knowledge of history, Gilles Bloch for ideas, code, and the multi-media interface, and Roger Schank, for his guidance, maxims, and teaching expertise.

References

[CARB79] Carbonell, J.G. Subjective Understanding: Computer Models of Belief Systems. Technical Report #150, Yale University Dept. of C.S., 1979

[COLL80] Collins, A. and Stevens, A.L. Goals and Strategies of Interactive Teachers. Technical Report #4345, Bolt Barnek and Newman, Inc., March 1980

[HAMM86J Hammond, K. Case-based Planning: An Integrated Theory. Technical Report #488, Yale University Dept. of C.S., October 1986

[NEUS86] Neustadt, R. and May, E. Thinking in Time. Harvard Univeristy Press, 1986

[SCHA86] Schank, R.C. Explanation Patterns. Lawrence Erlbaum Associates, Hillsdale, New Jersey, 1986.

[SUSS75] Sussman, G.J. A Computer Model of Skill Acquisition. American Elsevier, New York, 1975.