Toward a Model of Student Education in Microworlds[1]

Cristina Conati
Intelligent Systems Program
University of Pittsburgh
Pittsburgh, PA 15213
conati@cs.cmu.edu

Jill Fain Lehman
School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213-3890
jef@cs.cmu.edu

Abstract

Microworlds are educational environments intended to support the student in the active exploration of a subject-matter domain. We present preliminary work whose goal is to attain a better understanding of the educational effectiveness of microworlds through an examination of the learning processes that they exploit. The learning processes are made explicit within a computational model of the interaction between a student and a microworld for simple electrostatics. We focus, in particular, on the implementation of an episodic memory mechanism that gives insight into the processes involved in learning from incorrect behavior.

Introduction

In the educational community a large amount of enthusiasism has been engendered for the highly interactive microworld which behaves according to the laws and constraints of some subject-matter domain and permits the student to experience the nature of that domain through free or guided exploration. This enthusiasm is predicated largely on the belief that the student's activities in the microworld produce or foster education about the subject-matter domain. There is essentially no empirical evidence of substantive learning from mere interaction with a microworld, i.e., without an associated curriculum systematically designed to enable the learner to develop a set of well-defined skills (Carver, 1986; White88, 1988). Moreover, even when such a curriculum exists and experimental effect is evident, there is still little understanding of the processes through which such education occurs or through which the microworld experiences make their contribution. Thus, the goal of our work is to produce a process model account of education in microworlds. The research described here is preliminary, demonstrating a problem-solving organization that models one student's behavior in microworld interactions.

The Task: Electric Field Hockey

Our subject-matter domain is the simple electrostatics of discrete charged particles with unit mass. Our interactive microworld is Electric Field Hockey (hereafter, simply Hockey) (Chabay & Sherwood, 1989; Sherwood & Chabay, 1991). With respect to electrostatics, Hockey involves determining the trajectory of a unit-charge particle (the puck) from a given initial position, around a given set of obstacles, to a fixed final position (the net) by placing a number of additional unit-charge particles (see Figure 1). The motion of the puck along its trajectory is shown when the GO button is clicked on; the puck's velocity is recorded statically by the spacing of the dots.

The student has a limited number of options in controlling exploration in Hockey: velocity can be replaced with a vector representation of acceleration, and any of six levels of increasingly difficult play may be chosen. Increasing difficulty is achieved by variation along four dimensions: the number and configuration of obstacles, the starting position of the puck with respect to the boundaries of the playing field and the obstacles, the existence of other, unmoveable charged particles in the field, and the number of charged particles available to maneuver the puck into the net.

[1] This research was sponsored by the Markle Foundation. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Markle Foundation.

We have collected a nearly complete video record of two students' education in the domain. The completeness of this record was made possible by the creation of a new curriculum at Carnegie Mellon University that teaches students a purely qualitative model of electrical and magnetic phenomena as a theoretical basis for later, standard quantitative material. Our two students worked together through the notebook of desk-top experiments and questions that constitute the course material, then interacted separately with Hockey as a class assignment. The purpose of the assignment was to provide "an excellent physical feel for the extreme distance dependence of the Coulomb interaction" (Sherwood & Chabay, 1991). It is apparent from the video record that both students had at least a rudimentary understanding of the requisite physics concepts prior to their game interactions.

The Models: EFH-Soar

Our models are implemented within the Soar architecture (Laird, Newell & Rosenbloom, 1987; Newell, 1990), although due to space our description will be given at a functional level rather than in terms of problem spaces, operators and impasses. There are currently two models: EFH-Soar, a model based purely on a task analysis of Hockey which contains enough game knowledge and physics knowledge to play without error, and EFH-Soar2, a model of Student 2 based on her protocol. Both models possess the basic functionality for interacting with Hockey as a piece of external software, for re-encoding Hockey's spatial representation in symbolic terms, for reasoning about the placement of charges to create a trajectory that manuevers around or through an obstacle, for interpreting and evaluating a trajectory in terms of the forces that determine it, and for modifying the position of previously placed charges on the basis of the resulting trajectory. Both models play only the first three levels of the game. We concentrate here on EFH-Soar2 (hereafter, simply EFH2), which acquires some of the knowledge hand-coded into EFH-Soar.

Figure 1 shows the steps EFH2 takes in playing the first level of the game. Like Student 2, the system works its way from left to right across the screen, interweaving the positioning of each charge with an examination of the microworld's feedback on the effect of the placement in achieving the goal or subgoal.

One of the most important aspects of interaction with Hockey is the processing of spatial information displayed on the screen. The essential feedback from Hockey takes the form of continuous trajectories of the puck as it moves under the influence of the fixed charges. Although all the information available from the Hockey software is in terms of x-y coordinates, it is clear that students do not reason at that level. Based on prior work by Ward (Ward, 1991), EFH2 uses a highly approximate, qualitative spatial model which depends on continuous re-perception of the actual external world (the source of high-quality knowledge) to update and correct the low-quality internal representation. Figure 1, (a) shows a simplified version of the encoding of the representation of the screen available from Hockey into the features of EFH2's spatial model. These features include: a focus of attention mechanism, variable access to details about an object's properties depending upon whether or not it is in focus, and a qualitative representation of relative positions among objects. The qualitative representation encodes broad spatial relations among objects not in focus (e.g. left, right, above, below, and their compounds). When objects are brought into focus, the representation is augmented with distance and angular relations in terms of the system's internal resolution factor.[2] When it begins to play Level 1, the puck is the focus of attention.

Reasoning from the initial spatial model (a), EFH2 proposes clearing the obstacle as the first subgoal for sending the puck into the net. The process to accomplish the subgoal is divided into two phases: spatial reasoning (panel (b)) and physics reasoning (panel (c)). The spatial reasoning phase defines the trajectory that the puck must follow to accomplish the current subgoal. In the example, EFH2 performs shifts of focus until its focus of attention includes both the obstacle and the corner chosen as a clearing point. By bringing the corner into focus, the spatial model is augmented with a relative distance and direction between the corner and the puck. The appropriate path is then defined as the direction between the puck and the corner slightly decremented to allow the puck to clear the obstacle.

Once a path has been chosen, the system uses its knowledge of physics to find a strategy for moving the puck along the path (panel (c)). During this phase EFH2 decides, for example, whether to use an attracting or repelling force, how many charges to place, and where to place them. The strategy is then expressed in

[2]Although we believe the internal resolution can vary in grain size, this has not yet been implemented. The current resolution is 10 for distance and 32 for angular relation, laid out in the Cartesian plane.

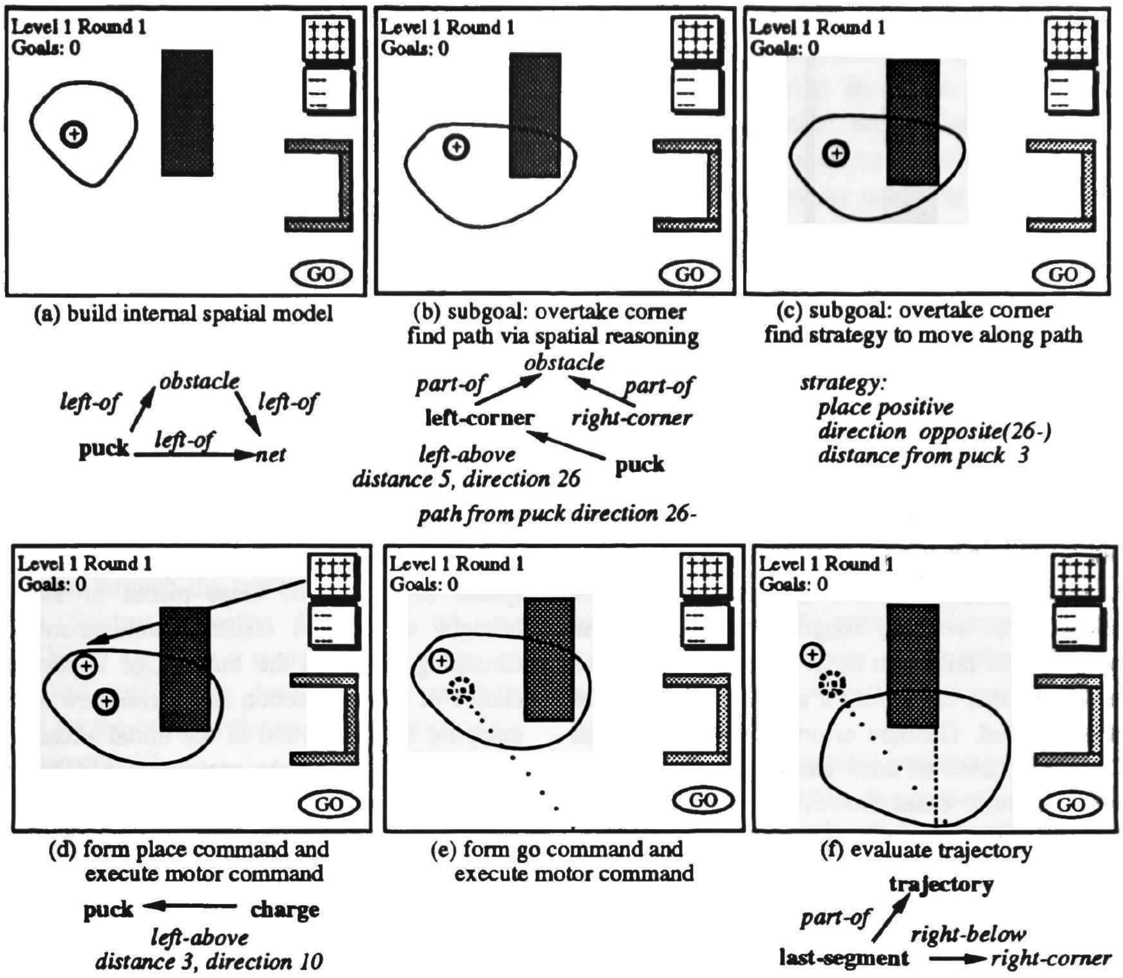

Figure 1: One iteration through EFH2's basic problem-solving loop
(cumulative state information in italics, focus objects in boldface).

terms of game actions — here, the placement of a positive charge with a particular angular relation and distance from the puck.

Once the action has been sent to and performed by Hockey, the new screen is perceived to modify the internal spatial model (panel d), adding the newly placed charge to the focus objects. The system is now ready to issue the GO command, which causes Hockey to display the motion of the puck subject to the electrostatic force of the placed charge (panel e). EFH2 encodes in the spatial model a qualitative representation of the trajectory, enabling the system to evaluate the result of the place action. To accomplish the evaluation, the system shifts its focus to the trajectory and the lower part of the obstacle (panel f). The position of the last part of the trajectory with respect to the obstacle's lower corners shows that the subtask of clearing the obstacle has been accomplished.

Since there are no other obstacles, EFH2 proposes the new subgoal of sending the puck directly into the net. The spatial and physics reasoning phases are repeated. This time the system moves its focus to define a path that starts inside the net and runs perpendicular to the trajectory. The strategy to move the puck along this path is, again, to use a repelling charge. Since the first placed charge is not close to the current focused area, its effect is ignored and the new charge is placed along the defined direction, close to the trajectory. The result of this second placement eventually sends the puck into the net.

The model's performance, like that of our student, is not flawless. In particular, since EFH2 reasons with a spatial representation which has a grain size much larger than the actual number of unique locations in the field represented inside Hockey, the system's intended location corresponds to a region inside the microworld.

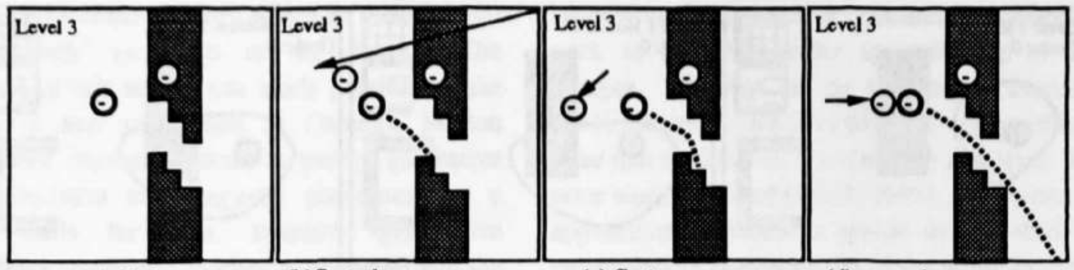

(a) initial state (b)first placement + go (c) first move + go (d) second move + go

Figure 2: An episode of learning from Student 2's protocol.

As a result, positioning of a charge can be correct with respect to physics knowledge but still fail to achieve the subgoal. In case of such a failure, the trajectory evaluation phase annotates the strategy with a description of the failure encountered (e.g. the puck is against the obstacle). When a failure occurs, the system proposes an error recovery subgoal that analyzes the description of the failure in order to find an action to eliminate it. In this case, only a simple adjustment of location is needed. Genuine errors of knowledge that result in failure cause an error recovery subgoal in the same way. In those cases, however, a failure triggers physics knowledge in the subspace that the student failed to bring to bear during the original problem solving process, i.e. knowledge that the student did not consider relevant to the current situation, but that she must re-evaluate to discover a solution (an example will be described in the next session).

What the Model Learns

The term *learning* refers to a change in long-term memory that results in a change in behavior. In ascertaining whether learning occurs in microworld interactions, the advantage of a process model over, for example, an empirical study of student test performance stems from the ability to actually examine changes in the model's long-term memory. In Soar, all changes to long-term memory, all episodes of learning, have at their basis the architectural mechanism of chunking (Newell, 1990). Chunking is a uniform, ubiquitous mechanism that adds knowledge to long-term memory throughout problem solving. It is a compilational, or integrational, technique that, in its simplest form, combines existing knowledge into new associations.

In the situation in Figure 1, for example, EFH2 begins with separate pieces of knowledge for deciding to overtake the corner first, for reasoning spatially about the path the puck should follow, for deciding whether to attract or repel the puck along the path, and for mapping the game concept of placing a charge into motor commands for interacting with Hockey. Between the time EFH2 perceives the initial situation (panel (a)) and the time it performs a motor action in the world (panel (d)), each of these pieces of knowledge is brought to bear in discrete problem-solving steps. Chunking takes all the knowledge implicated in the choice of a motor action and forms new associations, mapping the perception of the initial situation directly into a proposal for the motor action. The next time EFH2 is in a similar situation, it will not problem solve — it will simply recognize the situation and act in the way its prior experience dictates. In this way, the model becomes a more practiced Hockey player, able to find solutions to the problems in the game more quickly over time. Note that the quality of the solution does not improve via simple chunking, just the speed of attaining the solution. This process is not one by which the model, or the student, learns more physics, but one by which they continually recast the physics knowledge they do have into game-relevant form.

If the student has both perfect knowledge of physics and perfect knowledge of the game, then speed-up learning is the only type of learning that will occur. Of course, most students have neither. Consider the episode from Student 2's protocol shown in Figure 2. In panel (b), her initial placement of the repelling charge is inadequate to overcome the effects of the charge glued to the obstacle. Her next move seems to be an attempt to compensate for the direction of the force produced by the glued charge but is still too far away to produce the desired outcome. Finally, in panel (d), the forces combine correctly to achieve her subgoal of pushing the puck through the hole.

What did Student 2 learn from these interactions? We know she learned something, because when faced with an analogous situation at Level 4 of the game, she does not repeat this series of actions but places the initial repelling charge directly at the position analogous to the one in panel (d). Transfer occurs despite the fact that the situation in Level 4 differs in significant ways from the situation in Level 3 (the glued particle has a charge opposite that of the puck, the glued particle is on the lower of the two obstacles, and there is another obstacle between the puck and the net). Because the video record prior to her interaction with Hockey indicates that she has an abstract understanding of the relationship between distance and force, we conjecture that in this episode her abstract understanding is recast to fit the scale of the game (this conjecture is further supported by her verbal protocol at Level 4 when she says (her emphasis), "so we're gonna need a negative charge <u>close</u> that will hopefully umm take away some of the effects of that positive charge"). In other words, through the sequence of moves in (b) through (d), Student 2 makes the notion of *near* concrete. This is certainly a component in achieving the pedagogical goal of the designer ("an excellent physical feel for the extreme distance dependence of the Coulomb interaction").

It is straightforward to show that the model outlined above is inadequate to produce the transfer shown by Student 2. In the initial situation on Level 3, EFH2, like the student, ignores the effects of the glued particle, placing a charge aligned with a path through the opening. It chooses to place the charge at a distance from the puck that was adequate to achieve her goal at Level 1. Thus, chunks are formed that map the situation in panel (a) directly into that placement (call them PlaceChunks). EFH2 then issues the GO command, evaluates the resulting trajectory, and proposes error recovery (producing GoChunks, EvalChunks etc). The first error recovery action is to adjust for the direction of the force. So Move1Chunks are built that map the situation in panel (b) (the puck and charge as positioned and the current type of failure) directly into an action that moves the charge to its position in (c). After another GO command results in the failure in (c), a second error recovery moves the charge closer to the puck. Thus, Move2Chunks are built that map the situation in panel (c) (the puck and charge as positioned and the corresponding failure) to an action that moves the charge to its position in (d). When faced with a situation similar to (a) in the future, these chunks will be triggered and the system we have described, unlike Student 2, will go through the same sequence of motor actions seen in (b) through (d), rather than directly to the single, correct placement.

Our preliminary solution to this dilemma lies in giving EFH2 an episodic memory that is created and used by processes that allow the system to reconstruct past problem solving in order to avoid repeating mistakes. Assimilation, recognition, and recall are the three processes that must be coordinated (Lewis, 1992). The assimilation process results in chunks that encode situations the system has been in, where a situation is defined by the objects in focus, the current subgoal, the proposed action, and when available, the outcome of the action (success or failure). During the sequence in Figure 2, for example, assimilation chunks are formed that record the puck, glued charge and obstacles in the relation in (a) and the proposed placement action. Other assimilation chunks record that the outcome of this action was a particular kind of failure. Others record the configuration of the puck, placed charge, glued charge and obstacles in (b), the failure of the place action, and the proposed move. Still others encode the failure of the move action in (c). This assimilation process is repeated throughout the action sequence.

The episodic memories are used by the recognition process to notice when the system is about to do something it has already tried. Consider the next time EFH2 is in a situation like (a). As noted above, the PlaceChunks will fire, proposing the placement action. Now, however, assimilation chunks will fire signaling that this is a familiar situation. Familiarity invokes the recall process which tries to remember what the outcome of the proposed action was. Recall also uses assimilation chunks; it tries each of a small number of augmentations to the system's internal state waiting for recognition to occur. In other words, the system asks itself the question, "Do I have a memory in which this placement succeeded?" "One in which it failed because of this or that reason?"

Once a failure has been recalled, EFH2 must reconstruct its prior problem solving to find the actions that led to success. To accomplish this, the system imagines the spatial model that would result from applying the currently proposed action. "Imagining" means simply that a new charge and its spatial attributes are added to the spatial model without actually being present on the Hockey screen. Using this simulated spatial model, EFH2 tries the error recovery subgoal, triggering Move1Chunks that suggest the same move action performed originally in (b). Now the process repeats: the system recalls the outcome of the first move by generating failure types until it triggers an assimilation chunk. Then it imagines the outcome of the move in the spatial model and uses an error recovery subgoal to trigger the Move2Chunks. Finally, it recalls/recognizes that this move led to success. The system uses the distance and direction of the charged particle with respect to the puck in (c) to propose an alternative to the place action suggested by the original PlaceChunks. Proposing an alternative leads to Place2Chunks whose action is preferred to the original.

The chain of recalls and recognitions outlined above is performed the second time EFH2 is in situation (a). By reconstructing its past problem solving, the system circumvents the automatic replay of motor actions which simple chunking would have produced. In terms of observable behavior, EFH2 simply places the first charged particle correctly. Moreover, in future situations like (a), reconstruction via recall is not necessary — Place2Chunks map (a) into the correct placement directly.

Conclusions

The system outlined above is both too strong and too weak. It is too strong because the mechanism for assimilation, recognition, and recall is capable of reconstructing chains of memories of arbitrary length and involving the imagination of an arbitrary number of changes to the spatial model. The mechanism is too weak because it cannot, by itself, result in the transfer shown by Student 2 at Level 4. The flaws in the mechanism are the result of two simplifying assumptions: first, that assimilation captures all and only the necessary details of the situation; and second, that the assimilation, recognition, and recall processes are automatic, i.e. unmediated by other, more proactive processes.

The modified form of the mechanism we envision relaxes these assumptions. In it, each episodic memory is constructed by a potentially inaccurate assimilation process that can be automatic or mediated by reasoning. If assimilation is no longer guaranteed to capture the complete and correct details of the situation, then both overspecific and overgeneral episodes are possible. The former may result in the breakdown of the recognition and recall processes during reconstruction while the latter may serve as a foundation for single instance generalization. We also conjecture that when recall occurs upon re-encountering a familiar situation, transfer appears analogical, but recall employed at the moment of success or failure leads to transfer that appears to be a result of self-explanation (Chi, 1989).

Assimilation and recall are, by themselves, inadequate to explain or predict the full range of behavior we see in our students. Still, our work has led us to consider that conventional learning mechanisms — analogy and self-explanation, for example — may be understandable as variations of a single underlying process. These conventional learning mechanisms have been clearly implicated in successful education. Thus, our future work is directed at demonstrating how assimilation and recall, as component processes underlying conventional learning mechanisms, play a crucial role in acquiring subject-matter knowledge.

References

Carver, S. M. and Klahr, D. 1986. Assessing children's LOGO debugging skills with a formal model. *Journal of Educational Computing Research*, 2(4), 487-525.

Chabay, R. and Sherwood, B. 1989. *The Electricity Project and the cT Programming Language* (Tech. Rep. No. 89-10). Center for Design of Educational Computing, Carnegie Mellon University.

Chi, M. T. H., Bassock, M., Lewis, M., Reimann, P., and Glaser R. 1989. Self-explanation: How students study and use examples in learning to solve problems. *Cognitive Science*, *13*, 145-182.

Laird, J. E., Newell, A., & Rosenbloom, P. S. 1987. Soar: An architecture for general intelligence. *Artificial Intelligence*, *33*, 1-64.

Lewis, R. 1992. A Technical Thought on Soar and Short-term Memory. Slides from Soar Workshop 10, University of Michigan at Ann Arbor, February 1992.

Newell, A. 1990. *Unified Theories of Cognition.* Cambridge, Massachusetts: Harvard Press.

Sherwood, B. and Chabay, R. 1991. Electrical Interactions and the Atomic Structure of Matter: Adding Qualitative Reasoning to a Calculus-Based Electricity and Magnetism Course. *Proceedings of the NATO Advanced Research Workshop on Learning Electricity or Electronics with Advanced Educational Technology.* Berlin, Germany: Springer-Verlag.

Ward, B. 1991. *ET-Soar: Toward an ITS for Theory-based Representations.* Doctoral dissertation, School of Computer Science, Carnegie Mellon University,

White, B. Y. 1988. *ThinkerTools: Causal Models, Conceptual Change and Science Education* (Tech. Rep. No. 6873). BBN Laboratories, Inc.