Romoser, M. E., Woolf, B. P., Bergeron, D., & Fisher, D. "The Mental Rotation Tutors: A Flexible, Computer-based Tutoring Model for Intelligent Problem Selection." Conference of Human Factors & Ergonomics Society, New Orleans, LA. 2004.

The Mental Rotation Tutors: A Flexible, Computer-based Tutoring Model for Intelligent Problem Selection

Matthew R. E. Romoser[1], Beverly P. Woolf[2], Dan Bergeron[2], and Donald L. Fisher[1]
[1] Department of Mechanical & Industrial Engineering
[2] Computer Science Department
University of Massachusetts at Amherst
Amherst, MA 01003

1.0 Introduction

Intelligent Tutoring Systems (ITS) are altering the landscape of how course instruction, workplace training, and curriculum development are conducted. These systems vary in scale, scope and size. An ITS can be as simple as a twenty minute instruction module, or as complex and ambitious as a related series of semester-long courses. The ITS designer can choose from among a large number of development platforms which might include design software such as Macromedia Director or Flash, object-oriented programming languages and several different delivery modes such as CD-ROM, web-based or HTML systems or some combination of all of them. The instructor can choose from an equally large array of teaching strategies (i.e., apprenticeship, mastery learning, inquiry learning, or simulation) and ITS architectures (i.e., hierarchical, linear, exploratory, or node-based). As a result, research in ITS leads to all kinds of different approaches to computer-based intelligent tutoring.

This poses a challenging and potentially intimidating predicament for the everyday practitioner – designers and instructors who wish to find the best way to move their courses online but may not have a lot of experience in the field of ITS. In their "rush to the web", many instructors end up creating online lectures which function like web-based textbooks by saving their lecture notes in HTML format and perhaps adding a few multimedia bells and whistles. The student ultimately ends up reading these online and "turns the page" by hitting the *NEXT PAGE* button. While inexpensive to implement, these button-driven solutions or *frame-based* teaching systems (Woolf, 2002), deny the student an active learning experience. Constructing and testing hypotheses, receiving targeted feedback, and adjusting problem solving strategies during instruction are critical aspects of an active learning experience that are denied the student in frame-based learning systems (Knowles, et al., 1998; Woolf, 2002). Intelligent tutoring strategies dictate that the tutor must provide active problem solving activities that give immediate, customized feedback. The tutor should also dynamically choose problems that are appropriate to the student's current skill level and measure the student's overall progress against a predefined mastery criterion. The objective of this paper is to introduce a flexible model for intelligent problem selection and then to provide an example of its implementation. Here we describe the Mental Rotation tutors and demonstrate how the model is especially powerful when applied to domains such as mental rotation in which it is sometimes ambiguous how a student arrived at the incorrect answer. The novel model utilized in the tutors represents a marriage of Mastery Learning theory and our intelligent problem selection model.

Figure 1 contains a conceptual flowchart of a tutor using a Mastery Learning instruction strategy (Woolf, 2002, Siemer & Angelides 1998). The Mental Rotation tutors utilized this type of strategy to track and coordinate the student's progress with the material.

With Mastery Learning, course material is broken up into relatively small learning units. That way, the material in each unit can be matched with the instructional method that best suits it. When mistakes are made, appropriate feedback is given and the student is given more diagnostic problems to solve. This process is repeated until the student has mastered the material in the unit. The time spent in the learning situation varies with the student's aptitude (Woolf, 2002).

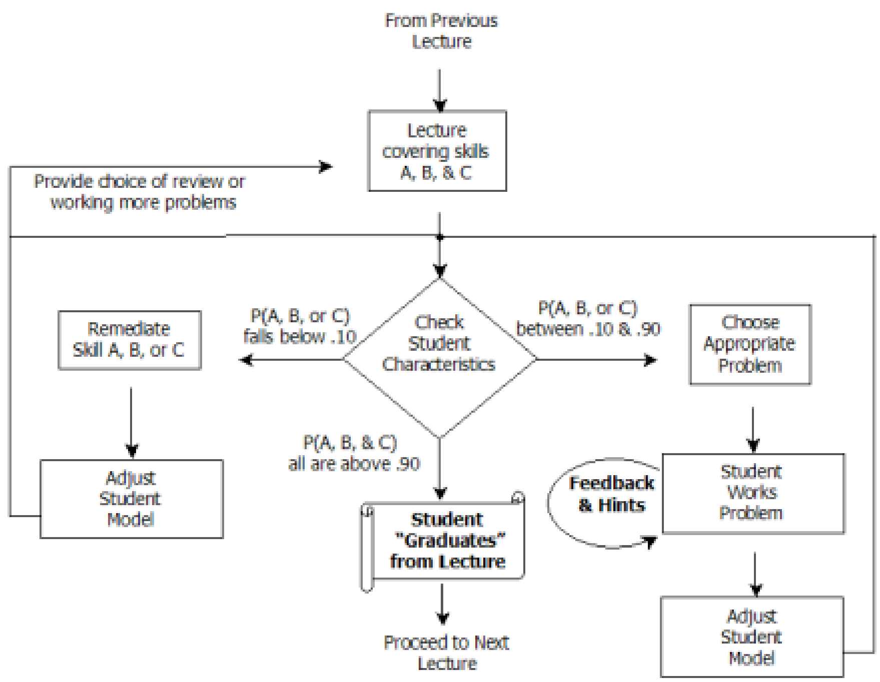

Figure 1 – Conceptual model for a mastery learning ITS strategy

2.0 The Mental Rotation Tutors

The Mental Rotation Tutors were designed for use in concert with 3-D modeling and design courses at the University of Massachusetts to help students better visualize the parts that they are designing. Previous research in our lab has revealed that students struggle with both inferring a rotation from a "before" and "after" view of an object and applying a known rotation to an object and selecting the appropriate final view (Woolf, et al., 2003; Romoser & Fisher, 2001; Romoser, et al., 2001). The tutor breaks down the process of mental rotation into two functional areas. Each functional area of the tutor was developed and evaluated separately.

In the first functional area (Phase I), the student must demonstrate an ability to infer what rotation is taking place to rotate an object from an initial isometric view to some secondary view. Figure 2 contains a screen capture of a typical Phase I problem. The student would respond by specifying the axis of rotation (red, blue, or gold), direction (CW, or CCW) and the number of degrees (90, 180, or 270) through selecting the appropriate icons and clicking "submit". The student could then select as many additional rotations as he or she felt were necessary to rotate the object from the initial to the final view. Students are given up to four tries to answer a problem correctly. With each incorrect response, feedback becomes more specific and targeted to the specific skill set represented by the problem. Figure 3 shows a few examples of the four levels of graduated hints that a student might receive.

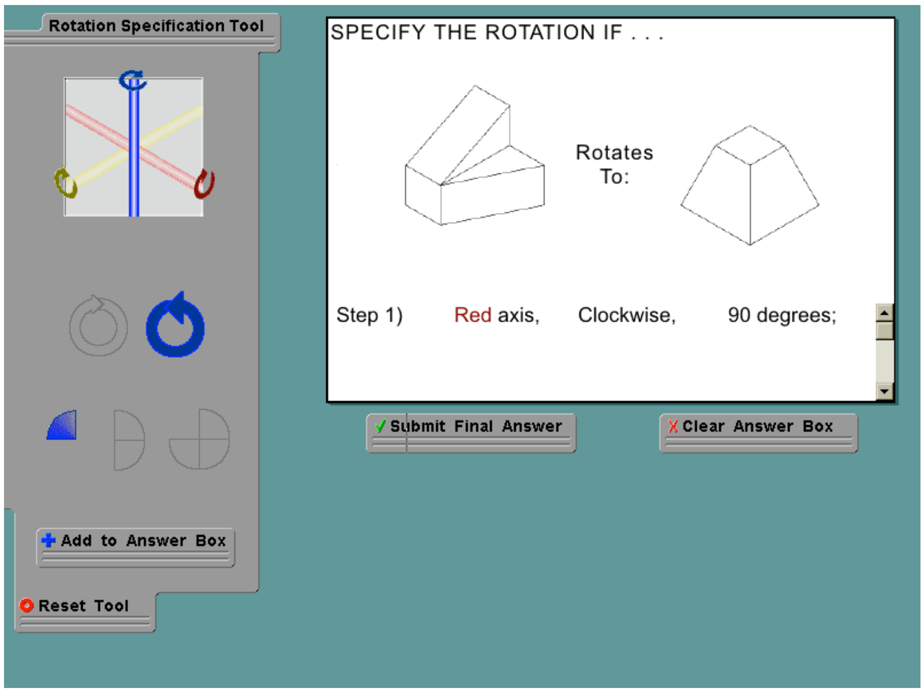

Figure 2 – Screen capture of Phase I tutor (infer rotation). Student specifies rotation(s) that will result in the final view on the right by using the graphical "rotation specification tool" – selecting axis, direction, then number of degrees, then adding that to the list of steps. (Red clockwise 90^0, blue counterclockwise 90^0 is one possible correct sequence.)

In the second functional area (Phase II) of the tutor, the student must demonstrate an ability to apply a known rotation to an object. Figure 4 contains a screen capture of a typical Phase II problem. The student is provided with an initial view of an object, one or two rotation descriptions, and a set of possible final views. The student responds by clicking on the final view he feels will result after the rotations are applied to the object. Again, students are given up to four tries to answer a problem correctly, with the feedback being adjusted appropriately at each phase.

The first step in the development process was to identify the skill sets. The tutors contained approximately forty objects, each with twenty-four unique rotations that could be executed. As a result, we had a library of approximately 960 problems to choose from. Each problem was rated for difficulty on a four-point scale on seven different parameters. The first four parameters were related to object complexity: faces, edges, protrusions, and notches. The last three parameters were related to rotation complexity:

axis (if it was aligned with the principle axis of the object), degrees of rotation, and direction.

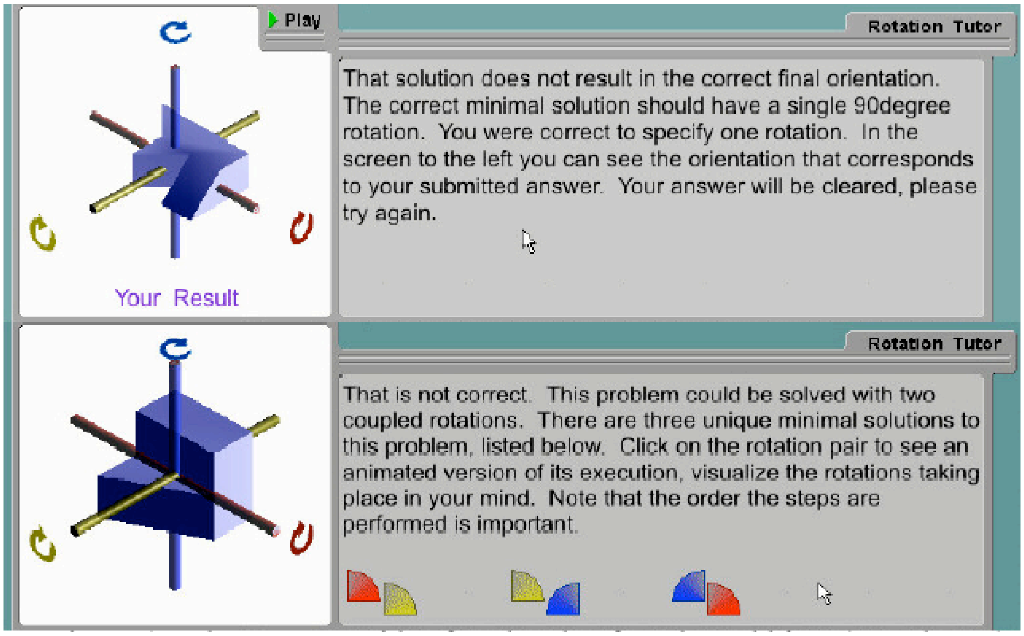

Figure 3 - The tutor provides four levels of graduated hints (two shown) including an animated version of the participant's proposed rotation steps ("play" button, left) as well as alternative correct solutions.

A student's progress with these skills was tracked and problems selected using our variation of a probabilistic strategy first introduced by Shute (1995). With this probabilistic model, a probability $P(S_i)$ is assigned to each skill *i* representing the model's confidence that a student has mastered it. Initially, $P(S_i)$ is set equal to .50. When a problem is completed, depending upon the number of hints required to arrive at the correct answer, the probability of skill mastery is increased in roughly inverse proportion to the number of hints required to arrive at the correct answer. A student is said to have "mastered" a skill if the probability is greater than some mastery criterion, $P(S_m)$, set by the domain expert who created the tutor. One of our significant contributions to this model was the introduction of a remediation criterion, $P(S_r)$. If $P(S_i)$ falls below $P(S_r)$, then the student is redirected to a special lesson dedicated to the development of that specific skill. These probabilities made up the student model portion of our ITS. The methodology is demonstrated in Figure 5.

In our case, since we were dealing with multiple skills per problem, the difficulty rating for each skill was used as a weighting factor in adjusting the probability. Hence, if the student answered correctly on the first try, then the probability increased more for those skills with a high difficulty rating than for those with a low (or no) difficulty rating. These new probabilities were then used by an expert system, which we referred to as the *domain reasoner,* which selected the next problem. The $P(S_i)$ for the four object complexity skills and the three rotation complexity skills were each averaged. This process repeated until these two average probabilities exceeded .90 (our graduation criterion). The result of this approach was that students received problems appropriate to their progress and spent as much time as was necessary with the material to learn all of the skills. For instance, if they had mastered the "faces" skill, but were struggling with "degrees of rotation", then the next problem selected would have a difficult face

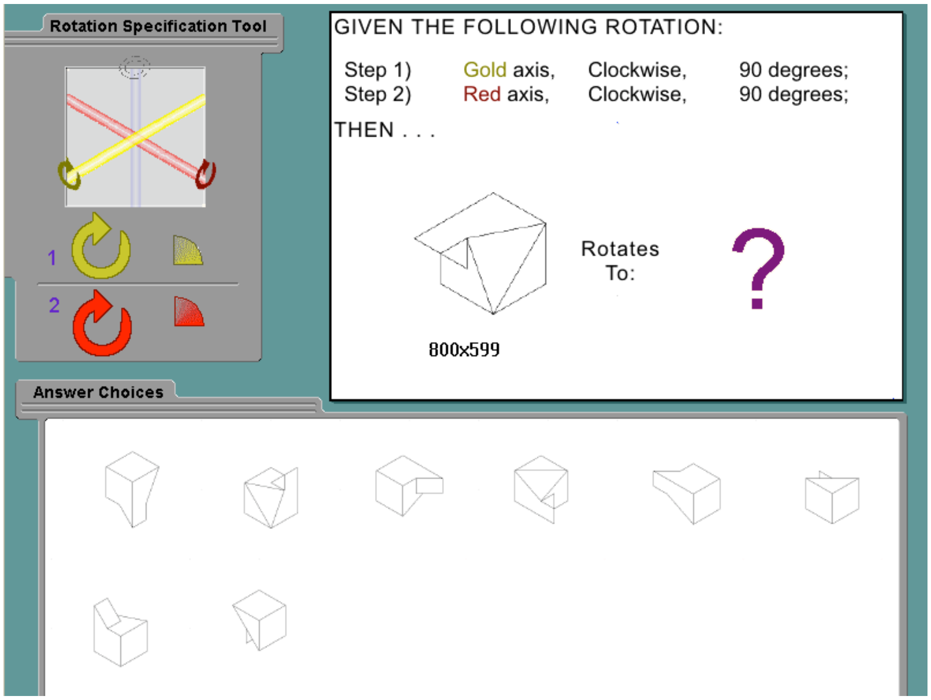

Figure 4 – Screen capture of Phase II tutor (apply rotation). The rotation to be executed is specified in both graphical (left) and text (top) format. The student selects the correct final view from the views at bottom. (The third picture from the left in the top row is the correct answer.)

arrangement (i.e., several non-normal to the major axes), but the assigned rotation would be simple (i.e., single 90 degree rotation). Another positive result of this approach was that the students who learned quickly graduated from the lesson with fewer problems than those who struggled. Therefore, the tutor "got out of the way" of the more advanced students who had demonstrated proficiency with the skills and prevented boredom from setting in (Woolf, et al., 2003).

3.0 Evaluation

The two sections of the mental rotation tutor were developed several months apart and each was evaluated separately. In both phases (Phase I – infer the rotation; Phase II – apply the rotation), each tutor was evaluated using twenty-eight participants from the schools of engineering and science for a total of fifty-six students. Participants were given the Shepard & Metzler's M.R.T. Test of Spatial Ability (MRT; Shepard & Metzler, 1971) as a screen for spatial ability (40 point max score). As a pretest, participants were randomly assigned to receive either the odd or even numbered problems from the Purdue Spatial Visualization Test (PSVT; Guay, 1980)(60 point max score). The top ten participants on the MRT were classified as "high-spatial" and the ten lowest as "low-spatial". The remaining participants were classified as "neutral". In each phase, after interacting with the tutor, the students were given the remaining problems from the PSVT that they did not receive as a pretest and took another MRT test of similar difficulty. There were also twenty control group participants who received no tutor, but took the pre- and post-tests one week apart from each other.

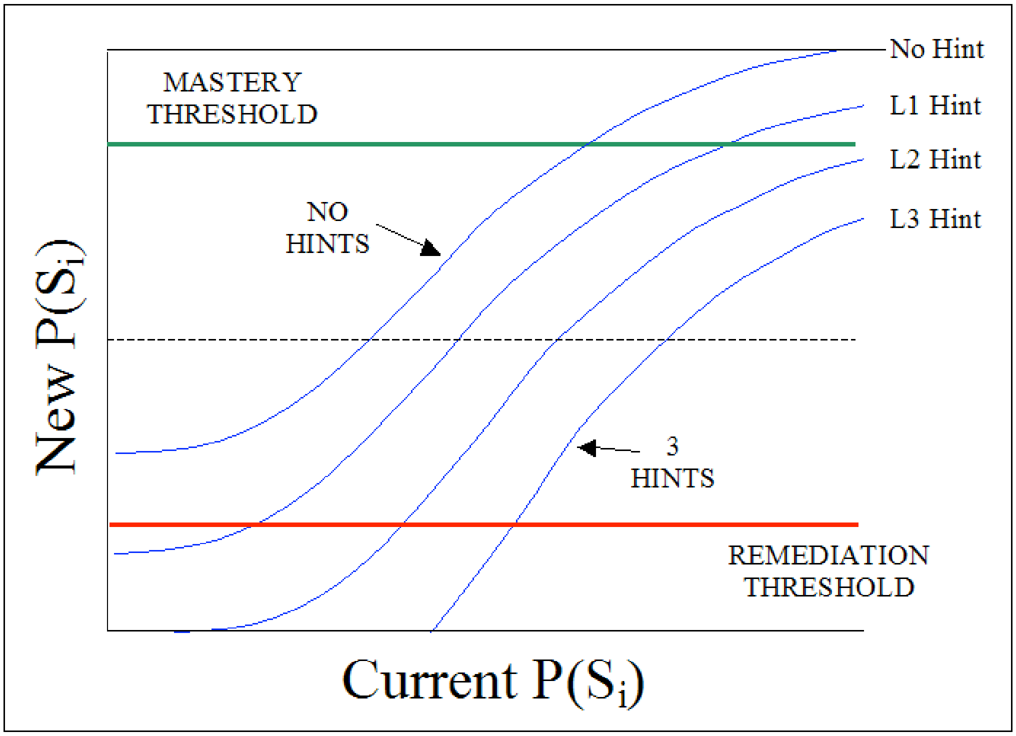

Figure 5 – Probability curves for adjusting P(S_i) – based upon Shute (1995).

The most interesting results came from comparing the high and low-spatial students. A paired-sample t-test was used to determine significant score increases. In Phase I (inferring the rotation), high-spatials increased their MRT score by only 1.10 points (35.0 to 36.1) while low-spatials increased their score by 8.4 points (13.0 to 21.4). The low-spatial increase was statistically significant, $t_{(9)} = -3.934$, $p < .005$. There were also significant increases on PSVT scores for both the high-spatials (41.5 pre-test; 51.5 post-test; 10.0 point change), $t_{(9)} = -2.739$, $p < .05$, and low-spatials (37.8 pre-test; 45.0 post-test; 7.2 point change), $t_{(9)} = -2.714$, $p < .05$. The high-spatials also completed significantly fewer problems on average than the low-spatials (9.7 versus 17.9, respectively), $F_{(1,18)} = 6.717$, $p < .01$.

In Phase II (applying the rotation), the increase in average MRT was 2.6 for high-spatials (32.3 to 34.8) and 8.0 for low-spatials (11.7 to 19.7). The increase for both low ($t_{(9)} = -7.303$, $p < .001$) and high-spatials ($t_{(9)} = -4.333$, $p < .01$) was significant. Unlike in Phase I, the lower spatials had a much larger average PSVT increase of 14.0 points (32.5 to 46.5) than high-spatials who increased by an average of 4.3 points (49.3 to 53.5). Only the increase for low-spatials was statistically significant, $t_{(9)} = -7.878$, $p < .001$.

Again, the high-spatials required significantly fewer problems to meet the mastery criterion than did low-spatials (9.4 versus 19.7, respectively), $F_{(1,18)} = 10.75$, $p < .005$.

The control group had smaller gains overall than participants in both Phases I and II. On the MRT, control group participants had an average increase of only 2.7 (15.0 to 17.7) for low-spatials and 2.1 (32.1 to 34.2) for high-spatials. On the PSVT, the average increase was 2.5 (30.8) for low-spatials and 4.3 (49.3 to 53.6) for high-spatials. Not one of these increases was statistically significant.

4.0 Discussion

Clearly, it was the low-spatial students who derived the most benefit from interacting with the tutors. This was to be expected. The tutors were designed to evaluate the student and get "out of the way" as quickly as possible if the student proved to be proficient with the skills early on. In the case of the high-spatials, the tutor's role shifted from *teaching* to *evaluating* and graduated the participants quickly. The fact that high-spatials required roughly half as many problems as low-spatials demonstrated that the tutors performed as designed in that regard. The control group data demonstrated that there were some practice effects from pre- to post-test. However, this effect was much smaller than the gains realized by the low-spatials who used the tutor. Gains for high-spatials in the control group were similar to those of the high-spatials who received the tutors – i.e. students who were already proficient in mental rotation did not get much benefit from the tutoring, which is not surprising.

From a development perspective, our strategy proved effective for selecting problems that were appropriate to the student's level of progress. ITS developers might benefit from adopting a strategy similar to the one outlined in this paper. When training soft skills – such as mental rotation – via computer, the instructor does not have the luxury of probing the student for the source of his or her confusion. Our approach could help in domains such as statistics, programming, or driving skills where problems to be solved by the student have multiple dimensions of difficulty. By developing a library of problems that combinations of difficulty on multiple skills, then it is possible to diagnose specifically which skills a student is having problems with. The computer can then seek out those problems that have the appropriate level of difficulty for each skill according to the student's progress.

REFERENCES

Guay, R. B. (1980). Spatial ability measurement: A critique and an alternative. A paper presented at the 1980 Annual Meeting of the American Educational Research Association, April, Boston, MA. (ERIC Document Reproduction Service No. ED189166).

Knowles, M. S., Holton, E. F. III, and Swanson, R. A. (1998). *The adult learner: Fifth edition.* Woburn, MA: Butterworth-Heinemann.

Romoser, M. R. E., & Fisher, D. L. (2001). Maintaining Kinematic Constraints When Performing Mental Rotations About a Fixed Axis: Implications for Instruction and Displays. *Proceedings of the Human Factors and Ergonomics Society 45th Annual Meeting – 2001,* 1400-1403.

Romoser, M. R. E., Katsikopoulos, K. V., Sathyan, D., & Fisher, D. L. (2001). Mental rotation as a function of kinematic constraints, *Society of Mathematical Psychology, 34th Annual Meeting, Brown University – 27-29 July, 2001, Refereed Presentation.*

Shepard, R.N. & Metzler, J. (1971). Mental rotation of three-dimensional objects. Science 171, 701-703.

Shute, V. J. (1995). SMART evaluation: Cognitive diagnosis, mastery learning & remediation, *Artificial Intelligence in Education Conference (AI-ED) 1995,* 123-130.

Siemer, J. and Angelides, M. C. (1998). Towards an intelligent tutoring system architecture that supports remedial tutoring. *Artificial Inelligence Review, 12,* 469-511.

Woolf, B. P. (2002). *Building intelligent tutors.* To be published. Amherst, MA: University of Massachsetts, Amherst.

Woolf, B. P., Romoser, M., Bergeron, D., & Fisher, D. L. (2003). Tutoring 3-Dimensional Visual Skills: Dynamic Adaptation to Cognitive Level, *International Conference on Artificial Intelligence and Education, Sidney Australia, July 2003.*