The AnimalWatch project: Providing adaptive teaching for students and adaptive support for teachers

Carole R. Beal
Dept. of Psychology
University of Massachusetts-Amherst

Ivon Arroyo
Dept. of Computer Science
University of Massachusetts-Amherst

Abstract

Smart teaching systems offer a way for schools to meet the challenges of providing instruction to students with a diverse range of abilities, interests and needs. AnimalWatch is an intelligent computer based tutor that provides individualized math instruction for 9-12 year olds. Math word problems are integrated into narratives about endangered species, to engage student interest and help students appreciate the value of learning math. Student motivation and self confidence in math ability are fostered through the selection of problems that provide challenge and the availability of help features that are tailored to students' individual learning styles. An accompanying web site allows teachers to download all or parts of the tutor, and provides curriculum information and discussion boards. Authoring tools available at the web site allow teachers to contribute their own word problems and to create their own AnimalWatch adventures.

Introduction

The chapters by Wright et al. and Greenfield et al. in this volume provide new evidence regarding the rapidly increasing exposure that American children today have to digital media, with the majority having access to computers and the internet at home. The same situation holds when children go to school: Almost 90% of American schools are connected to the Internet, and most classrooms have at least one computer. Yet the mere presence of computers in schools does not always mean that their promise is fulfilled. In fact, many teachers report that their computers are often underutilized. It can be difficult to find software that is easy to use, effective, and that will engage students' interest, and that meshes well with the existing curriculum. All too often, the powerful computers often end up being used for game playing or Internet surfing as a reward to students when their academic assignments are completed, or even as expensive plant holders or doorstops.

At the same time that powerful computer technology is becoming more readily available to schools, there is an urgent need for new approaches to instruction that will help schools meet growing challenges. These challenges include the need to serve an increasingly diverse student population. Students also now come to school from a wide range of home backgrounds; some have considerable prior experience with books and learning tools, whereas others need much more support in the transition to school. The needs of students with special learning characteristics and disabilities must also be considered. In addition, educators must also cope with rising expectations on the part of parents, politicians and others about educational standards, and with the growing pressure for accountability in the form of high stakes assessments and frequent achievement tests even for elementary school students.

Computers and education, revisited

One possibility to be considered is that computer based teaching systems may help educators to address more effectively the needs of all their students, through adaptive instruction. Computer based tutoring systems are not new; in fact, they have been around for decades. However, their potential in terms of actual classroom use has not yet been fully realized, for several reasons: First, adaptive teaching systems have traditionally been very expensive to develop, and they tend to be brittle, meaning that a computer tutor that has been developed for one particular domain cannot be easily updated or modified by the teacher for another purpose. Thus, if the teacher's needs or goals change, the tutor becomes useless. Second, most smart teaching systems have not been developed by educators, so they tend to be "stand alone" applications that do not consider teachers' goals, and thus do not become well integrated into the ongoing classroom curriculum. Third, the need to provide support for users has not always been carefully considered. Few classroom teachers have the time or experience to debug a piece of software that does not work the first time, or to persist through arcane installation procedures. Finally, until quite recently, there simply were not enough computers in most schools for the idea of providing regular instruction via computers to be plausible. Rather, a classroom might have had one machine that could serve one, two or a small group of students at a time, making it difficult to develop whole class activities around the tutor.

Although some of the factors that have limited the wide-spread use of smart teaching systems persist today, others are rapidly changing: In particular, the increasing availability of hardware in the form of computer labs designed to serve entire classes has made the need more evident for software that can provide individualized instruction to many students at a time. Developments in the area of artificial intelligence and machine learning have made it possible to design systems that can be used more flexibly. For example, the intelligent teaching component of the tutor can now be more clearly distinguished from the domain content, allowing teachers to add their own material, and also customizing the pedagogy for different groups of students (Ainsworth, Underwood & Grimshaw, 2000). The increased access to the Internet has made it possible for systems to be updated frequently, and for teachers to use design and authoring tools to customize the tutoring system for their own needs (Ritter, Anderson, Cytrynowicz, & Medvedeva, 1998). Thus, the failure of the past in terms of the impact of computer-based instruction does not necessarily have predictive power for the future. Rather, it is time for a new look at how intelligent computer based tutoring systems might be able to help schools meet the needs of a diverse student population (Andriessen & Sandberg, 1999; Beck & Stern, 1999; Bruer, 1997; Fletcher-Flynn & Gravatt, 1995).

The AnimalWatch project

The goal of the AnimalWatch project has been to develop a computer based mathematics tutoring program for 10-12 year olds that can be integrated into the classroom curriculum and customized by teachers through an accompanying web site. The project involves an interdisciplinary team of computer scientists, developmental psychologists, and classroom educators. Most of the team members also had interdisciplinary training and interests. For example, the computer science graduate students who were central to the project also earnedgraduate degrees in education. Classroom teachers were an integral part of the project, from the initial design meetings to the dissemination stage of making professional presentations at national conferences about the project. The overlapping backgrounds of the team partners helped to ground the project in the needs of the classroom rather than the laboratory.

AnimalWatch provides instruction in the mathematics topics most often taught in American fifth and sixth grades: whole number operations (addition, subtraction, multiplication and division), introduction to fractions, and addition and subtraction of like and unlike fractions and mixed numbers. The project's focus on mathematics as the target domain for smart teaching came about for several reasons: First, it is an area of the curriculum in which American schools tend to falter. International comparisons consistently indicate that students from the United States perform at best only on an average level in math, and often, far below students in other countries. At the same time, there is an urgent need for a technically skilled workforce, a need that is now being met in large part by the importation of mathematically proficient workers from other countries.

A second reason for the focus on mathematics is that teachers frequently reported that they would appreciate support in this area of the curriculum. Many teachers at the elementary level, most of whom are female, feel under-prepared to teach math; at all levels there is a shortage of well trained teachers of mathematics. The response in teacher education programs has not always been productive. For example, one of our project members gave a guest lecture at an institution that "prepared" student teachers for math lesson planning by having the students write poetry about their feelings about math.

The third reason for our focus on mathematics was that there are important issues of equity that remain to be addressed in this domain (Beal, 1994; Eccles & Jacobs, 1986). Specifically, although female students now enroll in almost as many math courses as males, the long standing gender gap in math achievement tests persists. There is also reason to be concerned about female students' level of competence in math, particularly when they must work under time pressure or transfer skills to problems not previously seen (Willingham & Cole, 1997). A related concern is that enrollment by female students in careers that require mathematical proficiency remains discouragingly low. Educators, researchers and funding agencies have called for new approaches to mathematics learning that can engage more students, particularly girls and women.

Theoretical framework

The overall design of the AnimalWatch tutor reflects the theoretical framework of Eccles, Wigfield and their colleagues (Eccles, Wigfield, Harold, & Blumenfeld (1993). Specifically, students' motivation and in turn their achievement in a particular academic subject such as math is thought to reflect their attitudes about the domain, including their liking of the subject, their belief that it is important and valuable to learn, and their concept of their own ability to master the subject. These constructs may vary independently. For example, a male student might dislike math but feel that it is important to learn and believe that he can succeed with effort, whereas a female classmate might dislike math, and feel that it is not important for her but believe that she could master the material if she tried. Extensive work by Eccles and colleagues has shown that students in elementary school and beyond can accurately report their attitudes, and that their attitudes are related to motivation and achievement in math as well as other academic domains (Beal, Beck & Woolf, 1998).

One implication of the Eccles et al. theoretical framework is that computer based teaching should help students see the value of learning mathematics. To accomplish this, math problem solving in the AnimalWatch tutor is linked with real world information and problems are embedded in real world contexts (Cordova & Lepper, 1996). One of our partner teachers, a master teacher of fifth graders known for her work on environmental science, first suggested the idea of creating math problems about endangered species (hence the name, "AnimalWatch"). This was an inspired suggestion because, of the sciences, environmental biology is equally appealing to both male and female students. We began by creating mathematics word problems about the Atlantic Right Whale, then added problems about the Giant Panda, and the Takhi Wild Horse (also known as the Przewalksi Wild Horse).

Students log onto AnimalWatch, choose a species, and are greeted with a personalized letter inviting them to serve as a wildlife biology consultant who will help to assess the status of the species. Students then are presented with a series of math word problems about the species. For example, a student working on the Right Whale adventure might view a problem about the average distance traveled per day on the whales' annual migration down the Atlantic coast. Each problem is accompanied by an image that serves as an illustration, or that presents information needed for the problem such as a graph, chart or other data sources that the student must examine to find relevant information. The content information represented in the word problems is as accurate as possible: the distances are factually accurate; sizes and weights of the animals are within known ranges, related historical and geographical information is correct, and so on.

Adaptive teaching in narratives

Students reported that they enjoyed working on math about the endangered species and that they felt that it was a good way to learn mathematics. However, in early field trials with the AnimalWatch tutor, we found that simply working on an indefinite series of math word problems was not sufficiently engaging for students; they needed to have some sense that they were making progress as they worked. We did not want to fall into the usual trap in software development of making the program into a competitive game, because this tends to appeal much more to boys than girls. Rather, we decided to merge the word problems into a narrative through which students would progress as they worked on math. Each narrative currently has four sections or "contexts". For example, the Takhi Wild Horse story begins with word problems concerning background information about the species, its history, and its distinctive characteristics. Then the student moves to a section on Mongolia, the original home of the Takhi, learning about its history, geography and culture. The third context focuses on how the Takhi horse was saved from complete extinction through the efforts of zoos around the world. In the final segment, the student takes a virtual trip to accompany a group of wild horses that are being returned to a nature reserve in Mongolia. Each context has a distinctive background color, and the transitions from one context to another are explicitly marked by text screens congratulating the student on his or her work and introducing the next section of the adventure.

Although we found that the addition of narrative to the tutor was very effective, one challenge presented by this strategy is how to move all student users through the story even though they may be working on different areas of mathematics. That is, a student who is still struggling with multiplication should ideally reach the final segment of a story in roughly the same amount of class time as a classmate who has moved on to fractions. Or, as we say in our project, "everyone has to get on the boat", referring to the Right Whale narrative which culminates in the student going on a virtual whale watching trip. Our reasoning was that it would not be helpful to students' confidence and motivation to be stuck in the first context while a student in the adjacent seat had completed the narrative. This was of particular concern because we had observed that students did frequently lean over to the next computer to see what classmates were working on.

In order to be able to adapt the narrative to students working on different math topics, we had to make the problem set space deep as well as wide. That is, word problems were required not only for all the math topics in our target domain (addition through mixed numbers), but each topic had to have 4-6 word problems for every section of the story, to ensure that a student who had reached the end of the story could still be solving addition problems if necessary. Sometimes, it was a challenge to find real world content that would support this problem set depth; to put it bluntly, there are only so many math problems about whales that one can think up. One strategy was to include problems about related topics such as history (e.g., some Takhi Wild Horses were captured by the Nazis in the Second World War), geography and culture (e.g., nomads of Mongolia), and other species (e.g., how the loss of the Mongolian steppes has also affected the Snow Leopard) in order to broaden the scope of our problem set.

A second strategy was to use problem templates that could be filled in "on the fly" with specific numbers, rather than relying on a fixed set of canned problems. AnimalWatch uses hundreds of problem templates, each with a range of possible numerical values that are realistic given the problem content. When a problem of a specific level of difficulty is needed for a particular student, the system searches a candidate space of problem templates (e.g., the set of two digit multiplication problems available for that particular part of the story), selects one, and fills it with specific numerical values that are appropriate for the student's current level of mastery. Creating novel problems as needed increases the flexibility of the system and helps to ensure that there are enough problems available at different levels of difficulty.

Adaptive teaching strategies and mechanisms

The theoretical framework guiding the AnimalWatch project implies that teaching methods should encourage students to believe that they can learn difficult academic subjects, given sufficient effort. This is especially important in the area of mathematics, because American students tend to believe that math success is due primarily to natural ability and that if they make errors, they must not be smart enough to understand the material. AnimalWatch addresses this issue by presenting mathematics problems that are targeted to the edge of the student's expertise, while also providing substantial help and problem solving support as needed to ensure success. The ability of the computer to assess and adapt its instruction on a minute by minute basis is a major advantage compared to traditional classroom instruction with one teacher working with groups of students. Studies in other domains suggest that the personalized instruction provided by an intelligent tutor can dramatically accelerate learning, because the student works on areas at the edge of his or her understanding, whereas traditional instruction moves either too quickly or too slowly.

The selection of appropriate math problems in AnimalWatch is determined by its student model function: an artificial intelligence module that "knows" the mathematics domain and creates an estimate of what the student understands of the domain. This estimate is continually updated as the student works, relying on errors made during problem solving and latency to enter a response via the keyboard. If a student makes many errors on a particular problem, AnimalWatch will present another problem of about the same difficulty, or even move backwards through the curriculum to something easier, both to check that the student actually has the prerequisite knowledge and on the assumption that a bit of success in problem solving will booster confidence. If a student solves a couple of similar problems quickly and without errors, AnimalWatch estimates that the student is ready to move on to problems involving a new mathematics operation or a new level of difficulty.

The ability to adapt instruction depends heavily on a good estimate of the student's comprehension. Yet determining what a student actually understands is technically quite difficult. Whereas a human teacher would be able to consider behavioral observations such as expressions of puzzlement and direct questions from the student, the computer tutor only has access to fairly crude indirect measures such as the latency of response and number and type of errors in order to create its estimate of student competency. These data are noisy; there are considerable individual differences in how quickly and accurately students work. For example, some students navigate the keyboard more expertly than others. Also, a small computational mistake at the beginning of the problem solving process can drastically impact the time required to solve a problem, but this does not necessarily imply that the student is seriously confused about the mathematical concept. Thus, AnimalWatch must continually evaluate its predictions about the student's behavior on a particular problem against what the student actually does and adjust its estimates accordingly.

Adaptive help

AnimalWatch provides help when students make errors, in the form of a simple text hint on the screen, a demonstration of a procedure such as digit carrying, or a series of instructional screens to prompt the student in the right direction. If the student still does not arrive at the correct answer, increasingly specific help is offered until the student gets the answer, even to the point of presenting the answer in the hint area of the screen and requiring the student to reenter it in the problem solution area. Thus, the student never fails to solve a problem when working with AnimalWatch; what varies is the amount of help and scaffolding that is needed.

Types of help in AnimalWatch

Responding to student problem solving errors offers another chance for AnimalWatch to adapt its instruction. The tutor has a battery of available hint types which range from simple message hints to highly structured and interactive hints that walk the student through a procedure such as finding the least common multiple for unlike fractions. The choice of which hint to present for a particular student on a particular problem is guided by the student modeling function, subject to some broad constraints. First, when the student makes an error, he or she first receives a simple message hint such as "Try again!" or "Are you sure you are trying to add (problem quantity 1) and (problem quantity 2)?" These simple message hints may not appear to be particularly helpful at first glance; however, we have found that almost half of the errors made by students are corrected on the next step after a simple message hint. Minor calculation errors or the choice of the wrong operation can be corrected easily by the student, and we have some reason to suspect that extensive help presented too quickly may undermine students' math confidence. Therefore, more detailed hints are reserved for problems on which the student has made several mistakes, with hint specificity increasing as needed.

The more extensive hints used in AnimalWatch range from fairly concrete, highly structured and interactive (meaning that the student is required to make many small steps) to more abstract, procedural hints. For example, a student who has made errors on an addition problem might see a concrete hint screen in which the two quantities are displayed in the form of screen objects such as blocks that can be dragged and dropped together into groups of ten that in turn become rod units. Dragging and dropping the screen objects parallels the grouping of concrete manipulatives in the classroom, and helps the student to observe directly the physical regrouping that is represented by numerical digit carrying and borrowing (or trading). In another case, the student who is struggling with a multi digit subtraction problem may receive a procedural hint in which the borrowing algorithm is displayed step by step. Generally speaking, AnimalWatch has a bias to first draw on concrete hints to help the student, and gradually shift to more abstract, procedural help as a topic (e.g., multi digit addition) is mastered. However, this is somewhat constrained by the particular topic and the level of problem difficulty. For example, if the student is working on a four digit addition problem, requiring the student to drag and group thousands of blocks would be impractical given limits on time and screen size, and so the abstract-procedural hints must be used instead.

Adapting to student learning styles

Much of our research with AnimalWatch has focused on how the tutor can adapt its instruction to students on the basis of gender and cognitive developmental stage (Arroyo, Beck, Schultz & Woolf, 1999; Arroyo, Beck, Woolf, Beal, & Schultz, 2000). With regard to gender, other researchers have reported differences in the strategies that male and female students prefer. Even though both groups may be equally likely to arrive at the correct answer, the approach adopted and the time taken may be quite different. In general, it seems as though female students tend to rely on more concrete and complete strategies, such as counting on fingers, or setting up equations and working out full solutions step by step (Carr & Davis, in press; Fennema, Carpenter, Jacobs, Franke, & Levi, 1998). These approaches tend to be accurate but are often slower than guessing and estimating, retrieving answers from memory, or reasoning on the basis of past examples of similar problems, all strategies that are more commonly observed among male students (Mills, Ablard, & Stumpf, 1993; Royer, Tronsky, Chan, Jackson, & Merchant, 1999). Therefore, what constitutes effective instruction in math may be somewhat different for girls and boys. For example, hints requiring the student to perform multiple problem solving steps without guidance may be less effective for girls than the more concrete, structured hints that involve manipulating screen objects.

Another potential predictor of response to hints of different types is the student's level of cognitive development. This is particularly important at the transition to adolescence: in a class of 11-12 year olds, some students may still be at the concrete operational stage, reasoning about problems on the basis of direct observation, whereas others may have moved into early formal operations and be capable of more abstract, hypothetical reasoning. The variation in cognitive developmental level can present a challenge to teachers in terms of selecting the right form of instruction in math. Students who are still reasoning with concrete examples may benefit more from screen manipulatives, whereas their formal operational classmates may find dragging and grouping virtual blocks to be tedious and less helpful than a dynamic screen demonstration of an algorithm for carrying.

To investigate the impact of different hints as a function of gender and cognitive developmental level, we analyzed student problem solving data that are automatically collected by the computer as students work. The primary data include time required to solve problems, the number of errors made, and the student's response to hints, meaning the probability that a hint on one problem was followed by problem solving success on a subsequent problem of similar difficulty. In one study, fifth grade students worked with AnimalWatch for three sessions over the course of one week. AnimalWatch was configured with a strong hint selection bias either towards hints that were highly structured and concrete, or hints that were more abstract and procedural in nature. Hints also varied in the degree of interactivity required from the student: hints that are highly interactive are also highly structured and walk the student through the solution process in incremental steps. We hypothesized these would be more helpful to girls.

Hint effectiveness was assessed by comparing the number of errors made on subsequent problems of the same type, the idea being that if a particular hint is helpful then the student should be able to move on to solve a similar problem with significantly fewer errors. Students' cognitive developmental stage was assessed via a computer presented battery of Piagetian reasoning problems. At the end of the activity, students completed a survey about their AnimalWatch experience, including questions in which they were asked to rate the helpfulness of the different types of hints that they saw.

The results indicated that male and female students did have different responses to the hint conditions. Specifically, girls made better progress with hints that were highly structured and interactive, in terms of a lower error rate on subsequent problems. In contrast, boys tended to prefer and make more progress when the hints were less structured. The students' cognitive developmental stage was also important in determining response to hints. Not surprisingly, both boys and girls of lower cognitive development needed more hints to solve the problems than students who had advanced to the stage of formal operations. The concrete operational students also benefited more from more concrete and structured hints, whereas formal operational students performed better with more abstract, procedural hints.

Interestingly, students' perceptions of what they found to be helpful or not helpful were fairly accurate. Students were shown screen shots of various hints and were asked if they had seen such hints, and if so, to rate how helpful they had been. There was a strong effect for girls between their cognitive developmental stage and their views of how helpful the different hints were. Girls who were still at the concrete operational stage rated the more concrete, highly structured hints as being much more helpful than the more abstract feedback, and as noted above, the structured hints did in fact work better for these students. In contrast, girls who were at the formal cognitive stage were more likely to think that the more abstract hints were most helpful. This was particularly true for hints that focused on multiplication and division.

Overall, the results indicated that not only is adaptive feedback especially important for girls, certain specific types of feedback are preferred by girls. This pattern paralleled an earlier study in which the fully adaptive version of AnimalWatch was compared to a no-help version in which problem solving errors were greeted with the response, "Try again". In this case, girls' performance faltered considerably in the no-help version, whereas the impact on boys was much less dramatic. In current studies, we are investigating whether information about the student user's gender and assessments of his or her cognitive developmental level can be directly integrated into AnimalWatch in order to guide its hint selection.

Using past performance as a guide to adaptive teaching

Another way to improve the adaptivity of intelligent tutors such as AnimalWatch is to use data from past users as an additional guide to optimize teaching decisions. That is, the problem solving performance of hundreds of students who have used AnimalWatch in the past can serve as a database resource for the student model function: When an estimate is made of a student's knowledge and a problem is selected, AnimalWatch can use its "knowledge" of past students with similar estimated knowledge to gauge the likely outcome of the problem solving process and make adjustments as needed. Thus, as more students use the system, its teaching performance can be continually refined and improved; it becomes increasingly likely that the system will be able to find a past user who has a problem solving profile similar to the current user and to evaluate the probable impact of a teaching decision based on past outcomes.

One constraint on this approach is that as the number of stored problem solving records increases, the time and resources required to search the space of problems expands dramatically. Given that AnimalWatch is being used on fairly low-end classroom computers and must reason in real time while the student user is waiting for the next problem, the search must be pruned. Techniques such as rollouts can be used to limit the search space while maintaining the predictive power of the past-performance database (Beck, Woolf, & Beal, 2000).

Adaptive support for teachers

Smart teaching systems will only have wide impact in the classroom if they can be used flexibly and can help teachers achieve their own goals. One strategy integrated into AnimalWatch is the ability to adapt the story line to the amount of time available for the student to work. In our field trials, students have typically worked with AnimalWatch for several sessions, the length of which varies across grades, schools, and even within the school day. To increase flexibility, the duration of each context can be adjusted by specifying the number of minutes or number of problems in each, so that the adventure can adapt to fit the anticipated time available. Also, when students have completed one narrative, they can choose another species and begin another narrative. The system will keep them at their current level of mathematics problem solving; they do not "start over" with simple addition in the new story but rather are presented with problems similar in difficulty to those that they worked on at the end of the previous story.

One long term goal of the project is to increase teachers' ability to adjust AnimalWatch's teaching without having to become computer programmers. Ideally, teachers could specify a goal and allow the tutor to figure out how to accomplish the goal on its own. As an initial step towards this goal, Beck designed a machine learning agent that would accept a specific teaching goal and then use stored performance data from hundreds of past users to run simulations that allowed the tutor to train itself to make appropriate teaching decisions. Beck found that the machine learning version of AnimalWatch could significantly reduce the time needed for students to progress through the mathematics curriculum (Beck et al., 2000). Although not currently practical for a classroom based tutor due to the intense processing demands, this approach should eventually allow teachers to enter a teaching goal and allow AnimalWatch to evaluate and modify its own performance accordingly.

Another strategy to promote adaptive use by teachers is the availability of authoring tools and support through the AnimalWatch web site. Here, teachers are provided with options and strategies for customizing the program and extending it to meet their specific needs. The web site provides information about the curriculum and options for downloading the program (although the very large number of images makes distribution via CD-ROM more practical for most users). A bulletin board allows teachers to share ideas about ways to use the tutor and to discuss what needs to be fixed or improved, and what new features would be worth adding.

Most importantly, teachers can propose new word problems and select images or upload their own graphics to accompany the problems. The web site is designed to make the word problem authoring tool easy to use, and the overall philosophy is open source, meaning that although researchers provided the initial content, the larger community of users will expand the content. An important part of the open source philosophy is that users should evaluate the development of AnimalWatch, judging what problems are the best to download, what people are the most reliable creators of word problems, and which adventures are most effective for use with students. To accomplish this, rating mechanisms are provided, along with a discussion forum in which teachers can comment on the features of the system and contribute ideas about how to use it effectively in their classrooms.

Thanks to the independence of the intelligent teaching component from the stories, teachers can create their own endangered species adventure. These functions allow teachers to customize AnimalWatch to fit their curriculum. For example, if a teacher is developing a unit on Borneo, he or she might create an adventure about the orangutan, including maps and background information about the country and its people. Although the original focus of AnimalWatch was on mathematics tutoring, the inclusion of real world content would support interdisciplinary projects in the classroom, including email contact with experts, environmentalists, and students in countries that are home to the endangered species under study.

Summary

Cognitive strengths, motivation, and interests vary dramatically from one student to another, even if they are the same age, and these characteristics affect the success of classroom instruction. The notion that instruction should be adapted to individuals, or at least to groups of students with similar characteristics, has long been valued in education (Snow, 1977). Increasingly, computer based tutoring systems offer a way to accomplish this goal. Smart teaching systems have long been able to adjust teaching on the basis of student performance, for example, by adjusting the difficulty of problems if necessary, or trying a different type of hint, or returning to an earlier topic to review a critical concept. In addition, new approaches in intelligent teaching systems are leading to the development of tutors like AnimalWatch that can consider the learning style and cognitive stage of the student as a guide to teaching decisions, and that can be customized and extended by teachers. Although it does not seem likely that computer based systems will ever entirely replace human teachers, it does seem probable that such systems will become increasingly important as tools to enhance education and support teachers.

Author Note

The AnimalWatch project was supported by the National Science Foundation Program in Gender Equity (HRD-9714757). The views, opinions and recommendations expressed here are those of the authors and do not necessarily reflect the views of the granting agency. We would like to acknowledge our project colleagues Beverly Woolf, David Hart, Joseph Beck, and Rachel Wing. We would also like to express our deep appreciation to Charlene Galenski, Diana Campbell, and Amy Ryan for their enthusiastic support of our field work, as well as to Douglas Tierney, principal of the Deerfield Elementary School in South Deerfield, Massachusetts, and Dr. Mario Cirillo, principal of the Chestnut Accelerated Middle School in Springfield, Massachusetts. Correspondence concerning the AnimalWatch project may be directed to Carole R. Beal, Dept. of

Psychology, Tobin Hall Box 37710, University of Massachusetts Amherst, Amherst MA 01003, or via electronic mail at < cbeal@psych.umass.edu >.

References

Ainsworth, S., Underwood, J., & Grimshaw, S. (2000). Using an ITS authoring tool to explore educators' use of instructional strategies. Proceedings of the 5th International Conference on Intelligent Tutoring Systems.

Andriessen, J., & Sandberg, J. (1999). Where is education heading and how about AI? International Journal of Artificial Intelligence in Education, 10, 130-150.

Arroyo, I., Beck, J. E., Schultz, K., & Woolf, B. P. (1999). Piagetian psychology in intelligent tutoring systems. Proceedings of the 9th International Conference on Artificial Intelligence in Education, 600-602.

Arroyo, I., Beck, J. E., Woolf, B. P., Beal, C. R., & Schultz, K. (2000). Microadapting AnimalWatch to gender and cognitive differences with respect to hint interactivity and symbolism. Proceedings of the 5th International Conference on Intelligent Tutoring Systems.

Arroyo, I., Schapira, A., ., & Woolf, B. P. (2001) . Authoring and sharing word problems with AWE. Proceedings of the Tenth International Conference on Artificial Intelligence in Education, IOS Press.

Beal, C. R. (1994). Boys and girls: The development of gender roles. New York: McGraw Hill.

Beal, C. R., Beck, J. E., & Woolf, B. P. (1998, April). Impact of intelligent computer instruction on girls' math self concept and beliefs in the value of math. Poster presented at the annual meeting of the American Educational Research Association, San Diego.

Beck, J. E., & Stern, M. (1999). Bringing back the AI to AI & Ed. Proceedings of the 9th International Conference on Artificial Intelligence in Education, 233-240.

Beck, J. E., Woolf, B. P., & Beal, C. R. (2000). ADVISOR: A machine learning architecture for intelligent tutor construction. Proceedings of the 17th American Association of Artificial Intelligence Conference.

Bruer, J. T. (1997). Schools for thought. MIT Press.

Carr, M., & Davis, H. (in press). Gender differences in strategy use: A function of skill and preference. Contemporary Educational Psychology.

Cordova, D. I., & Lepper, M. R. (1996). Intrinsic motivation and the process of learning: Beneficial effects of contextualization, personalization, and choice. Journal of Educational Psychology, 88, 715-730.

Eccles, J. S., & Jacobs, J. E. (1986). Social forces shape math attitudes and performance. Signs, 11, 367-380.

Eccles, J. S., Wigfield, A., Harold, R. D., & Blumenfeld, P. (1993). Age and gender differences in children's self and task perceptions during elementary school. Child Development, 64, 830-847.

Fennema, E., Carpenter, T. P., Jacobs, V. R., Franke, M. L., & Levi, L. W. (1993, June-July). A longitudinal study of gender differences in young children's mathematical thinking. Educational Researcher, 27, 6-11.

Fletcher-Flynn, C. M., & Gravatt, B. (1995). The efficiency of computer assisted instruction (CAI): A meta analysis. Journal of Educational Computing Research, 12, 219-241.

Mills, C. J., Ablard, K.. E., & Stumpf, H. (1993). Gender differences in academically talented young students' mathematical reasoning: Patterns across age and subskills. Journal of Educational Psychology, 85, 340-346.

Ritter, S., Anderson, J., Cytrynowicz, M., & Medvedeva, O. (1998). Authoring content in the PAL algebra tutor. Journal of Interactive Media in Education, 98, 3.

Royer, J. M., Tronsky, L. N., Chan, Y., Jackson, S. J., & Merchant, H. (1999). Math fact retrieval as the cognitive mechanism underlying gender differences in math test performance. Contemporary Educational Psychology, 24, 181-266.

Snow, R. E. (1977). Individual differences and instructional theory. Educational Researcher, 6, 11-15.

Willingham, W. W., & Cole, N. S. (1997). Gender and fair assessment. Mahwah NJ: Erlbaum.