

Intelligent modeling of the user in interactive entertainment

Carole R. Beal¹, Joseph Beck², David Westbrook², Marc Atkin², & Paul Cohen²

¹Psychology; ²Computer Science

University of Massachusetts at Amherst

Amherst MA 01003

cbeal@psych.umass.edu; beck, westy, atkin, cohen@cs.umass.edu

Abstract

A theme of the symposium is to explore ways to employ AI to make games more appealing to people who do not enjoy current genres, and to expand the market for interactive entertainment beyond the traditional niche of young male players. We suggest that AI techniques employed in the world of intelligent tutoring to model the user and adjust instruction, help and content could be fruitfully adapted to interactive entertainment. In computer-based educational tutoring, adaptivity to user behaviors and characteristics such as gender and cognitive developmental level have been shown to increase learner motivation, engagement and achievement in the area of mathematics learning. Similarly, utilizing data regarding player behaviors such as latency and errors to construct a model of the player would allow for more adaptive game play, which in turn would increase the appeal of computer games to a wider audience.

Intelligent modeling of the user in interactive entertainment

Increasing the appeal of interactive entertainment

From one perspective, the world of computer games is a major economic success story, with sales now beyond those of more traditional media such as films, and with more power to attract users' leisure time than television or books (Jennings, 2001). However, from another perspective, the picture is not so bright: although large, the interactive entertainment market is not expanding in any significant way beyond the original target market, i.e., adolescent and young adult males, by attracting new fans from other demographic niches. In particular, girls and young women are still only one third as likely to be consumers of interactive entertainment as their male peers (Jennings, 2001). Popular explanations for this gender gap typically focus on girls' supposed dislike of fantasy, magic and action genres, or on their rejection of sexist portrayals of female characters (cf., Gailey, 1992; Rabasca, 2000). Yet these explanations are not supported by research in adolescent development (Beal, 1994). For example, busty, scantily clad female game characters (e.g., Lara Croft), map directly onto the role models that are prevalent in

popular media consumed by young females, as a glance at the cover of any issue of *Cosmopolitan* -- the magazine most widely read by young women in North America-- will readily confirm. Television programming focusing on themes of female-employed magic and science fiction has been extremely popular with female viewers since the 1970s (e.g., from "Bewitched" and "I dream of Jeannie" to today's "Buffy", "Charmed" and "Witchblade") (Sternglanz & Serbin, 1974). And, if anything, female characters in many computer games are more powerful and exciting role models than are found in the world of TV and print media. So why isn't the world of interactive entertainment more broadly appealing?

ITS and interactive entertainment

One approach to address this question is to ask what characteristics of current interactive entertainment might be less appealing to some prospective players, and how AI techniques and approaches might address this problem and in turn bring more players into the market (Rubin, Murray, O'Neil, & Ashley, 1997). Our suggestion is that a number of techniques and approaches that are already well established in the area of intelligent tutoring systems (ITS) and user modeling could be of great potential benefit to the interactive entertainment community. In the past, these areas have not talked to one another, and the strengths of each area are not duplicated in the other. Yet many of the issues and concerns are similar, and research in the ITS community is already adopting approaches popularized in the world of interactive entertainment. Thus, a tighter connection between the two research areas may be productive for both.

For example, intelligent tutors (including our own NSF-supported tutor for middle school mathematics, AnimalWatch), typically do an excellent job of effective teaching (Fletcher-Flynn & Gravatt, 1995). Our users (students) make rapid progress, achieve specified learning goals, and show increased motivation. The intelligent tutor's excellent teaching performance is made possible by modeling each user's abilities and adjusting its teaching to each student. Yet intelligent tutors generally lack appeal; although they are highly interactive, they typically lack the fast pace and exciting graphics that are now standard in most games. This is a growing problem given that younger users now often have considerable media experience and correspondingly high expectations for the look and feel of

interactive software and other media. The objective of our current NSF-supported project, Wayang Outpost, is to embed an intelligent tutor for SAT-Math test preparation within a narrative adventure set in Kalimantan, employing multimedia techniques that have to date been found primarily in interactive entertainment.

In turn, the interactive entertainment community may be able to bootstrap new approaches from lessons learned in the world of interactive education. Most contemporary interactive games are much more visually and viscerally engaging than intelligent tutors. However, by comparison with tutors, most games provide very little user support, particularly in the initial levels or stages. Although some games provide tutorials, many do not, and some are not particularly helpful. For example, the tutorial in "Star Trek:Elite Force" leaves the player stranded at the top of a ladder, without any guidance about the next place to navigate to continue play. Others simply assume that there will be considerable knowledge transfer from previous games. But if the goal is to attract nontraditional players, then introductions and tutorials need to become more detailed. As it stands now, the learning curve in many interactive games is far too steep to win over less traditional players. Most new users do not get very far before losing and giving up, and early defeat is highly predictive of abandoning the game entirely. In addition, the time investment-return ratio is very different for non traditional players. An experienced player may consider it reasonable to devote many hours to acquiring game strategies; players from less traditional markets are unlikely to feel the same way. Rather than berate potential users for their lack of initiative and commitment, why not configure the game to infer intelligently how to keep them interested and engaged?

From the perspective of intelligent tutoring, it seems clear that adding an intelligent user modeling component would widen many games' appeal considerably by fitting the level of play to individual users. One thing that makes games popular (and therefore sells many copies) is play that is difficult enough to give users a sense of accomplishment, but not so difficult as to make them quit in frustration. Challenge plus accomplishment makes users want to play the game again, or recommend it to a friend, or buy the sequel to the game. Yet a sizeable proportion of players never reach the highest levels of most games, meaning that a great deal of expensive development effort goes under utilized. The existence of and extensive reliance on cheats is another indication that games do not serve many users (players) well. And, it is our contention that a failure to adapt intelligently to the user is a major barrier to the industry-wide goal of attracting more players who are not part of the traditional, core user base of adolescent and young adult males who are willing to devote hours to cracking a difficult game. In short, the "one size fits all players" approach to game development overlooks the accumulating evidence in the ITS community that adaptivity is critical to a good user response -- and increased sales to a wider market.

Making interactive entertainment more adaptive

We suggest an approach that we believe will offer an avenue to expand the user base into nontraditional areas, and at the same time dramatically expand the level of AI incorporated into the game. Specifically, our goal is to make the entire game intelligent. The game should observe the player's behavior, create a model of what aspects of the game the player knows and does not know (yet), and offer choices to the player that are designed to keep him or her motivated and engaged. We envision a game that provides hints and shapes the options that are available to the player so that he or she remains engaged with the game, particularly in the early levels when players can become easily discouraged and frustrated. As the game continues and the player acquires more expertise, the pace and difficulty can ramp up accordingly. With a game that incorporates a player model, replay opportunities could be expanded so that the "same" path through the game will appear different and be experienced differently, in terms of challenge, as a function of what the game estimates the player is ready for.

Fortunately, many of the AI techniques and approaches that are necessary to make games (not just characters) intelligent already exist in the world of intelligent tutoring and user modeling. In our earlier work with our intelligent tutor for mathematics, AnimalWatch, we discovered that both micro and macro adaptation of system performance to the user are important. At the micro level, the use of a student model to guide instruction through intelligent problem selection and help provision was associated with higher achievement and motivation. Macro adaptation refers to the use of more general characteristics about the user, such as gender and cognitive style, to direct the choice of problems and the type of instruction provided. We have found that students with different characteristics respond best provided with hints and help that fit their preferred learning style. For example, girls make faster progress through the math curriculum when provided with hints that contain many structured interactions, whereas boys progress faster with relatively short text-based hints. Concrete examples are especially helpful for students at a less advanced stage of cognitive development, relative to those who are capable of abstract reasoning. Similarly, we expect that certain users of games will perform better under different circumstances or, as a more relevant measure, will play the game more and enjoy themselves more.

A system's ability to adapt to the user is important for motivation and engagement, in addition to achievement. In extensive field evaluation studies, we have found that students who work with the adaptive version of our mathematics tutor are significantly more engaged and provide higher ratings of enjoyment than students who work with a less adaptive version (Beck, Arroyo, Woolf, & Beal, 1999). Most importantly, the impact of adaptivity on performance, motivation and enjoyment is especially strong for female users of the AnimalWatch math tutor. We have found that there are striking gender differences in

response to various types of hints and help provided (Arroyo, Beck, Woolf, Beal, & Schultz, 2000; Beck et al., 1999). Male users prefer and perform well with hints that are brief, abstract, or that illustrate an algorithm (e.g., trading in long division). In contrast, female students prefer and perform best with hints and help that offer concrete illustrations, virtual manipulatives (e.g., rods and blocks that can be dragged and dropped to solve a division problem), and that involve considerable interactivity. There are no gender differences in progress through the curriculum or in the difficulty of problems solved, but there are gender differences in how males and females prefer to solve challenging math problems and in the degree of tutor support and structure that is provided. Our findings in the world of intelligent tutoring are reinforced by research into more traditional instructional approaches also indicating gender differences in help effectiveness (Zambo & Follman, 1994). This gender difference supports our suggestion that nontraditional game players (more often, females) will respond more positively to adaptive games.

To clarify, our proposal differs from past efforts to "add more AI" to games. As it is currently defined, AI in interactive entertainment is often focused primarily on relatively minor adaptation of the computer's play in response to a human player's strategy or accumulation of resources. Advances in the latter case have been fairly modest, and play quality is often quite poor, as assessed by elite players. Also, because character behavior is fairly constrained to begin with in many cases, the absolute scope for improvement to game quality by adding character AI is necessarily limited.

Modeling the user/player

The major research challenge involved in integrating user modeling techniques effectively into games will be to find ways to identify what to model and how to model it. In intelligent tutoring systems, the domain to be taught is usually well analyzed and structured by domain experts. In the area of mathematics problem solving, relevant student behaviors include latency to solve the problems and errors made on the problems. The impact of the help behavior of the tutor is assessed by the student's performance on subsequent problems of the same type and difficulty. Because mathematics is a well structured domain, and the skills required to solve various types of problems are well documented in the field of education, creating a student model to guide instruction and identifying the features of student behavior that the tutor should track was straightforward.

We recognize that the task of user modeling in the games world may be more challenging than in the ITS world. Although the game domain is also highly constrained in many cases, our sense is that the range and type of player behaviors to track may be less predictable than is the case for many tutoring systems. Even so, a number of candidates could be suggested that might generalize across different games. First, some games provide hints to users during tutorials (although many do

not, leaving the novice user to struggle on his or her own). This implies that some thought about coherent player strategies has been done. Also, the computer AI has a set of strategies that it follows. Both of these can serve as a first attempt at defining what to model about the user. Drawing on the expertise of elite players would also be a promising approach, using classic knowledge engineering techniques to elicit heuristics and strategies.

Certain player behaviors can be identified that would have predictive significance for game progress and user engagement. For example, quit behavior may be a strong indicator of user frustration; analyzing the precursors to "quit" would be revealing. In evaluation studies with our AnimalWatch middle school math tutor, we learned that rapid reentry of an answer within a brief time interval was indicative of student frustration and boredom; this behavior served as a signal that the problems being presented were too repetitive and that the student was ready for a new math topic.

Information about other player behaviors could also be utilized, such as latency to respond. Latency data are easy to collect -- in fact, in some sense they come for free in an interactive system -- and they could be analyzed in the early stages of play to determine the user's average response time and adjust the pace of play to optimize that user's experience. This would involve the entire process, not merely having more things happen at the same time, as is often the case now. For example, a shot might move more slowly through the game world space for an older user than for a 17 year old. Calibrating the pace of processes as a function of response time could well make games more appealing to nontraditional markets. In the case of female players, there are no gender differences in perceptual processing speed or in absolute reaction times. However, in some timed problem solving situations, such as the SAT exam, females seem to prefer a more careful approach, which necessarily takes longer (Willingham & Cole, 1997). In mental rotation tasks, the gender difference in accuracy is minimal, but males achieve the answer much more quickly than females. Similarly, in studies of math fact retrieval, there are no gender differences in accuracy, but the fastest responders are typically males (Royer, Tronsky, Chan, Jackson & Merchant, 1999). Thus, one possibility is that the pace of many games is not calibrated appropriately for many female players. Similarly, it is well documented that cognitive processing speed declines from the early twenties on; a game should recognize the optimal pace of play for the older player and adjust accordingly. Of course, latency data are typically very noisy. In evaluation studies of our AnimalWatch math tutor, we utilized performance on several problems as an index of student understanding, because a long delay in solving one problem might simply reflect the student's having stared out the window, chatted with a classmate, or gone to the restroom, rather than confusion about the math concept involved in the problem. Even so, averaging responses over the first few minutes of play would readily address the noise issue.

Another example of a player behavior that could be employed by an intelligent system would be a player's failure to use a strategy or resource when it would be appropriate to do so, suggesting that the player is unaware of the options. Bad moves can be diagnosed and remediated by the structure of the game in subsequent play, just as math errors are followed by help and the selection of additional problems that will help the student master the target skills.

A critically important component of intelligent tutors is the design of a battery of hints that will meet the needs of users with different learning styles and domain expertise (Gertner, Conati, & Van Lehn, 1998). Hints have not traditionally been part of games design, yet they play an important part not only in user progress but also in motivation, particularly for girls and women. There is a great potential role for AI research here in terms of decisions about how to alter a game to facilitate a player's progress, for example, by highlighting a key object that would otherwise be overlooked, or moving a critical doorway or gate, or bringing forward a character who offers advice or counsel at the right time, or limiting the number of choices that can be made just enough so that the novice player has a chance to process the information. Simply providing help or hints when needed is not sufficient; as we have learned in mathematics tutoring, hints must be matched to the student's characteristics and preferences to be effective (Arroyo et al., 2000).

Adding adaptivity via machine learning

Traditionally, considerable research would be required to identify what works best for users with various profiles: in terms of help and hints, i.e., macroadaptation. However, another promising approach to the identification of relevant player behaviors needed to make games adaptive would be to use machine learning techniques. In intelligent tutoring, the data from hundreds of previous student users can be used to help the system adapt its instruction. In intelligent tutoring, the data from hundreds of previous student users can be used to help the system adapt its instruction. For example, the AnimalWatch system gathered data from all of the users to create a model of the population (Beck, Woolf, & Beal, 2000). This provided the tutor with an initial set of standard behaviors by our users, and reduced the amount of time it must act with an individual user before it could make inferences about appropriate teaching decisions. With on-line games, in which data can be collected automatically, it should be possible to use these large data sets to identify the actions that are most predictive of player progress and repeat play. It would also be possible to identify "matches" of players at similar levels of expertise and route game moves accordingly.

Conclusions

The mantra of games design seems to be that what makes games sell well is that they are "fun" -- a quality that is generally undefined. One concern might be that making

games more adaptive to the user might simply make them too easy and therefore less "fun". We would argue in response that an intelligent game should be able to adjust its difficulty several steps ahead of the user so that it is never too easy. But at the same time, it should never be too hard (never getting beyond Level 1 is not fun!). The important point is that what is easy or hard will vary for individual users. We have shown in the user modeling world that adding adaptivity to an intelligent tutor increases its effectiveness both in terms of achievement and motivation. Similarly, the ability to assess the user in an ongoing manner, to model his or her understanding of the game, and to continually utilize this information to direct the course of game play may well make gaming more attractive to a considerably wider market.

Acknowledgments. The AnimalWatch and Wayang Outpost intelligent tutoring projects are supported by grants to the first author from the National Science Foundation. The views, opinions and recommendations expressed here are those of the authors and do not necessarily reflect the views of the granting agency.

References

- 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*, Montreal.
- Beal, C. R. (1994). *Boys and girls: The development of gender roles*. New York: McGraw Hill.
- Beck, J. E., Arroyo, I., Woolf, B. P., & Beal, C. R. (1999). An ablative evaluation. In *Proceedings of Artificial Intelligence in Education*.
- 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*.
- 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.
- Gailey, C. (1992). Mediated messages: Gender, class and cosmos in home video games. *Journal of Popular Culture*, 15, 5-25.
- Gertner, A. S., Conati, C., & Van Lehn, K. (1998). Procedural help in ANDES: Generating hints using a Bayesian network student model. In *Proceedings of the 15th National Conference on Artificial Intelligence*, 106-111.

Jennings, N. (2001, April). *Children's use of technology in multiple settings*. Paper presented at the biennial meeting of the Society for Research in Child Development, Minneapolis.

Kafai, Y. B. (1996). Electronic play worlds: Gender differences in children's construction of video games. In Y. B. Kafai & M. Resnick (Eds.), *Constructivism in practice: Designing, thinking and learning in a digital world* (pp. 97-123). Mahwah NJ: Erlbaum.

Rabasca, L. (2000, Oct.). The Internet and computer games reinforce the gender gap. *Monitor on Psychology*, 31.

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.

Rubin, A., Murray, M., O'Neil, K., & Ashley, J. (1997, March). *What kind of educational computer games would girls like?* Paper presented at the annual meeting of the American Educational Research Association, Chicago.

Sternglanz, S. H., & Serbin, L. A. (1974). Sex role stereotyping in children's television programs. *Developmental Psychology*, 10, 710-715.

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

Zambo, R., & Follman, J. (1994). Gender-related differences in problem solving at the 6th and 8th grade levels. *Focus on Learning Problems in Mathematics*, 16, 20-38.

