An Adaptive Learning Architecture for Next Generation Simulation Training Systems

Jeff Sestokas, Christopher Burns
Applied Research Associates, Inc.
Klein Associates Division
Fairborn, OH
jsestokas@ara.com, cburns@ara.com

Thomas Forbes Worth II
Department of Germanic Languages and Literatures, Harvard University
Cambridge, MA
tworth@fas.harvard.edu

ABSTRACT

In traditional scenario-based serious games, players can choose a skill level and then proceed to game play whereby they are exposed to challenging, ambiguous situations that mimic real-world reality through intelligent agents, 3-D environments, and other computer-based stimuli. During game play, users are presented learning activities which involve complex decision-making and problem solving in order to advance to future time horizons and complete the scenario. However, some scenario-based serious games on the market today don't adapt to the player's skill level during game play. Moreover, most serious games don't offer opportunities for players to learn about their strengths and weaknesses, receive real-time in-game assessment feedback on their performance, and share diverse solutions and strategies during and after game play in order to update, adapt, and advance their understanding of the knowledge being presented to them.

In this paper, we will overview several types of adaptive learning environments and present a conceptual model for integrating an intelligent adaptive learning architecture to an existing government-owned, web-based simulation gaming platform. Additionally, we will discuss how the use of dynamic pathing and sequencing in simulation design can allow for players to assess their initial knowledge and skill level prior to and after game play through pre and post-testing. Also, in this paper we will review a conceptual design of an After-Action Review that can show traceability of user performance in an adaptive simulation system and provide remedial training recommendations to the player.

ABOUT THE AUTHORS

Jeff Sestokas is an Instructional Systems Technologist and Designer in the Skilled, Performance, & Expertise Group of Klein Associates Division, specializing in the application and diffusion of instructional technologies to military personnel and corporate firms. As a member of Klein Associates, he has worked on projects developing training for complex, multi-agency and military command environments with a focus on designing scenarios, evaluation and assessment tools within computerized and live environments. He holds a B.A. in English specializing in secondary integrated language arts and a M.A.Ed. in instructional technology from Baldwin-Wallace College (Berea, Ohio).

Christopher Burns is a Senior Instructional Systems Technologist and Designer at Applied Research Associates, Inc., Klein Associates Division. He works in the Skilled Performance and Expertise Group and has an extensive background in software development, CBT authoring, graphics and video production. His research activities revolve around instructional design, learning theory, and knowledge management, with an emphasis in multimedia, simulations, and instructional product development. Mr. Burns is a retired Army Chief Warrant Officer and holds a B.S. in Computer and Information Science from University of Maryland, a M.S. in Instructional Media and Information Technology from University of Central Arkansas, and a M.S. in Instructional Systems Technology from Indiana University.

Thomas Forbes Worth II holds an Armstrong Fellowship as a doctoral candidate in Older Germanic Languages, Literatures, & Cultures at Harvard University, specializing in heroic epic. He has extensive experience with the role-playing genre of games and endeavors to incorporate simulation and role play into the courses which he teaches, such as Viking literature (Scandinavian 160), in order to enhance students' cultural awareness of medieval societies. He holds a B.A. in Latin Language & Literature and Classical Archaeology from the University of Michigan, an M.Phil. in European Literature from Oxford University, and expects his Ph.D. in 2011.

An Adaptive Learning Architecture for Next Generation Simulation Training Systems

Jeff Sestokas, Christopher Burns
Applied Research Associates, Inc.
Fairborn, OH
jsestokas@ara.com, cburns@ara.com

Thomas Forbes Worth II
Department of Germanic Languages and
Cambridge, MA
tworth@fas.harvard.edu

INTRODUCTION

In recent years there has been an increased use of serious games or immersive learning simulations to measure high-level learning processes and accelerate the transition toward expertise by motivating learners to acquire knowledge and skills by presenting instruction in dynamic, entertaining, and creative ways. As serious games have evolved, so has their design to accommodate user needs and simulate real world experience through high-fidelity visuals and interactions (Zielke, et al., 2009). Although user needs and hi-end visuals are important to user experience during game play, game design should concentrate on providing a contextualized, adaptive learning experience with detailed feedback given to a player based on his performance level.

To address this issue, we have designed a theoretical framework that focuses on providing an adaptive learning experience whereby a user can progress or regress to appropriate skill levels in the simulation based on his response or non-response to well-crafted stimuli. The model is based on research examining how learners acquire new skills and make decisions in their natural settings under high stakes, time-pressured conditions. In addition, we have applied this model as an enhancement to an existing government–owned, web-based platform that supports and trains decision-making at both operational and tactical levels.

In 2008, the use of serious games as an effective training tool for learning and skill development of today's workforce was predicted to reach 1.5 billion dollars in global market sales (Derryberry, 2007). Serious games are considered different from traditional video games because their intended objective is to educate military and corporate organizations with measurable and sustained outcomes in order to improve performance and behavior instead of just for recreation or personal use (Zyda, 2005).

In this paper, we will describe several types of adaptive learning environments and present a theoretical design framework that illustrates user progression in a multi-tiered scenario-driven simulation environment. Also, we will discuss the components that constitute scenario-driven adaptive training systems, describe how the use of dynamic pathing and active sequencing in simulation design can allow players to assess and reflect upon their initial knowledge and skill level prior to and after game play, and then examine the criteria for advancement of performance levels. Finally, we will review a conceptual design of an After-Action Review that can show traceability of user performance in an adaptive simulation system and provide remedial training recommendations to the learner.

ADAPTIVE LEARNING ENVIRONMENTS

The emergence of adaptive learning environments for web-based education has been in existence since the early days of the internet (Brusilovsky, 1999). According to Brusilovsky, Pesin, and Zyryanov, (1993), adaptive learning environments are intelligent educational systems that combine features of traditional Intelligent Tutoring Systems (ITS) with hypermedia components which adapt and record learner performance as players interact with the system. Over the years, these system environments have manifested into multiple categories or types. Brusilvosky & Paylo, (2003) and Specht & Burgos (2006) categorized adaptive learning environments into seven areas:

- Interface-based: This environment focuses on the user interface elements such as color, size, and shadow to support adaptive learning.

- Learning-centered: This environment focuses on the learning process by dynamically adapting the content to the instructional sequence the content in various ways. The learning path changes every time the learner starts.

- Content-based: This environment focuses on the presentation of content as it adapts to the learner's experience. A content-based interface may include adaptive navigational support to guide learners towards relevant and interesting information. For example, content-based adaptive learning environments may provide direct guidance, adaptive link sorting/hiding/annotating/generation.

- Problem solving support: This environment guides the user on the appropriate steps in order to solve problems. The guidance could come from an intelligent mentor programmed by a predefined set of rules.

- Adaptive information filtering: This environment focuses on providing information that is relevant and sequenced by the performance of the learner.

- Adaptive evaluation: This environment focuses on the system or assessment content changing based on the performance of the learner to simulated injects and/or guidance of an intelligent mentor.

- Adaptive timing**:** This environment focuses on the ability for a system to modify or adapt a path based on the learner's performance within a designated time allotment.

Types of Adaptive Feedback

Within adaptive learning environments, simulation feedback plays a vital role in assisting learner performance. Traditional simulation feedback allows a learner to reflect upon his performance after game play in order to think through the decisions and problems made during the simulation. However, in an adaptive learning environment simulation feedback may or may not assist a learner after an event while still in game mode to solve problems and realize workable solutions. During and after game play, feedback may take on many roles to inform, correct, and disrupt learner performance. For example, a mentoring agent could provide critical facts about the situation in order to inform a player. However, if a player disregards or does not have time to read the information, he may not possess the data needed to make an informed decision, thus potentially resulting in a negative consequence.

The second type of adaptive feedback is corrective. Corrective feedback presents information in order to help improve or re-adjust a player's performance. For instance, if a player is continuously making bad decisions, a mentoring agent might take over and pause the game to have the player participate in a remediation assessment in order to teach the proper knowledge and skills needed to continue with game play. Another example of corrective feedback could be a simulated inject designed to scold a player for not taking the proper action or inaction.

The third and final type of adaptive feedback is disruptive. The purpose of disruptive feedback is to confuse, interrupt, or lure a player from taking action on a specific task or a set of tasks. For example, a player is tasked to write an email warning of a pending disaster to a specific simulated character. As soon as the player begins to write the email, he is disrupted by a telephone call, then an important email message, then a video, and so on. The disruptive feedback could be one or a series of lures to distract the player for accomplishing the task or objective.

ELEMENTS OF SCENARIO DRIVEN ADAPTIVE TRAINING SYTEMS

For the purposes of our adaptive framework discussed in this paper and the model of the technology platform used, there are several design elements which must be acted on in order to develop an instructionally sound adaptive learning experience. The narrative, timeline, and injects are critical design elements which drive immersive learning experiences.

The first design element is the narrative. The narrative is the storyline which chronicles past and future events and provides context to the learner during game play. In order to fully engage learners, it is important that the narrative design account for immersion. For immersion to take place within a narrative, the words must express a captivating story with a dilemma or problem in which players must engage and make decisions in the situation. Additionally, the narrative should start off by presenting general information of an event and provide details about the environment. As the situation unfolds, information and events reveal greater information, which enables learners to construct meaning and begin to make sense of the story. Real-world situations place severe time-pressure on decision makers, and so should simulated events. Adding disruptions or information overflow can increase the feeling of time-pressure. However, if too much time pressure exists, learners may make only quick reactions that prevent them from assessing situations and thinking through responses.

Another element used in our theoretical adaptive framework is the development of multiple timelines along which a learner can progress through the game. Timelines provide an instructional designer with a structure around which he can organize detailed simulated injects. Every major event poses an opportunity for learners to encounter new challenges in which they must make decision(s) and take action(s) which will affect exercise outcomes.

The third and final design element to make up a scenario-driven adaptive training system are simulated injects. Injects are the driving components of serious gaming simulations, allowing learners to take action and perform tasks. Once an instructional designer establishes the narrative and timelines, he begins to create injects which adhere to the narrative and timeline. Several elements are important to consider when developing injects for adaptive serious gaming environments:

- *Content*: body of information conveyed to the learner
- *Type*: inject form (phone call, live conversation, video, 3D interaction, etc.)
- *Time*: inject occurrence on the timeline
- *Simulated Entities*: characters to and from whom the information is conveyed
- *Feedback*: standardized responses to the injects/stimuli
- *Consequence*: positive and negative outcomes of responding to an inject
- *Advancement Criteria*: measures used to indicate whether or not a learner progresses to higher level challenges. Criteria may take the form of a correct or incorrect response, time used to take action or inaction, etc.

Realistic, well designed injects motivate learners to engage in game play to make decisions and take actions. For instance, if the learner is presented with a live conversation inject and dialogue content is too ambiguous or does not use proper domain verbiage, the learner may or may not respond to the stimulus, thus compromising its original intent. Associating negative or positive consequences to simulated injects is also important. For example, in our adaptive learning model if the learner does not perform or respond to a task and there are no repercussions (e.g. make a follow-up telephone call) then the learner will fail to recognize the purpose of the task, may or may not be demoted to lower performance level, and will not complete the overall task objective.

Injects, along with their associated actions, should allow learners to gradually gain insight into the nature of the scenario (Davis & Kahn, 2007). Injects should provide learners with the opportunity to expand their experience base and ability to structure information as it is collected. In addition, well developed injects allow learners to focus and comprehend the situation and enable them to assess events and reflect on actions.

A CONCEPTUAL MODEL OF AN INTEGRATED ADAPTIVE LEARNING ENVIRONMENT

The Importance Pre and Post-Testing

Pre- and post testing can provide a measurement device for learners to assess their initial knowledge and skill levels in a particular domain or subject area and acts as a baseline assessment for them to select their appropriate skill level prior to game play. Pre and post assessments help an individual to practice the mission and think through potential problems and possible solution strategies. Additionally, pre-assessments can help activate already learned knowledge causing the player to have a potential advantage during game play because his mind is focused and primed on the tasks awaiting them. Post-assessments help learners retain the knowledge and skills learned during the game by providing an opportunity for rehearsal. The more learners rehearse and practice with the appropriate instructional technique, the greater the likelihood they will be able to transfer the knowledge and skills into long-term memory (Clark, 2008).

Dynamic Pathing/ Active Sequencing

Dynamic pathing is the concept of taking different paths through a simulation based on the decisions made and actions taken within a defined time frame or what we call a “time horizon”. The decisions and actions selected by the learner are evaluated by the simulation portal in order to provide a customized learning experience tailored to his performance levels and abilities.

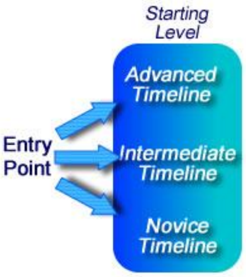

Figure 1. Simulation Starting Point

Prior to game play the user selects an initial skill level or entry point into the simulation and then proceeds to play the immersive learning simulation (see Figure 1). At the end of the first time horizon or timeline segment, the system analyzes user performance based on the specific actions taken and decisions made in order to determine if the user should remain on the current timeline or be shifted to a alternate timeline based on his performance. In this way, the user's experience is dynamically customized to provide an appropriate amount of challenge while playing the simulation. Figure 2 depicts a user who began playing the simulation on the "Novice" timeline and was shifted up to the "Intermediate" timeline at the end of "Time Horizon 1" based on good decisions and actions to simulation injects/stimuli that were presented within the timeline segment. Later in the simulation, difficulty levels for the learner were then subsequently increased further to the "Advanced" timeline and then back down to the "Intermediate" timeline at time horizons 2 and 3 respectively. This path through the simulation depicts a user who performed well and was promoted to more challenging timelines within the simulation. However, after reaching the "Advanced" timeline, the learner's performance began to degrade during time horizon 3 and he was moved back down to the "Intermediate" timeline for the duration of the simulation (see Figure 2). The concept of dynamic pathing can deliver many potential user experiences depending on what skill level is initially selected and what timeline paths the user follows through the simulation based on user performance.

The notion of Active Sequencing provides for further customization of the learning experience within a timeline segment based on user performance. Based on the accuracy of decisions and/or actions made by the user, certain injects may be triggered or suppressed from being spawned within a time horizon which dynamically changing the user's simulation experience. This feature helps the learner experience the consequences of his actions and decisions. By making good decisions, the learner can stave off potentially bad things from happening within the scenario. Conversely, making poor decisions or failing to take action can cause bad things to happen within the simulation similar to what happens in real life.

Theoretical Design Approach

Our theoretical adaptive framework is based on a model of performance and skill acquisition developed by Dreyfus and Dreyfus (1986). Their model describes a five-stage approach for skill and performance development ranging from novice to expert. Currently, the government-owned simulation platform used in conjunction with this conceptual model delivers game play at three varying levels of difficulty ranging from novice, intermediate, and advanced. These levels are selected by the learner at the beginning of the simulation. Our enhancement to the existing serious gaming platform is to continue to give learners the option to select their initial competency levels, while allowing the simulation to adjust the difficulty level as the game progresses based upon the performance of its players.

There are many games which currently offer their players a variety of novice, intermediate, and advanced modes of play. Some, like Tetris, require an advancing player to perform essentially the same function at each stage, simply placing ever greater constraints on the time allotted to making decisions. Others, such as most racing games, present players with increasingly skilled opponents and require them to refine the subtleties of their own performance in order to achieve their objectives, while gradually limiting the margin for error. Games of this sort continually adapt in order to keep a player challenged, and can be highly effective for training a person to carry out specific tasks. Games in this vain are not, however, very useful instructional tools for teaching people how to effectively acquire and act upon information gathered from a variety of sources. To this point, few recreational nor serious games have yet been presented with environments which adapt to the ability of a player to manage dynamic and changing situations requiring more than the fulfillment of a single task.

A serious game which teaches its players to manage complex situations cannot adapt itself simply in speed of decision making or by limiting the margin for error. On the contrary, we propose that an effective training tool would in fact increase the possible margin for

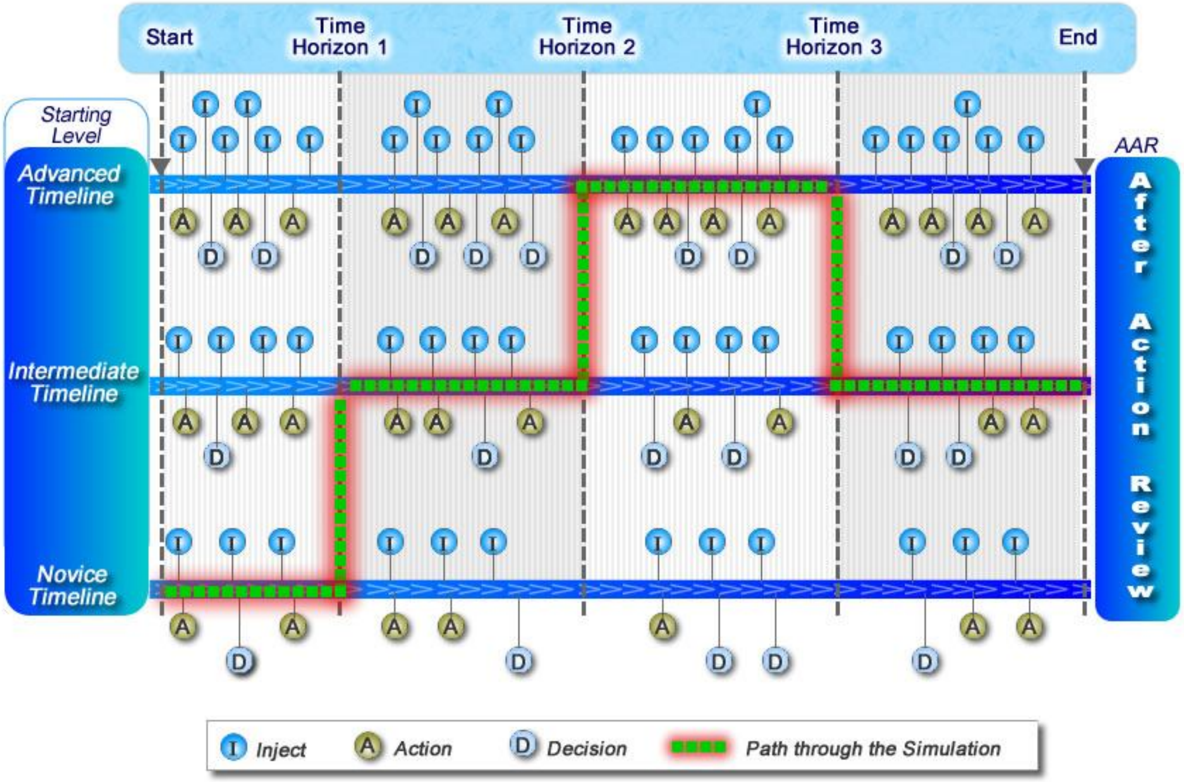

Figure 2. Conceptual Adaptive Framework Using Multiple Timelines

error by a rather large extent. The fundamental suggestion we are putting forth is that an integrated adaptive learning environment of this type is most effective when it offers its players an increasing level of freedom in choosing which actions to take and when to take them in response to the stimuli presented.

Imagine a game in which a player has an ultimate objective at the end of a narrative, with a number of minor objectives set up along the timeline between the start and the finish, the resolution of which all have an effect on the success of the ultimate objective, but none of which can be singularly responsible for its success or failure. A novice player would be prompted by simulated entities with specific information and a choice of several pre-determined actions possible in response at set points within the timeline. He does not have a choice in when he receives the information, its source, or when to take action. His only freedom is to choose between several possible actions. This format is actually employed already by a large number of games in order to familiarize the player with the interface. The idea to constrain the novice player's freedom to make decisions is based on his limited experience and "context-free" rules in situations characteristic of their domain (Dreyfus and Dreyfus, 1986).

Once the player has gained enough familiarity with the interface to know where information can be gathered and what sorts of actions are available within the simulated environment, however, this sort of guidance or mentoring should no longer be needed. At this point, a player reaches an intermediate stage. Obviously there can be several different levels of intermediate skill which a game can address. It is not important how many intermediate levels are available. The significant distinction between a novice and an intermediate player in the type of serious game we envision, however, is the freedom afforded the intermediate player in seeking information and choosing to act upon it. An intermediate player could be prompted with information by a simulated entity, for example, but not given any list of possible actions to take. He would have to decide upon his own action within the constraints of the simulated environment. At another point within the narrative, a player could be presented with a minor objective, and a specific set of actions which could resolve the matter, but no information regarding which option would be most appropriate. It would fall upon the player himself to determine which of the different simulated entities within the game would provide the most useful information for achieving the objective. He could also be prompted with all of the information necessary to take a specific action, but given freedom to implement the action at a time of his own choosing, considering such variables as the changing resources available to him. Essentially the intermediate player is given a less rigid framework for achieving an objective than a novice player, but does not have to be completely self-reliant. Also, he can choose where to find information, or what action to take, or when to take it, but not all at the same time.

Because experts seem to have a highly tuned intuition that enables them to perceive situations and respond quickly and accurately (Klein, 1998), the advanced player, on the other hand, would be given wide freedom of choice in all of these matters. When an objective is presented, the player is responsible for consulting the appropriate simulated entities on his own in order to find out everything he needs to know in order to be able to choose a suitable action to take at an appropriate time. As much freedom as possible will be given to the player in order to accomplish each objective based on a variety of ways and varying degrees of success.

The potential for advancement between the levels will be determined by an assessment of each minor objective within the timeline. A novice should be advanced to the (first) intermediate level simply by completing his first objective. Effort should be made by the creators of the game to provide the player with as much exposure to the interface as possible within the framework of the first objective in order to assure a smooth transition. At the intermediate level, however, advancement should be granted only for a high degree of success in accomplishing a specific objective. This success can be measured in a variety of ways.

A player who achieves a lower degree of success in completing an objective should be dropped down a level, where the increased level of prompting by simulated entities and constraint in available actions should help him focus more clearly upon his specific objective. Furthermore, we suggest that a player be given the option to switch between his current level and novice level at any point, in case he needs the mentoring or interface clarification supplied there.

Assessment of a player’s performance at the intermediate and advanced stages requires first that an objective be broken down into constituent components, such as information gathering, type of action, time of action, etc. If, for example, the game prompts the player with an objective, and the designers of the game have written the narrative such that there are four pieces of information held by the simulated entities which would be useful to the player in making a decision about how to act, a player would ideally take an action immediately after receiving the fourth piece of information. Few players could be expected to act so precisely, however. If information gathering were plotted on a bell curve, with four pieces of information at the apex, then a player who took action after receiving the fourth piece would be awarded a maximum score. If he acts too soon, with only three pieces of information, his score will be slightly less, but not significantly so, and when the objective is completed, the game will prompt him with feedback including the piece of information which was missed and how it applied to the decision to act. If the player continues to seek information after receiving the fourth piece, perhaps not realizing that he already possesses all of the knowledge needed to take responsible action, the assessment will account for each additional attempt to acquire information from simulated entities. Prompting one or two entities for information will only have a slight effect on his overall score, but continued inquiry will have increasingly negative effects as the slope of the bell curve becomes steeper. Similarly, if he acts before acquiring much information at all, his score will be more highly impacted than if he had missed only one relevant piece of information. It is a simple matter to include in the assessment of each objective a clear explanation of why each of these pieces of information is important within the context of the objective.

If the variable is time, an in-game clock could be started at a specific point along a timeline. The player is expected to take an action at a time of his choosing between the start and the stop of the clock. An ideal time for action must be chosen by the game designers. A bell curve could be drawn over that period of time. If the player acts too soon, or too late past the ideal time for action, he would receive a much lower score than if he acted within a standard deviation of the ideal moment. When the objective is completed, a less successful player could be prompted with specific reasons as to why he should have acted sooner or later than he chose to act, and how that choice affected the resolution of the objective.

If we imagine a situation which combined variables in both time and information gathering, it is possible that the player finds himself unable to acquire all of the relevant information by the time the ideal moment for action in the timeline arrives. If he takes action at this point, he should be given a full score for the time of the action, and a reduced score for the amount of information he had when he acted. The prompt after the resolution of the objective should explain to him that the time to act was appropriate, but reveal to him the significance of the information which he failed to acquire, and advise him where to locate similar pieces of information when encountered with future objectives in the same vain. The more advanced a player is, the more variables an adaptive learning simulation would present him with.

Figure 3 details a situation where a player has the freedom to choose between a number of possible actions, and to decide for himself the best time to take those actions. The combination of the three graphs shows how the bell-curve assessment strategy may be applied to similar situations, where freedom is afforded to a player within multiple variables in a simulation. The action assessment graph, on the left, is tilted ninety degrees from customary orientation for convenience. In this simulation, there are seven actions labeled A through G that are available to a player. Action D is theoretically the best solution to the scenario presented, and a player who chooses action D would receive the highest possible score from this graph for that choice. The time assessment graph shows that the best time for a player to take an action in order to resolve this scenario is at time signature 3. Acting within a standard deviation of time signature 3 still has a profoundly positive effect on the resolution, but the further from time signature 3 that a player takes action along this timeline, the less effective his action will be in resolving the situation. Nevertheless, taking any sort of action in order to resolve the

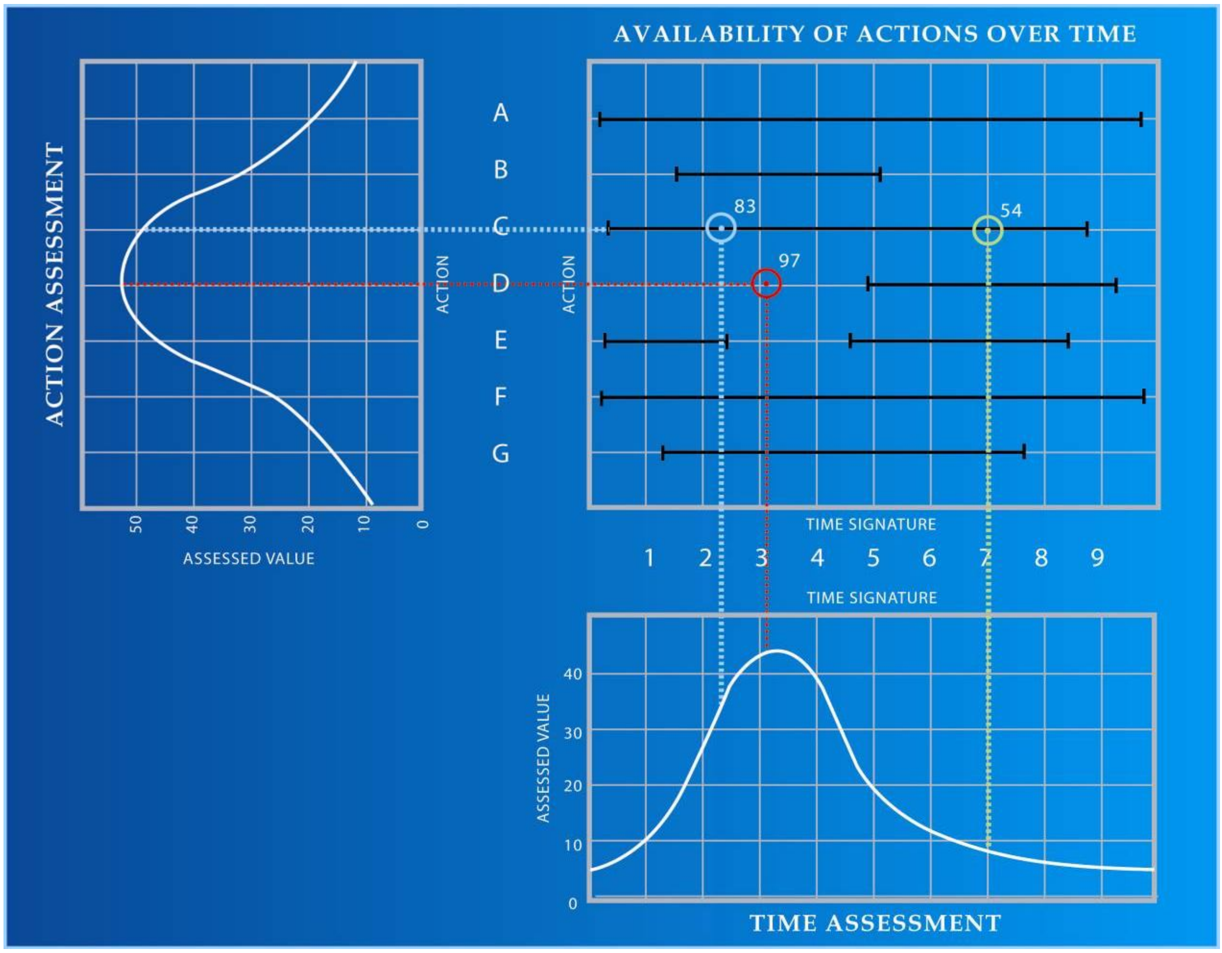

Figure 3. Availability of Actions Over Time Assessment

situation, even if grossly late, still has some value. It is for this reason that the bell curve is particularly suitable to these assessments, rather than a straight line.

The graph showing the availability of actions over time, however, complicates the matter. Rarely, in a real world situation, are all actions available to a person at all times. Actions A and F, for example, are available essentially for the duration of the simulation. Action E, however, is available only at the beginning and end, and action D only after about time signature 4.8. As the simulation plays out, the actions available to a player at any particular moment can change. Theoretically, the best action to take and the best time to take it would be action D at time signature 3. This theoretical maximum is represented by the red dot. As one can see, in this particular simulation, a player does not have action D available at time signature 3, and is therefore forced to make a choice about taking less-than-ideal actions at less-than-ideal times. The blue dot represents a relatively successful choice by a player, who takes a very good action, even if slightly early. The green dot represents a player who takes the exact same action, but far too late for it to resolve the scenario in an effective manner. The numbers next to each dot are the sum of the assessed actions and time signatures from their respective graphs. The assessment numbers in this specific graph, however, are quite arbitrary, and can be adjusted to fit any specific simulation. Here, they are merely intended to illustrate a relative difference between possible choices. As one can see, the freedom afforded to an advanced player in such an environment greatly increases the margin for error from an ideal. For that reason, it is important that the feedback which a game provides to a player after the completion of an objective take these matters into consideration. The player represented by the blue dot should be praised for choosing a suitable action, even if it is not the ideal solution to the situation, because he acted in a timely manner. The player represented by the green dot, on the other hand, should not only be corrected for delaying his action for such a long time, but also for choosing a less-than-ideal action when the ideal action was available to him.

A less advanced player presented with the exact same scenario would be limited with respect to one of the variables. He could, for example, be prompted by the game and told that he must take action E. He would then have the freedom to choose when to take that action, but nothing else. Conversely, a less advanced player could be prompted to take action at time signature 4. He would have the free choice between the five possible actions available at that time, but nothing else. A more advanced player, on the other hand, would be presented with the freedom of choice in increasingly greater numbers of variables, such as the gathering of information or the tapping of resources. No matter how many variables are included, however, assigning specific values to the choices available to a player, as we do here, makes it a simple task for the designer of a game to set limits above and below which the level of a game will automatically change. Perhaps our player represented by the blue dot will advance to the next level of difficulty because he scored above a 75, and our player represented by the green dot will regress for scoring below 55. The specific numbers are unimportant, and depend upon the discretion of the game developer. Bell curve assessments, however, are a simple way of quantifying the results of decisions made by players with respect to multiple variables.

CONCLUSION

The purpose of our theoretical adaptive framework discussed in this paper is to demonstrate a blueprint toward developing an adaptive serious gaming system that could provide a greater opportunity for an organic learning experience whereby individuals can be consistently challenged in a meaningful way and granted freedoms to face challenges as they arise. This framework is meant to be a starting point for modeling, simulation, and educational professionals to continuously refine and advance as adaptive learning technologies are developed and implemented.

ACKNOWLEDGEMENTS

We acknowledge the contributions and support of our colleagues Dr. Laura Zimmerman, Dr. Gary Klein, Dr. Beth Veinott, Dr. Shane Mueller, Matt Kraus, Bernadette Becerra, and Mary Ann Pigora of ECS.

REFERENCES

Brusilovsky, P., Pesin, L., & Zyryanov, M. (1993). *Lecture Notes in Computer Science. In G. Goos & J. Hartmanis (Eds.),* Towards an adaptive hypermedia component for an intelligent learning environment (Vol. 753, p. 348-358). Berlin: Springer Berlin / Heidelberg.

Brusilovsky, P. (1999). Adaptive and Intelligent *Technologies for Web-based Education. In C. Rollinger & C. Peylo (Eds.)* Künstliche Intelligenz 4, Special Issue on Intelligent Systems and Teleteaching, 19-25.

Brusilvovsky, P., & Paylo, C. (2003). *Adaptive and Intelligent Web-based Educational Systems.* International Journal of Artificial Intelligence and Education, 13, 156-169.

Clark, R. C. (2008). *Building Expertise: Cognitive Methods for Training and Performance Improvement*. (Vol. 3rd Edition). San Francisco, CA: Pfeiffer & Company.

Davis, P. K. & Kahan, J. P. (2007). *Theory and methods for supporting high level military decision making.* (Technical Report prepared under contracts F49642-01-C-0003 and FA7014-06-C-01). Santa Monica, CA: Rand Corporation

Derryberry, D., I'm Serious.net. (2007, Sept. 26). *Serious Games: online games for learning* San Jose, CA: Adobe Systems Incorporated. Retrieved June15,2009,from http://seriousgames.ning.com/forum/topic/show?id=63075

Dreyfus, H.L., & Dreyfus, S.E., (1986). *Mind over machine: The power of the human intuitive expertise in the era of the computer*. New York: The Free Press.

Klein, G. (1998). *Sources of power: How people make decisions*. Cambridge, MA: MIT Press.

Specht, M. & Burgos, D. (2006). Implementing *Adaptive Educational Methods with IMS Learning Design*. In ADALE Workshop Proceedings of AH2006 Conference.

Zielke, M. A., Dufour, F. A., Christopher, T. V., *Donahue, J. K., Jennings, E. B., et al. (2009). Serious Games for Immersive Cultural Training: Creating a Living World.* IEEE Computer Graphics and Applications, 29(2), 49-60.

Zyda, M. (2005). *From visual simulation to virtual reality to games.* IEEE Computer Society, 38(9), 25-32.