Customizing Interactive Training Through Individualized Content and Increased Engagement

Brian Magerko
Michigan State University
East Lansing, MI
brian@magerko.org

Robert E. Wray, Lisa Scott Holt, Brian Stensrud
Soar Technology, Inc.
Ann Arbor, MI
wray, lholt, stensrud@soartech.com

ABSTRACT

Simulation-based training offers the potential for relatively low-cost training available at any time and almost any duty station. However, a main drawback of simulation-based training is the lack of oversight in the training process. Simulations often depend on a fixed number of pre-defined training scenarios that are designed to test training objectives but not to deliver a training experience customized to the specific trainee's current level of skill and understanding. In this paper, we introduce the Interactive Storytelling Architecture for Training (ISAT), which uses an intelligent agent, the *director*, to assemble training scenarios that test the skill level of individual trainees. The director also provides indirect feedback about trainee actions during the execution of a training scenario, subtly adapting the training environment to stress unmastered skills and suggest remediation. This approach results in a training experience that is specialized to the trainee's individual needs and potentially more engaging, resulting in faster development of trainee proficiency.

ABOUT THE AUTHORS

Brian Magerko is an Assistant Professor of Educational Gaming at Michigan State University. He received a Ph.D. in Computer Science and Engineering from the University of Michigan. His doctoral work focused on building a general interactive drama architecture and exploring the benefits of predictive modeling for drama management. Dr. Magerko's research interests include interactive storytelling, believable agents, knowledge-based systems, cognitive architectures, user modeling, and machine learning.

Robert E. Wray is Chief Scientist at Soar Technology, Inc. He received a Ph.D. in Computer Science and Engineering from the University of Michigan. His doctoral work focused on maintaining logical consistency in agent reasoning systems and his innovations were incorporated in the Soar architecture. Dr. Wray's research encompasses many areas of artificial intelligence, including agent-based systems and agent architectures, machine learning, cognitive science, and knowledge representation and ontology.

Lisa Scott Holt is a Research Scientist at Soar Technology, Inc. Dr. Holt spent 9 years at the Center for Innovation in Learning at Carnegie Mellon University. During her time there, she received her Ph.D. in Cognitive Studies in Education from the University of Pittsburgh in 2001. For her dissertation she designed, implemented and evaluated an interactive tutoring environment for introductory physics. Dr. Holt's current research focuses on the design of technology tools to support learning and the design of interfaces based on an information-processing model of human cognition. She is also interested in methods to assess what learners actually know and can do with their knowledge.

Brian Stensrud is a Research Scientist at Soar Technology, Inc. He received his Ph.D. in Computer Engineering from the University of Central Florida (2005), and also holds B.S. degrees in Mathematics and Electrical Engineering from the University of Florida (2001). His doctoral dissertation dealt with learning high-level tactical behavior using a modified neural network architecture. Prior to joining Soar Technology, Dr. Stensrud was involved in a variety of research projects as a member of UCF's Intelligent Systems Laboratory (ISL) involving human behavior modeling, neural network and evolutionary programming applications, affective reasoning, and robotic applications of agent-based systems.

Customizing Interactive Training Through Individualized Content and Increased Engagement

Brian Magerko
Michigan State University
East Lansing, MI
magerko@msu.edu

Robert E. Wray , Lisa Scott Holt, Brian Stensrud
Soar Technology, Inc.
Ann Arbor, MI
wray, lholt, stensrud@soartech.com

INTRODUCTION

Advanced Distance Learning (ADL) technologies are critical for training war fighters in distant locales where schoolhouse access is limited and resources are scarce. Immersive computer games are increasingly used as an ADL technology because they provide compelling, engaging learning experiences that have been shown to reduce training time and increase training persistence. Game technology also represents a significant step forward in the "train as we fight" vision of effective, anytime, anywhere training. In virtual environments, the realism of the training experience helps trainees obtain a better sense of "how" and "why" they are learning to perform some task, whether it's an urban room clearing exercise, fire training in a shipboard engine room, or an aircraft maintenance scenario.

Human-controlled systems where a trainer actively directs a trainee's experience are usually the preferred training environments. However, human trainers are a very constrained resource and are usually only available at prescribed times and duty stations. They are often quite constrained in providing individualized training. For example, in the U.S. Army's 91W10 combat medic training, a schoolhouse instructor demonstrates medical procedures (such as applying a tourniquet) to trainees using a dummy to represent an injured person. The instructor typically guides one student through the process hands-on, while the other students in the class observe the target performance. Giving each trainee hands-on experience in each medical procedure would maximize training effectiveness and persistence. In a schoolhouse situation with limited time and resources, however, hands-on training for all is usually not practical.

Automated systems, such as virtual reality training, address some of the limitations of human-mediated training. They may cost less to operate and can offer individual trainees the experience of performing the trained skills. Computer games are increasingly used as very low-cost virtual training platforms with a computational profile that allows them to be used for ADL, as well as in more direct training situations. Computer games provide somewhat adaptive experiences by reacting to trainee actions, but, in general, game behaviors are typically only loosely or indirectly coupled to training goals. Generic training games can result in poor training or training experiences that must be repeated many times to achieve a desired effect. Further, skill assessment (based on behavior) can be difficult and may also be less precise than in most web-based learning environments where skill is evaluated based on responses to specific questions. Without appropriate oversight, computer games can result in negative training. Unlike the schoolhouse, a human trainer is rarely present to guide the trainee and provide the most appropriate experiences. Within the freedom of the game environment, the trainee is able to perform many actions, some of which are appropriate for the training and some of which are not. Thus, while computer games offer "fun" experiences for the trainee, without careful design, the training experience itself can either be lost in the game play or fail to take advantage of the freedom offered by the game environment.

Our approach to interactive training, the Interactive Storytelling Architecture for Training (ISAT), is designed to address the limitations of computer games for ADL outlined above and to fully realize the potential of games as engaging, individualized learning environments. To accomplish these goals, ISAT introduces an intelligent agent into the software environment. This agent, which we call the *director*, serves a role similar to the schoolhouse trainer. The director chooses and customizes training scenarios based on a trainee's skill level and previous experiences. The director also subtly guides the trainee through a scenario by dynamically adapting the environment to the dramatic needs of each scene and the learning needs of trainee. Unlike the schoolhouse trainer, the director acts "behind the scenes" and is imperceptible to the trainee, which allows the training experience to remain immersive and engaging. This approach applies earlier work on interactive storytelling, the Interactive Drama Architecture (IDA) (Magerko, 2005; Magerko and Laird, 2004), to the interactive training domain. ISAT extends IDA to the requirements of training systems, including a skill model that the director uses to evaluate the proficiency of a trainee with respect to training objectives and to tailor training scenarios to exercise specific skills. This paper will present ISAT's design, both the design implemented in a prototype system and the approaches to be explored in a complete implementation.

Requirements for an Interactive Trainer

Before describing ISAT's design and initial implementation, we first address the requirements guiding the design and development of an interactive training architecture.

1. *Training Effectiveness*. A training system should effectively train the desired skill set. This requirement subsumes all others in that the end measure of a training system is how accurately and efficiently trainees acquire the targeted skills.

2. *Engaging Experiences*. Engaging systems have a greater likelihood of providing persistent and efficient training. A system is more engaging if it provides some secondary motivation for a trainee (such as interesting gameplay or a compelling narrative) in addition to being "required training." An important requirement for engagement is that the training situation be believable and realistic. Believability does not necessarily imply high fidelity models of the environment and the actors in it (Wray & Laird, 2003), but rather trainee experiences should be consistent with what one would encounter in a real situation.

3. *Individualized Training*. In order to maximize effectiveness, training should adapt to the particular needs of individual trainees. Computer-based training technologies should scaffold (provide the trainee with needed guidance and support) and fade (gradually remove or reduce scaffolding to promote independent performance) (e.g., Collins, et al, 1989; Kintsch, 1993; Vygotsky, 1978). The system's guidance and feedback is tailored in accordance with the trainee's performance.

4. *Generality*. Any architectural approach to interactive training should be one that could be applied across domains and training environments. A specialized technology that performs well in a single domain but not others is certainly useful, but does not have the utility of an approach that can be applied in many different domains.

Previous Approaches

In recent years, pedagogical software agents have been implemented in interactive learning environments to individualize the learning experience, providing learners with needed support (scaffolding). Users learn and practice new skills in a virtual world, and the instructional system adapts to individual learners through the use of artificial intelligence. Intelligent agents can interact with the learner directly (e.g., as a coach or learning companion) or indirectly by altering the environment.

Software agents have been used to realize a range of support and guidance in educational technology systems. A software agent is simply a software system designed to interact (perceive and act) within an environment, including physical environments (e.g., a robot) or virtual environments (e.g., a computer game "bot"). Pedagogical agents have been developed to provide explicit tutorial guidance in training systems (Rickel and Johnson, 1997), and these developments have demonstrated the power of using software agents in knowledge-based learning environments (Johnson, Rickel, and Lester, 2000; Moreno, Mayer, and Lester, 2000; Lester, et al, 1997a; and Lester, et al, 1997b). In most cases, pedagogical agents have been represented explicitly in the simulation environment as an animated virtual instructor that is always present. The explicit representation of the instructor may be appropriate for some training, but it also has the serious drawback of compromising the realism of the training scenario (see Requirement 2).

ISAT builds on the notion of a pedagogical agent that guides user experience but is implicit rather than explicit. An implicit pedagogical agent does not provide direct guidance but rather structures the environment to give the trainee appropriate experiences (i.e., tailored to needs and skill level). Because the agent is not represented explicitly, the trainee's sense of engagement and responsibility for making decisions is naturally and effectively supported. As the trainee's skills improve, the pedagogical agent makes fewer and less frequent remediation actions, providing a natural implementation of fading. Pedagogical agents that directly support remediation and storytelling should efficiently support a large numbers of users, individualize training, and provide simulated experiences that transfer to improved performance in the real world.

Our approach to altering the trainee's experience is similar to that used in intelligent tutoring systems (ITSs). The defining feature of an ITS is that it carefully oversees a learner's work to provide needed guidance. ITSs incorporate a rule-based expert model of the target skill that is used to monitor and guide novice learners as they engage in the new activity. The intent of an ITS is to model the actions and interventions of a human tutor, which is the most effective means of instruction (Bloom, 1984). ITSs identify the need for instructional interventions by comparing a model of expert performance with a model of the learner's performance (Corbett, 2001). However, the pedagogical nature of these interventions is quite limited (see Requirement 3). Model-tracing methods force the learner to proceed in steps of a specified grain size corresponding to the underlying rules, which constrains the progression of the learner's actions. If the learner makes a recognizable error that has been pre-programmed as a buggy rule in the system, an error message is presented. If the learner asks for help, a help message is presented to guide the student toward the correct solution. These error and help messages are very context specific, as they are generated by matching the learner's solution with the underlying model of expert solution (Anderson, Corbett, Koedinger, & Pelletier, 1995). This information processing approach is useful where the representation of domain knowledge has a significant rule component such as algebra, physics, or language (Johnson, et al, 2004) (see Requirements 1 and 4), but not as useful in unstructured learning domains.

THE INTERACTIVE STORYTELLING ARCHITECTURE FOR TRAINING

ISAT is designed to provide individualized training through the *real-time adaptation of stories.* We are developing general techniques and components (an architecture) that will enable existing and new training systems to dynamically adapt content to support training goals and increase trainee engagement. The ISAT architecture is focused on an internal and imperceptible agent, the *director*, in a constructive simulation environment. To explain how ISAT works and to test its operation, we will describe ISAT in the context of a specific example training system. However, ISAT is designed as a general capability, rather than a component in a single training application (see Requirement 4). Our future plans include integration of ISAT with additional training domains. Multiple, successful integrations will provide empirical evidence of what is now only a claim with respect to generality in design.

Architectural Design

ISAT is heavily influenced by the aforementioned Interactive Drama Architecture (IDA), which employed an intelligent director agent to mediate between story content, synthetic character behavior, and the player's actions in the game environment (Magerko, 2005; Magerko and Laird, 2004). ISAT's design, shown in Figure 1, spans the virtual environment, the trainee, a human trainer who authors training content, the synthetic or "non-player characters" (NPCs) characters that populate the game environment, and the director agent. The trainer authors story content with both dramatic and teaching principles in mind. Future work will investigate authoring tools suitable for this task. The director is responsible for managing the trainee's experience in response to the training scenario content, the trainee's actions in the game world, and the skill model of the trainee, which tracks his mastery across a set of domain skills. The director can influence the trainee's experience by giving synthetic characters certain actions, by altering the environment (e.g., spawning characters or environmental sounds), or through selection and instantiation of scenario content.

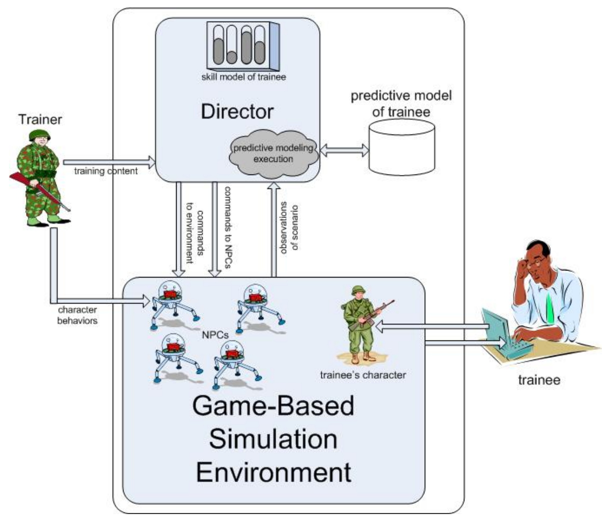

Figure 1. Components of the Interactive Storytelling Architecture for Training.

Environment

Our current prototype system, shown in Figure 2 below, is implemented in the 91W10 Tactical Combat Casualty Care Simulation (TC3), a combat medic training system under development by Engineering and Computing Simulations, Inc. (ECS). The TC3 trainer is an immersive 3D simulation that provides the knowledge, skills and practice necessary to significantly improve the performance of combat medics. The courseware is based on the same systematic approach for Basic and Advanced Trauma Life Support that is used to train Emergency Medical Technicians (EMTs). However, it requires students to demonstrate an ability to perform these tasks in tactical situations that are dictated by the principles of Tactical Combat Casualty Care.

The existing TC3 proof-of-concept implementation highlights "Tactical Field Care," which means that the medic is out of direct fire, but the threat has not been defeated (versus "Care Under Fire," where the medic is under direct fire). After an initial introduction, the trainee sees the simulated world from a first-person point of view, holding his rifle. The trainee may then survey the situation by moving around in a 100-ft. radius circle of terrain representing a courtyard in an Afghani village. Trainee actions include slinging the rifle, manipulating casualties, and accessing items in the medical bag to assess casualties, prioritize treatment, and treat injuries.

Figure 2. Screenshot from prototype ISAT implementation in 91W10 TC3 Trainer.

Skill Model

The director keeps track of the trainee's mastery of target skills being trained. As the trainee executes actions in the environment, the director relies on knowledge that indicates when a trainee is performing well or poorly in a particular skill. For example, if the trainee applies a tourniquet too far from the wound to be effective, then the director would downgrade the "apply tourniquet" skill. The set of target skills used in the TC3 trainer is shown in Table 1. The current ISAT prototype employs a simple, qualitative skill model. Developing more fine-grained, multi-dimensional skill models (such as using *recency* and *quantity* of training on a given skill) will be explored in future work. This will entail modeling player skill across time, including how skills should decay and the relationship between frequency of testing and scaffolding.

Director Actions

The director is responsible for managing the training environment (e.g., coordinating synthetic characters), hypothesizing about the trainee's mastery of the targeted skills, and selecting training content. The director can affect the world in accordance with these roles by executing various actions in the environment.

Table 1. Target skills used by ISAT director in the 91W10 TC3 Trainer.

Skills	Description
Care under fire	In a hot zone situation, return fire and help secure the area before attempting treatment or extraction.
Find temporary fighting positions	Minimize exposure to fire by staying low to ground and using objects in the environment.
Prioritize casualties	Decide the order in which patients should receive care.
Secure casualty	Make sure the casualty poses/can pose no threat.
Check vitals	Check casualty's breathing, heart beat, and skin color.
Apply tourniquet	Decide when a casualty/situation requires a tourniquet and apply the tourniquet.
Manage airway	Open airway if the patient's breathing is labored.
Manage chest wound	Take correct steps to treat a casualty with a chest wound.
Monitor	Monitor casualty until evacuation
Extraction	Properly extract a casualty from a hot zone.

Scene Choice

A scenario is defined as a partially-ordered collection of scenes. Each scenario is authored by a human trainer and encapsulates a single training experience. Scenes in a scenario are temporally-ordered to maintain consistency in the training content, akin to the temporal ordering used in planning languages. Once all of a given scene's parents have been performed, then that scene can also be considered for performance. We use standard languages and representations developed by artificial intelligence researchers to represent dependencies and interactions among scenes.

A scene is defined as a set of skills associated with abstract plot content. Each scene has a set of skills that it could possibly test. The abstract content authored in each scene is intended to describe the basic situation, leaving the specifics to be instantiated by the director at run-time. This allows the trainer to specify scenario content at an abstract level, offering more possible game experiences and training scenarios that are more tailored to the trainee's needs (see Requirement 3).

The director is capable of selecting the next scene based on both pedagogical and dramatic criteria (see Requirements 2 and 3). The director compares the trainee's current skill model with the available scenes (and the skills associated with those scenes) to decide which scene should come next. The skill model represents the set of skills that the system is supposed to be training as well as the trainee's demonstrated proficiency in those skills. As the trainee progresses, the director updates the skill model, noting at which skills the player is good or bad. After the completion of a scene, the director compares the trainee's skill model with the scenes that could logically come next. The director looks at skills at which the trainee is bad or untested and chooses the scene that would best test those skills. Next, the director instantiates the abstract scene in the context of both the current state of the world and the subset of skills that specifically applies to the skill model's current state.

This scene selection method is based purely on pedagogical motivations and does not address engagement. A modification we are considering to this approach, which will be explored in the full TC3 trainer implementation, is the use of a scene choice heuristic that takes account of both the skill model and a model of dramatic progression. This approach is akin to the use of heuristics for plot content choice in interactive storytelling systems (Mateas and Stern, 2002; Weyhrauch, 1997; Magerko and Laird, 2004). This method requires the annotation of scenes along a series of author-defined dimensions (e.g., "tension" or "interest"). As the trainee experiences the scenario, a set of scenes would be initially selected based on the skills tested, as described above. The director would choose the next scene from this candidate set based on the scenes' dramatic ratings. A sensible, straightforward heuristic might be to choose scenes in terms of tension along a traditional Aristotelian arc — introduction, dramatic tension rising to a climax, followed by a resolution. At the beginning of the scenario, less tense scenes would be chosen, followed by more tense scenes, up to a climax, etc. This approach provides a two-tiered method of scene selection that puts a premium on training value during selection, but does so with dramatic and narrative goals considered as well.

Story Direction

We define *direction* to be any action that the director executes in the training environment. This may include instantiating synthetic characters or objects, giving commands to synthetic characters (e.g., new goals or knowledge), and altering the environment itself (e.g., creating an ambient sound). *Story direction* is specifically an action taken by the director agent to fulfill plot content. For example, if the ambush scene involves the trainee's platoon leader yelling "Take cover!" the director would send a command to the platoon leader character to perform the dialog. ISAT does not require strongly autonomous agents that rely on their own goals while interacting with the trainee & the environment to populate the world (see Mateas and Stern, 2002, for a discussion of the weaknesses of strongly autonomous agents in interactive storytelling systems).

Skill Direction

S*kill direction* indirectly guides the trainee in performing a particular targeted skill. If a trainee has not been tested in a given skill, such as "prioritize casualties," or has already shown to be particularly bad at this skill, then the director will give explicit in-game advice to the trainee, such as directing the CO to order to the trainee to "Be sure to attend to the casualties in critical condition first!" As the trainee shows aptitude at a given skill, the director's guidance will fade into progressively less explicit strategies, such as directing the most critical patient to gasp several times to get the trainee's attention. In this manner, ISAT directly supports the notion of fading. Fading of guidance in skill direction is a way for the system to adapt the training world to the trainee's particular needs across the duration of the training (see Requirements 1 and 3).

Reactive Direction

Reactive direction is a technique used in interactive storytelling systems to prevent player actions from moving the experience outside of what is covered by the story content (Magerko and Laird, 2004; Young, et al, 2004). Authors for interactive stories create "story spaces" (Magerko, 2005) when authoring plot content for an interactive story, such as character behaviors, rules for generating dramatic situations, or plot points. If a player executes an action that takes the story outside of that story space, then the system executes reactive direction to modify the world to keep the player's experience within the story space. For example, if the player tries to shoot a character that is important to future plot developments, the system makes the gun misfire or jam.

ISAT applies this approach to the highly interactive environments for which it is designed. The more interactive a world is (i.e., the more possible actions in the world), the more possible it is that the trainee could execute an action that prevents the logical continuation of the story. For example, getting lost in the environment is completely irrelevant to the training of a combat medic and would harm the trainee's progression through the training scenario. It is the director's role to execute reactive direction in the world, such as noticing when the trainee is getting away from his squad and directing a fellow soldier to fall back to help guide him. This approach helps address the problems that can arise in an interactive environment without relying on the standard, less believable techniques used in commercial computer games, such as disallowing any action that isn't the "right" one (see Requirement 2).

STORYBOARD

This section describes an example storyboard to be used for our future implementation in the 91W10 TC3 trainer. Rather than convey specific details, the goal of this section is to illustrate the kinds of scenarios and actions ISAT will be able to support by the conclusion of implementation. We will focus on showing the director's decisions and actions with respect to the proficiency of a skill. Initially, we have chosen a simple, qualitative model of proficiency: untested, bad, fair, neutral, good and very good. In future work, we will develop richer models of proficiency. For example, if a trainee demonstrated "very good" proficiency at a skill in a training session over a month ago, then the director might choose to test that skill's persistence in the trainee. The skill model uses the skills listed previously in Table 1. (*Director actions are parenthesized and in italics.)*

Scene 1: Bombing in a Marketplace

The initial scene takes place in a village market in Afghanistan. The trainee arrives on the scene of a bomb explosion in the marketplace as part of a quick response force (QRF). There are only civilian casualties, mainly in a bazaar open to the market. The rest of the squad quickly secures a perimeter as the trainee is ordered to tend to the wounded in the bazaar. There are four casualties, and two that are obviously alive and moving. It is only a matter of time before an ambush could happen. The trainee needs to quickly prioritize, diagnose, treat, and then secure the patients.

Once in the bazaar, the trainee quickly checks the pulse and breathing of all of the casualties. He checks the worst-looking casualty first, which yields no pulse or breath (*upgrade check vitals*, *downgrade secure casualty*). He moves on to the next casualty, who has a leg blown away, without doing any more (*upgrade prioritize*). He checks this casualty's vitals (*upgrade check vitals*, *downgrade secure casualty*) and sees that he is still alive. He quickly applies a tourniquet, but too high up from the wound and without removing minimal clothing (*downgrade apply tourniquet*).

He moves on to the next casualty, who has a chest wound from shrapnel (*upgrade prioritize*). He checks for weapons, then checks the casualty's breathing, which is labored but existent (*upgrade check vitals; upgrade secure casualty*). The trainee removes clothing to bandage the entry wound (*upgrade manage chest wound*), sits the casualty up (*upgrade manage chest wound*), and bandages the exit wound (*upgrade manage chest wound*).

The trainee moves on to the last casualty, who is already dead from loss of limb (*upgrade prioritize*). He forgets to check for vital signs (*downgrade check vitals*). As he begins treatment, the director moves on to a new scene. The first scene ends either when 12 minutes have passed or when the trainee is attending to the last casualty (*begin next scene*). Scene 2 is authored to be an ambush on the quick response team. The particular details of the attack depend on the trainee's skill model. At this point, the director's assessment of the trainee's skills would be:

- *check vitals*: fair
- *prioritize*: very good
- *stay covered*: untested
- *secure casualty*: neutral
- *apply tourniquet*: fair
- *extraction*: untested
- *manage airway*: untested
- *manage chest wound*: very good
- *monitor*: untested
- *soldier first*: untested

Scene 2: Ambush of Quick Reaction Force

The director will have the ability to map a scene description to an instantiation that matches this particular trainee's needs and test untested or less than good skills. Therefore, in preparing the next scene, the director will:

- Choose a scene that can test *stay covered*, *extraction*, *monitor*, *soldier first*, *secure casualty*, *apply tourniquet*, and *check vitals*.
- Spawn the ambushers with an RPG and small weapons to test *stay covered* and *soldier first*.
- Give the enemy a goal of firing on the trainee when he appears from the bazaar to test *stay covered*.
- Give the enemy a goal of hitting at least one soldier with the RPG (creating a wound to test *apply tourniquet*), and wounding at least three people total to test *check vitals* and possibly *monitor* and *extraction*.
- Give the injured soldier the feature of being hysterical after being hit to test *check secured*.

While the trainee is attending to the third patient in the bazaar, shooting begins outside and there is a loud explosion. The trainee continues checking the dead amputee for vitals (*downgrade soldier first*). A US soldier yells "medic!" The trainee rushes outside and then pauses to take in the situation (*downgrade stay covered*). The ambushers direct fire toward the trainee, which causes him to crouch low and move toward the burning Humvee parked near the bazaar. The trainee surveys the situation as weapons fire from several directions across the square outside the bombed bazaar. Within a few yards are two wounded soldiers. One is lying motionless on the ground, while the other is screaming for a medic. The trainee moves immediately to the screaming soldier (*downgrade soldier first*). The director reactively directs the platoon leader to yell at the medic: "We need some more firepower!" The trainee unslings his weapon and begins returning fire.

After a few minutes, the shooting subsides. The response team moves out in a mop up operation. There are two casualties visible to the trainee, the two wounded soldiers. There is also a third casualty that the trainee cannot see, on the other side of the Humvee. The trainee moves to the soldier who had been screaming and checks his vitals. The soldier has received multiple gun shot wounds and is hysterical, but his vitals are stable. The trainee begins to apply pressure to the wounds to stem the bleeding without having first safetied the hysterical soldier's weapon (*downgrade prioritize*, *downgrade secure casualty*).

The trainee moves next to the unmoving soldier. His vitals are erratic and barely perceptible (*upgrade check vitals*). The trainee suspects a chest wound and begins to remove clothing from the soldier. He finds the wound and immediately begins to bandage it without testing for an exit wound (*downgrade sucking chest wound*).

The trainee goes behind the Humvee and finds another US casualty. He checks vitals, which are reasonably good (*upgrade check-vitals*). The soldier has lost the lower part of his left forearm in the RPG explosion. He cuts the soldier's uniform away from the wound and applies a tourniquet close to the wound (*upgrade apply-tourniquet*).

With the situation secure and the patients mobilized, the trainee calls for an evac to extract the wounded soldiers (*upgrade extraction*). He continues to monitor the vitals of each patient periodically while awaiting the ambulance (*upgrade monitoring*).

At the conclusion of Scene 2, the skill model would approximately be:

- *check vitals*: very good
- *prioritize*: fair
- *stay covered*: bad
- *secure casualty*: bad
- *apply tourniquet*: fair
- *extraction*: good
- *manage airway*: untested
- *manage chest wound*: fair
- *monitor*: good
- *soldier first*: bad

When the director considers how to construct Scene 3, *soldier first* and *secure casualty* will be high priority skills. The director also notes the trainee has not yet been tested in *manage airway*, and will select an injury to test that skill.

Scene 3: Offensive Infantry Operations

Based on the current skill assessment, for Scene 3 the director will:

- Choose a firefight scene that can possibly test *stay covered*, *soldier first*, *manage airway*, *manage chest wound*, *apply tourniquet*, *secure casualty*, and *prioritize.*
- Create the enemy squad in the firefight sequence to use rifles.
- Make sure one of the soldiers is shot in the throat to test *manage airway*.
- Leave behind a civilian casualty who is booby-trapped to test *secure casualty*. Because the trainee has repeatedly failed to demonstrate mastery of *secure casualty*, the director chooses an action that will directly demonstrate why this action is important.

DISCUSSION

ISAT addresses the requirements laid out for an interactive training system.

- Requirement 1: *Effectiveness* should be at least as good as the underlying training system. The suggestion from training literature that both engagement and individualization improve training effectiveness indirectly suggests that ISAT should improve overall effectiveness, given that it incorporates both of these properties (Lepper and Henderlong, 2000). However, the actual impact of ISAT on training effectiveness has not yet been investigated. As ISAT is more fully instantiated in the TC3 simulation, we plan to perform user tests to evaluate the impact of ISAT on training effectiveness.

- Requirement 2: *Engagement* is addressed by the use of an immersive, 3D game environment, as well as by the heuristic selection of training content.
- Requirement 3: *Individualization* of the training experience is achieved through the various roles played by the director agent, such as story direction, skill direction, and reactive direction.
- Requirement 4: *Generality* has been a guiding concern throughout ISAT's design, but has yet to be shown. ISAT, because of its attention to the larger-grained size of real-world actions, seems specifically suitable for more skill-based domains, such as training medics, versus more cognitive-based domains, such as language training (Johnson, et al, 2004). Future work will entail using ISAT in several environments to show the applicability of this approach across domains. However, there is no aspect of ISAT that specifically depends on the domain that we have selected as our first environment.

ISAT's design, while implemented in a prototype system, has several facets that have yet to fully mature and will be refined in future design. The next implementation will extend the director's role in the 91W10 TC3 trainer to support the capabilities described in the storyboard.

Skill Model Refinement

The skill model used for the ISAT TC3 prototype is simple. While it does demonstrate the benefits of using a skill model for skill direction and scene selection, it does not detail how the skill model should be precisely defined in full implementation. How to represent the decay of skills over time, to interpret how often a skill has been tested, and to encode "getting better" or "getting worse" at a skill will all be examined in our future research and implementation efforts.

Scene Instantiation

After selecting a scene based on its pedagogical and dramatic content, the director must instantiate the scene's abstract content. As mentioned above, a scene is partially defined by an abstract description of the content, such as "the trainee's squad is ambushed by local insurgents." How this content is instantiated by the director depends again on the trainee skill model. The director applies a transition function that translates between the abstract content, the current state of the game, and the trainee's skill model. If the trainee is particularly bad at applying tourniquets, for example, then this ambush scene would be instantiated to include an enemy fighter with a rocket-propelled grenade, a weapon that can cause injuries requiring the application of a tourniquet. This transition function must include mappings from the skill model to instantiated content, as well as from abstract content to the simulation environment. The instantiation of scene content is a part of ISAT's design not yet implemented in the prototype. Our future work will be heavily focused on refining this approach.

Predictive Direction

One role of the director that has yet to be explored is suggested by work done in IDA (Magerko and Laird, 2004). The use of *predictive direction* has been shown to be a more subtle and believable alternative to reactive direction. As opposed to waiting for problematic actions to occur, the director attempts to predict the player's actions in the near future. If that prediction reliably shows a conflict with story content, then the director can employ strategies to redirect the scene. Future research will focus on how to incorporate this approach into ISAT, as well as what defines a good predictive model for a given training domain.

SCORM Compliance

Conformance to the Sharable Content Object Reference Model (SCORM) 2004 will promote reusable and interoperable shared content objects (SCO) across multiple learning management systems (LMS). In SCORM, interactive training content currently can be described, sequenced, tracked, and delivered like all other content. However, interactive training often has more complex requirements for data tracking. Different interactive training models will result in different solution approaches.

A mechanism for assessment needs to be developed that allows student performance data to be extracted from the TC3 and communicated to the LMS. LMSs are very limited in how they can interact with learning content; consequently, the LMS is unable to "watch" what is happening inside a given SCO and take action.

Prototypes exist for SCORM-compliant LMSs that can dynamically sequence SCOs on the basis of user performance (Morales & Aguera, 2004). The ISAT TC3 prototype dynamically sequences within the SCO itself. We plan to compare these sequencing methods and their underlying data representations with the goal of making ISAT SCORM compliant and enhancing SCORM. This contribution has potential to advance the DOD's ADL vision by making interactive learning environments part of the reusable, transferable pool of learning resources.

Evaluation

Our evaluation of the ISAT prototype will examine the following aspects of ISAT:

Training context. The training context (i.e., the environment, scenarios, scenes, and NPC actions) is intended to be rich, authentic, dynamic and believable. The director is supposed to act in the background to guide the experience imperceptibly.

Trainee learning. The ultimate goal of the project is to improve the effectiveness of the training through use of the ISAT approach and director's scenario management. The training system was explicitly designed around specific learning objectives. These learning objectives were also used in the design of the director. For the training system to be judged effective, the trainees must acquire various skills and knowledge.

CONCLUSIONS

We have developed and refined the conceptual framework for ISAT and applied its concepts to a specific training system, the 91W10 Tactical Combat Casualty Care Simulation. ISAT suggests engagement and pedagogically-motivated direction can improve training experiences. The provided storyboard details an example that shows the benefits of our design.

ISAT offers a number of distinct advantages over traditional intelligent tutoring systems and constructive simulation technology. Most importantly, ISAT integrates the advantages of these previously distinct approaches to educational technology. ITSs offer explicit pedagogical guidance, building on a cognitive user model to understand and guide a trainee's learning (although ISAT uses a simpler behavioral skill model to provide guidance). Virtual and constructive training offer the experience of learning by doing, of being in the actual environment (or a close approximation of it) in which real-world performance will occur. ISAT includes pedagogical support but offers it in the context of the actual environment. Rather than explicit pedagogical direction, the trainee receives subtle cues and feedback that do not break the "spell" of engagement in the experience. This approach is best-suited for *unstructured teaching domains*, meaning that it should be more effective for domains that involve physical skills and actions as opposed to mainly mental operations, which are more suited for intelligent tutoring applications. Our evaluation phase will help better understand ISAT's flexibility.

ISAT also offers the potential to increase training effectiveness by drawing on important elements of interactive storytelling—instantiation of abstract plot content and heuristic scene choice. Because the training experience is designed to be engaging, training goals are reinforced and those skills that most require attention drive scene selection and instantiation. The result should be that trainees get more direct exposure to the skills that need the most work, naturally speeding the training process. Because the experience is engaging and (as a result) enjoyable, trainees may participate in training more often and voluntarily, further accelerating training. Engagement also potentially leads to memorable experiences, facilitating in-the-field recall of prior training experiences and thus possibly leading to greater persistence in the training experiences.

Another potential benefit of ISAT is that it supports scaffolding and fading in a realistic setting. The director will help a student execute a task by providing feedback (e.g., in the form of utterances by NPCs). However, because the trainee is not aware of the director during execution of the training scenario, the director can easily "disappear" as the trainee becomes more skilled and requires less feedback. The director's guidance strategies will become more and more subtle until the director is not intervening for pedagogical purposes at all. While there is growing consensus in the educational technology field to use these technologies, scaffolding and fading were not explicit goals of ISAT. Instead, they result from the primary goals of providing realistic, engaging experiences tailored to the needs and requirements of individual trainees. Having completed a small initial feasibility demonstration, we are continuing to test these ideas via further exploration and development in the 91W10 Tactical Combat Casualty Care Simulation. Ultimately, our goal is to show the generality of ISAT as an architectural plug-in for training systems and to provide it as a cost-effective component for more effective trainers.

ACKNOWLEDGEMENTS

This work was funded by a Small Business Innovation Research grant, contract number N61339-05-C-0027, sponsored by US Army RDECOM-STTC and issued by NAVAIR ORLANDO TSD. We thank Dr. Amy Henninger for her preliminary design work on ISAT, Engineering and Computing Simulations, Inc. for providing the TC3 training simulation environment and their collaboration efforts, and Mr. Bill Pike, RDECOM-STTC, for valuable contributions and guidance.

REFERENCES

Anderson, J. R., Corbett, A. T., Koedinger, K. R. and Pelletier, R. (1995). Cognitive tutors: Lessons learned. *The Journal of the Learning Sciences,* 4(2): 167-207.

Bloom, B. S. (1984). The 2 sigma problem: The search for methods of group instruction as effective as one-to-one tutoring. *Educational Researcher*, 13 (6): 4-16.

Collins, A., Brown, J. S., & Newman, S. E. (1989). Cognitive Apprenticeship: Teaching the craft of reading, writing, and mathematics. In L. B. Resnick (Ed.), Knowing, learning, and instruction: Essays in honor of Robert Glaser (pp. 453-494). Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.

Corbett, A. T., Koedinger, K. R. & Hadley, W. H. (2001). Cognitive tutors: From the research classroom to all classrooms. In P. S. Goodman (Ed.), *Technology Enhanced learning: Opportunities for Change*. Lawrence Erlbaum Associates, Mahwah, NJ.

Johnson, W. L., Marsella, S., Vilhjálmsson, H. (2004). The DARWARS Tactical Language Training System. *Interservice / Industry Training, Simulation, and Education Conference.*

Johnson, W. L., Rickel, J., and Lester, J. (2000). Animated Pedagogical Agents: Face-to-Face Interaction in Interactive Learning Environments *The International Journal of Artificial Intelligence in Education.* Vol. (11): 47-78.

Kintsch, E. (1993). Principles of instruction from research on human cognition. In J. M. Spector, M. C. Polson, & D. J. Muraida (Eds.), Automating Instructional Design: Concepts and Issues (pp. 23-42). Englewood Cliffs, NJ: Educational Technology Publications

Lester, J., Converse, S., Kahler, S., Barlow, T., Stone, B., and Bhogl, R. (1997a). The Persona Effect: Affective Impact of Animated Pedagogical Agents. *In Proceedings of CHI '97,* Atlanta, GA., p.359-366.

Lepper, M. R. and Henderlong, J. (2000). Turning "play" into "work" and "work" into "play:" Twenty-five years of research in intrinsic vs. extrinsic motivation. In C. Sansone and J. M. Harackiewicz (Eds), *Intrinsic and Extrinsic Motivation: The Search for Optimal Motivation and Performance*, 257-307. San Diego: Academic Press.

Lester, J., Converse, S., Stone, B., Kahler, S., and Barlow, T. (1997b). Animated Pedagogical Agents and Problem-Solving Effectiveness: A Large-Scale Empirical Evaluation. *In Proceedings of the Eighth World Conference on Artificial Intelligence in Education.* Kobe, Japan, p.23-30.

Marsella, S., Johnson, W. L., LaBore, C.. (2000). Interactive Pedagogical Drama. *4th International Conference on Autonomous Agents*. Barcelona, Spain, p.301-308.

Magerko, B. (2005). Story Representation and Interactive Drama. *1st Artificial Intelligence and Interactive Digital Entertainment Conference.* Marina Del Rey, CA.

Magerko, B. and Laird, J.E. (2004). Mediating the Tension between Plot and Interaction. *AAAI Workshop Series: Challenges in Game Artificial Intelligence*. San Jose, California, p.108-112.

Mateas, M. and Stern, A. (2002). A Behavior Language for Story-Based Believable Agents. *AAAI Spring Symposium Series: Artificial Intelligence and Interactive Entertainment*.

Morales, R. & Aguera, A. S. (2004). Dynamic Sequencing of Learning Objects. Proceedings of IEEE ICALT 2004.

Moreno, R., Mayer, R., and Lester, J. (2000). Life-Like Pedagogical Agents in Constructivist Multimedia Environments: Cognitive Consequences of Their Interaction. *In Proceedings of the World Conference on Educational Multimedia, Hypermedia, and Telecommunications (ED-MEDIA).* Montreal, Canada, p.741-746.

Rickel, J. and W. L. Johnson (1997). Intelligent Tutoring in Virtual Reality: A Preliminary Report. *International Conference on Artificial Intelligence in Education*.

Weyhrauch, P. 1997. *Guiding Interactive Drama*. PhD Thesis, Carnegie Mellon.

Wray, R. E. and J. E. Laird. (2003). Variability in Human Behavior Modeling for Military Simulations. *Conference on Behavior Representation in Modeling and Simulation*, Scottsdale, AZ.

Young, R. M., Riedl, M. O., Branly, M. Jhala, A., Martin, R.J., Saretto, C.J. (2004). An Architecture for Integrating Plan-based Behavior Generation with Interactive Game Environments. *Journal of Game Development* (1): 51-70.

Vygotsky, L. S. (1978). Mind in society: The development of higher psychological processes. Cambridge, MA: Harvard University Press.