Forging a New Simulation Technology at the ICT

At USC's Institute for Creative Technologies, computer scientists, US Army personnel, and entertainment professionals are developing the foundation for an experiential learning system.

Richard Lindheim
William Swartout
University of Southern California

Throughout history, artists from all disciplines have provided scientists with inspiration, supplying the unbridled speculations that researchers, through hard work and perseverance, have made real. We can trace this unspoken partnership back at least as far as the sixteenth century, to the imaginative drawings of Leonardo da Vinci. In the late nineteenth and early twentieth centuries, Jules Verne and H.G. Wells extrapolated from the inventions of the Industrial Revolution to foreshadow such modern developments as submarines, airships, and space travel. Imagination's pace accelerated from the 1920s onward, with pulp science fiction writers envisioning atomic bombs, flying cars, and robots. Later, Arthur C. Clarke predicted global communications via satellite, and the television series *Star Trek* introduced the handheld communicator, a concept embodied in today's cellular telephone.

Given the rich harvest this informal cross-pollination of the creative and technological fields has provided, it seems obvious in retrospect that a deliberate commingling of these two cultures could offer both immediate and long-term benefits. Yet no one attempted to create such a collaboration deliberately focused on simulation and learning until the founding of the Institute for Creative Technologies at the University of Southern California in 1999. For a detailed account of the factors involved in creating the ICT, see the "ICT Genesis" sidebar.

The US Army's need for highly realistic training drove this experiment in collaborative engineering, which gathered officers from the Army's Simulation, Training, and Instrumentation Command, computer research scientists at USC, and professionals from the Hollywood-based film community and the Silicon-Valley-based electronic-games community. Initially, the Mission Rehearsal Exercise (MRE) provided the most significant challenge and the most ambitious project for this diverse team.

MRE PROJECT OVERVIEW

The MRE project seeks to create a virtual reality training environment in which soldiers will confront dilemmas that force them to make decisions in real time under stressful and conflicting circumstances. By allowing soldiers to see the consequences of their actions and decision-making skills in a simulator, the Army expects to better prepare its troops for dealing with similar dilemmas in the real world.

Since the end of the Cold War, the need for this kind of training has grown more acute as the variety of US military operations has expanded enormously. In addition to conventional combat operations, US military personnel frequently undertake a broad spectrum of missions that include peacekeeping operations, disaster relief efforts, and noncombatant evacuations. Because these actions may require troops to deploy to virtually any location across the globe, providing advance training tailored to a wide variety of specific situations presents a daunting challenge.

These operations often involve close interaction between soldiers and the local populace, so troops must understand the local culture and anticipate how an area's inhabitants will behave. Further, such operations often occur in the context of a bewildering array of political and logistical issues that seriously test soldiers' decision-making skills. If soldiers fail to accurately assess and respond to the situation, they risk triggering potentially disastrous misunderstandings that could have international repercussions.

Assembling a technical and creative team

Building the MRE system required assembling a diverse group of individuals and organizations with a broad range of talents. On the technical side, AI researchers from USC/Information Sciences Institute and USC/ICT created the automated reasoning and emotion modeling technology for the virtual humans.

ICT Genesis

Michael Macedonia
US Army Stricom

The computer and Internet revolutions have substantially changed the direction of entertainment from delivery in a mass medium such as television to a mass customized experience via the Web and PC. However, the art of entertainment still requires stories, characters, and direction to make the experience meaningful and enjoyable.

Leveraging Commercial R&D

The US Army faces the same challenge of adapting to the changes brought about through the mass marketing of supercomputing via platforms such as the PlayStation 2 and Xbox and low-cost 3D graphics.

The entertainment industry has in many ways grown far beyond its military counterpart in influence, capabilities, and investments. The Interactive Digital Software Association estimates that, in 1998, interactive entertainment businesses invested approximately $2 billion in new technology R&D, an increase of more than 20 percent over the previous year.[1] This investment far outweighs current US Army R&D for training and simulation technology. Moreover, we cannot ignore industry advances. Witness the rapid pace of graphics systems' development for game consoles and personal computers, which nearly double their performance every nine months.

The Human Touch

There is also an urgent requirement for representing the human dimensions of war and conflict to provide training for the truly difficult decision-making problems our soldiers must face. Our experience in Kosovo and Bosnia has shown that we need troop leaders who can handle the dilemmas posed by ethnic and social strife.

To provide such expertise, we must develop interface technologies such as natural-dialogue systems and intelligent agents that can simulate real-world problems. A 1997 study initiated by Anita Jones, US Deputy Director of Defense Research and Engineering, documented this need.[2] Early in 1999, US Army leaders recognized a need for a major transformation of our forces to overcome the limitations of our current simulation technologies. Effecting this transformation requires developing new training and simulation systems for dealing with future conflicts that leverage the capabilities of both the entertainment industry and academia.

The US Army and the Department of Defense selected the University of Southern California as a strategic partner in the development of the Institute for Creative Technologies because of USC's unique confluence of scientific capabilities and entertainment-industry relationships, which the Army deemed necessary for simulation leadership.

ICT's Prime Objective

The prime objective, as reaffirmed by Michael Andrews, US Deputy Assistant Secretary of the Army for Research and Technology, is to build a special partnership with the entertainment industry and academia.

Some of USC's unique qualifications arise from its location in Los Angeles, at the hub of both the entertainment and aerospace industries; some arise from its standing as a leading private research university; and some arise from the capabilities and stature of its component units and the working relationships they have developed with industry. For example, USC's top-ranked School of Cinema-Television grew side-by-side with the entertainment industry and continues to maintain close ties with it.

USC established the Institute for Creative Technologies, shown in Figure A, under the auspices of the US Army Simulation, Training, and Instrumentation Command (Stricom) to focus on developing the art and technology for providing synthetic experiences so compelling that participants react as if the simulations are real. That is, ICT will bring verisimilitude—the quality or state of appearing to be true—to synthetic experiences. Beyond the military, ICT will also advance a compelling new medium for entertainment, education, the arts, travel, and other potential applications.

The establishment of the ICT is just one of many steps needed to provide the essence of verisimilitude in training and virtual reality systems. The US Army will explore all avenues of entertainment technology to keep pace with the challenges presented to it, whether in distributed learning applications or embedded training systems. Ultimately, we want to prepare our soldiers for the future by experiencing it.

References

1. M. Macedonia, "Why Digital Entertainment Drives the Need for Speed," *Computer*, Feb. 2000, pp. 124-127.
2. M. Zyda, ed., *Modeling and Simulation: Linking Entertainment and Defense*, National Research Council, National Academy of Sciences Press, Washington, D.C., 1997.

Figure A. Offices of the Institute for Creative Technologies. Tilted wall paneling forces sound upward and downward, providing privacy without the constriction of a cube environment. Open work areas facilitate collaboration while small offices that can be closed off with curtains provide privacy and bunk beds for rest during extended working hours.

Figure 1. The Mission Rehearsal Exercise presents officer trainees with situations that test their decision-making and leadership abilities.

Researchers from Boston Dynamics developed the virtual humans' animated bodies, with Haptek providing expressive faces. Audio researchers from USC's Integrated Media Systems Center created the immersive sound system and mixed and synchronized the effects and background sounds. Researchers from USC/ISI worked with AT&T's Next-Gen TTS speech synthesis system to create the most natural-sounding output. Finally, programmers and system developers experienced in creating real-time graphics tweaked the graphics system to provide acceptable performance.

To create the content for the MRE system, we needed an art director to design the environment's overall look, actors to serve as models for the virtual humans, and artists to model the animated characters, buildings, and environmental details. This effort—our first attempt to combine the expertise of the entertainment community with that of the technical and military simulation communities—would ultimately prove far more resource-intensive than we imagined. Coordinating all the teams involved and integrating the various systems and subsystems turned out to be a major task.

Hollywood brings story and character to simulations

As the team worked at conceptualizing the MRE simulation, a core divergence emerged between how the Army, computer scientists, and entertainment people viewed the project. The entertainment people usually took an approach diametrically opposed to that of established scientific and military procedures.

Entertainment people who work on simulations focus on the project's concept, theme, and story, shaping these elements to create the simulation's desired impact. This focus on the project's net result differs markedly from the scientific approach to simulation, which usually starts with a definition of the realism level required.

The entertainment community uses story and character to create emotional involvement for two reasons. First, writers know that a work's ability to engage the emotions determines its impact. We readily recall great moments in motion pictures and television because they evoked a strong emotional response. Films as varied as *Jaws*, *Casablanca*, *The Sixth Sense*, and *Saving Private Ryan* became popular classics because of the emotional responses they generated. Yet films that lack such an emotional punch—even big-budget special-effects films such as *Judge Dredd* and *Mission to Mars*—quickly fade from memory. For entertainment people, then, emotional impact delivered through story and character provide the key to connecting with the audience.

Entertainment people also use emotional involvement because it's expedient. The industry has learned that viewers will ignore a film's flaws and inconsistencies if it emotionally involves them. Applying this knowledge to simulations, we know that once participants become emotionally engaged we can relax the requirements to render every detail at the highest resolution.

The Army's concept of story differs considerably from the entertainment industry's. What the Army considers a story, Hollywood labels an event list. A sequence of events does not itself create a story; a story requires linking events in a way that builds to a dramatic climax.

To date, simulations have, at best, made rudimentary use of character despite the critical importance entertainment veterans place on this component. Yet any simulation for training people to work with one another in decision-making tasks must, by definition, place a premium on realistically depicting how individuals—even simulated ones—react to each other. Given the technical challenges involved and the high priority placed on implementing the capability, integrating character into the MRE simulation proved to be the project team's most difficult task.

Making it concrete: The MRE story

In the MRE, we seek Hollywood's influence most strongly in the rich story structure that guides, but does not completely determine, how our simulation unfolds. A good story sequences events so that emotions and tensions build and ebb. Plot twists and surprises, such as the scene shown in Figure 1, maintain interest and involvement. The MRE uses story structure to build toward the dilemmas the trainee must resolve, offering different paths through the structure to reflect the different options the trainee can choose, thereby making it possible to see vividly the consequences of each decision.

For example, in our demonstration scenario the computer generates all the scenario's characters except for the trainee, a young lieutenant. The lieutenant has been instructed to rendezvous with his troops at a staging point before proceeding to help quell civil unrest occurring in a small town. The action unfolds as follows:

- *Surprise*. When the lieutenant arrives at the staging area, his platoon sergeant informs him that one of his Humvees has collided with a local civilian car. The lieutenant sees a small boy on the ground, seriously injured, the boy's frantic mother kneeling beside her injured child.

- *Dilemma.* Should the lieutenant continue with his mission or stop, render aid, and arrange for a medevac?
- *Complication.* A TV cameraman shows up and starts filming. Any mistake the lieutenant makes could appear on the international news.
- *Complication.* If the lieutenant decides to arrange for a medevac via helicopter, the MRE may challenge him by relaying a radio call from troops already in town that reports intensifying unrest, shots fired, and a request for assistance.
- *Dilemma.* Should the lieutenant split his forces, sending some ahead while keeping others behind to help with the medevac, or should he keep his unit intact?

When confronted with these choices, the trainee may receive assistance from the virtual platoon sergeant. Given that sergeants usually possess substantially more field experience than new lieutenants, the Army teaches its entry-level officers to listen to the advice their sergeants offer. The virtual sergeant thus embodies US Army doctrine and coaches the lieutenant toward the most appropriate course of action.

Technology goals

When we began developing the MRE system, we knew we would need to push the technology's boundaries in several areas to attain the kind of compelling, immersive experience we desired. Although we sought to achieve advances in particular areas, we also recognized that much of the experience's immersiveness would come from combining components that had thus far never been integrated. Further, we felt that integrating these components would grant us a better understanding of the technology and how the basic research should proceed in individual areas.

Thus, we wanted to make the MRE system fully integrated as early as possible while continuing to enhance the capabilities of some components. These concurrent goals presented a dilemma: You can't integrate components still in development. We resolved this dilemma by adopting a hybrid approach to system development. Doing so allowed us to create an integrated system that used rudimentary versions of some capabilities, which acted as placeholders until more sophisticated capabilities became available.

MRE'S HYBRID APPROACH

Because entertainment people focus on the simulation's ultimate objective, they readily adapt any technology at hand if it does the job. Thus, in motion pictures and television, visual-effects experts commonly combine varying technologies, a hybrid approach that can encompass a huge gamut of technologies. When creating visual effects, these experts readily mix live photography, trick effects, models, and computer-generated images.

Using this eclectic approach as a template, we adopted a hybrid approach to constructing the MRE system. The behavior of the virtual humans in the scenario could be controlled by scripts, an AI reasoner, or an AI reasoner augmented by an emotional model. We found that scripted control worked well for bit players in the simulation, such as people in a crowd and some of the soldiers. AI reasoning proved useful for characters whose many interactions required them to exhibit a broad range of behaviors. Finally, we used an emotional model for the one character in our scenario who really needed to show emotion, the mother.

Our integrated system used rudimentary versions of some capabilities as placeholders until more sophisticated capabilities became available.

For speech, we used both prerecorded dialogue from actors and text-to-speech synthesis. We found that the prerecorded speech worked best for characters with small parts because they had comparatively little to say. Prerecorded speech also was useful for characters who needed to express a lot of emotion—difficult to achieve with current text-to-speech systems.

For the sergeant, who had a large part, text-to-speech provided the best solution, especially since his calm, advisory role did not demand that his voice express a wide range of emotion. Further, as our scenario expands in scope the sergeant will need to deal with an increasing number of situations. For this character, the advantages of dynamically generating the speech needed outweighs the expressiveness we give up by using text-to-speech.

Taking a hybrid approach also allowed us to introduce new technologies in the places where they would make the most difference, while saving us effort in areas where simple techniques would suffice. Hybrid solutions also let us create a complete mission rehearsal scenario so that we could assess how it works as a whole, without having to solve all the subproblems in the most general way. We expect to incrementally improve the MRE system over time by replacing simple solutions for control, speech, and sound with more sophisticated techniques as our research in these areas progresses.

TECHNICAL RESEARCH AREAS

Successful development of the MRE simulation required technical research into two major areas:

- achieving immersion through graphics and sound, and
- creating believable virtual humans.

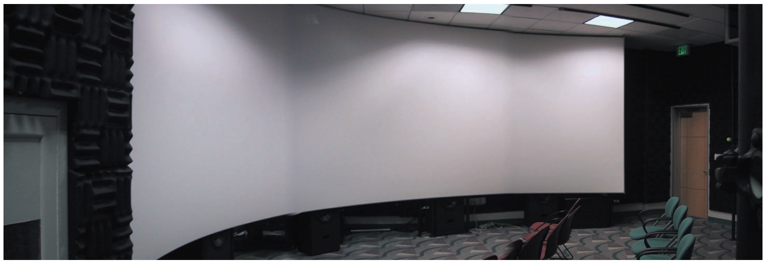

Figure 2. The Mission Rehearsal Exercise simulation room features cutting-edge sound and graphics. The simulation's action, which features life-sized computer-generated and -directed characters, unfolds on an 8-foot-high by 30-foot-wide curved screen.

We integrated a variety of technologies from several vendors to achieve our goals in these areas.

Immersive graphics and sound

We use high-resolution graphics and spatialized sound to immerse the participants in each simulated setting. An SGI Onyx Reality Monster IR3 generates the MRE images, which three Barco 909 projectors project onto a curved screen, shown in Figure 2, that is approximately 8 feet high and 30 feet wide. Multigen-Paradigm's Vega simulation development environment drives the simulation.

The immersive sound system, developed by Chris Kyriakakis, employs 10 audio and two subwoofer channels.[1] The audio channels feed speakers arrayed not only around the listener but also at different heights, making it possible to create effects unavailable with current home or theatrical sound systems. The system localizes sounds along vertical and horizontal axes to increase immersion. Our initial demonstration features the creaks, rattles, and engine noise of driving a Humvee, the murmur of a crowd gathered in a small village, and the almost deafening thwack-thwack of a helicopter passing overhead.

Integrating the sound system with the graphics proved challenging. Our current version of the demonstration system spatializes and renders the sounds offline, then plays them back on an Apple Macintosh when simulation events trigger them. The timing must be precise, because even a small discrepancy between the events onscreen and the sound can be distracting, as anyone who has watched a badly dubbed movie can confirm. We resolved this issue by generating SMPTE time codes on the SGI that the Mac interprets to trigger the sounds. Spatializing the sounds offline in advance presented a more serious limitation. For example, the path the helicopter takes must be fixed and could not change dynamically during the simulation. This requirement in turn limited the simulation's interactivity and variability. To overcome this limitation, we're currently researching ways to spatialize sound dynamically in real time.

Virtual humans

We populate our simulated worlds with computer-generated characters—virtual humans. They play the roles of local citizens, friendly and hostile forces, and other team members. At times, the characters may also act as coaches, giving advice to the trainee.

A script triggered by events in the simulation controls the simplest characters. We use artificial intelligence reasoning to determine the behavior of more sophisticated characters as events unfold.[2-4] Using AI reduces the need to script these characters and allows them to react to situations as they unfold—even situations that arise from unanticipated events.

In addition to AI-based reasoning, we can further augment characters with emotional models that guide their reactions to events.[5,6] Expressing such emotions makes virtual humans more believable. People expect other people to react emotionally to events, so virtual humans that lack emotions seem flat and unconvincing.

Characters with emotion models have goals and values. When these goals are supported or threatened, the characters react emotionally. For example, one MRE simulation character, a mother, must deal with a situation in which her son has been injured in an automobile accident. Naturally, she concerns herself first with getting help for her son, a goal that she believes some of the soldiers in the simulation can help her achieve. When she sees some of them leave, she becomes visibly upset. Although we have only begun emotional modeling, we can see already how depicting emotional behavior in simulations makes virtual humans more believable.

Boston Dynamics created bodies for the virtual humans, developing a variety of body types and a library of motions for the various roles characters might play in a scenario: sergeant, infantryman, onlooker, medic, and so on. Developers used motion capture data from actual actors who performed typical movements to create the motions for each character type, then they used the Boston Dynamics software to help smooth the transitions between different motions.

Motion capture provides natural-looking animations, but it limits the set of motions possible. This restriction sometimes imposed limits on the kinds of behaviors developers could incorporate into the scenario because they could not depict certain motions. Expanding the library of motions and using animation techniques such as inverse kinematics that do not require motion capture data can reduce these limitations.

Haptek developed expressive faces for some of the virtual humans. The system developers can control the faces so that they smile, frown, or show other expressions. In addition, the faces support lip sync so that the lips move realistically when the character speaks.

We took two approaches to controlling the characters, using either

- a tool to create a precise script of motions for the character to follow or

ICT Project List

In addition to the Mission Rehearsal Exercise, the ICT has undertaken several other projects that explore the synthesis of immersive multimedia and AI to create effective simulation, learning, and entertainment experiences.

Previsualization Project—Future Combat Systems

The Army's Future Combat Systems, an effort that will potentially redefine the nature and capability of US ground forces, is a "system of systems" that emphasizes rapid deployability, survivability, and lethality. The Army has sought a ground-up conceptualization with a three-year development period for the FCS project, departing radically from existing methods and hardware.

Designed to deal with the kinds of "pop-up" crises that dominate today's headlines, FCS is intended to be deployable in only 96 hours, thanks in large part to a 20-ton weight limitation. Nonetheless, the Army expects it to offer the safety and reliability of today's defense hardware.

The ICT contributed to FCS concept development by engaging the creative energies of the film-entertainment community and the computer game industry in an interdisciplinary effort. Social scientists, human factor specialists, and military subject-matter experts complemented the ensemble. The group worked to imagine FCS capabilities and the circumstances in which it might one day be used.

The process began with orientation and training to provide the context necessary to understand the kinds of challenges and missions US forces will likely face in the future. The group traveled to military installations and participated in numerous Tactical Decision Exercises. Briefings and selected readings rounded out this subject-matter immersion.

Next, group members developed future US Army deployment scenarios that covered a broad range of scope and mission categories. Finally, the participants explored and assessed actual conceptual solutions in a series of guided discussions. The ICT FCS project used the Institute's previsualization tool to foster the design process.

As the concept development phase draws to a close, the ICT continues its collaboration with Army development experts and end users to perfect FCS components and systems. With the immersive previsualization tool at its disposal, the ICT can help accomplish an exciting and innovative early step—in essence, using FCS before it is built.

Experience Learning System

Current simulation training systems have an acceptable-to-good ability to present physical objects and terrain. However, these systems suffer from several severe limitations. With the Experience Learning System (ELS), we plan to develop for the Army a simulation training system that can overcome current simulation shortcomings while playing a complementary role with its legacy simulation systems. Ultimately, the Army seeks to implement a networked, completely immersive system that benefits from the creative input of entertainment industry artists and virtual reality technology.

State-of-the-art simulators do not yet provide many of the capabilities necessary to achieve this goal, so the ELS requires basic research along several vectors, including immersion, story, artificial-intelligence characters, direction, setup, architecture, and networking.

To realize the ELS, we intend to investigate five fundamentally different approaches to simulation:

- *Virtual motion.* A VR environment primarily intended for vehicle simulation, virtual motion uses a nearly 360-degree view of a simulated world around a vehicle and could include a motion platform.
- *Virtual world (FlatWorld).* A simulation intended for dismounted infantry operating in a small cluster of structures, or any situation in which participants need to walk around in a simulation, this approach uses Hollywood, stage, motion-picture production, and theme-park design techniques. The system will manipulate visual surfaces to create the experience of being in a specific location. The MRE system provides a first step in this direction.
- *Virtual environment.* A simulation that uses augmented reality to place synthetic objects in a real environment, virtual environment addresses situations in which simulated and real objects must be mixed, or where the training exercise involves a space so large and detailed that simulating all of it becomes impractical.
- *Ubiquitous simulation.* The virtual-motion, virtual-world, and virtual-environment simulation approaches —such as the Star Trek Holodeck— cannot by themselves fulfill the Army's desired learning goals. Instead, these approaches must be coupled with newly emerging high-end video-game technology such as the Sony Playstation 2 and the Microsoft Xbox. The program would provide participants with specific games that both reinforce training objectives and offer corrective learning for their specific areas of weakness.
- *Advanced Leadership Training Simulator.* ALTSIM uses networking capability to bring individuals together in a command-and-control environment for leadership training. A PC-based simulation, the project can run on any computer. It uses high-end graphics to act as a "window" into the command center, along with Internet technology and immersive sound to create a realistic simulation environment.

The integrated ELS can be configured for any application, from training Army personnel for new assignments in remote locations around the world to preparing for disaster relief to studying countermeasures for potential terrorist activities. It can also be used as a previsualization tool. The ELS has far-reaching applications well beyond the Army. One day, this immersive distance learning system might change the entire spectrum of education.

For the training scenarios we envision, the most appropriate way for the characters to interact is using natural language.

- an interface that allowed those characters with AI reasoning to select motions dynamically from a library of possible movements during the simulation.

For the training scenarios we envision, the most appropriate way for the characters to interact is using natural language. For input, we used a commercial speech recognition system, developed by Entropic, to recognize what the trainee said. For output, we used two approaches: Next-Gen TTS, a text-to-speech system, and prerecorded speech.

Developed by AT&T, Next-Gen TTS produces sounds that are more natural than most synthesis systems, but it has a limited range of expression and number of voices. When necessary, we used prerecorded speech to bypass these limitations. Beforehand, actors recorded a library of verbal phrases that the characters might say. During the simulation, the computer selected the most appropriate dialogue sample for playback. We found that the flexibility of the text-to-speech system made it most appropriate for the sergeant in our scenario, who had the most to say. The greater emotional range possible with prerecorded speech, on the other hand, made it appropriate for characters who had less to say but needed to impart more expression and emotion in their speech.

SURMOUNTING TECHNICAL LIMITATIONS

Creating a dynamic scene with important story elements, adding characters that will utter the dialogue, and providing appropriate stage directions normally requires the efforts of a seasoned Hollywood screenwriter. Yet scripting the MRE presented a considerably more difficult challenge given the technical limitations imposed upon the process.

First, given that the story process was itself interactive, normal linear storytelling procedures did not apply. Writers involved in the project had to learn a new way of telling stories. Fortunately, chief writer Larry Tuch has considerable experience in developing interactive materials and viewed the story issues as challenges rather than problems.

Second, to use Tuch's nomenclature, "dumb" actors would play the characters in our scenario. Even the dullest live actors can learn lines, know where to look and place themselves, and speak when they are supposed to during a scene. Because AI characters in the MRE do not possess these attributes, each element of their behavior must be programmed separately.

Third, the writer had to work closely with the computer scientists to confirm what the simulation's actors could and could not do. For example, Tuch wanted some of the soldiers to get into the Humvee and drive away. Many technical limitations, however, complicated this otherwise simple goal, including extensive motion-capture animation for the characters, who would still look stiff and unrealistic when entering the vehicle. The solution involved a workaround: The characters walk offstage, exiting one edge of the screen, then the Humvee drives away. Viewers automatically fill in the "missing" frames, assuming that the soldiers entered the vehicle from the off-screen side.

The most critical collaboration between writer and technical people centered on the core emotional moment in the MRE: the automobile accident that injures the mother's young son. To advance the story aspect of the simulations, some actions critical to the training exercise trigger peaks in the mother's emotions. When she sees some soldiers run to their Humvees and leave the accident, she gets angry and frustrated, reminding the officer in charge that her son needs help.

To have the desired impact, the mother's emotional progression and expression required a mix of body language and verbalization. An easy task for conventional film or television to accomplish, this behavior proved a major technological hurdle for the MRE team. Three basic problems challenged us:

- None of the characters could touch one another. The program we used lacks boundary capability, which means that if one character touches another an overlap results, completely destroying any appearance of reality.
- Although we used motion-capture technology to approximate body language, the technology lacks the precision needed to express realistically the various emotional states the writer desired.
- Most of the motion-capture animation shows the mother operating at an energy level too high to effectively show her changing emotional state and to believably portray her building anxiety.

Solving these problems required intensive collaboration between Tuch and computer scientists Stacy Marsella and Jon Gratch. First, Tuch revised the script to retain the emotional impact even though characters never touch one another. This restriction proved particularly difficult because it prevented the mother from holding or even touching her injured son.

Tuch strove to overcome this limitation by creating dialogue specific to each stage of the mother's emotional progression. To create more precise body movements, Marsella and Gratch selected clips from the Boston Dynamics motion-capture sequences and used them as building blocks to construct new, more intricate movements.

The Boston Dynamics team also added a parameter that let Marsella and Gratch slow down certain body movements. This had the effect of creating more subtle effects that let the team modify the mother's energy level to better match her emotional state.

Sound also played an important part in creating emotional impact for characters with limited mobility and imprecise body movements. The team recorded the dialogue in Serbo-Croatian as discrete sound clips. By combining the sound clips with the motion-capture adjustments, the team could adjust the mood and intensity of the mother's emotional states.

The Mission Rehearsal Exercise simulation's effectiveness owes much to techniques perfected by the entertainment industry. While we knew well the technology's limitations and the compromises agreed to by the project's creative and technical people, these factors proved inconsequential to audiences, whose enthusiastic reactions exceeded our expectations.

For MRE participants, the situation and CGI characters became real. When they exited the simulation, Army officers said, "I've been there. This is almost too real." Those who had encountered similar incidents in real life became emotionally moved. Their reactions provided a powerful lesson in the impact that story and character can have when we integrate them into the simulation experience.

An exciting side benefit of the MRE is its integration with other elements of the Experience Learning System. In the demonstration, we took care to coordinate the MRE's story and situation with the Advance Leadership Training Simulation (ALTSIM). When run, both simulations interact and share significant events. Next, we intend to add newly developed computer games into the system, with our ultimate goal the seamless integration of all component systems.

The MRE's positive results have generated support for our research efforts. The team will move forward rapidly now, with a focus on developing other scenarios and improving the technology. The next MRE version will solve many of the problems we've described. Other challenges will certainly emerge, however, as the ICT moves forward in its effort to build more effective education and training simulation environments. ✷

Acknowledgments

The authors thank the MRE team for their hard work, intelligence, and dedication. Jonathan Gratch, Randall Hill, Lewis Johnson, Stacy Marsella, Jeff Rickel, and Richard Whitney developed the artificial intelligence and emotion models and integrated the natural language support for the virtual humans. Marc Raibert, Adam Crane, Whitney Crane, and Marlon Veal developed the animated bodies and expressive faces for the virtual humans. Chris Kyriakakis developed the immersive audio technology used in the demonstration, and David Miraglia produced the sound content. Larry Tuch wrote the scenario. Jacki Morie provided art direction. Michael Murguia and Erika Sass developed the 3D graphic models for the simulation. Kate LaBore scripted the behaviors of the non-AI virtual humans. Marcus Thiebaux and Ben Moore integrated the various software systems and made everything work. Jay Douglas operated the demonstration system, and Trevor Hawkins played the role of the lieutenant trainee. Ramon Gonzalez provided computer support and kept the machines running. Finally, Brigadier General Pat O'Neal (Ret.) and Elke Hutto provided important military knowledge to help structure the scenario.

References

1. C. Kyriakakis, P. Tsakalides, and T. Holman, "Surrounded by Sound: Immersive Audio Acquisition and Rendering Methods," *IEEE Signal Processing Magazine*, vol. 16, no. 1, 1999, pp. 55-66.
2. J. Rickel and W.L. Johnson, "Animated Agents for Procedural Training in Virtual Reality: Perception, Cognition, and Motor Control," *Applied Artificial Intelligence*, vol. 13, 1999, pp. 343-382.
3. J. Rickel and W.L. Johnson, "Virtual Humans for Team Training in Virtual Reality," *Proc. 9th Int'l Conf. Artificial Intelligence in Education*, IOS Press, Burke, Va., 1999, pp. 578-585.
4. J. Rickel and W.L. Johnson, "Task-Oriented Collaboration with Embodied Agents in Virtual Worlds," *Embodied Conversational Agents*, J. Cassell et al., MIT Press, Cambridge, Mass., 2000.
5. J.E. Gratch, "Marshalling Passions in Training and Education," *Proc. 4th Int'l Conf. Autonomous Agents*, ACM Press, New York, 2000, pp. 325-332.
6. S. Marsella, W.L. Johnson, and C. LaBore, "Interactive Pedagogical Drama," *Proc. 4th Int'l Conf. Autonomous Agents*, ACM Press, New York, 2000, pp. 301-308.

Richard Lindheim *is executive director of the University of Southern California Institute for Creative Technologies. His research interests include computer graphics and 3D animation, as well as large-screen immersive display technologies. He received a BS in electronic engineering from the University of Redlands and did graduate work at the University of Southern California, specializing in computer animation. Contact him at lindheim@ict.usc.edu.*

William Swartout *is director of technology at the University of Southern California Institute for Creative Technologies. His interests include artificial intelligence, virtual humans, agent architectures, and natural-language processing. He received a PhD in computer science from the Massachusetts Institute of Technology. Contact him at swartout@ict.usc.edu.*