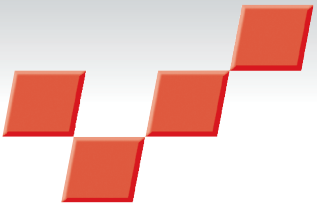

Serious Games for Immersive Cultural Training: Creating a Living World

Marjorie A. Zielke, Monica J. Evans, Frank Dufour, Timothy V. Christopher, Jumanne K. Donahue, Phillip Johnson, Erin B. Jennings, Brent S. Friedman, Phonesury L. Ounekeo, and Ricardo Flores ▪ *University of Texas at Dallas*

In Rockstar Games' Grand Theft Auto IV (GTA IV), the in-world juke box plays songs remixed for Liberty City, the fictional, yet very real feeling, place where protagonist Niko Bellic pursues ethically questionable missions with lingering moral angst. Virtual nonplayer characters (NPCs) use umbrellas when it rains. Low-income neighborhoods have cracked and broken streets with eroding asphalt. Various brands of cars handle differently, and the driver's body reacts accordingly when slamming on the brakes. The publisher uses its own RAGE (Rockstar Advanced Game Engine) software and incorporates the Euphoria engine from NaturalMotion to create a living world where gamers care not only about the character but also about Liberty City.[1,2]

As a major strategy of the game, players spend time learning Liberty City's culture. Does this mean that virtual Liberty City has a culture? Culture can be defined as socially transmitted information that shapes and regulates human behavior so that people can survive and reproduce. Furthermore, culture has mental, behavioral, and material aspects and provides a model for proper behavior.[3] Indeed, GTA IV's Liberty City is a living world with a culture.

Although fictional, GTA IV represents a real issue in 21st-century society—the need to understand culture to survive and be successful. From an American perspective, the 21st century has been largely characterized by military and political conflicts in the Middle East, where marked differences in philosophy make diplomatic and military strategies difficult to formulate and implement. Stakes are high in dealing with cultures that are unknown. Frustration permeates, and dollars and lives can be lost when distinctly different cultures interact blindly.

Living worlds offer a nonlinear, unscripted process for experiencing and safely learning the cognitive complexity and nuance of culture through emergent high-fidelity simulation. The 3D Asymmetric Domain Analysis and Training model uses visual, auditory, behavioral, and cultural models for immersive cultural training using the living-world construct.

Supported by funding from US government sponsors, the Institute for Interactive Arts and Engineering (IIAE) at the University of Texas at Dallas (UT Dallas) has created a serious game that lets players increase their cultural expertise in simulated Afghan rural and urban environments. Toward that end, we developed the 3D Asymmetric Domain Analysis and Training (3D ADAT) model, a recursive platform for the development and visualization of dynamic sociocultural models. This model integrates visualization, sound design, and behavioral and cultural modeling with recursive assessment tools to create a living world that is sensory and culturally realistic. Figure 1 (next page) illustrates the rich detail and character development possible with the 3D ADAT model.

Serious Games for Learning Culture

Video game technology can be a successful bridge to cultural understanding. One reason, of course, is the veracity of the 3D environment, with fuller sensory representation and immersion than is

Gaming Development for Cultural Training

For the past several years, the US military has recognized the value of using video games for training, to accommodate the media preferences of many military personnel and to leverage the fidelity of the technology. Middle Eastern cultural training is a key component of this initiative. The Institute for Creative Technology (ICT) at the University of Southern California (USC) is a major developer of interactive military training modules. Certainly the need is obvious for technological training solutions to accommodate an emergent and asymmetric theater among a little-understood culture. For example, the Enhanced Learning Environments with Creative Technologies (Elect) suite of products utilizes the PsychSim social simulation system, which seeks to understand social interactions.[1] Elect employs an intelligent coach and tutor to provide the student with pedagogical feedback about social and cultural issues. Current applications are located in a virtual Iraq, but plans exist to generalize the application to other cultures.

Also from USC, the Densely Populated Urban Environments (DPUE) module explores modeling of large-scale (40,000 to 50,000 people) densely populated urban locations in noncombat situations at multiple levels of detail.[2] Emphasis is on the development of a large-scale modeling framework and the results obtained under specific infrastructure and power struggle constraints.

Another entry, Global Conflicts: Palestine, by Serious Games Interactive, supplies a firm storyline. The game allows the player to be a journalist with the goal of getting a "solid story" among "troubling issues of human rights, checkpoints, settlements, martyrs and suicide bombs that haunt the region."[3]

References

1. "Elect," Inst. for Creative Technology, Univ. of Southern California, 2007; ict.usc.edu/projects/print/447.
2. "Densely Populated Urban Environments (DPUE)," Inst. for Creative Technology, Univ. of Southern California, 2007; ict.usc.edu/projects/print/473.
3. "Global Conflicts: Palestine," Serious Games Interactive, 2008; seriousgames.dk/new/products.php.

Figure 1. A screenshot of the 3D Asymmetric Domain Analysis and Training (3D ADAT) model. This model represents behavior and cultural models in a visualized, cognitively complex interactive environment.

possible with text-driven, noninteractive, or nonvisual learning tools. Serious games allow increasingly high-fidelity visual representation, integration of behavioral and cultural models that drive group and individual-agent AI, and ambient and event-driven sound. They also allow the assumption of different player roles and asynchronous learning environments that enable a self-directed training experience. All of this leads to a compelling and motivating learning environment.

Networked serious games allow for multiplayer involvement, training without physical presence, and an ability to update the virtual culture with new information, characters, and interior and exterior physical environments. Players can contribute to the knowledge base of the simulated society. Serious games can also be multitiered and take advantage of secondary asset sources such as Google Earth. Such games leverage the multimedia gaming environment familiar to recreational players. They are simulations with more sensory and cognitive complexity, and can readily imitate the dynamic and nonlinear nature of life itself from a behavioral and cultural perspective. Serious games can be unscripted and behave asymmetrically—similar to the cultures they attempt to model. They represent a "safe" way to make mistakes and to appreciate the multilevel implications of on-target or misguided cultural interactions and interpretations. Such games facilitate data capture, analysis, and assessment along with recursive dynamics. A serious game can be designed as a living world—representing the physical, auditory, cultural, behavioral, and emotional characteristics in increasingly realistic ways.

Living Worlds Teach Culture

The goal of many of the recent deeply immersive and interactive training applications is to create a safe way to learn culture in as realistic an environment as possible. (For some examples, see the "Gaming Development for Cultural Training" sidebar.) Living worlds are the next evolution toward this objective.

Living worlds go beyond linear game play and therefore are more directly similar to real life, which, of course, is not scripted. In real life, when visitors go to a new culture, what do they do? They explore. They interact with inhabitants. They determine who in a society is worth meeting to accomplish a goal, and who is of little apparent consequence or current value. Furthermore, as illustrated in Figure 2, by observing interpersonal interactions, visitors learn about the culture. Nonverbal signals such as facial expressions also offer cultural clues.

Furthermore, interactions in real life are not isolated or without consequence. If we meet someone and somehow our interaction is not satisfactory, that person will likely gossip about us, and thereby affect the attitude of the community. Offend a person who has either much influence or many friends and acquaintances, and the negative viral affect can be far-reaching. However, if a powerful person with influence is favorably impressed with an interaction, then positive effects can also grow exponentially. In short, we have an impact on the place we visit. A living world should reflect these dynamics.

The 3D ADAT Model

The model focuses on the living-world construct to create an agent environment at the individual, group, and broader geographic level. There are 200 independent NPCs in the environment—each with a history, an agenda, and the ability to form an emergent opinion. The 3D ADAT model focuses on understanding culture and creating technology that evolves and interacts in asymmetric, nonlinear ways and mirrors the unpredictable behavior of real-life settings.[4–6] The experience in the environment is unscripted and dependent on actions of the player. The feedback is not by a tutor, per se, but rather an analysis of the unfolding of the interactions within the community. Furthermore, the emphasis of the 3D ADAT model is on high-level fidelity of the visual environment, auditory models, and behavior of the NPCs.

In addition, the model relies on not only the individual actions and reactions of the NPCs but also the behavior of the group. This leads to accuracy in situation assessment. The echoing effect of the players' actions lingers throughout dynamic game play. The game's "viral" nature allows the player to listen to community buzz. The multilayered audio structure of the 3D ADAT model adds the final element of cognitive complexity to the environment, simulating global, regional, and local sound that supports the living-world construct.

Figure 2. A depiction of social behavior and facial expression in the 3D ADAT model. Understanding proper cultural nuance in handshake; personal space; and the differences in greeting men, male elders, and women of different ages are critical for acceptance in new cultures. Facial expressions vary by culture and can also indicate emotion or stress. Making these observations helps us learn a culture.

Although the challenges addressed by the 3D ADAT model have been daunting for quite some time, we believe we have raised the threshold of realism within our environment to achieve a living world, reminiscent of the most popular recent commercial game releases (for some examples, see the "Living Worlds in Commercial Games" sidebar on the next page), that acts very closely to the environment it mirrors.

Research and Validation

The research for the visualization, auditory component, and behavioral and cultural models spanned a variety of sources and included information provided by the sponsors and developed internally at UT Dallas. Research included unclassified human-terrain (that is, cultural landscape) team reports, review by subject matter experts throughout the military community, social-network-based photo-sharing Web sites such as Flickr, and even a doctoral dissertation on Afghani architecture.

Literature and commentary supplied by subject matter experts validated several of our assumptions about Afghani culture. The traditional Afghani salutation displays, greetings, and farewells were confirmed as accurate. Several fine-grain aspects of Afghani etiquette were validated. These included close interpersonal standing distance, small talk preceding serious conversations, the topic of women being taboo, and indirect communication being preferred over directness. We did discover that our understanding of the structure and nature of Afghani names needed slight adjustment.

Living Worlds in Commercial Games

Living worlds are an increasingly important design construct in commercial games. Popular titles have made important technology breakthroughs in the creation of dynamic, reactive, and, above all, convincingly realistic worlds. Rather than create simple game AI where nonplayer characters (NPCs) move along predictable, prescribed paths, titles such as Big Huge Games' The Elder Scrolls IV: Oblivion and Rockstar Games' Grand Theft Auto IV (GTA IV) use complicated character simulations to more fully realize a lifelike, albeit fictional, game space.

New game engines and middleware are providing the required depth of realism. Many games have focused on providing the player with a complex, emotionally realistic AI companion, such as the scientist Alyx in the Half-Life 2 series or the player's faithful dog in Fable 2. Others have focused on creating smarter NPCs that can adapt to their environment. The first-person shooter F.E.A.R. generates context-sensitive behaviors on the fly for a variety of enemies to provide greater combat and strategic challenges for the player.[1]

In relation to our work in cultural training, the most interesting games are those attempting to realize true living worlds through dynamic virtual populations. The Elder Scrolls IV: Oblivion uses the Radiant AI system, in which "a character's awareness isn't strictly limited to a few hard-scripted objects or activities. It radiates out into the surrounding environment."[2] NPCs in the Oblivion game space can choose to interact with nearly anything they come into contact with, including the player, other NPCs, and environmental objects such as chairs, books, farm implements, and weapons. These interactions are shaped on the basis of multiple parameters, including the NPCs' individual statistics (affecting such things as combat skill and intelligence) and a list of internal goals each NPC must carry out.

The intent of this system is to create populations of game characters with realistic behavioral patterns and structures, to alleviate any feeling that the player is wandering through a static, unchangeable, and therefore less-convincing world.

Part of the success of Radiant AI, according to producer Gavin Carter, is that the AI creation system is very closely linked to the scripting and quest development systems, making it easier for developers to create immersive, complex scenarios.[2] The Oblivion team also developed tools to link a robust conversation and dialogue system to numerous changeable conditions in the world space, so that NPCs can dynamically comment on everything from individual NPC and player names to the state of the weather. These tools allow developers to create a world that truly feels living through generated dialogue, rather than by writing specific dialogue that quickly becomes stale and breaks a player's immersion in the game space.

Other games have looked at different constructs to create believable worlds. Creatures, an artificial-life program created by Steve Grand, allows players to breed small furry creatures called Norns that can be taught to speak in simple English, feed themselves, and eventually protect themselves against outside threats.[3] Rather than creating parameters for the creatures, the game codes Norns from the genetic level upward, each one with a distinct DNA code, biochemistry, and neural-network brain. The game is notable as a breakthrough in alife (the process of simulating biological phenomena in an artificial environment), in part because Norns behave and evolve similarly to real living creatures, comparable in intelligence to small domestic animals such as cats.

These living-world systems can have important value for teaching culture through serious games because they present players with unpredictable, real-life situations that can be quickly recreated, and to which there are no simple, black-and-white problems or solutions. Furthermore, although the games described here are uniformly single-player (although players of Creatures can trade the DNA-files of their respective creatures with others), the potential for implementation in multiplayer and massively multiplayer titles is immense.

References

1. A. Champandard, "Assaulting F.E.A.R.'s AI: 29 Tricks to Arm Your Game," 2007; aigamedev.com/reviews/fear-ai.
2. B. Stein, "Radiant A.I.," Microsoft, 2005; xbox.com/en-us/games/t/theelderscrollsivoblivion/20051208-radiantai.htm.
3. S. Grand, "What Is Creatures Labs' Cyberlife?" Gameware Development, 2008; gamewaredevelopment.co.uk/creatures_more.php?id=520_0_6_0_M27.

Of course, the research-and-validation process is ongoing.

Cultural Knowledge Engineering

Like other living worlds, the 3D ADAT model makes a deep effort to simulate the physical space it represents through cultural knowledge engineering.[7] Before a simulated self-driven cultural environment can be created, the target culture must be understood in terms of daily errands, conversational topics and etiquette, the culturally appropriate response to community stresses, and proper conduct of negotiation. Furthermore, these constructs must be understood at multiple levels within the environment. After this research phase is complete, the second step of the cultural-knowledge-engineering process begins—translating the essential elements of the culture into programmable behaviors and artifacts.[7]

As we mentioned before, current applications of the 3D ADAT model examine urban and rural Afghan culture.[6] At the highest level, Afghan culture can be described as patriarchal, roughly egalitarian, and tribal. In tribal cultures, social power is organized by kinship relations.[3] In Afghan culture, older men have great influence over younger men, women, and children through local traditions and Islamic law, yet kinship obligations constrain the power of the individual and steer behavior toward benefitting the family network.[8]

Ideologically, the guiding principles of Afghan culture are a sense of familial and tribal honor, gender segregation, and indirect communication. The resulting network of interdependent connections is especially important given the traditional lack of a strong central government and responsive institutions.

Game Play

The two key elements of game play are winning over individuals, factions, and the entire living world, and characterizing the NPCs.

Winning over the living world. Ultimately, in the game based on the 3D ADAT model, the player wins by successfully interpreting the environment and achieving the desired living-world attitude toward that player. The entire living-world game space is fueled by the knowledge-engineering process we just described, and winning requires successfully navigating cultural moves in the game space.

NPCs who represent typical inhabitants of the Afghani city or rural village meet the player. Each NPC comes with demographic and psychographic characteristics such as gender, age, community status, and culturally specific behavioral attitudes. Furthermore, the NPCs have factions, and within these groups, NPCs have a wide-ranging degree of influence dependent on their place in society. For example, elder men have much influence. Boys have some influence. Women and girls have little influence.

If the player wants the inhabitants of the living world to feel positively toward the interaction, the norms of the cultural environment must be followed. For example, if the player addresses a married woman of child-bearing age directly, this is not culturally acceptable. The woman will then, albeit demurely, gossip through the community that the player does not display socially acceptable behavior. This opinion is shared virally through the NPC community. Should the gossip reach an NPC with much influence, the damage of the inappropriate encounter is manifested. Eventually the entire community becomes a unified agent with an opinion of the player's handling of the society's culture.

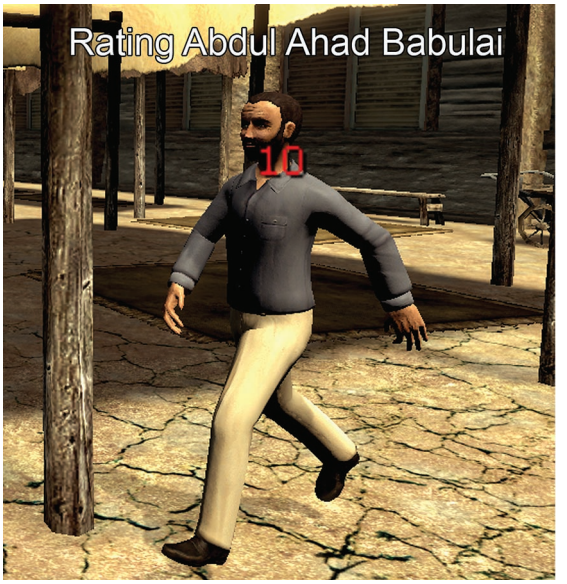

Figure 3. A rating for an NPC (nonplayer character) in the game based on the 3D ADAT model. Players rate the emotional mood of NPCs on the basis of visual, auditory, and behavioral clues in the environment.

Characterizing NPCs. As Figure 3 illustrates, the game lets the player rate an NPC's emotional mood—such as happy, sad, angry, or frustrated—on the basis of conditions in the world or as a result of the NPC's interactions with the player or other NPCs.

Visualization—the Living-World Environment

For the player to test his or her ability to succeed in the cultural environment, it is important that the living world represent the physical space (in this case, approximately three city blocks) as accurately as possible and present a full array of sensory clues. For example, the urban Afghan space models a three-tiered socioeconomic structure with corresponding virtual representation. Laundry hangs from the lower level of apartments because the middle-class inhabitants do not have electrical appliances. Vendor carts are inhabited during the day but become vacant at night as the residents in the lower economic strata return to their dwellings on the outskirts of town. The glass and bricks on the buildings are modeled after regional photographs. Interiors of typical regional edifices such as mosques are also part of the game play.

Unreal Engine 3, one of the most popular engines for commercial shooter games, was chosen as the base for building the cultural simulations. As such, it provides both a powerful platform for rapid development and an interesting challenge for adaption. Our primary motivation for choosing this engine was its graphical capabilities. The Unreal Engine has built-in visual scripting languages that allow us to easily and quickly develop complex pixel shaders, animation blending algorithms, and particle systems—all vital to realism. This

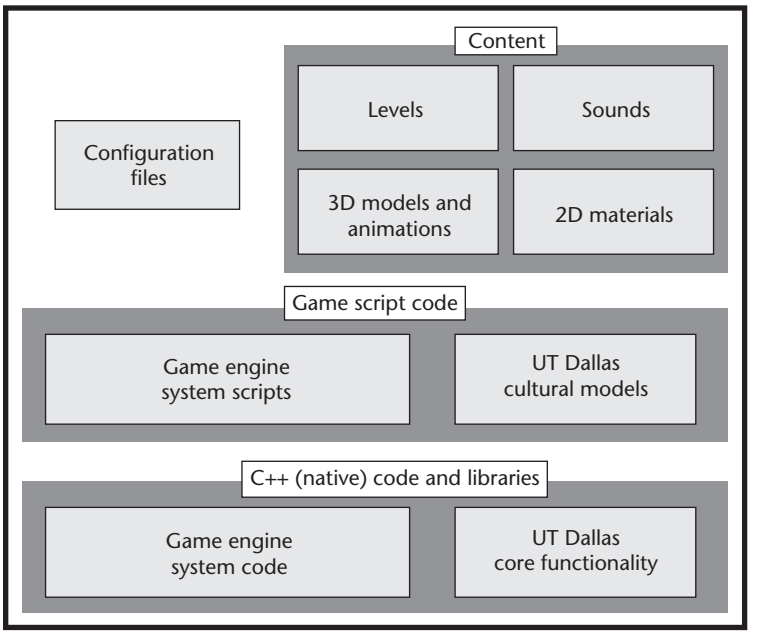

Figure 4. The architecture of the 3D ADAT model. The model utilizes a game engine, scripting, and C++ development.

Figure 5. The Random People Generator allows us to quickly create high-fidelity, physically diverse, and culturally accurate NPCs for the environment. This technique will also allow us to integrate multiple cultures into the living world as needed.

"technical art," which usually requires the attention of a skilled programmer, can instead be built by someone with little or no programming experience. This art content is hooked into our cultural models through scripting. To maintain flexibility, the simulations are being designed in a modular fashion that will accommodate additional candidate engines. Figure 4 illustrates the relationship between the game engine, scripting, and C++ in the development of the 3D ADAT model.

Visualization—the World's Inhabitants

Men, women, boys, and girls of all ages and economic strata make up the inhabitants of the city and act according to established social and cultural norms. Elders are represented in the world with white tips on their beards. Clothing typifies the appropriate affiliations and class. Some characters in the world act in an atypical manner, as would be expected in a large group of people that would likely include some deviants.

Generating characters. Our model stochastically describes the members of our virtual world by randomly generating families from demographic data. At the beginning of each game, our process

- generates one family for every home, with data for economic status, and
- uses those family attributes to generate men, women, and children based on demographic data for the target culture.

The result is that each game has completely different characters and families but that each instance of the game is still roughly representative of the actual demographic data.

Using UT Dallas's Random People Generator (RPG), as illustrated in Figure 5, we construct the physical materialization of an NPC from several components, such as hair, facial hair, eyes, hat, upper torso, lower torso, and foot. Components are organized in an assemblage of categorized meshes. An NPC is generated from a random combination of one resource from a category for each component, within biological and cultural conditions. Head and facial hair color, skin tone, and skin details such as pockmarks and wrinkles are given random numerical values in a range, allowing the procedural generation of an even more diverse population.

Motion capture. To create characters who look and act the way individuals would in an Afghan urban or rural setting, motion-capture animation is performed on individuals as similar to the actual inhabitants as possible. Our motion capture capability allows us to capture both facial and full body performances from real actors and apply them to virtual NPCs. The motion capture illustrates not only typical movements such as walking and sitting but also culturally endemic actions such as kneeling. Facial animations also offer culturally specific expressions.

We use the Vicon MX optical system with 16 high-speed cameras. Eight of our cameras are the

T13 model, which offers 1-megapixel resolution, captures 10-bit grayscale using 1,120 × 896 pixels, and runs at 120 frames per second. The other eight are the T40 model, which offers 4-megapixel resolution, captures 10-bit grayscale using 2,352 × 1,728 pixels, and runs at 120 fps.

As is standard in motion capture, the actor wears a suit with reflective markers placed on key joints to capture skeletal movement. As the actor begins his or her performance, the cameras constantly flash their LEDs and capture the light that bounces back from the reflective markers. The system can then triangulate a marker's position on the basis of the data sent by multiple cameras. Using this process, the system can then make a virtual representation of the markers. This representation runs in real time and appears as motion. We then apply the data to the virtual-character skeletons.

For facial motion capture, we place 3 mm reflective markers on the actor's face. These are placed according to major facial muscles, in order to capture expression. We can then process and place these data on different facial setups for virtual characters.

The primary reason to use motion capture is to represent highly nuanced cultural facial and body movement. Figure 6 illustrates the integration of motion capture to create gestures during conversation.

Figure 6. Motion capture converted into virtual-character behavior. Motion capture helps us develop culturally representative gestures.

Behavior in the Living World

In order to create a self-driven game world, a number of interrelated behavioral models were designed for the NPCs. These models control activities ranging from gossiping to daily behavior such as errands.

The gossip model. As Figure 7 illustrates, the game uses the gossip model to spread faction-relevant information. The gossip model represents NPCs' opinions of the player, each other, and abstract ideas, such as a religion or a particular faction. Opinions are modeled on a scale of +100 to −100, where 100 represents adoration, 0 represents apathy, and −100 represents abject loathing. As NPCs move around the game space and interact with each other, they might gradually shift the collective opinion of the game space.

An important facet of the gossip system is the way in which NPCs react to the player. As we mentioned before, the NPCs change their opinions of the player on the basis of the player's behavior in the living world—in essence, reacting to whether the player has been culturally polite or offensive.

The traffic/errand model. The traffic/errand system models what NPCs do over the course of a normal day, providing the opportunity to meet and exchange ideas. The system also creates a strong illusion of a bustling, dynamic town. Players can follow an NPC guided by this system throughout an entire day in which the NPC will go shopping, perform chores, eat meals, and eventually return home to sleep.

The traffic/errand system functions by assigning each NPC an appropriate list of tasks to complete during the day such as dressing, shopping, and working. This list is randomly generated and weighted on the basis of specific criteria for each NPC in the

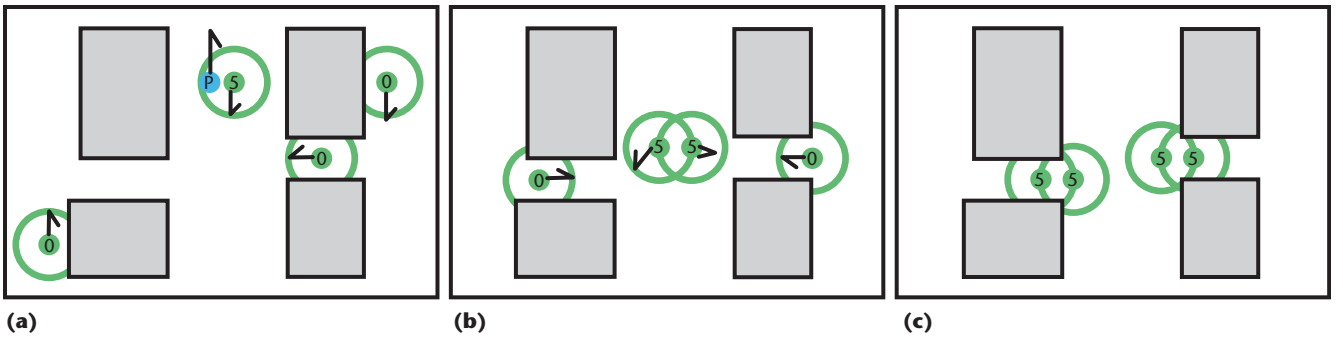

(a) (b) (c)

Figure 7. The gossip model: (a) In stage 1, the player (P) interacts with an NPC. In this example, the NPC increases his or her opinion of the player from 0 to 5 on the basis of the interaction. (b) In stage 2, NPCs continue to run errands. NPCs that have an opinion of the player begin to share them with NPCs that do not have an opinion. (c) In stage 3, eventually all NPCs in the area have formed an opinion of the player through word of mouth.

Figure 8. A shura is a Middle-Eastern dispute resolution custom. It can be dynamically generated within the living world if circumstances warrant.

game. These rules include what kind of job, if any, each NPC has, along with the NPC's social status.

Another aspect of navigation, the grouping model, causes the NPCs to gravitate toward other NPCs they consider friends or family. They also avoid NPCs they dislike.

The situational-awareness model. Players can detect that a situation or conflict might be imminent on the basis of the daily actions of NPCs. The situational-awareness tool weighs the resources of major factions or groups in the game space against their opinions of one another. As resource levels become increasingly imbalanced, the NPCs will automatically reduce their errand lists, dependent on priority. For example, in a high-threat situation, NPCs will still gather food but will spend little to no time socializing with their neighbors. NPCs will also tend to move slightly faster through the game space.

The resource variable is set manually for the different factions. By giving different factions significantly different levels of resources, the amount of threat can be increased in the simulation. This can also be used as one of several ways to set a difficulty level for the game. A setting with wildly differing resource amounts will be more unstable than one where most factions have equal resources.

The tactical, cultural, and faction model. Tactical and cultural refer to the types of actions that a player can take while in the living world. Tactical actions are specifically meant to help achieve the goals of identifying NPCs and rating them in regard to certain attitudes.

Cultural actions are those taken to maintain or increase the local opinion of the player or groups the player represents. These include things such as shaking hands at the start of a conversation, introducing yourself the first time you speak to an NPC, and any number of other polite local customs.

The Social Environment

Factions can be created by blood and marital family ties, school affiliations, nationality, and, of course, a variety of economic and religious influences. Factions also affect the degree of "viralness" in a community. Entire factions can turn against a player, greatly influencing game outcome.

Players also bring their own opinions and game play into the world. As we mentioned before, players who react positively or negatively with the community produce world attitudes. This can come from personal experiences or a player's own cultural and behavioral perspective. Analyzing results of players with different levels of cultural knowledge can feed back into the system to enhance the reality of the game.

Dynamically generated game play also occurs, such as NPCs pursuing expected chores or errands. The player presumes a certain semblance of normality within the living world, where specific errands are run by world inhabitants each day under a variety of conditions or game states such as friendly, nervous, happy, or angry. The NPCs choose to run their typical day, or not, depending on the mood in the environment.

One specific form of dynamically generated game play is a *shura* (see Figure 8). A shura is an informal type of Afghan dispute resolution and collective decision-making system. Shuras take place both in rural and urban contexts.[9] Theoretically, any member of the community can call a shura on any subject. However, only major leaders in the community do this without consulting others first. Shuras are exclusively male gatherings where many participants are the heads of households. Younger males are free to attend and speak but often will remain silent and defer to their elders. All male community members directly affected by the topic of the shura are expected to attend, because any decisions made are considered binding. Consultation, bargaining, and lobbying will continue without any strict time limits until a resolution is reached. Elder figures will typically try to guide the shura toward an honorable and satisfactory end in order to encourage peace and stability within the community.[10]

A shura can be generated independently and spontaneously by NPCs if conditions warrant. The NPCs fill three roles: the arbiter responsible for controlling the shura and the two opposing parties of the dispute. The shura is dynamically played out on the basis of each party's resources and the severity of the dispute. The goal of the arbiter is to encourage the parties to exchange sufficient resources for both to survive. This can be difficult if one party particularly dislikes the other. To assist in this goal, the arbiter's ability to influence both parties gains strength as the shura goes on.

Cognitive Complexity

Visualization with the 3D ADAT model offers environmental clues and representations. The behavioral models offer social and emotional cues. Auditory stimulus is also presented in the environment. Certain sounds are expected; others are not. Some sounds, such as call to prayer, might add to the cultural immersion and the necessary multisensory cognitive complexity. We also would expect to hear vendors in the street and children playing. However, if we don't hear these sounds or the NPCs are not where we expect them, this might be a sign that something is amiss. The auditory component also helps the player work through potential confusion that foreign sounds add when he or she enters a new environment.

The Auditory Component

Creating sound design for a living world that accurately reflects culture requires strategies beyond those used in the linear representation of cinema or the traditional interactivity of video games. In film, an auditory scene adds to immersion and realism and is indicative of subtle changes in moods and attitudes. However, sound in a traditional movie does not focus on a particular viewer or respond to that viewer's actions. Conversely, video games use the actions of the player to contribute to the actualization of the environment. However, sound in a traditional interactive approach is used primarily as an effect and as such represents the change in the environment directly resulting from a visible cause related to player action.

For a living world, the film and traditional interactive approaches need to be complemented by an ecological approach to analyzing and modeling the dynamic characteristics of the sonic environment. Sound must reflect the multisensory cognitive complexity of the environment and contribute to immersion by augmenting or complementing the consistent realism of the representation. Because the most relevant characteristic of a living world is its capacity to evolve and change, the sound has to support a consistent formalization with the behavioral and cultural models of the project, such as those representing errands or gossip. To achieve this effect, we've divided the systemic or ecological sounds of the living world into three spheres: global, regional, and local.

Global, regional, and local sound design. The global sphere contains all of the combined sounds of distant traffic, construction, weather conditions, and so on from any parts of the living world. The changes occurring in this sphere are equally perceived by all the NPCs and players. The global range is affected only by changes occurring in the environment and is the result of events such as the passing of time or the massive convergence of multiple actions.

The regional sphere contains all the sounds that can be heard only as independent elements within this sphere. Examples of isolated, local sounds are a radio playing in a room and construction noise on a street. Outside this sphere, they are masked by or assimilated to the background sound from the global sphere. In this sphere, the player and all the NPCs present perceive the changes. The regional range of perception is affected by actions performed by the player or NPCs in the environment, or by actions that the NPCs perform on objects.

Creating sound design for a living world that accurately reflects culture requires strategies beyond those used in the linear representation of cinema or video games.

The local sphere contains all the sounds of proximity: footsteps, voices, and activities of the closest NPCs. In this sphere, only the player perceives changes. The local range of perception is affected only by the presence, location, or displacement of an NPC in the environment and by the settings of the sounds decided by the user.

Technical implementation of sound. We use several methods to generate the three-tiered dynamic sound environment. Sound cues can combine several different wave files and randomly mix and modulate the sounds, creating a unique playback each time the cues are triggered. Sound cues can be placed throughout a level to emit ambient sounds. The cues can also be attached to NPCs so that they themselves emit sound as they traverse the environment. We then program NPCs to select from an array of different sound cues, on the basis of factors such as subculture and mood.

Through programming and scripting, we trigger sounds to play on the basis of special events, such as actions of players, reactions of NPCs, the time of day in the game, or speech during a conversation.

We can filter or modify different characteristics of each individual auditory component of an environment. For example, we can adjust reverberation to match different acoustic environments, such as alleyways; small, medium, or large interiors; cities; mountains; and plains. An area's ambient sound

Figure 9. A conversation between two NPCs. Conversations occur throughout the living world between the player and NPCs and between the NPCs themselves. NPCs can talk on a variety of topics such as small talk, sports, weather, and America and can spread gossip.

can also change on the basis of the number of given actors in that space. This means that the ambient sound in an area can be louder and more hectic as NPCs continue to gather in that location.

Conversations. To facilitate realistic conversations with our NPCs, a natural-language generator pulls in contextual data to randomly generate sentence responses. The process consists of a lookup table, which defines the general structure of responses, and topic data that "fills in the blanks." In order to get grammatically correct responses, our sentence structures include fields for demonstrative and nominative pronouns, gender, nouns, subject-verb agreement, and more.

For instance, the sentence structure

%subject @have@ made %me angry recently.
%Npronoun should just leave %me alone.

would manifest in many forms on the basis of the current topic of conversation. All of the following are valid sentences generated by this sentence structure:

- John has made me angry recently. He should just leave me alone.
- The Wilkinsons have made me angry recently. They should just leave me alone.
- I have made myself angry recently. I should just leave myself alone.

Designing these sentences obviously requires a strong understanding of the target grammar. A major benefit of the underlying system is that new languages can be easily added given an understanding of the target grammar.

Initially, conversations will occur as an interaction between two characters. When an NPC is activated by the player, the system determines which conversation to use and then activates the default topic, such as a greeting. The system then follows the scripted pathway made by the designers for that particular topic. Once this conversation path or "thread" has completed, the player may select a new topic. Until that point, the player is barred from changing the topic.

The player might be offered a set of options during a thread. This list of options replaces the list of topics in the interface, and the thread stops until the player clicks on one of the responses. This facilitates two-way dialogue and interaction between an NPC and the player beyond just simple interrogation.

Well-designed conversations could be intelligence-agnostic. That is, a well-designed conversation could work for both player-NPC interactions and NPC-NPC interactions. For instance, the conversation designer can put in tags to give hints to the AI about what topics are acceptable and what options are suitable for various NPC types.

Figure 9 shows an example conversation between two NPCs.

Serious about Assessment

The 3D ADAT model also allows for analysis of the cultural behavior exhibited by the player in the game. Conversations and interactions between the NPCs and the player are recorded through a text log to provide game performance analysis. For example, a log is kept of offensive actions such as rudely initiating or breaking off a conversation. Other offensive actions such as touching a female on the arm or gesturing with the left hand are also recorded. A log of conversations between NPCs is also kept, allowing a textual analysis of the viral communication of the NPCs.

A trainer also can assess the player's ability to perceive an NPC's emotional state compared to the actual value set by the living world. Furthermore, the assessment tool lists all possible choices for player behavior and conversation, highlighting both the player's choice and the most culturally appropriate or inoffensive response. The tool provides scores on the opinion of the player at the NPC, faction, and village level. Additional comments can be provided that highlight the player's weaknesses, explaining why a particular response is most appropriate. The player receives a score based

on these performance criteria. The goal of the assessment tool is to evaluate the player's in-game choices and provide meaningful feedback to improve future performance. Figures 10 illustrates the type of performance statistics that can be gathered, such as the number of conversations held and the number of people met, and the ability to tag an emotional state on individual NPCs and determine the overall attitude of the village toward the player.

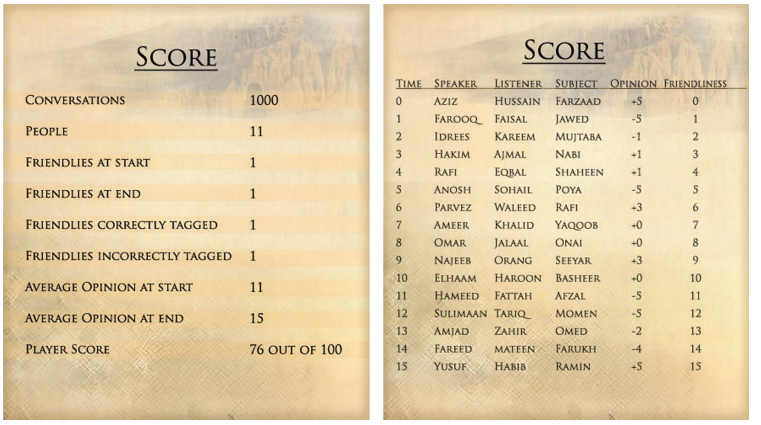

Figure 10. A display of statistics regarding the player's performance. The player wins the game by winning over the multilayered agents (for example, individual NPCs, factions, and the entire village or urban area) and correctly identifying the emotional state of the NPCs. Assessment tools also include the tracking of topics of conversation in the game.

Opportunities for Further Research

The 3D ADAT model is in the next stage of development for field rollout. We are expanding the geographical areas and cultures being simulated, as well as exploring medical applications and other sectors. Internal usability studies offer encouraging results regarding the ease of use of the 3D ADAT model for cultural training and understanding. We anticipate much more field data in the coming months to verify these initial findings.

Although this application of our model is specifically aimed at understanding correct interaction with Afghan culture, many other appropriate applications exist. A virtual cultural trainer such as the 3D ADAT model can capture multisensory, behavioral, emotional, and cultural nuance on both the macro and micro level. With the correct content, the platform could introduce players to serious game representations of any framework, be it global, national, corporate, or social. Training New York City police to work with fractional gang members on the street or in prison would be another powerful application.

Furthermore, this immersive, self-driven environment presents opportunities for more than just cultural acclimation. In fields where the random acquisition of knowledge is common, such as nursing or other medical fields, rare or unique training situations with cultural aspects could be presented with minimum risk.

Owing to the high profile of previously little-researched countries such as Afghanistan, culture is emerging as a critical area of study. Serious games, particularly those associated with cognitively complex living worlds, offer a promising area of research and create new thresholds for visual, auditory, behavioral, and emotional representations.

Acknowledgments

Many of the developments presented in this article rely heavily on four years of US Army-sponsored research at the University of Texas at Dallas Institute for Interactive Arts and Engineering. We acknowledge the work of Thomas Linehan, Andrew Blanchard, and the many graduate students, staff, and others from the institute.

References

1. H. Goldstein, "Grand Theft Auto IV Review," 2008; http://xbox360.ign.com/articles/869/869381p1.html.
2. B. Ortutay, "Watch Out: Game Characters Getting Smarter," msnbc.com, 13 July, 2008; www.msnbc.msn.com/id/25643688.
3. J.H. Bodley, *Cultural Anthropology: Tribes, States, and the Global System*, McGraw Hill, 2004.
4. T.E. Linehan and A.J. Blanchard, "Culture and Character: A Training Game," US Army Training & Doctrine Command Futures Center, 2006–2008.
5. T.E. Linehan and A.J. Blanchard, "HUMINT—War Gaming Trainer," US Joint Forces Command, 2006–2008.
6. M.A. Zielke and T.E. Linehan, "First-Person Cultural Trainer," US Army Training & Doctrine Command—G2, 2008–2009.
7. E. Feigenbaum and P. McCorduck, "The Fifth Generation: Artificial Intelligence and Japan's Computer Challenge to the World," Addison-Wesley, 1983.
8. R. Stewart, *The Places in Between*, Harcourt, 2006.
9. J. Beall and S. Schütte, "Urban Livelihoods in Afghanistan," Afghanistan Research and Evaluation Unit, 2006; www.lse.ac.uk/collections/DESTIN/pdf/AREU%20Urban%20Livelihoods%20SP%20FINAL%20PROOF%205Oct2006.pdf.
10. F. Barth, *Political Leadership among the Swat Pathans*, Athlone Press, 1965.

Marjorie A. Zielke *is an assistant professor of arts and technology at the University of Texas at Dallas*

and the principal investigator and project manager on several training and simulation project grants. Her primary research interest is cyberpsychology. Zielke received her PhD in aesthetic studies with an emphasis in arts and technology from UT Dallas. Contact her at margez@utdallas.edu.

Monica J. Evans *is an assistant professor of computer game design at the University of Texas at Dallas. Her research interests include game studies, serious games, and interactive narrative. She also coordinates the UT Dallas Computer Gaming Entrepreneurship Competition. Evans received her PhD in aesthetic studies, with an emphasis in arts and technology, from UT Dallas. Contact her at mevans@utdallas.edu.*

Frank Dufour *is an assistant professor of arts and technology at the University of Texas at Dallas. His interests include perception of music, auditory scene analysis, and acoustic ecology. Dufour received his PhD in communication sciences from the University of Paris 8. Contact him at fod051000@utdallas.edu.*

Timothy V. Christopher *is a lecturer in arts and technology and is earning his PhD in that area at the University of Texas at Dallas. His technical interests include digital and analog games, educational and experimental games, and game theory. Christopher received his MFA in arts and technology from UT Dallas. Contact him at khimbar@utdallas.edu.*

Jumanne K. Donahue *is a graduate research assistant in arts and technology and is earning his PhD in that area at the University of Texas at Dallas. His technical interests include cultural and behavioral simulation modeling. Donahue received his MFA in arts and technology from UT Dallas. Contact him at jkd033000@utdallas.edu.*

Phillip Johnson *is a graduate research assistant in arts and technology and is earning his MA in that area at the University of Texas at Dallas. His technical interests lie in producing compelling and explorative spaces using advanced graphics technologies. Johnson received his BA in arts and technology from UT Dallas. Contact him at pcj041000@utdallas.edu.*

Erin B. Jennings *is a graduate research assistant in arts and technology and is earning her PhD in that area at the University of Texas at Dallas. Her technical interests include instructional design and assessment, and serious game and simulation design. Jennings received her MS in education in instructional design for online learning from Capella University. Contact her at erin.jennings@utdallas.edu.*

Brent S. Friedman *is an undergraduate research assistant in arts and technology and is earning his BS in computer science from the University of Texas at Dallas. His technical interests include code optimization, language design, and metaprogramming. He's a member of the ACM and the International Game Developers Association. Contact him at bsf062000@utdallas.edu.*

Phonesury L. Ounekeo *is a graduate research assistant in arts and technology and is earning her MFA at the University of Texas at Dallas. Her technical background includes 3D modeling, texturing, level design, concept art, and creative writing. Ounekeo received her BA in arts and technology from UT Dallas. Contact her at plo041000@utdallas.edu.*

Ricardo Flores *is a research assistant at the University of Texas at Dallas. His technical interests are motion capture and game animation. Flores received his BA in arts and technology from UT Dallas. Contact him at ricardo.flores13@gmail.com.*