Virtual Humans as Participants vs. Virtual Humans as Actors

Andrew Gordon and Mike van Lent
The Institute for Creative Technologies
University of Southern California
13274 Fiji Way, Marina del Rey, CA 90292

Introduction

Should virtual humans be thought of as actors playing a role in a virtual play, or real participants living in a virtual world? Although the question is of philosophical interest, it also has implications in how the virtual human's knowledge, goals, communication and sensing are implemented. The question is also one aspect of the more general tension between realism and believability in virtual environments. In most cases, components of a virtual world (such as its physics) are made believable as a short cut when fully realistic components are more desirable but prohibitively expensive in development time and/or processing power. In the actor vs. participant case, however, an argument can be made that actors who behave believably, but have additional unrealistic knowledge (more knowledge than their characters would realistically know) and capabilities, are actually more useful than realistic participants. We propose the 2002 Spring Symposium on AI and Interactive Entertainment include a discussion on this subject. This proposal consists of an introduction to the topic and two one-page position papers to get the conversation started.

The choice of virtual human as participant vs. actor affects multiple aspects of the design and development of the virtual environment. Among the most critical issues are three that concern how the virtual characters themselves are modeled: what sort of perception capabilities do these virtual humans have, what are the goals that they are pursuing, and what means do they have to communicate with one another? Each of these areas are elaborated below:

Perception: If the virtual human is modeled as participant, sensing must be realistically limited to information that can be seen, heard, felt, smelled, or tasted from their vantage point in the virtual environment. Thus, sensing for participants needs to take into account the virtual human's field of view, lighting of the environment, background noise and other sensory limitations and interference. A participant's behavior is realistic in that it is based solely on the information available to his realistic senses. Actors, on the other hand, may best be implemented such that they have full, omniscient sensing of everything in the environment. This is analogous to how a human actor in a play knows what is going on behind him on stage, even without looking, due to his knowledge of the script and the roles of the other characters. This does not mean that actors need not know the realistic limits of their perception, however, as they must pretend that they do not have this information in order to make their behavior believable to the audience.

Goals: For virtual humans modeled as actors, the additional knowledge that is provided through omniscient perception is employed in pursuit of narrative goals: to behave in a way that elicits a desired reaction from the audience (or a human participant). In a sense, actors pursue their characters' goals in a play along with certain meta-goals of achieving particular effects in the minds of the audience. To support the pursuit of these additional meta-goals, virtual human

actors need to be constructed to include additional procedural knowledge about the ways in which entertaining, educational, or emotional responses can be elicited – the sorts of techniques that might be taught to humans in acting schools. In addition, certain types of declarative knowledge are needed to support these techniques, including knowledge of future events in the storyline, knowledge of the full context of the situation, and knowledge of the internal mental states of other characters. In effect, this knowledge is about the story being presented to the human participants, which enables the use of narrative mechanisms such as foreshadowing or sucker-punch endings. For modeling virtual humans as participants, no such meta-goals are employed: the characters simply pursue the goals that humans would realistically pursue if the virtual environment were real. The participant's procedural knowledge need only include methods for achieving their immediate goals, and their declarative knowledge is based only on background information and the evidence gained through their perceptions.

Communication: A final implementation issue that is affected by the choice of participant or actor virtual humans is that of communication between characters. For virtual humans modeled as participants, communication between characters, either between virtual humans or between virtual and real humans, must be limited to ensure realism. To successfully communicate, characters must share a common language and be in close proximity to one another or employ a communication device. Virtual actors, on the other hand, need not adhere to these restrictions when communicating with each other as long as their behaviors remain believable to the audience. In fact, virtual actors can conceivably have capabilities that even their human actor counterparts do not have, such as direct actor-to-actor collaboration through some sort of virtual extra-sensory-perception (ESP) - where human participants would not be able to perceive that a communication had even taken place. This sort of communication could be especially useful in support of dramatic effects that rely on the synchronization of actors' behavior.

In the following two sections, opposing views on the choice between virtual humans as actors vs. participants are presented.

Virtual Humans as Actors

Andrew S. Gordon

While computer game artificial intelligence (AI) has not traditionally been seen as an arena for deep philosophical debate, there is one question that is increasingly raised that has significant practical importance: should we model the computer-controlled agents in games as *actors* that adopt the role of non-player characters or as *simulated people* in a simulated world. Fortunately, this is a debate that proponents on either side, and is one that invites computer game AI researchers to address questions that are central to the entire AI field. The comments presented here are made in support of the former position, computers as role-players.

To understand the opposing positions, imagine that it were possible to explore the personal thoughts of a computer-controlled agent in any one of the adventure games that are on the market, e.g. a computer-controlled teammate in id Software's Quake series or Epic Megagame's Unreal series. If we model these non-player characters as simulated people in a simulated world, then we might imagine thoughts like the following: "It looks like those blue team villains are on the run. I think it is safe to come out of hiding now. I'll just wait for my red team leader to make

the first move, and I'll follow him with my other comrades to glorious victory!" In contrast, if we model these agents as actors that play the role of these non-player characters, the inner dialogue may be somewhat more complicated: "The user doesn't seem to realize that his character's team has the advantage in the game. If the user decides to advance now, he will be in for an exhilarating and entertaining battle. I'll make it clear to the user that my character is waiting for his character to make the first move, and that even though my character will probably be killed at a crucial moment for dramatic effect, I follow his character into this glorious combat!"

Given the current state of computer game AI, neither of these two inner dialogues is remotely feasible. However, even here in these fictional sentences, one of the key differences between these two approaches is apparent: by modeling computer agents as actors, they become explicitly concerned with the entertainment (and perhaps also the educational) value of the game to the user. Analogous to theater actors on stage or film actors in movies, we would like our computer-controlled characters to behave differently than the characters they are playing would actually behave if the situation were real to support entertainment goals. Just as theater actors will project their voices much louder than their character ever would, or wait until the audience's laughter subsides until delivering their next line of dialogue, we would like our computer-controlled characters to be primarily responsive and reactive to the entertainment needs of the user, not to the mission needs user's character. Ideally, every amount of development work that goes into the production of a computer game should increase its value as entertainment, and this is especially true in deciding how the virtual characters should be modeled.

Given the alternate position, where non-player characters are themselves computer simulations of people, entertainment itself becomes some sort of emergent effect – a by-product of the simulated environment that can be influenced either by carefully tuning the capabilities and dispositions of these simulated characters, or by reserving some degree of supervision and control of these agents by some central, fun-preserving game-master algorithm. This approach has proven successful in many past computer games, such as Maxis's line of simulations of cities, planets, and household (SimCity, SimEarth, and The Sims, respectively), where the entertainment value seems to draw from the concept of emergent behavior itself. One severe downside of this approach is apparent in both these games as well as in the new crop of massively multiplayer adventure games, in that any notion of a *storyline* is one that is imagined by the users rather than authored by game developers. By choosing to model virtual humans as actors, designers have more direct control of narrative flow, enabling designers to inject a much-needed dose of real storytelling prowess into otherwise unstructured games.

From a practical point of view, the temptation is adopt the simulated person approach on the grounds of simplicity. This position is certainly clear when the non-player characters are exceedingly stupid: it is much easier to model the thought processes of a mindless zombie than to model the thought processes of an intelligent actor playing the role of a mindless zombie. Indeed, mindless zombies are one of the few sorts of truly believable non-player characters in current computer games. Presumably, there is some hope among advocates of this position that progress in AI will yield smarter and smarter computer-controlled agents for use in games over time. Today's mindless zombies will incrementally be replaced by slightly more intelligent zombies, then very stupid humans, and ultimately by non-player characters that have the intelligence of

average people. The obvious fallacy of this view is that it is unimaginable that computer game users will ever be able to go up against ruthless masterminds of incredible brilliance, mutant renegades with inexplicable mental powers, or superior alien intelligences – all of which are characters that have become commonplace in the cinema despite the cognitive limitations of film actors and actresses. A more appropriate vision of the contribution of AI progress will be in developing more intelligent models of actors, and in the short run we'll have to be content with systems that play the roles of ruthless masterminds with the intelligence of mindless zombies.

Virtual Humans as Participants

Mike van Lent

Both virtual actors and virtual participants are capable of populating virtual environments and interacting with the avatars of humans exploring those environments. In choosing between implementing a virtual actor and a virtual participant it is necessary to consider the believability of the agent and the cost of implementing the agent. In many domains a virtual actor is better suited to eliciting the desired reactions in the human audience, however, virtual actors are generally more costly and less believable. In these cases, the advantages and additional costs of a virtual actor must be weighted. Frequently, a virtual participant will be cheaper, behave more realistically, and be almost as effective.

In many computer games the AI agents "cheat" through the use of omniscient sensing, magically appearing resources, and blatant disregard for the rules imposed on the human player. The designers of these games are using these cheats in an attempt to create challenging, but still believable, behavior in the easiest way possible. Although the resulting AI may be challenging it almost never turns out to be believable. This is evidenced by game players frequent complaints about "AIs that cheat." Virtual actors face exactly the same problem, taking advantage of unrealistic capabilities without the knowledge of the player. Virtual humans as participants, on the other hand, have no unrealistic capabilities and can't cheat. Although creating virtual humans that fill the role of challenging opponent is more difficult, the resulting opponents are more realistic and can provide a more satisfying gaming experience.

In some domains, a level of realism is required that virtual actors don't provide. For example, some training simulations require the virtual participants to behave very realistically to insure that the training will transfer to real situations with real humans. Other examples include cognitive modeling applications and simulations used to evaluate strategy and tactics. For these domains believable behavior isn't sufficient. The behavior must be as realistic as possible. One requirement realistic behavior is realistically modeling the knowledge and sensory information available to the agent. Limiting a virtual actor to realistic knowledge and information effectively turns it into a virtual participant.

As described in the introduction, virtual actors have additional knowledge and additional sensory information designed to reason about how the actor's behavior affects the human audience. In many domains this can make the human's experience in the virtual world more enjoyable and valuable. In a computer game, for example, a virtual actor can tailor its behavior to make the game more fun for the human player. This already happens, at a primitive level, with AI opponents in strategy games that are programmed to adjust their attack to lose to the human after

a close battle. In training simulations a virtual actor can tailor the training lessons to the student. Fast learners can be presented with more challenging lessons and frequent new concepts while slow learners can be given easier lessons and frequent reviews. I concede that there can be an advantage to virtual actors for many applications.

However, this advantage comes at a cost. As is made clear in the introduction, and supported in the previous paragraph, virtual actors have additional costs that virtual participants do not. In addition to the knowledge required to pursue their character's goals in the environment, virtual actors must also have the knowledge necessary to pursue the meta-goal of tailoring their behavior to achieve the desired effect in the environment. A good actor playing a soldier needs to know more than a real soldier. That actor needs to know what a soldier knows and needs to know how to act. This extra level of knowledge requires additional implementation time, to program the knowledge into the agent. It also requires additional processing time and memory to allow the agent to pursue the meta-goal of acting in addition to the believable pursuit of the goals inherent in the role being played.

The tradeoff between cheaper, more efficient virtual participants and additional functionality of virtual actors needs to be evaluated on a case-by-case basis. However, the success of sophisticated virtual participants in many applications suggests that virtual actors are not often, or even usually, the best choice.

Finally, virtual participants play a role in hybrid approaches, approaches that include both virtual actors and virtual participants. One example of a hybrid system is a virtual director who reasons about eliciting the desired effect in the audience and sends directions to a number of virtual participants who's only additional knowledge comes from the director's directions. Another example might be a computer game with virtual participants playing teammates of the human player and virtual actors playing the NPCs who control the player's progress through the story. These hybrid approaches provide the advantage of virtual actors where necessary while retaining the simpler virtual participants where they are sufficient.

The Debate Continues

The comments presented here mention only a few of the various pros and cons of either approach to modeling virtual humans, and they may be other approaches that need to be considered as well. The opinions of the community of developers and researchers that participate in the AAAI Spring Symposium on Artificial Intelligence and Interactive Entertainment will have a significant role in guiding future work in this area, and everyone is encouraged to join the debate.