IEEE Intelligent Systems
www.computer.org/intelligent

Game AI Is Dead. Long Live Game AI!

Alexander Nareyek

Vol. 22, No. 1
January/February 2007

Editor: Alum Preece
University of Aberdeen
apreece@csd.abdn.ac.uk

Game AI Is Dead. Long Live Game AI!

Alexander Nareyek, *National University of Singapore*

Nowadays, many games use basic game AI techniques such as pathfinding, steering, and finite state machines.[1] And, the number of middleware companies that focus on game AI technology is growing. Initiatives such as the International Game Developers Association's Artificial Intelligence Interface Standards Committee, which I set up and chair, help to standardize and improve the applicability of middleware gaming solutions. Meanwhile, most of academia seems trapped tackling problems of limited use,[2] such as increasing some remotely relevant algorithms' performance by a few percentage points or producing yet another agent architecture.

With industry addressing most game AI problems in seclusion and academia examining these problems halfheartedly, what hope is there for a future impact of academic contributions?

For techniques such as pathfinding and finite state machines, I'm afraid the light we see at the end of the tunnel is an oncoming train. Why compete with a specialized middleware industry on some minor efficiency gains? However, the increasing computational power now available to game AI unlocks some great opportunities for academia.

At the National University of Singapore, in our Interactive Intelligence Labs and multidisciplinary Games Lab, we're pushing ahead in several areas.

Automated content generation

With the ongoing increase in realism and high fidelity in virtual game worlds, art- and content-production costs are getting out of hand. For example, you can no longer have one artist paint sprites for an in-game object or create a simple 3D model with a color texture. Today, for the same in-game object, you usually have an artist who paints conceptual pictures; a modeler who designs the corresponding 3D model, rigging it and producing normal maps; an artist who provides various textures (such as diffuse, specular, or transparent); and an animator who designs the 3D model's animation, perhaps using additional expensive resources such as motion capturing. This trend of increasing complexity won't end anytime soon. For example, it continues to expand because of the need to add physics properties to the models.

At the same time, players want ever-increasing free-roaming environments. The number of necessary assets and the additional time for design, along with increasing art and content complexity, creates an enormous cost explosion. It's no wonder that each new generation of games has higher production costs.

Consequently, *procedural* techniques are gaining popularity—that is, techniques that generate art, environments, and other content automatically, reducing the amount of handcrafted art. Figure 1 shows an example of a game that liberally applies procedural technology for its large forest areas.

Large-scale automated art and content generation practically screams for AI technology. Sure, you can generate the content without AI, but AI lets designers and artists do their work more efficiently and increases the resulting quality. To do this, we need to address such questions as

- How do we design exciting vistas, natural-looking terrain, interesting city layouts, and stylish textures and animations?
- How can we generate art and content that fulfills the game's design vision, art direction, and technical requirements?
- Designers and artists will use our tool; how does this affect concepts, algorithms, and interfaces?

Advanced techniques might even generate new content in real time during gameplay.

However, let me insert a word of caution on *unpredictable* outcomes in procedural generation, which some view as a desirable feature (especially those with a machine learning background). From a development perspective, outcomes that the designers didn't foresee can often produce undesired results or break players' suspension of disbelief. The problem is that a game contains all kinds of implicitly represented mechanics and knowledge, such as which elements provide player entertainment, what game objects are for, or what graphical

1541-1672/07/$25.00 © 2007 IEEE
Published by the IEEE Computer Society

Figure 1. A screenshot of the game The Elder Scrolls IV: Oblivion, which applies procedural technology for forest areas.

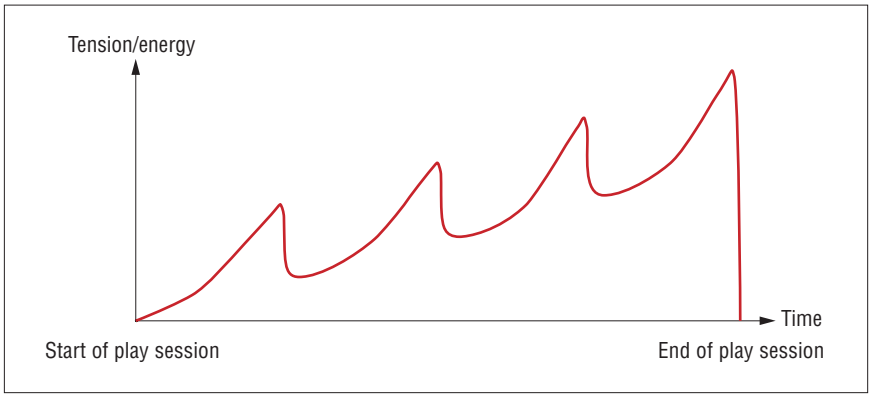

Figure 2. Applying a storytelling rule to gaming: the player's excitement should steadily increase with spikes in excitement every seven to 10 minutes.

elements mean. So, unpredictability could lead to automatically placing a house on a "river" object or developing a very effective, but boring, new strategy for player enemies. Despite marketing slogans to the contrary, hardly any game will thus feature this kind of fully unpredictable AI anytime soon.

Automated storytelling

Depending on the game genre, having a story isn't always important, but for a highly emotional and meaningful experience, stories play a major role.

Most games feature a main storyline, perhaps with some limited branching—for example, the player can choose a good or evil path. This limited number of player options is inevitable because the storyline is handcrafted and thus an expensive component.

If we want to automatically generate a story in real time during gameplay, the previous warning about unpredictable outcomes applies. Like content generation, stories draw on an incredible amount of implicit knowledge that we can't fully model. So, while fully variable storytelling is far off in the future, we don't need to reach the Holy Grail in just one leap.

For storytelling, reactive approaches, like in The Sims, which simply trigger new character actions according to the current game situation, aren't particularly useful because they don't consider the overall player experience. Much more useful is overall story planning, in which we generate a storyline that matches the play experience that we have in mind for the overall play session and update it if the player does something unexpected.

For example, say our AI plans a story in which the player meets a nice princess who is then kidnapped by an evil mage, and the player must rescue her. However, during an initial scene in which the player is supposed to develop some feelings for the princess, he might instead run away with the maid. Our AI must then replan the remaining story, paying attention to an overall cohesive story experience and structure.

Many rules exist for creating a good structure for a storyline in literature and film, and many carry over to interactive games. For example, one rule is that the player's excitement should steadily increase throughout the game session and include "whammies" every seven to 10 minutes, depending on the player's age and demographics (see figure 2). So, if the player runs away with the maid, in replanning the story, the AI must stick to the same tension curve—for example, changing the high-energy point of the princess's kidnapping to a situation where the maid's husband comes after the player.

I won't get into further details, but depending on which variety level you're aiming for, you can add procedural set generation, automated dialogue generation, and much more. Automated storytelling for games is a huge task and will present research challenges for many years to come.

Virtual actors

Of course, there's no story without people (or orcs, or aliens, or whatever). And again, I'm not talking about simple finite-state machines or reactively triggered scripts for realizing characters' behaviors. Nonplayer characters in computer games should ideally be much more than the simple reactive punching bags in most of today's games.

As a simple example, say we want a monster to consider what action to take to get through a door that the player is blocking. In terms of AI, this is a relatively easy mechanical reasoning task. However, it takes much more than such mechanical reasoning to realize nonplayer characters, which is why we call such characters *virtual actors*. They not only need to act in a goal-driven way that makes sense to the user but must also display convincing emotions and personality, fill a specific character role, and actively drive the

story progression. The need for virtual actors also applies to a much wider area than games—for example, for applications that include virtual assistants, teachers, tour guides, or even friends.

Figure 3 shows a concept study for a forthcoming Interactive Intelligence Labs project, showcasing a virtual actor as a pedagogical agent. In this art lesson, the user learns about cubism and some background on Picasso and gets to draw a cubist-style picture using a separate art tool. The virtual actor doubles as tutor and posing model for the picture. She communicates the background knowledge; responds to the user's request on how to pose; comments on issues in his or her drawing style, such as picture composition and color palette; and generally establishes a social binding and motivation for the user. In essence, a kind of story planning is again at work—this time focused on selecting the virtual actor's actions.

Figure 3. A virtual actor in an art lesson.

Some people argue that with the rise of multiplayer environments, there's less need for intelligent nonplayer characters because players can interact with other players. I can see why people jump to this conclusion, because most massively multiplayer online games don't yet feature meaningful stories. However, a good story requires much more than spontaneous or random interactions with other players.

Adapting to the player

Unlike noninteractive forms of entertainment such as films, we don't need to optimize toward an abstract demographic audience; instead, we can learn more about our specific player. The more we know about the player's preferences and mental state, the more we can provide a stronger emotional impact and more meaningful experiences.

To realize a game that adapts to the player, we need to answer questions such as these:

- What constitutes meaningful and emotional experiences in a game?
- How can we automatically measure and assess the player's relevant mental and emotional state during the game?
- How can we manipulate the player to induce specific beliefs, desires, and intentions and lead him or her to specific emotional states?

Much research needs to be done in cooperation with social psychologists and neurologists. In fact, the progression of game AI as a whole requires a multidisciplinary approach, including participation from the social sciences and the arts as well as technology. This "new" game AI isn't about AI researchers working alone in secluded research labs.

The research areas I've described here have little to do with the "old" view of game AI. There's nothing bad about finite-state machines, pathfinding, and so on. Indeed, our higher-level AI will use these techniques at a lower level. But from an academic perspective, we shouldn't focus on small quantitative improvements, but rather on areas that prepare qualitative jumps. I've described some exciting areas, but opportunities for other interesting and relevant research abound. An increasing number of researchers are entering the game AI research arena, and I'm confident that we're in for some great advances in the future.

Because electronic gaming is slowly turning into a mainstream market, the term "game AI" will probably disappear from public perception. Marketing will most certainly not label their products as having AI, to avoid the technical and nonhuman image and to better position their products as "alive and emotional" instead. Whatever the public perception will be, it's AI that powers it, so I see a bright future for this technology.

In this article, I've hardly touched the topic of the underlying AI technology for the described challenges. Our approaches at the Interactive Intelligence Labs and the Games Lab are mainly focused on goal-directed action planning. A primary research focus for us is adapting this technology to the requirements of dynamic real-time environments, such as virtual gaming worlds. In this respect, we're following the direction of the EXCALIBUR project (www.ai-center.com/projects/excalibur/publications.html) and taking it to the next level.

Game AI is a great field of research, and I can hardly imagine a more exciting and rewarding research area. If you share my passion, drop me an email; we're staffing up strongly! ⬒

Related Web Sites

Interactive Intelligence Labs:
www.ai-center.com/ii-labs
AI SIG page of the International Game Developers Association:
www.igda.org/ai
AI SIG Game AI Newsletter:
http://five.pairlist.net/mailman/listinfo/gameai_news

References

1. A. Nareyek, "Artificial Intelligence in Computer Games—State of the Art and Future Directions," *ACM Queue*, vol. 10, 2004, pp. 58–65; www.ai-center.com/references/nareyek-04-gameai.html.

2. A. Nareyek, "Computer Games—Boon or Bane for AI Research?" *Künstliche Intelligenz*, vol. 18, no. 1, 2004, pp. 43–44; www.ai-center.com/references/nareyek-04-gameairesearch.html.

Alexander Nareyek is director of the Interactive Intelligence Labs at the National University of Singapore, and he heads the university's multidisciplinary Games Lab. Contact him at alex@ai-center.com.