



In the

# NEWS

ALSO FEATURED

THE YEAR THE POWER GRID  
GOT SMARTER

Features Editor: Dennis Taylor  
dtaylor@computer.org

## Toward More Human Video Game Characters

Mark Ingebretsen

**T**hey wield deadly swords, trash planets, and fly spacecraft at the speed of light. But the animated video game bots that engage players for hours really aren't that smart, at least when it comes to their AI capabilities. The behaviors of video games' nonplayer characters (NPCs) largely have been scripted by developers using rule-based systems that leave little room for autonomous decision making.

"A lot of the early work in game AI has been so-called path planning, getting a character to go from point A to point B," says John Laird, professor of engineering at the University of Michigan, who has designed games for research purposes and for the military.

In the past, game consoles limited AI's role. "AI has always suffered at the expense of graphics and game physics," says Penny Baillie-de Byl, associate professor of serious games at the Breda University of Applied Sciences, who wrote *Programming Believable Characters for Computer Games* (Charles River Media, 2004). "About 10 percent of processing was available for NPC AI processing," she says.

Even though current-generation consoles are fast enough to permit far more robust AI behaviors, Laird and other academics who are looking to boost AI's presence in video games say it will be a while before the commercial gaming community taps the full potential. That's because AI-enabled NPCs

that can learn and alter their strategies run the risk of being too adept for their human counterparts. "The objective of successful games is to ensure the player progresses and that nothing blocks their progress to a point where they just give up," says Baillie-de Byl. To keep human players engaged, developers have even been forced to dumb down NPCs so that they wouldn't overwhelm their human opponents.

As a result, serious AI game character development has been relegated to academia. Free from the market demands game companies face, university researchers are creating characters able to communicate with their digital colleagues and even exhibit something approaching true emotions. At the same time, two Massachusetts Institute of Technology (MIT) researchers are working on NPCs that incorporate enough knowledge about particular situations that their ability to respond intelligently to player input approaches that of humans.

### Dinner conversations

"In the world of machine learning, there is no such thing as too much data," says Deb Roy, an associate professor at MIT and director of the Cognitive Machines Group at the school's Media Lab. Roy has gone to great lengths to follow up on that belief. The March 2008 *Scientific American* chronicled his and his wife's efforts to record every sound and movement their infant child made—all part of an experiment aimed at collecting data that might help enable robots to talk.

Working with Jeff Orkin, his doctoral student, Roy has helped fashion an equally ambitious video-game-like experiment called the Restaurant Game (<http://therestaurantgame.net>). The game puts players in a restaurant setting, where they can adopt one of two familiar roles, each with a simple objective. As Orkin explains: "You're a waitress, try to earn some money. You're a customer, try to have dinner." Orkin's prior experience was as a senior engineer with Monolith Productions, developing AI for the games F.E.A.R. (First Encounter Assault Recon) and No One Lives Forever 2.

The experiment aims to build a data set incorporating as many variations of the dialogue and physical behavior involved in the server-patron interaction as possible. When finished, it will be similar to *n*-gram models used in natural language processing, which denote the conditional probabilities of the many ways individual words can be sequenced in a language. Instead of only capturing words, Roy and Orkin record *interaction traces*, which track the series of linked behaviors and utterances extracted

## How to Reach Us

### Writers

For detailed information on submitting articles, write for our Editorial Guidelines ([isystems@computer.org](mailto:isystems@computer.org)) or access [www.computer.org/intelligent/author.htm](http://www.computer.org/intelligent/author.htm).

### Letters to the Editor

Send letters to

Dennis Taylor, Lead Editor  
*IEEE Intelligent Systems*  
10662 Los Vaqueros Circle  
Los Alamitos, CA 90720  
[dtaylor@computer.org](mailto:dtaylor@computer.org)

Please provide an email address or daytime phone number with your letter.

### On the Web

Access [www.computer.org/intelligent](http://www.computer.org/intelligent) for information about *IEEE Intelligent Systems*.

### Subscription Change of Address

Send change-of-address requests for magazine subscriptions to [address.change@ieee.org](mailto:address.change@ieee.org). Be sure to specify *IEEE Intelligent Systems*.

### Membership Change of Address

Send change-of-address requests for the membership directory to [directory.updates@computer.org](mailto:directory.updates@computer.org).

### Missing or Damaged Copies

If you are missing an issue or you received a damaged copy, contact [membership@computer.org](mailto:membership@computer.org).

### Reprints of Articles

For price information or to order reprints, email [isystems@computer.org](mailto:isystems@computer.org) or fax +1 714 821 4010.

### Reprint Permission

To obtain permission to reprint an article, contact William Hagen, IEEE Copyrights and Trademarks Manager, at [copyrights@ieee.org](mailto:copyrights@ieee.org).

from each player session. For example, a trace might capture a customer leaving a big tip for an especially good steak that the waitress delivers quickly. Traces are combined to form a statistical model that can be thought of as a generalized script of typical restaurant behavior and dialogue.

Once thousands of these traces have been recorded, an unsupervised, bottom-up learning system generates a statistical model. NPCs can access this model to guide their behavior in new interactions. For example, suppose a customer expresses satisfaction over a well-prepared steak. The AI-endowed waitress NPC that's programmed to maximize her income can sift through thousands of previous scripts and determine there's a high probability she'll receive a big tip. She can also note what follow-up behaviors will most likely ensure she'll get that tip.

Roy and Orkin chose a game format to collect their data to exploit the opportunity provided by online games to collect a large amount of data in a small amount of time. And in fact, this simple game has been played over 8,000 times by over 10,000 players. "People do all sorts of things because we intentionally gave them an open-ended scenario," explains Orkin. A game similar to the restaurant scenario might teach NPCs in commercial video games how to successfully fight human players or guide bots that help healthcare workers diagnose a patient's ailment.

Roy and Orkin say their approach, which builds on '70s-era research by AI pioneer Roger Schank, provides a methodology for generating more robust scripts than previous approaches that tried to anticipate and hand-code every behavior variation in a scenario. "There have been repeated attempts using different contexts to try to create intelligent characters by essentially encoding common sense or all this background knowledge," says Roy. "It turned out to be impossible for human programmers to anticipate all the rich variation of behavior."

### Thanks for the memory

Another approach that shows promise as a training tool for game characters has been around for some time. John Laird, Allen Newell, and Paul Rosenbloom developed the Soar open-source cognitive architecture in the '80s while they were at Carnegie Mellon University.

Before applications such as Soar emerged, "agents were created using finite-state ma-

chines that incorporated knowledge through rules," Laird explains. He says that Soar simplifies the task by making it easier for developers to insert the varied pieces of knowledge required to direct a character. In addition, it "allows the interactions to arise naturally as the system is executing," he says.

For example, suppose an NPC hiding in a room is tasked with shooting a human player's avatar when it enters the room. To make the NPC's job easier, it might be fed knowledge about the room's dimensions, its entry and exit points, and the like. So, when the NPC "saw a player enter a room but did not have a good shot," Laird explains, it would wait until that avatar moved to a more advantageous position before firing.

Soar also lets NPCs store information they've detected during game play and use this memory to guide their actions. For example, if an avatar enters a room and disappears behind some furniture, Soar's memory would enable the NPC opponent to remain ready to fire when the avatar reemerged. Alternatively, the NPC could communicate with another NPC, directing it to attack the avatar. In contrast, if the NPC was acting from a list of rules, it might disengage the moment the avatar ducked behind the furniture.

### I feel therefore I am

With game research funding in short supply, Laird says his current work with Soar has focused on its application to robotics—particularly for the military, where robots with NPC-like qualities could autonomously locate enemy troops. Laird's other research centers on giving Soar the ability to program emotions into NPCs, an interest he shares with Baillie-de Byl.

Whereas some see emotions as distractions from logical thought, Baillie-de Byl and her colleagues conjecture that they act as a filter to guide decision making. Synthesizing emotions in NPCs not only lets researchers better understand how they affect human behavior but also can make the game characters much more realistic.

To prove her point, Baillie-de Byl trained a neural network using a data set taken from human subjects. The data set was originally compiled by behavioral researchers Craig Smith and Phoebe Ellsworth, who developed a gridlike representation of how people experience feelings. The grid mapped 15 human emotions, ranging from guilt to

## IEEE Computer Society Publications Office

10662 Los Vaqueros Circle, PO Box 3014  
Los Alamitos, CA 90720-1314

### Staff

Lead Editor

**Dennis Taylor**

dtaylor@computer.org

Senior Editorial Services Manager

**Crystal R. Shif**

Magazine Editorial Manager

**Steve Woods**

Senior Editors

**Dale Strok and Linda World**

Assoc. Peer Review Manager

**Hilda Carman**

Publications Coordinator

**Alkenia Winston**

Contributing Editors

**Molly Gamborg, Rita Scanlan,**

**and Joan Taylor**

Production Editor

**Jennie Zhu**

Technical Illustrations

**Alex Torres**

Director Products & Services

**Evan Butterfield**

Digital Library Marketing Manager

**Georgann Carter**

Senior Business Development Manager

**Sandra Brown**

Senior Advertising Coordinator

**Marian Anderson**

**Submissions:** For detailed instructions and formatting, see the author guidelines at [www.computer.org/intelligent/author.htm](http://www.computer.org/intelligent/author.htm) or log onto *IEEE Intelligent Systems'* author center at Manuscript Central ([www.computer.org/mc/intelligent/author.htm](http://www.computer.org/mc/intelligent/author.htm)). Visit [www.computer.org/intelligent](http://www.computer.org/intelligent) for editorial guidelines.

**Editorial:** Unless otherwise stated, bylined articles as well as products and services reflect the author's or firm's opinion; inclusion does not necessarily constitute endorsement by the IEEE Computer Society or the IEEE.

boredom to fear, and how given situations affected them. The list of situations encompassed such things as a person's perceived degree of control or the required degree of effort. "For example, happiness is an emotion generated in a high-control, moderate-effort domain," Baillie-de Byl explains. She adds, "The neural network training resulted in a nonplayer character which was evaluated as being emotionally intelligent by humans who participated with it during experimental trials."

**A**dd the efforts of Baillie-de Byl, Laird, and Roy and Orkin to those of other researchers, and it's easy to imagine NPCs that can react with human-like adaptability to whatever situation they're thrust into. But such characters are likely still far off. Even if they existed, they still might not be the totally engaging opponents avid gamers dream about. The NPCs' behaviors do become predictable over time," says de Byl, "but so then does that of real human beings."

## The Year the Power Grid Got Smarter

Mark Ingebretsen

**W**hen they write the history of the intelligent electric-power grid, 2008 might be remembered as the tipping point. In May, Xcel Energy announced it had begun the test phase of its SmartGridCity project in Boulder, Colorado. As part of the project, the first of its kind in the US, Xcel will install more than 15,000 "smart meters" in residential

and commercial locations. The meters will provide Xcel with real-time information about energy demand, helping it optimize delivery and prevent service interruptions.

Additionally, customers will be able to monitor their energy use via a Web site. They'll also be able to choose whether they want to receive power from the region's wind and solar farms or purchase it from whatever source is cheapest. Eventually, SmartGridCity will equip homes with "smart" appliances that can be regulated by the grid itself to ease power consumption. Plug-in hybrid cars, when they become available, might also be enlisted as backup power suppliers when outages occur.

An article in the May 2008 *Forbes* called SmartGridCity "'Tomorrowland' for electricity." The article also said a plethora of "start-ups, big utilities and government have unleashed a wave of investment to make the electric grid smarter and more efficient," while cashing in on the "enormous" profits awaiting those able to reinvent how households receive power.

Indeed, the Boulder test comes at a time when America's power grid—like many throughout the world—is being stretched in two directions. On the one hand is the

rising demand from plasma TVs, computers, and an ever-growing number of other power-hungry household gadgets. On the other hand are the growing numbers of wind and solar farms. State laws require utilities to purchase power from these small-scale suppliers. But because wind and solar farms function at the whim of the weather, utilities have an increasingly hard time forecasting how much power must be brought online at any given moment.

Utilities got a taste of the problems intermittent resources can cause in Texas this spring, says Ralph Masiello, chief technology officer with KEMA, a global energy consulting firm involved with GridWise. "Wind was forecast to die off," says Masiello. "But it died off an hour or two before the forecast said, and the grid was barely ready." Although Masiello and other analysts say wind and solar assets are critical to meeting future energy demand, the question remains, he says, "When we get more of this stuff, what's it going to be like?"

Indeed, as the grid becomes increasingly complex, the once-a-month meter readings used to monitor it today will likely seem quaint compared to what's in store, says Patrick Mazza, research director with Cli-

mate Solutions, a regional environmental think tank in Olympia, Washington. He and other energy experts look to intelligent agents and their ability to adapt in real time to volatile conditions as perhaps the best means of controlling the grid. "Today's smart meters and advanced metering infrastructure are leading there by providing power user information at near or real time," he says.

Intelligent agents would join other forms of AI, specifically neural networks, that utilities have long been using. Trained using years of weather and energy-use data, neural nets routinely forecast demand a day or even hours in advance. AI's other long-established role in monitoring the grid consists of intelligent sensors that provide advance warning of malfunctions on everything from underground cables to large-scale transformers.

But smart agents will play a much larger role. As a first step, Masiello envisions a university campus grid laden with agents that continually communicate with one another. "The system would comprise neural nets and some kind of optimization program," he says.

In this scenario, the central control unit might continually update agents located throughout the campus. Those agents would then adjust the power usage of appliances, heating and cooling equipment, and the like, and notify the central unit about their actions. "Then the central operator might adapt and republish [information on its status] until it sees what it wants," Masiello explains.

Full-blown visions of the smart grid extend this design to residences and generating facilities. As Mazza describes it, in a decade or two, the entire grid will be populated with autonomously operating sensors and microprocessors. Agents in homes, for example, will negotiate with each other to divvy up the limited power available during brownouts, while giving grid operators updates on their actions. Meanwhile, agents at solar and wind farms will sell power to the highest bidder while updating everyone in real time on their own power output and perhaps even the amount of carbon they're offsetting. The result, says Mazza, will be "a collaborative network of almost biological complexity."

Ambitious as that sounds, Australia's Commonwealth Scientific and Industrial Research Organization (CSIRO) will test elements of just such a network later this year. Underlying that network, the Australian project will make "extensive use of machine learning and multiagent-systems techniques," explains Glenn Platt, a CSIRO engineer heading the development of the network's smart agents ([www.csiro.au/science/SmartAgents.html](http://www.csiro.au/science/SmartAgents.html)).

As part of the network, agents in homes will receive input from residents via a touch screen. Commercial customers will interact with the network via a more sophisticated interface.

The grid design also calls for agent control of wind and solar facilities that will use machine learning to predict their output on the basis of demand and weather data. These agents might also order that energy be stored in or retrieved from batteries as conditions warrant.

**D**espite their autonomy, the agents in the Australian project, like those planned for similar projects elsewhere, will solicit far more customer input than is the case today. Responsible energy-use decisions by consumers will be critical to managing growing demand. Yet consumers will always have differing, legitimate needs for power. "This is one of the main advantages of our technical approach," says Platt. "The use of multiagent and artificial intelligence techniques means that our system can be flexible enough to accommodate the needs of the end user but also provide the reliability the wider electrical network needs." ■



## RUNNING IN CIRCLES LOOKING FOR A GREAT COMPUTER JOB OR HIRE?

The IEEE Computer Society Career Center is the best niche employment source for computer science and engineering jobs, with hundreds of jobs viewed by thousands of the finest scientists each month in *Computer* magazine AND online!

- > Software Engineer
- > Member of Technical Staff
- > Computer Scientist
- > Dean/Professor/Instructor
- > Postdoctoral Researcher
- > Design Engineer
- > Consultant
- > Hardware Engineer

IEEE  
computer society

<http://careers.computer.org>

The IEEE Computer Society Career Center is part of the *Physics Today* Career Network, a niche job board network for the physical sciences and engineering disciplines. Jobs and resumes are shared with four partner job boards - *Physics Today* Jobs and the American Association of Physics Teachers (AAPT), American Physical Society (APS), and AVS: Science and Technology of Materials, Interfaces, and Processing Career Centers.