IEEE Intelligent Systems
www.computer.org/intelligent

In the News

Mark Ingebretsen, Maya Dollarhide

Vol. 22, No. 4
July/August 2007

In the News

ALSO FEATURED

The Rise of Self-Aware Robots

Features Editor: **Dennis Taylor**
dtaylor@computer.org

Intelligent Tutors Make the Grade

Mark Ingebretsen

The field of intelligent tutoring received a boost this May, when the US Department of Education awarded $2 million to Massachusetts' Worcester Polytechnic Institute and Carnegie Mellon University. The four-year grant will help researchers at the two schools refine Assistment (www.assistment.org), an intelligent tutor for middle-school mathematics.

Like other intelligent tutoring systems, Assistment tailors its instruction to individual students, using embedded AI functionality to provide feedback and assistance—all without human overseers. However, the federal government was perhaps equally interested in the software's other function, according to Assistment's developer, Neil Heffernan. "It was designed to give schools an accurate evaluation tool" of their students' math abilities, explains Heffernan (www.cs.wpi.edu/~nth), an assistant professor in WPI's Department of Computer Science.

Indeed, studies by Heffernan and his colleagues reveal that aggregate data provided by Assistment can predict how well students will perform on the state basic math skills tests mandated by the US No Child Left Behind Act. That predictive ability could make Assistment a godsend for schools struggling to comply with the 2001 federal law, which requires that students in school districts achieve basic competency levels. Assistment's potential importance harks to the seemingly growing value of creating programs that teach while providing human instructors with valuable feedback.

Three established intelligent tutors—Assistment; Tactical Iraqi, a military language-learning program; and Maria, an intelligent avatar—exemplify this combination of instruction and data capture. However, these three programs take differing approaches to interface design, using bare-bones text display, an immersive video game, or an interactive talking head.

Keeping it simple

"Interface design is not that important to results," Heffernan says, at least when it comes to math education. Assistment's simple design minimizes bandwidth, a major plus because the program resides on Internet servers at WPI. The program's on-demand (that is, server-side) model relieves schools of the hassle of installing it. But it also facilitates high-level data collecting by Heffernan and his colleagues.

To make Assistment work on the Web, Heffernan devised XTA (Extensible Tutor Architecture), an offshoot of XML. He and his colleagues also created a series of tags that track Assistment's 98 different eighth-grade math skills. The roughly 1,000 math questions in the program each receive a low-level tag such as "linear equations" or "Pythagorean theorem." So, school systems can better see areas requiring improvement, and teachers can focus on individual student problems.

According to Heffernan, Assistment uses Bayesian networks to track which of the 98 skills students have mastered. He says that Bayesian networks are the ideal method, because you can tag a single question with more than one of the 98 skills, and the networks let the computer guess which skills to address when a student makes an error. Heffernan says that by using these networks, Assistment can make statistically significant predictions of students' state exam scores.

While groups of students use the program in a school's computer lab, teachers monitor which students request help the most, a sure sign they're struggling.

So far, over 4,000 students near WPI and Carnegie Mellon have tried Assistment during the 2006–07 school year. Paul King heads the math department at Worcester's Forest Grove Middle School, where roughly 950 students use Assistment twice a week. He comments that teachers there generally give the program high marks for its ability to reveal where one-on-one instruction can do the most good. Students already proficient at math are helped the most by Assistment, he says, adding that less adept performers tend to "game" the system. That is, in lieu of working to solve the math problems themselves, "they ask the tutor a lot of questions," he says.

King says he'd like future versions of Assistment to make data on student performance more user friendly by using detailed graphs and the like. Heffernan is working on

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

an online grade book to go with the program and on a system that will automatically alert parents by phone or email if students don't complete their homework. In fact, homework could become intelligent tutoring's killer app, Heffernan believes. Students would receive AI-enabled help in the evenings, while teachers could make good use of the resulting extra classroom time.

Life-and-death learning

Math instruction seems easy when compared to learning another language, especially when that language is nuanced with colloquialisms and spoken in the emotionally charged atmosphere of a war zone. Tactical Iraqi meets that challenge by placing learners in a video game-like environment. The program has taught an estimated 10,000 American military personnel how to speak Iraqi Arabic and observe important local customs.

Game modules, which take 100 hours in all to complete, duplicate the situations soldiers in Iraq face, from house-to-house searches to learning the location of belligerents. To succeed, players must interact with locals, who are intelligent agents, impressing them with language and cultural skills. The better the players perform, the more cooperation they receive. What's more, the game modules are linked so that errors incurred in early modules continue to haunt players as they progress.

Such realistic simulations form an important component of military training, says Tactical Iraqi's lead developer, Lewis Johnson (www.tacticallanguage.com), a former linguistics professor at the University of Southern California. DARPA began funding a language-training simulation at USC's Information Sciences Institute in 2003, the year the Iraq war began.

That project, which eventually led to Tactical Iraqi, had its share of challenges—in particular, "how to employ speech technology to work robustly with learners' speech as well as to detect common errors in learners' speech," Johnson explains.

Johnson and his cohorts, who would eventually form the Tactical Language Training company, built their solution around Julius (http://julius.sourceforge.jp/en_index.php?q=en/index.html), an open-source speech-recognition toolkit. They brought in Iraqi cultural experts to help create the game's realistic scenarios. These scenarios are key to helping Julius recognize and respond to novice learners' frequent mispronunciations. As Johnson explains it, "We do a lot to use the context or current focus to restrict the size of the speech recognition grammar dynamically to what we're trying to do at a particular point in time. That helps with accuracy."

Accuracy should improve over time as well. Johnson and his 15-employee company collect data from students and use it to continually retrain the speech recognizer, he says.

To build a realistic Iraqi world, Johnson licensed the Unreal game development engine from Epic Games, the same engine used to create popular video game titles such as Shadow Ops and Harry Potter (www.unrealtechnology.com). Johnson then employed Bayesian statistical techniques to track students' progress in real time. When students working through the game modules demonstrate cultural awareness—say, by asking to speak to the head of the household during a weapons search—the program immediately updates their score. Students and instructors receive detailed progress reports at each module's conclusion.

Tactical Iraqi fills a particular need because it substitutes for Iraqi language instructors, who are in short supply. To Johnson, that meshes perfectly with what intelligent tutoring is supposed to do. "Part of it is to provide trainees with feedback, but also to inform instructors how students are doing. So, instructors can apply their time more effectively," he says.

Talking heads

Surveying student progress reports is helpful. But what if teachers could actually converse with an AI tutor—say, an avatar—in the same way students do? The teacher could learn about the student's progress and then supply the avatar with new information or instructions. In turn, the avatar could pass that information on to the student. And all the while, the avatar's knowledge of the student, the learning process, and the subject matter would increase.

That's what avatars produced by Shahin Maghsoud and Robot-Hosting (www.robot-hosting.com), his New Zealand- and British Columbia-based company, promise to do—and in a friendly, human-like manner. "[Our] robots show emotions that are a combination of several basic emotions," he explains. He adds that the avatars can generate 24 human emotions, including happiness, sadness, and anger. The avatars' intelligence and emotions are based on such approaches as case-based reasoning, propositional-logic reasoning,

IEEE Intelligent Systems

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

Group Managing Editor
Crystal R. Shif
cshif@computer.org

Senior Editors
Shani Murray, Dale Strok, and Linda World

Assistant Editor
Brooke Miner

Editorial Assistant
Molly Gamborg

Publications Coordinator
Hilda Carman

Contributing Editors
Keri Schreiner and Joan Taylor

Production Editor
Jennie Zhu

Technical Illustrations
Alex Torres

Associate Publisher
Dick Price

Membership/Circulation Marketing Manager
Georgann Carter

Business Development Manager
Sandra Brown

Senior Production Coordinator
Marian Anderson

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

IEEE Intelligent Systems

How to Reach Us

Writers

For detailed information on submitting articles, write for our Editorial Guidelines (isystems@computer.org) or access 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

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

On the Web

Access 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. 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.

Missing or Damaged Copies

If you are missing an issue or you received a damaged copy, contact membership@computer.org.

Reprints of Articles

For price information or to order reprints, email 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.

commonsense reasoning, and epistemic-logic reasoning.

Maria, one of Robot-Hosting's avatars, holds the title of assistant lecturer at both the University of Auckland's Information Systems Department and the University of Arizona's electrical engineering department. Others provide customer service support for companies or serve as online entertainers. Peter Burggraaff, former IT Manager at Farmers Trading Co., a New Zealand retail chain that used Robot-Hosting's avatars for customer support, noted that the "key benefit the customers identified was fast and factual information" when they didn't want to speak with a live person because it might take too long.

Regardless of their mission, the cartoon-like Web-based avatars possess a 203,000-word vocabulary and 106,000 logical-reasoning rules. Each comes with an empty database that can be continually filled with information on a subject.

So, in essence, Maria serves as an intelligent querying agent that uses case-based reasoning and other AI techniques to fulfill data requests where no exact matches exist. But what separates Maria from run-of-the-mill search engines is her ability to gather information on the individual students and to detect and react to each student's emotions in real time.

Depending on the avatar's complexity, clients might pay anywhere from US$30,000 to $1.5 million. But, Maghsoud says, the AI characters pay for themselves because they're especially suited to handle repetitive student questions, thus saving teachers' time. Robot-Hosting plans to enhance student-avatar rapport by using neural networks to augment the avatar's reasoning abilities. "It is in our R&D plans to connect a camera and grab the emotional state of the user by seeing him and respond accordingly," Maghsoud says. "Of course, when we need to do more complex real-time processes, a faster CPU will help."

That extra processing power, once it becomes readily available, might lead to a convergence of the strengths of intelligent tutors such as Assistment, Tactical Iraqi, and Robot-Hosting's avatars. Thus, a future avatar might be able to simultaneously interact with thousands of users, providing its overseers with detailed micro- and macrodata on its global classroom's performance. Meanwhile, students or teachers might someday be able to place their enhanced avatars within increasingly detailed video game-like environments that might themselves continually adapt on the basis of the feedback the system receives.

The Rise of Self-Aware Robots

Maya Dollarhide

What has four legs, is shaped like a starfish, and practices at self-awareness? Ask Hod Lipson, assistant professor of mechanical and aerospace engineering at Cornell University; Josh Bongard, a former Cornell postdoctoral researcher now on the faculty at the University of Vermont; and Cornell graduate student Viktor Zykov. They have created a robot that learns about itself and uses that knowledge to operate and self-correct when injured.

Know thy self-model

The AI technology behind the Starfish robot is a computer algorithm that continuously creates, modifies, and discards self-models, explains Bongard. "For example, the robot starts with random self-models: it doesn't know that it has four legs, so the self-models may be virtual snakes, eight-legged creatures, or a creature with randomly attached legs. One after the other, the algorithm sends the same motor commands to the self-model as it did to the physical robot," he says.

The technology causes the self-model to generate virtual sensor data, which the algorithm then compares against the sensor data from the physical robot. The closer the match is, the more accurate the self-model must be. The algorithm discards models with low accuracy and makes modified copies of the better models. “By continuing this process, the algorithm eventually converges on a set of very accurate self-models: in effect, the robot ‘discovers’ that it has four legs, or, after damage, that it now has only three,” says Bongard.

The agony and thrill of learning

The project began in 2004, when Bongard and Lipson discussed the possibility of creating a robot prototype for NASA that could recover from damage on its own. Bongard and his team at the Cornell Computational Synthesis Lab have had many challenges along the way. Interestingly, the starfish shape wasn’t one of them. Bongard and Lipson both say the shape was accidental. The prototype was originally going to have six legs, but partway through the project the team realized they could demonstrate self-modeling on a simpler, four-legged robot. The real problems began when the robot was complete, the technology installed, and the simulation process up and running.

“Noise (especially biased noise) made training the simulator more difficult. We also found that certain chaotic aspects of the robot dynamics (for example, bifurcations due to symmetry) caused inherent unpredictability that baffled the internal simulator,” says Lipson. But they also experienced the joy of seeing their creation finally produce accurate self-models. “We watched it as it internally rehearsed moving patterns,” says Bongard. “Once it had a moving pattern that it was confident with, it tried it out in reality. Luckily, the physical robot moved very similarly to the way the self-model predicted it would. This last part only took 10 seconds, and it was a very thrilling moment. It was the first steps, if you like, of a new kind of robotics.”

Lipson says that the most interesting success occurred when the team removed one of the robot’s legs. “We did this essentially to test the algorithm, to see if the internal models would also lose a ‘virtual leg’ in response to our removal of the physical leg,” says Lipson. “Indeed this happened, and the robot began to limp, again without doing any locomotion trials. The gait it chose was rather surprising.” Instead of standing up, walking a distance, and then resting (like other legged animals), it “uses its back leg to throw its body weight forward to move,” says Bongard.

Starfish and Superbot

Although unique in many ways, the Starfish is kin to other robots with awareness capabilities. Wei-Min Shen, director of the University of Southern California’s Polymorphic Robotics Laboratory, says the Starfish is similar to robots created at USC. “We have been working on a similar project for a long time here at the USC Intelligent Systems Institute,” says Shen. “In 2002–2003, one of our robots called CONRO [*Con*figurable *Ro*bots] could already endure and be aware of the loss of its limbs and adapt its behaviors to compensate for the change of its own topology. Its current successor SuperBot [www.isi.edu/robots/superbot.htm] can also dynamically discover its own topology and detect changes when they happen.”

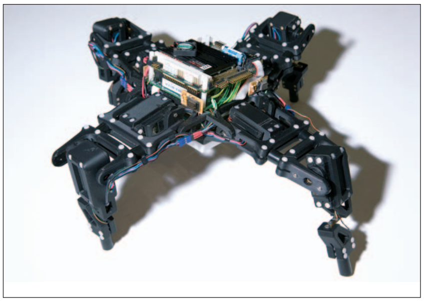

The Starfish robot. (photo courtesy of Cornell University.)

Shen says that the main difference between the Starfish and the USC robots is that the Starfish’s parts don’t communicate with each other, while the CONRO and SuperBot modules do. “So the discovery of a robot’s own (changing) topology is accomplished differently: Starfish uses the sensor feedback, while CONRO and SuperBot use communication between modules; the body of the Starfish contains its brain,” says Shen.

Peter Will, Shen’s colleague at USC, has been following Lipson’s progress. Will points out that Lipson’s inspiration might have been Karl Sims at the Massachusetts Institute of Technology. Sims was the first to look at how arbitrary collections of polyhedral modules might locomote.

“Sims used genetic algorithms in simulation to explore movement options and then computed the best motion according to some metric for a given configuration,” says Will. “Some of the motions were bizarre but still very effective ... such as standing fully erect and falling full-length forward and then doing it again. The important contribution from Lipson is that he and his group have done essentially the Karl Sims exploration or its equivalent in what appears to be real time.”

Curiouser and curiouser

Lipson and Bongard say their robot is a proof of concept and has a long way to go before becoming a part of field robotics, but there are many possibilities. “One is to use more elementary representations and see if a robot is able to learn not only its shape but also other aspects of the world like Newtonian physics,” says Lipson. “These would be even more basic and would allow it to predict the outcome of a broader range of actions.”

Both scientists are interested in using robot teams to cooperatively create a joint model of themselves and the world by sharing experiences. “The idea is that the multiple robots with the same body plan would trade self-models, integrating the most accurate self-models into its own-self models,” says Bongard. “The hope is that if this is done correctly, a second robot may recover from damage more quickly than the first one, if it can draw on the experience of that robot.”

And perhaps the robots will one day teach us more about our own human nature. “This project suggests something about the nature of curiosity,” says Bongard, who has coauthored a book on the mind-body connection called *How the Body Shapes the Way We Think* (MIT Press, 2006).

“In order to build accurate self-models as quickly as possible, the robot does not simply move randomly but tries to move in a lot of different ways and to extract more information from the real world. Again, this suggests why humans are so curious: there may be a link between curiosity, self-awareness, and intelligence,” Bongard says. “I think robots will start appearing that can work together to share their experiences and learn how to act in an increasingly wide range of unanticipated situations, such as learning to use a tool it has never seen before, generating innovative solutions to complex problems, and working alongside humans.” ■