IN THE NEWS

Features Editor: **Lee Garber**, lgarber@computer.org

Programmable AI Applications: Putting Users in Control

Mark Ingebretsen

Search engines help us navigate through vast seas of online information, but they give users few options for managing how that information is obtained. Lars Hard, founder and CEO of the Swedish firm ExpertMaker, wants to change that with something he calls *microsearch engines*.

As Hard explained it, microsearch engines are topic focused so they can deliver results with more precision. Rather than relying on encyclopedic Web indexing, as traditional search engines do, microsearch engines use reasoning techniques such as probabilistic inference, deterministic inference, optimization, estimation, and configuration to fine-tune their results.

However, what's really unique about ExpertMaker—it's currently in beta testing—is that users themselves can program it to their needs. Hard said a large part of the aim with ExpertMaker is to provide an easy-to-use tool that gives untrained users access to powerful AI technologies. "Our goal is that [microsearch engines] should be like a paint program," Hard said. "But instead of shapes and colors, we paint in knowledge with structured data and relations."

That's a goal shared by other experimental AI-enhanced applications. Researchers in Australia are developing an intelligent grid system that lets homeowners decide how to optimize their energy use. User input might also help email users separate junk from important communications. Videogame nonplayer characters might even reach a point where they require human players to collaborate with them to achieve the game's objectives. All these apps offer a hint at what the ultimate user interface could look like someday.

Beyond the Algorithm

Utilizing AI as a tool specifically to facilitate human-machine cooperation is distinct from other specialties within the field. Rather than focus on traditionally difficult goals such as object or speech recognition or complex robotic movements, the collaborative paradigm centers on the idea that humans and intelligent machines work best when each leverages the other's strengths.

The paradigm could be seen as a logical step in the evolution of machine intelligence. A few years ago, processing power vastly amped up the performance of everything from toys to sports car transmissions. That power let users program their devices to perform sequences of actions impossible to do mechanically. The next step up the evolutionary ladder was adding intelligence to devices. So rather than simply being programmed, the devices seek to learn and adapt to user behavior. TiVo's ability to automatically record shows based on users' past preferences and Amazon's ability to do the same with products are the early stage results of this evolution.

Overall, this approach has led to a plethora of useful algorithms, although users rarely interact with

1541-1672/10/$26.00 © 2010 IEEE
Published by the IEEE Computer Society

them directly. And while these algorithms have given applications unprecedented functionality, some like Hard say simply creating more algorithms shouldn't be the end-all goal. "It is often more important to remove the barriers stopping more people from accessing AI technology than trying to develop new algorithms," he said. "Imagine allowing ordinary people access to case-based reasoning, tree induction, image recognition, neural networks built by genetic algorithms, and so on. Then we can start talking about change."

A Stable Analogy

The collaborative paradigm attempts to do just that, creating a new human-machine relationship in the process. In the white paper "NASA Personal Air Transportation Technologies," NASA researcher Mark D. Moore described aircraft in which AI handles much of the flying, using a fitting analogy to describe this new human-machine collaboration: A personal air transportation vehicle's design "begins to simulate an intelligent transportation device that we have depended on for hundreds of years," he noted—namely, a horse.

Like intelligent devices, horses possess enough smarts to allow humans to train them to suit their needs. When the training is complete and the horse and its rider are on the trail, the horse's intelligence helps it sense the environment and navigate to avoid hitting obstacles, including other horses and their human riders. Horses can also communicate how they feel to their riders.

Riders, meanwhile, can control a horse's movements with relatively simple commands. Moreover, the user isn't "required to instruct the horse along a specific path." Because horses typically follow trails on their own, the human rider's attention is free to focus on the overall course and decide how best to perform higher-level tasks that the horse is incapable of performing.

Note too that the horse's role requires a number of differing kinds of intelligence, object recognition, motion control, audible communication, and so forth. All these traits have analogies in machine intelligence, of course. In fact, some new-generation intelligent applications likewise rely on an assortment of AI techniques.

> As programs become more generic, the ability of an AI interface to control them could well become a key differentiator.

"Robustness and relevance come from multiple AI approaches," said Hard at ExpertMaker. Those approaches include "extensive data mining for insights and powerful desktop tools based on machine learning that allow us to model different chunks of knowledge rapidly."

Reliance on multiple approaches allows the company's microsearch engines to work within many disciplines. An example might be helping people find products they need, based on predefined qualitative parameters, Hard said. "By mirroring the work of a salesperson, we can reach much further than a simple social filtering." Thus, where social filtering might determine that "if you like *x* then you probably like *y*," the varying AI used in the company's microsearch engines can evaluate far more parameters. This, in turn, allows the engines to work in more difficult venues such as medical diagnosis, Hard said.

Power and Control

Search isn't the only application where users have been given control over AI technologies. Australia's Commonwealth Scientific and Industrial Research Organization (CSIRO), in partnership with a number of government and private entities, is preparing to test an extensive smart grid system to monitor and control energy distribution in several New South Wales cities. The system, which will use machine-learning and multiagent systems to analyze data from a multitude of sensors in user homes and throughout the grid, is tasked with better controlling the rise in peak demand for energy resulting from air conditioning and other appliances.

"Our agents cooperate with each other to bring global benefits to the agent system," explained Glenn Platt, a CSIRO engineer heading the development of the smart agent network. Agents can to learn the usage patterns of individual consumers and adapt their performance to this environment.

But user control is also key. Participating homes will receive touch screens that let them input their preferences and give feedback on how the system is performing. As an example, Platt said, "our air-conditioning agents consider human comfort factors based on feedback such as 'too hot,' 'warm,' 'too cold,' so the agents will learn a comfort model for their local customer, striving to keep them comfortable."

The Adaptive Inbox

Relying on the electrical grid for the most part, untold millions of us use email on a daily basis. Margaret Burnett, computer science professor at Oregon State University, has been developing an intelligent email program that gives users enviable control over how messages are filed. The prototype user interface she and her colleagues OSU Assistant Professor Weng-Keen Wong and Simone Stumpf, a lecturer at the City University London, created relies on machine learning to let users actually ask the program why it took certain actions, such as filing an important memorandum in the junk folder. When the program identifies the embedded rule that caused its action, users are prompted to amend the rule if they wish.

The process of tweaking the rules to suit an individual user is never ending, Burnett noted. That's because users' needs continually change and the new emails that arrive each day likewise might require new rules.

The researcher's program is written in Java as an overlay to the root application's language, meaning the team's AI interface can be adapted to other commonly used programs, such as word processing, as well. The only requirement is that the root application already contain the needed machine-learning algorithms so the user interface can tap into it.

In fact, as programs become more generic, the ability of an AI interface to control them could well become a key differentiator. "We think user adjustment and debugging of mass-produced software is the way of the future, to increase user trust in these systems and to provide faster adaptation to users," Burnett explained.

If the group's approach catches on and we see a growing number of programs that give users AI-based control over their functionality, the very definition of programming itself could change.

"In the end, users' preferences are the 'gold standard,'" Burnett said, "since the learned program is customizing itself to that particular end user. Therefore, the end user is also a 'programmer'—namely, the programmer who is doing the debugging."

The Interface Is Your Friend

Even as programmers become software authors and users become debuggers under Burnett's scenario, the intelligent interface that lets users control a program's underlying AI will also no doubt evolve. And a clue to what that interface might look like can be found with videogames. Penny de Byl, associate professor of serious games at NHTV Breda University in the Netherlands, envisions an animated game character "that can speak fluently to the player on almost any topic, exhibit and act on emotions, plan and adapt to a changing environment." What's needed, she said, "is an NPC whom the player must rely on for their own success in the game."

As in games, so in software—and life.

AI on the Team Bench

Mark Ingebretsen

High-end security systems used in airports and other public gathering spots employ AI to compare a visual database of suspect behavior with what live cameras actually see. Researchers at the Universidad Carlos III de Madrid (UC3M) were working on just such a system when they realized that the same technique could be applied to team sports.

"Our objective was to analyze the tactics of the two opposing teams in a game in order to obtain indicators that determined what types of plays or strategies would be optimal for each," explained Principal Investigator Miguel Ángel Patricio, who is head of research at the Artificial Intelligence Group at UC3M. He and his fellow researchers were joined in the research by colleagues in the Sport and Performance group at the Universidad Politécnica de Madrid (UPM).

AI-based security systems work by creating a visual data set containing examples of suspicious behavior—people loitering in groups, for example. In team sports, players position themselves at key locations on the playing field in order to put a scoring or defensive tactic in motion. If the researchers could assemble a data set with known examples of that strategy created, say, from an opponent's previous game films or better still from multiple cameras, the computer might be able to identify that strategy when it occurred in an ongoing match-up. "This information will help the trainer to determine the strategy to take at any moment during the course of the game," Patricio said

Their work called upon a laundry list of AI methodologies, including machine-learning techniques, statistics, databases, knowledge-representation formalisms, case-based reasoning, approximate reasoning, knowledge acquisition, neural networks, and visualization data. It was tested using student basketball club teams in the area.

This isn't the first time AI has been applied to sports. In the 1990s, IBM developed Advanced Scout, a data-mining program used by more than two dozen National Basketball Association teams. Robert Schumaker, an information systems expert at Iona College's Hagan School of Business who studied how to apply machine learning to greyhound racing, noted that sports is a ripe area for AI research in part because sporting competitions "can be thought of as a well-defined sequence of events that are conditionally based on game-based actions." Thus, a game's overall structure can be readily partitioned into sequences of video action.

He uses baseball as an example of how sequences in a game are relatively easy for AI systems to identify

> A game's overall structure can be readily partitioned into sequences of video action.

and tag. "When transitioning from one batter to the next," Schumaker explained, "the first batter will either get on base or head for the dugout." Either action would constitute a separate event, which would be easy to identify and tag.

In addition to identifying the tactics of competing teams, Schumacher sees applications for the work of Patricio and his colleagues in sports medicine and athlete training. "By building a baseline of player actions, teams would have the ability to detect deviations that could indicate the onset of injury," he said. "Likewise, a baseline of movements by successful players could help younger players adapt quicker and become more successful themselves. Just like a machine can analyze your swing in golf and suggest improvements, researchers in biomechanics could soon adapt this practice to a wide range of sports."

Patricio believes his technology needs more testing before it can be found on the sidelines at professional sporting events. However, with athletes paid more than the heads of large corporations these days, teams might readily spring for the technology in order to protect their valuable assets. And while Patricio has looked at other sports such as squash, he thinks North American football, with its elaborately scripted plays and lengthy list of rules, might prove particularly interesting.

Front Seat versus Back Seat Driving

Mark Ingebretsen

This October, Google's announcement that its driverless Toyota Prius prototypes had logged 140,000 miles caused people to wonder if the world was indeed ready for a driverless car. Some pointed to a driver-free car's potential liability issues. Others expressed alarm that we might one day share the road with cars operating on autopilot.

In truth, humans were always present inside Google's test vehicles, at the ready in case of a software malfunction. Although somewhat coy about the specific AI technologies that enabled the company's test fleet to wheel around sans human control, Google did say the project leveraged its hefty investment in road-mapping technology along with the formidable number-crunching capacity of its data centers.

The project organizers, some of whom were enlisted from DARPA's 2007 Urban Challenge driverless vehicle competition, hope their work will allow vehicle passengers to focus on more enjoyable things while being ferried about. "The U.S. Department of Transportation estimates that people spend on average 52 minutes each working day commuting,"

IEEE Intelligent Systems

IEEE Computer Society Publications Office
10662 Los Vaqueros Circle, PO Box 3014
Los Alamitos, CA 90720-1314

noted a Google blog post (http://googleblog.blogspot.com/2010/10/what-were-driving-at.html). "Imagine being able to spend that time more productively."

Roadways might also become safer if people let machines do the driving, Google noted, citing World Health Organization statistics that "more than 1.2 million lives are lost every year in road traffic accidents."

Reducing those numbers is also a focus of a European initiative called DRIVSCO. Rather than developing cars that drive themselves, DRIVSCO aims to create a vehicle that can learn its driver's habits and intercede when necessary.

> After driving along a road once, DRIVSCO will remember all its twists and turns forever and then guide you through them day or night.

"We deliberately use a very simple form of learning by demonstration," explained Florentin Wörgötter, the project's principal investigator and a professor at Bernstein Center for Computational Neuroscience (BCCN) Göttingen, Germany. During the trial runs, cameras aboard the test vehicle, a Volkswagen Passat, recorded the route ahead. "Learning consists basically of saving such snapshots in a database," along with data on how the driver negotiated the stretch of road, Wörgötter explained.

The next time the driver took the same road, the camera and sensors enabled the vehicle to compare the current images to those in the database. That way, the system can alert the driver about approaching curves and note if the driver is prone to approaching them too fast.

The driver-assist approach is part of a trend within AI that emphasizes collaboration between intelligent machines and their human operators (see main story). It also seems to be what auto manufacturers are actively touting today. Cars that parallel park themselves are already old hat. More recently, the Mercedes' commercials for the German automaker's E-class line, for example, focus on the car's lane-change warning and collision-avoidance technology.

Although the DRIVSCO system is more advanced than what Mercedes and other manufacturers currently have in their showrooms, it still has certain limitations. If a new situation occurs, the system can't predict or prescribe the driver's actions with sufficient certainty, so it won't issue a warning. "DRIVSCO is based on predicting the driver's behavior according to previously observed driving examples," Wörgötter explained. Hence, if a tree falls across the road or a bridge suddenly collapses, any defensive response is up to the driver.

Fortunately, such situations are rare. Moreover, DRIVSCO is a fast learner. Drive along a road just once, it will remember all its twists and turns forever and then guide you through them day or night. "But the more samples available the better," said Wörgötter. "Our idea is that DRIVSCO should be a 'life-long' learner, constantly refining its experience." ■