



Intelligence, Designed

COMPUTER SCIENCE, RESEARCH, FEATURE

By Steve Franklin, Illustrated by Jenna Luecke



UT Austin Villa soccer players criss-cross the field, kicking, scoring, colliding with opponents and working as a team. Their coaches get players to work seamlessly together, and they operate like a well-oiled machine—or a few such machines, really. These players, along with the students and professor who created them, took first place in the 2014 RoboCup Soccer 3D Simulation League, a sports match for robots.

UPDATE: Listen to our podcast episode about how insights from neuroscience could help make the robots more nimble (Apr. 2017)

Welcome to the artificial intelligence age. These days, you'll find AI in the workplace, the home, even on a sports pitch. From hospitals to highways, artificial intelligence offers new solutions to real-world problems.

Scientists at UT Austin, one of the world's top computer science programs, say artificial intelligence is enabling breakthroughs that will improve everyday life.

But teaching an artificial system to be intelligent isn't as easy as the movies make it look. Lots of time and energy go into training computers on seemingly simple concepts.

"You have to first appreciate the enormous gulf between computers and people. Computers are profoundly stupid machines," says Bruce Porter, Chair of the Department of Computer Science. "If you can get a computer to do something smart, like IBM's Watson on *Jeopardy*, you have to respect how much work went into that accomplishment. Because computers as they come out of the box are so useless, so stupid, to get them to do something intelligent like these systems have done is a tour de force."

Teaching Intelligence

Consider the amount of computation your brain does for something as simple as catching a ball. It recognizes the object, calculates speed and direction, and computes where to intercept the ball. These are easy tasks for humans,





but teaching them to a computer means tackling each piece separately.

Kristen Grauman, head of the Computer Vision Group at UT Austin, teaches computers to understand what they see. For this, machines need to identify categories of objects or activities and communicate with humans.

Part of teaching computers to recognize objects is getting them to decipher descriptive terms like *furry*, *metallic*, or *formal*. Imagine trying to explain *furry* to someone who has never held or seen an animal, and you get a sense of the challenge. To address the gap, Grauman is developing learning algorithms that can estimate the presence of descriptive attributes in images. Such predictions pave the way for clear communication between human

“teachers” and visual recognition systems. A person can tell the system that dogs are furry, four-legged animals, and the system can conjure a model even before it’s shown actual image examples.


“If you can get a computer to do something smart, like IBM’s Watson on Jeopardy, you have to respect how much work went into that accomplishment.”

Nuanced Learning

An image retrieval project, WhittleSearch, builds on this technology to let people make relative comparisons between images, allowing users to fine-tune their search. “Suppose you have in mind what something looks like: your ideal handbag or pair of shoes, or the criminal you witnessed commit a crime. Trying to express that to a computer in words alone is quite limiting,” Grauman says. Instead, WhittleSearch allows users to give comparative feedback, like that the pair of shoes should be *more formal*, or the criminal had a *broader nose*. This way, the machine targets the right image from its vast database.

Some of Porter’s work focuses on teaching human language. In Project Halo, he and his team successfully built a computer that could learn a science subject well enough to pass an Advanced Placement (AP) exam. This involved machines answering questions they had to work out the answers to, not ones for which they could simply retrieve information or perform a basic operation.

Porter and his lab were also among the scientists who contributed expertise to the construction of IBM’s automated question-answering system, Watson. Whereas search engines yield results from a keyword search, Watson actually answers questions posed to it, enabling the computer to beat two *Jeopardy* champions in 2011. In an upper-division Automated Question Answering class taught by Porter and colleagues at UT Austin and six other universities, students build applications that use Watson to answer complex questions.



Many artificial intelligence problems come down to getting machines to master similar higher-level processes that mirror learning and reasoning.


“We’re never in exactly the same situation twice,” Peter Stone, director of UT’s Learning Agents Research Group explains. “If you’re a computer in a new situation you have to decide, not what action worked when I was exactly in this situation, but which actions in similar situations were good and which were bad.”

Programming computers to recognize those nuances happens sometimes with a technique called reinforcement learning. Other approaches include evolutionary computation—which runs neural networks through a survival-of-the-fittest gauntlet modeled on the biological process of evolution, helping systems learn new behaviors—and multi-agent systems, a tactic that involves programming components to work together in ways that make the most of each part’s unique abilities.

Ease and Safety

Researchers see the potential for artificial intelligence to bring real societal benefits. Techniques that Stone and his team used in creating the robot soccer team could someday be employed to develop enhanced disaster-response robots. Some of Grauman’s artificial intelligence systems can capture and summarize video in ways that could prove useful for everything from national security to search and rescue efforts to helping people with memory impairment.

Finding new ways to apply the technology to daily life, so people can go about their business more easily and safely, excites AI innovators. Stone and students in a Freshman Research Initiative course have worked to build that icon of artificial intelligence in popular culture: a helper robot. Students in the Autonomous Intelligent Robotics freshman research course work towards building robots that can navigate campus buildings and may one day make deliveries, answer questions and



even act as tour guides for visitors.

Risto Miikkulainen, director of UT’s Neural Networks Research Group, also points to the healing potential of some AI systems. He is working with colleagues at Yale on an AI strategy that models how schizophrenia forms in the brain. Researchers hope to devise new treatments for the complex disease based on what they learn.

Inevitably, artificial intelligence will also help us understand ourselves better.

The Future

Artificial intelligence experts foresee a future where AI unleashes new creativity and knowledge, not unlike what the Internet has done.



Self-driving Cars

Peter Stone taught a Freshman Research Initiative course on autonomous vehicles where researchers worked with Austin Robot Technology to create an autonomous car named Marvin. The vehicle could drive in real-world settings while following the rules of the road. A future of self-driving vehicles and autonomous traffic management may not be far off.

“Computers can already fly a passenger jet much like a trained human pilot,” Stone notes. “Soon vehicles will be able to handle most of the driving tasks themselves, but once they’re popular, we’ll need to coordinate those vehicles on the streets.”

To that end, Stone and his students developed virtual intersection systems to make auto travel safer and faster. In models, self-driving cars were able to move through intersections with little or no need to stop.

[Read an op-ed by Stone about why driverless cars will mean safety, convenience, and social change.](#)

“Creative problem-solving is a computational technique that’s just now coming of age because of the increase in computational power,” Miikkulainen says. “Once we are able to automate discovery, we will have solutions that we could not come up with as humans because the systems are too complex.”

He posits that automated creativity could be used to design new kinds of robots, new kinds of cars and games that are more challenging.

Porter envisions “an artificially intelligent science oracle—a computer-based system that we would consider to be the all-knowing authority in some area of science.” This oracle would be able to read every scientific article written in a certain field and could connect the dots between similar papers, passing that knowledge on to human researchers.

Inevitably, artificial intelligence will also help us understand ourselves better. There are insights into human cognition, complexity and intelligence that come from teaching computers to emulate us. This is just one way that, Porter says, “Technology is making the world a better place.”

UPDATE: [Where is Artificial Intelligence Headed?, a Q&A with Peter Stone \(Sep. 2016\)](#)

[ABOUT](#) [CONTACT](#) [SUPPORT CNS](#)

