


ARTIFICIAL INTELLIGENCE



POPULARIZED IN science fiction, AI is no longer merely a figment of our imaginations. In the past few years, the technology has moved from research labs into our everyday lives. It's being

used in medical devices, smart-home systems, and video games—to say nothing of robots and autonomous cars. And AI has started to do what many people have feared: outsmart humans.


In this special report, *The Institute* describes one of the latest AI developments: deep learning, in which neural networks teach themselves and make decisions on their own. Google, Microsoft, and other tech companies are racing to improve deep learning and apply it to new applications. Noteworthy AI applications already in place include computer systems writing news articles and preventing credit card fraud [p. 6].

Companies can't find enough engineers qualified in machine learning and other AI applications, so we provide tips for breaking into the growing field [p. 8]. And you'll find IEEE products and services [p. 15] as well as conferences and books [p. 16] on AI and related topics.

We profile IEEE Fellow Fatih Porikli about his work in computer vision and deep learning [p. 17]. We also feature an IEEE member who helped found Clearpath Robotics, a leader in self-driving vehicles for warehouses [p. 18]. And don't miss our Q&A with 2017 IEEE president-elect candidates IEEE Life Senior Member Jim Jefferies and IEEE Fellow Wanda Reder [p. 10].

To let us know about your work in AI or to comment on what you've read in this special report, email the editors: institute@ieee.org.

—Monica Rozenfeld, associate editor




Mastering Deep Learning

*The next step for AI is machines that get smarter,
on their own* BY MONICA ROZENFELD

ERIC FROMMELT





and other tech companies are in a race to apply deep learning to make machines intelligent without much help from the programmers.

Take that voice-activation service, for example. Say a user with an accent is dictating a message amid background noise. A deep-learning machine could detect and process those factors to interpret what the person is saying. And this is just a starting point. Applications for deep learning are endless.

"Deep learning can be applied horizontally across many fields and applications," says IEEE Member Rajan Goyal, a distinguished engineer at Cavium, in San Jose, Calif., who works on next-generation accelerators. He is exploring silicon designs for deep learning. "The power of deep learning is that it can solve many problems that have been impossible to solve to date."

THE MIND OF A MACHINE

Deep learning is a relatively new form of artificial intelligence that gives an old technology—a neural network—a twist made possible by big data, supercomputing, and advanced algorithms. Data lines possessed by each neuron of the network communicate with one another.

It would be impossible to write code for an unlimited number of situations. And without correct code, a machine would not know what to do. With deep learning, however, the system is able to figure things out on its own. The technique lets the network form neural relationships most relevant to each new situation. First, however, the machine needs to learn how to learn.

One of the first deep-learning exercises was carried out by Yann LeCun, director of AI research at Facebook. LeCun taught a computer system how to recognize the differences between images of dogs and cats. When the system chose incorrectly, he would correct it until the program figured out the reason it was wrong, such as not, for example, having considered the shape of a nose or an ear. Eventually, the machine started to distinguish a cat from a dog nearly every time.

"Just like human beings grow up learning about the environment around them and reacting to it,

machines can also learn complex tasks this way," Goyal says. Basically, the machines learn from their own mistakes.

Now add millions of inputs beyond those for cats and dogs, and you start to have a highly intelligent system. Because tomorrow's intelligent systems will have so much knowledge, they aren't likely to make the mistake of reading a word incorrectly just because the handwriting is poor or a letter is missing. Instead, the machine would, for example, consider the context of the word, such as how it is written, and on what medium—on a billboard, say, or in a newspaper—as well as what words or images surround it.

"Instead of teaching the machine everything, let it teach itself," Goyal says. "By creating a system that can learn on its own, the time to develop it is drastically reduced."

The more an intelligent machine knows, the faster it can pick up new information. Eventually, humans might not be able to teach it much at all. One of the latest examples is the AlphaGo program, which defeated Go champion Lee Sedol. Go, an ancient board game renowned for its intuitive strategy, has more possible combinations than the number of atoms in the universe.

In response to the victory, Demis Hassabis, cofounder of DeepMind, the company that developed AlphaGo and has since been acquired by Google, said it demonstrates that AI could be used to solve problems that confound humans.

And then there's Google Brain. For each image of a location collected for Google Maps, a team of employees had the tedious task of clicking yes or no for whether a photo was of an actual address and not, for example, of an empty stretch of street or unoccupied woods. Then the company's engineers trained its computer system to handle the task. And using deep-learning image recognition, Google's machines were able to identify with street addresses all the homes and buildings in France in less than an hour.

IEEE Fellow Li Deng, chief scientist of AI at Microsoft Applications and Services Group and research manager at Microsoft Research, pioneered research and applications

in deep learning speech recognition. With colleagues at Microsoft Research he explored multimodal intelligence involving images and natural language for computers to communicate like humans. Deng received the IEEE Signal Processing Society's 2015 Technical Achievement Award for contributions to deep learning and to automatic speech recognition.

When a deep-learning system views an image or video, it can describe what it sees. The system identifies visual cues—such as woman, camera, flowers, and purple—then uses natural-language models to generate many possible sentences describing the scene. The system can then quickly determine what it understands to be the most sensible description: A woman is taking photos of purple flowers.

The opposite is also true. If you type in a descriptive sentence, the machine can bring up the most relevant media it finds on the Web.

That ability to understand an image's content and express it in natural language, Deng says, will be useful for a wide variety of applications. At a Microsoft conference on 30 March, a blind software developer showcased how he was able to "see" using a deep learning-enabled headset. The user tapped a button on the headset to take a snapshot of the scene in front of him, and the system explained what was in the photo. The system could even describe facial expressions and detect if a person in the image looked happy or confused, for example.

THE FUTURE OF AI

More precise than humans, intelligent machines will be able to pick up subtle cues, such as differentiating fake smiles from real ones—which is often difficult for a person to discern—Goyal says. Deep learning also will enable machines to predict a person's needs, Deng adds. For example, in the future when you text a friend that you plan to see a movie, Uber or a similar car service could be automatically programmed to pick you up if you do not have your own car. Intelligent systems, he says, will be able to make decisions far more accurately and faster than humans can, and we're already starting to see that happen. ♦

HAVE YOU EVER USED


a voice-activated service such as Apple's Siri only to find it completely missed what you were saying? Or played a game against a computer and felt it didn't even put up a fight? That's about to change with advances in deep learning, which improves computers' ability to process information and make decisions—like people do, and sometimes even better.

Deep-learning techniques allow a computer system to connect the dots from many different areas of knowledge, similar to how the human brain works, to make the best decision possible. Facebook, Google, Microsoft,

AI Is All Arc

Machines are writing news articles, composing music, and protecting our money BY AMANDA DAVIS

ARTIFICIAL intelligence has crept into our everyday lives, even though we're not always aware of it. Here are some examples of AI's many applications.



MAKING HEADLINES

Journalists might not care for this, but AI programs are becoming smart enough to compile bits of information and turn them into articles. Although long-form features and investigative pieces are still being left to reporters, the Associated Press, Fox News, Yahoo, and other outlets are using AI to fill in the blanks of simple, data-driven stories, like financial summaries and sports score recaps.

Last year Automated Insights, a software company in Durham, N.C., launched the news-writing bot Wordsmith, which fills in the blanks of an article template with simple words and information such as stock market figures or baseball stats.

Wired magazine really put Wordsmith to the test by using it to write an obituary for IEEE Life Fellow Marvin Minsky, the AI pioneer who died in January. The bot produced a 150-word "just the facts" obituary with such information as Minsky's name and age, the cause of death, and names of family members—which could be

useful if a news outlet wants to be the first to publish the story.

There's also an AI program to write clickbait headlines—the kind designed to drive online readers to an article. The program uses recurrent neural networks (RNNs), which form connections based on the data they receive.

After uploading and practicing on several million articles from BuzzFeed, Gawker, The Huffington Post, and other sites, the software has produced a number of grammatically correct—yet offbeat—headlines such as "How to Get Your Kids to See the Light" and "This Guy Thinks His Cat Was Drunk for Five Years."


which adds headlines to short articles—all assembled by the RNN program.

LETTING GO OF THE WHEEL

Self-driving cars might be a few years from the showroom floor, but several manufacturers have already incorporated AI in their cars.

This year BMW introduced its 750i xDrive model, the first car that can park itself with no one behind the wheel. Simply press a button on the car's remote and the full-size sedan backs out of a driveway or maneuvers itself into a tight parking space.

The 2015 Infiniti Q50S and the 2015 Mercedes-Benz S65 AMG have systems of sensors



As its creator, Lars Eidnes, explained in a blog post, "We can show an RNN a bunch of sentences and get it to predict the next word, given the previous words. So, given a string of words like *Which Disney Character Are*, we want the program to produce a reasonable guess, like *You*, rather than, say, *Spreadsheet*."

Eidnes has launched an auto-generated news site, Click-o-Tron,

that engage the brakes when the car comes too close to the vehicle in front, an object on the road, or a pedestrian. Such models also have a lane assist feature, which adjusts the car's course when it begins to drift.

In January, Toyota hired IEEE Member Gill Pratt, a robotics expert, as CEO of its research institute, and it plans to spend US \$50 million to

support collaborative AI research at Stanford and MIT. The company is working on a camera that can read road signs and determine the colors of traffic signals—which could help vehicles get through intersections safely.

INTELLIGENT PLAY


AI programs can keep gamers on their toes as their computer-based opponents become less predictable. First-person shooter games, like *Call of Duty*, make significant use of AI to enable the background characters to analyze their surroundings to find objects or take actions that might be crucial to their survival. The characters take cover, investigate sounds, and communicate with other AI characters to increase their chances of victory.

Another game, *Left 4 Dead*, features an AI "director." The director isn't an onscreen character but an invisible force that increases the number of opponents to beat as the game progresses based on each player's situation, status, skill level, and location.

In fact, AI has been part of video games for more than 60 years. An early example was Nimrod, a computer designed for the 1951 Festival of Britain exhibition by IEEE Life Member John Bennett and built by engineer Raymond Stuart-Williams. Exhibition attendees could play Nim, a mathematical game of strategy, against the computer. The player would make moves by pressing buttons on a panel, with each button corresponding to a light on the machine; the computer would run through calculations to make its next move based on the player's actions.

The 1978 game *Space Invaders* introduced AI into computer-generated opponents, whose movements depended on the player's input. *The Sims*, one of the best-selling computer games of all time and still on the market,

and Us



remiered in 2000 and took AI step further. Players create characters, build houses for them, and assign them careers, but the characters exhibit a certain degree of free will, reacting to situations (such as meeting new neighbors or dealing with kitchen fires) in somewhat unpredictable ways.

OR YOUR ENTERTAINMENT

Are you bored with your current music playlist or can't decide which TV show to watch? I might be able to help. Netflix, Pandora, Spotify, and other streaming services use the technology to recommend movies and music based on their customers' past selections. By applying deep-learning algorithms, the services make recommendations that their subscribers are likely to enjoy.

To make accurate choices, the companies' programmers manually tag songs, movies, and TV shows with certain traits. A song on Pandora, for example, might have deep bass and dynamic vocals listed as traits. If you listen to such a song all the way through, AI selects another one with similar attributes.

Not only are computers now selecting music, but they're also composing it. Researchers at the Sony Computer Science Lab, in Paris, are working on algorithms to let computers produce original symphonies and other music.

A symphony could be in the style of composer Johann Sebastian Bach, for example, or the AI could produce riffs that sound like, say, jazz musician John Coltrane's, or original melodies that conjure up a particular pop star. The jazz bot can even come up with songs

that combine the attributes of two musicians, such as a tune that sounds like it was written by composer and conductor Pierre Boulez and played by saxophonist Charlie Parker.

"The commercial applications of such efforts may include endless streams of original music in shopping malls that can respond to crying babies with soothing harmonies, as well as time-saving tools for busy composers," says William Hochberg, who wrote about the technology in *The Atlantic*.


But can computer-composed music move an audience the way live music can? Is it able to replicate nuances, like a singer's imperfect vibrato or a blues musician's soulful strum of a bass guitar? AI has a way to go, according to Hochberg. Oftentimes "ham-fisted dynamics and pointless melodies" make it clear, he says, that the song wasn't composed or performed by a human.

FRAUD PREVENTION

It's a scenario becoming familiar to more and more people: You go to a café to grab a cup of coffee, only to have your credit card declined. Minutes later, you receive notice that your card has been suspended because a large


purchase was made by someone at a store on the other side of the globe. Your account is frozen, and a new card arrives in the mail several days later.

AI and data-analysis programs work together to stop thieves. AI systems are fed a large sample of normal and fraudulent transactions so they can learn what each type looks like. If, for example, a client typically makes most of her credit card purchases locally, and then buys a high-priced item in another state or country, the system flags the transaction.



Similar technology has been in place since the late 1990s, but as IEEE Member Jungwoo Ryoo says in an article on *The Conversation*, an online research publication, the process has now become almost instantaneous.

The algorithms in use today can handle more data and do it faster, making the job of fraud detection not only less labor-intensive but also more accurate, says Ryoo, associate professor of information sciences and technology at Pennsylvania State University, in State College. That is important for companies such as PayPal, which processes some 1.1 petabytes of data at any given moment. Because AI improves through practice, every new piece of information makes the system that much smarter and more efficient. ♦





Landing a Job in Artificial Intelligence

Companies can't hire skilled professionals fast enough BY JOHN R. PLATT

THOUSANDS OF openings in artificial intelligence and machine learning posted on job boards are going unfilled. In fact, though AI is one of the fastest-growing areas for high-tech professionals, according to a recent *Kiplinger* report, there are too few qualified engineers. "Supply is far lower than demand," says Boris Babenko, a machine vision engineer at Orbital Insight, a company in Palo Alto, Calif., that uses AI to make sense of data gathered from satellite images. That's true of all software engineering, but AI is a niche on top of that.

The need for AI specialists exists in just about every field as companies seek to give computers the ability to think, learn, and adapt.

"If you look hard enough, any industry you can think of has a need for AI and machine learning," says Geoff Gordon, acting head of the Machine Learning Department at Carnegie Mellon University, in Pittsburgh.

Many open positions exist in Silicon Valley, of course, but AI jobs can be found in Boston, New York, Pittsburgh, London, Hong Kong, and just about any city, Gordon says.

A recent report from Shivon Zilis, a founding partner at the investment fund Bloomberg Beta, points to Apple, Google, and IBM as AI's three top hiring companies. But those three firms are far from alone. Zilis's report shows that AI is used in many fields including advertising,

agriculture, health care, and transportation.

GETTING IN

How do you get a job in AI? According to Gordon, some workers start in software engineering or a data-heavy field such as physics.

"Others might come from a field like biology," he says. "Machine learning becomes an important part of what they do, and they end up loving it." He says a lot of his Ph.D. students have returned to school to study AI after a few years in industry.

A background in software engineering, experts agree, is a must-have.

"We assume that when people first come in they have not only formal thinking ability but also the know-how to code and work with computers," Gordon says. The exact programming lan-

guage doesn't matter; most students know several.

"We love seeing candidates who have had some open-source projects," Babenko says, "so we can look at the code they've written."

Beyond technical skills, AI requires an innate sense of curiosity and a drive for problem solving.

"We're trying to train people who can take on the impossible problems and solve them," Gordon says. "Someone once described our students as elite machine-learning ninjas who would get dropped in by black helicopters to solve all your problems."

A combination of analytical ability and creativity also matters, according to Matthew Michelson, chief scientist at InferLink, an AI firm in El Segundo, Calif.

"This is a difficult combination to find, but you need to be analytical to understand the data and to craft algorithms," Michelson says. Creativity is important, he adds, because "the problems are often new and require new solutions." He looks at candidates' hobbies—he's partial to those who developed games—when considering how they might handle problems.

STAYING CURRENT

As for education, jobs exist for those with a master's degree, and there are plenty of lower-level positions as well. Employers hiring in AI value Ph.D. candidates for their depth of education and the work they produce during their doctoral program.

Attending conferences [like those on p. 15] is a good way to keep your AI knowledge up to date—vital with the field evolving so rapidly—and to find job leads.

"My advice to those interested in working in AI is to network, attend events, and follow industry news closely—become part of the industry conversation," says Jana Eggers, CEO of Nara Logics, a synaptic intelligence

*Just about
ever
industry
needs
employee
with AI
skill*

company in Cambridge Mass., that combines neuroscience and computer science. "It is the best way for you to assess your fit with a company well as to learn of professional opportunities."

Babenko praises competitions such as those by Kaggle, which styles itself as "the world's largest community of data scientists." The competitions can be great for networking, he points out.

WHERE TO LOOK

The big tech companies are all hiring, and paying top dollar for talent. The deep-pocket companies might, in fact, be pricing some of the smaller one out of the market.

"I don't think Orbital Insight can compete with larger companies," Babenko says of his employer.

But well-equipped candidates should consider their own personalities, adds: "Google has a lot of cool projects, but you're drop in the bucket there whereas a smaller company, like Orbital, provides him with a feeling that he's doing more."

The job boom, for the most part, has been in industry. Gordon has not seen a corresponding increase in academic jobs, although he expects that to change.

And the number of AI jobs will only keep increasing. "The field will continue to be hot," Gordon says. ♦