

MACHINE LEARNING,

BIG

UND

The ancient Chinese game of Go has simple rules, yet is extremely sophisticated. With a large board and few restrictions, the game is said to be a googol (10 to the hundredth power) times more complex than chess. There are more possible positions in Go than there are atoms in the universe.

The game, in which opponents take turns placing black or white stones on a board, is played largely through intuition. Players understand that moves made early in the game can shape the match dozens of plays later. Go's subtleties, patterns, and elegance have captivated players, scholars, and mathematicians for millennia.



HOW WILL THE IC HARNESS ADVANCEMENTS
IN ARTIFICIAL INTELLIGENCE?

ERSTANDING

BY MELANIE D.G. KAPLAN



“WHETHER WE CALL IT AI or machine learning, they’re different methodologies to try and replicate the human thought process.”

— Dr. Colleen McCue, data scientist, CACI

According to Google DeepMind, which created the computer program AlphaGo, Go has long been viewed as a pinnacle against which to measure artificial intelligence (AI). AlphaGo uses a form of AI called deep neural networks to mimic expert players and achieve a 99.8 percent success rate against other

programs designed to take on Go. In October 2015, AlphaGo became the first computer program to beat a professional Go player, and in March, AlphaGo took on Go world champion Lee Sedol, beating him in four matches out of five.

Industry leaders and academics often use games as a testing ground for algorithms. In 1997, IBM supercomputer Deep Blue defeated the world chess champion, and in

2011, the company’s Watson supercomputer beat two champions at Jeopardy. Two years ago, DeepMind algorithms learned to play dozens of Atari games.

The Intelligence Community (IC) has discussed AI for decades, and in the months after AlphaGo’s success, a wave of surprise swept through the community. For many, the human-like intuition required to be successful at Go represents the progression of AI to the next level.

“Many experts thought this wouldn’t happen for another 10 years,” said Michael Laielli, a research and development technologist at the National Geospatial-Intelligence Agency (NGA). “We’re very interested in how [Google] used deep learning in just the right way, combined with other approaches, to solve a very hard problem.”

AlphaGo’s victory was a wake up call for the IC—not only is machine learning advancing at a breakneck speed, it also demands changes to a long-established culture.

Experts debate the precise definitions of AI, machine learning, and deep learning as well as their relationships to one

another. In 1959, Arthur Samuel defined machine learning as a field of study that gives computers the ability to learn without being explicitly programmed. Deep learning is commonly considered to be a more in-depth subset of machine learning. Regardless of terminology, experts agree these technologies promise a sea change in problem solving.

Terry Busch, chief of the Defense Intelligence Agency’s (DIA) integrated analysis and methodologies division, believes machine learning and automation will allow his experts to be just that, experts, rather than data managers. Finding new ways to interrogate data will benefit the tradecraft and yield powerful results, Busch said. But he’s also aware of the challenges.

“We’re in about third grade,” Busch said of DIA’s machine learning capabilities in April at USGIF’s second Data Analytics Forum in Herndon, Va. “And we need to get to graduate school—fast.”

BIG UNDERSTANDING

It wasn’t long ago that the phrase “big data” created a deafening buzz in the IC. Now, the focus has shifted—in a subtle but profound way—from collecting information to empowering it. How can we derive meaning from big data? How can we use big data to change outcomes? Today’s buzz is “big understanding,” for which intelligent machines are essential.

“Whether we call it AI or machine learning, they’re different methodologies to try and replicate the human thought process,” explained Dr. Colleen “Kelly” McCue, a data scientist at CACI and author of *Data Mining and Predictive Analysis: Intelligence Gathering and Crime Analysis*. “We use them to confirm what we know or think we know and to discover new patterns, relationships, and entities.”



BUZZWORDS

Although proper definitions are still being debated among experts and scholars, the following are generally accepted definitions to help understand the nuances among these increasingly common phrases.

ARTIFICIAL INTELLIGENCE: A branch of computer science focused on the development of technologies that simulate human cognition and general intelligence.

MACHINE LEARNING: A multidisciplinary approach allowing computers to adapt without the need for explicit programming, i.e., the machine may modify its own rule set to learn new rules for future problems. Machine learning includes two training methods: supervised learning, in which known events are used to create models; and unsupervised learning, in which the machine learns from naturally occurring trends, patterns, sequences, and associations. The two methods can be used alone or in parallel. A common example of a machine learning paradigm is an artificial neural network mimicking a biological one.

DEEP LEARNING: A branch of machine learning that attempts to mimic the architecture of the human brain by using multiple levels of highly connected simple models, i.e., a multilayer neural network. Deep learning systems can infer new insights from a large set of provided data or inputs.

McCue said AI is particularly helpful in achieving two goals: performing tasks humans cannot (such as processing massive amounts of information); and performing tasks humans should not (such as those that are highly repetitive).

“Whether you’re trying to optimize sales in your store or figure out where the bad guy will go next, trying to process that much data is not a good use of human time,” McCue said.

Esri President Jack Dangermond said machine learning is about finding the needle in the haystack. He is especially fond of the phrase, “burn the haystack down.” Esri’s capabilities now allow users to find the proverbial needle in the haystack in minutes, not weeks—whether it’s discovering incongruities, recognizing patterns, or revealing fraud. The key is integrating big data and AI into a GIS platform, which allows the user to detect links between people and places, changes in imagery, or interactions between humans, events, and locations.

“In the intel space, this may mean finding bad guys based upon their spatial data anomalies discovered out of billions of human observations,” Dangermond said. “We are not as interested in big data and smart computations as we are in creating big understanding and being able to drive that understanding and insight into the actual workflows or right down to the tactical operations people making decisions.”

In automating some tedious tasks typically performed by humans, such as manually scanning an image, AI frees us to think creatively, “one of the best attributes of the human brain,” noted NGA’s Laielli. Many in the IC say machine learning is, at the very least, a workforce multiplier—a powerful tool for finding the most valuable bits among vast amounts of data.

But to create truly intelligent machines, Laielli said, the computer must be able to connect those bits of data, think critically, and produce cohesive analysis as a result.

DEEP LEARNING

The ImageNet Large Scale Visual Recognition Challenge began in 2010



WHAT'S NEXT?

JEFF JONAS ON SELECTIVE CURIOSITY

Standing on the grounds of the U.S. Capitol, one hears birds, buses, and the chattering of tourists. Off in the distance is a siren, the overhead whir of a helicopter, the clanging of a construction site. In the nation’s capital, these sounds are typical, and they generally don’t elicit inquisitiveness. Throw the crack of a gunshot into the mix, and one instantly becomes quite curious.

This is what IBM Fellow and USGIF board member Jeff Jonas has coined “Selective Curiosity.”

“You can’t be curious about everything,” Jonas said, noting that the out-of-the-ordinary cannot always invoke our curiosity either. “There are simply not enough free computational cycles to investigate everything.”

In his lab, Jonas is developing Selective Curiosity features, in which a machine will not only determine it needs more information to reach a conclusion, but independently seek the information from sources such as Google or LexisNexis.

In a 2014 blog, Jonas wrote that his G2 Selective Curiosity tool—which he hopes to reveal later this year—“finds a maybe that would matter.” In the Intelligence Community, machines routinely produce “maybes.” In such situations, this new technology aims to help a machine determine what it wishes it knew (i.e., an individual’s employer or home address) and figure out where it should go to find the answer.

“In big data there are millions of maybes,” Jonas said. He added that tipping and cueing—which tips analysts or an algorithm toward a particular person or object in the data—is riddled with false positives. “Fixing these accuracy rates will not come from more algorithm work, it will come from more data.”

But which data? Jonas said Selective Curiosity will allow the machine to figure out—and then search for—exactly the type of data it needs, yielding significantly more accurate decision-making than existing methods can deliver.

and has since become the benchmark for large-scale object recognition. Competitors in industry and academia train their algorithms using more than one million labeled images in more than 1,000 categories. During competition, their software is tasked to name unlabeled images, and in 2012 the event hit a milestone: For the first time, the winning team used deep learning to improve accuracy.

Loosely based on the brain’s neural network—yet not nearly as intricate—deep learning feeds off big data. Because it is designed to work like a brain, a deep learning enabled machine doesn’t

need to understand (or be explicitly programmed for) the universe’s rules for every possible scenario. Rather, its algorithm learns. In essence, developers teach and guide the algorithm. For example: “This is a cat; go identify other cats.”

Machines that learn from experience are now tasked with intuitive problems, such as recognizing speech, faces, and preferences. Considered a branch of machine learning, deep learning already touches human lives in countless ways: DeepMind helps doctors detect cases of acute kidney injury; Netflix recommends movies for you to watch;



Microsoft's CaptionBot describes photos; Google's spam filter uses an artificial neural network to intercept 99.9 percent of unwanted emails; Facebook's DeepFace recognizes your friends; Google and Chrysler are building driverless minivans that can recognize pedestrians; and Amazon's Alexa keeps track of your grocery list. Some say it's not long before a machine will review inbox messages, filter out the noise, and show us only the emails that require our attention.

"You train the machine like a newborn," explained Dr. Dennis Bellafiore, who teaches in the geography department at Pennsylvania State University and studies machine learning. "But once you do, boy are they fast. And [the machine] will go beyond repetitive tasks. It will start to consider various scenarios, possibilities, and things that may be related from a geographic and temporal perspective."

Andrew Jenkins, a principal data scientist at DigitalGlobe, said perhaps one of the most significant tasks taken on by machine learning is feature extraction, which has traditionally relied on human expertise.

"What deep learning provides is an unsupervised process that allows the algorithm to learn and extract the best features—taking the human out of the loop," Jenkins said. "If you want the

machine to learn to detect a car in a picture, for example, you'd give it many samples of a car, and the algorithm would pick up on features such as distinct edges, color, and shapes. Then it could see a picture of another vehicle and say, 'I have 95 percent confidence this is a car.'"

Laielli said historically, visual recognition has been a hard problem for machines. "[Achieving this] has brought computers to a level arguably as good as humans on certain visual tasks."

He added, deep learning has given machines an important boost in the ability to learn from large amounts of data. Deep learning is not just hype, Laielli stressed. Academia, government, and industry are keeping an eye on how the other sectors are applying the technology, and some describe the effort to snatch up deep learning experts as a virtual arms race. NGA, for example, is paying attention to research being conducted by companies such as Uber, Facebook, Apple, and Google. The banking and finance sectors are investing heavily in this area. Advances in neuroscience continue to help humans understand how the brain works and how we learn, which in turn helps data scientists develop smarter computer programs.

Sarah Battersby, a senior research scientist at Tableau Software, said deep

learning is particularly helpful to the IC: "It would be impossible to have enough trained analysts and time to evaluate all the sensor data available," she said.

However, a well-trained machine learning system could sift through the data and then human analysts could evaluate the smaller set of more relevant data.

But these techniques are only as good as your data and training methods, Battersby added. "One should never implement a machine learning system to provide 'the answer' without evaluating the potential quality concerns of the result," she said.

Potential GEOINT applications for deep learning include crisis response management, unmanned vehicle operations, air traffic control, and industrial process control. Deep learning algorithms can keep up with events that happen in different times and places, look for objects in enormous geographic areas, or locate words in vast social media repositories. Using satellite imagery, weather analytics, and human activity, for example, deep learning can help predict famines and droughts that could lead to uprisings or make informed decisions about the best place to land a helicopter. The machine helps humans understand, reason, and learn; and the more it works, the smarter it becomes.

IBM WATSON is a technology platform that uses natural language processing and machine learning to reveal insights from large amounts of unstructured data.




PHOTO COURTESY OF IBM

MACHINES ON OUR TEAM

A critical ingredient to machine learning is providing machines with enough feedback so they know when they're right or wrong. Although experts agree humans need to stay in the loop, finding the sweet spot of human-computer balance is tricky at best. Concerns about humans being totally replaced by machines might be unfounded. Yet experts in many areas of AI still talk about the potential dangers of intelligent computers and express concern about technological singularity—when computer intelligence exceeds that of the human brain with dire consequences.

"The biggest issue for the 21st century, for individuals and organizations alike, is determining how to achieve the correct balance between humans and automation to optimize outcomes," said Richard Boyd, CEO of SZL, which offers a cognitive machine learning system by the same name.

As an example of this balance, Boyd described an object-scanning machine that is 85 percent confident in its assessment. A human follows behind, amends answers, and the machine learns. Those adopting AI look at machines as an integral team member and perceive the goal as synergy between machine processing and human attention and care.

"Whether you're trying to do calculations for space missions or work on the human genome, machines are good for some things, and humans are good at some things," Boyd said. "The combination of them working together is a powerful way to go after unstructured data—and separate the wheat from the chaff—from chatter on social media to satellite imagery."

Although machines can sometimes pick up subtleties that aren't apparent to humans, Patrick Biltgen, technical director for analytics at Vencore, emphasizes machines don't understand nuance. This is immediately clear when you upload slightly ambiguous photos into Microsoft's CaptionBot. The classic example is that a computer sees a picture of a cat—perhaps in an unusual position—and says, "I am not really confident, but I think it's a cat," whereas a 3-year-old human will know, unequivocally, that it's a feline.

Computers also have a difficult time understanding deception and sarcasm, which can lead to false alarms.

"If you train a computer to recognize what an invasion looks like, it might be tripped up by a fake invasion," Biltgen said. "There's promise with AI, but the journey is very frustrating, especially with GEOINT."

He cited an example of an airborne motion sensor that was confused by the movement of trees blowing in a strong wind. That created an alert that there were 50,000 "movers" in a city—and it took a human to understand this scenario was implausible in a city with fewer than 50,000 residents.

"So you go back and tweak the algorithm," Biltgen said. "[There's often] some fairly obvious false alarm that a human should have known to correct for."

The possibilities for error are boundless: A computer can report that a large number of vehicles uncharacteristically appeared at a location where it previously hadn't seen cars. What the machine doesn't know—until a human educates it—is that it's the first Sunday of football season, and the cars are gathering in a stadium parking lot.

"This is an example of the challenge of 'normalcy,'" Biltgen said. "Is the empty parking lot normal or is football season every fall normal?"

In another widely cited example, motion imagery systems may be calibrated to detect loitering and U-turns, under the assumption that anyone who loiters or U-turns is up to no good. But what about loitering taxis? Or mail trucks that often turn around? Machine learning systems require extensive tuning to eliminate false alarms.

Computers also struggle when it comes to judgment—which is why Predator UAVs don't decide where to fire and robots don't sit in a jury box. A machine may discover a city in Syria has new fortifications, or it may determine that a runway has doubled in size, but it can't tell us if either was built by friend or foe. A machine—at least for now—cannot discern whether an adversary is lying.

Gabe Chang, IBM's federal CTO architect, contributed to an article on deep learning for USGIF's 2016 State of GEOINT Report. Chang said

intelligence agencies still have a long road ahead before bringing the benefits of machine learning to analysts and warfighters.

"We need to find a faster way to transition from the innovative to the everyday," Chang said. He added that decision-makers, traditionally risk-averse, need to be "more willing in this cognitive world to make mistakes and learn from them."

Implementation challenges are considerable. As academia pursues ways for non-technical people to use data science tools in the workplace, industry leaders try to embed AI into the workflow in ways that are seamless, relevant, and actionable.

"If you present something that's overly complex or not timely, it's the analytic version of a tree falling in the woods," McCue said. "The end user needs to be able to understand the output, and you have to be able to analyze data in motion. You can't say, 'Just put the investigation on hold for three days while I work on these algorithms.'"

She added that because of the long-established "I'll know it when I see it" culture among analysts, it's important to incorporate AI slowly.

"Some of the models are so complicated, you have to just look at the metrics and say it works," McCue said. "That's not something analysts are comfortable with right away. You can earn the trust of the analyst by incorporating a little machine learning and then building up to some of the more computationally intensive capabilities."

At USGIF's April Data Analytics Forum, Skip McCormick, a CIA data scientist, said even the best algorithms don't stand a chance without ground truth against which to measure.

"If you don't have some sort of ground truth to measure your forecast against, you have no way to know whether your algorithm is any good and no way to convince an analyst they should pay attention to it," McCormick said.

Despite the cultural and logistical challenges, McCue said deep learning in the GEOINT space is a critical next step for the IC—and the time is right.

"Enormous content and high-performance computing and algorithms have all come together. The timing is beautiful," she said. "It's game-changing." ■

A MACHINE
may discover
a city in Syria
has new
fortifications,
or it may
determine
that a runway
has doubled
in size, but it
can't tell us
if either was
built by friend
or foe.
