

BRIEFING

IF 3816 August 2, 2017

almost anything bold about the

Who Owns the Future? (2013)

-Jaron Lanier

future, and you will almost certainly sound ridiculous to someone.

ARTIFICIAL INTELLIGENCE, NOT YOUR AVERAGE ALGORITHM: WHAT AI IS AND WHAT AI IS NOT

CONTEXT & DYNAMICS

"It didn't play Go as a human does." That is what the world champion of the game Go said after losing to the artificial intelligence (AI) application called AlphaGo, which we discussed last year (IF 3712). But since then, this mystery of AI applications and other less mysterious AI software have been providing critical insights to more and more industries and services. In that advancement, however, things are being called artificial intelligence that are actually advanced automation. To get a better understanding of this curious software approach, we look at several lessons learned from the AlphaGo experience: Artificial intelligence applications (1) program themselves; (2) are not just basic algorithms; (3) can be inscrutable; (4) are very narrow in range for now; and (5) have created a "Sputnik moment" for some world leaders. Money is flowing into this arena, and services and products are starting to emerge.

OPPORTUNITIES

• Disease diagnosis could become a more exact process, leading to longer life expectancy and better treatment results.

Say

- More and more efficiencies across industries as well as higher and higher levels of sophistication in the decision-making tree will expand profit margins.
- Companies providing AI applications, microprocessors and other pieces of the system benefit.

RISKS

- Advanced artificial intelligence could eliminate higher and higher levels of jobs, creating, as some have warned, a jobless (and, thus, salaryless) future.
- Given that their thinking processes cannot be explained, rogue programs could create unanticipated control problems.
- AI could become a particularly effective tool for hackers.
- Weapons operated via artificial intelligence, especially by rogue countries, could lead to warfare of a different sort than is known today.





It's Amazing, But Then....

Last year, DeepMind's AlphaGo defeated the South Korean champion of Go, the ancient East Asian board game with a mind-numbing range of possible moves and countermoves. The complex game has pathways of play that rise to roughly the number 2 followed by 170 zeroes, a figure that approaches the number of atoms in the observable universe. It can lead to nearly 400 billion unique games. Lee Se-Dol, the top Korean player, explained his amazement at how AlphaGo played at one point: "It was impossible to make such a move. It didn't play Go as a human does." Then this past May, DeepMind took its AlphaGo software to China, where it defeated the world champion of Go, Ke Jie, who added to Lee's assessment: "I was shocked. Many of the moves can never happen in human competition." (South China Post, 5/13/17; New York Times, 3/14/16; Discover, 4/17)

AlphaGo is an example of artificial intelligence, a machine-learning approach to software – that is, the machine learns from experience, changes its own program to make adjustments to a new reality presented by the data and produces results that can mystify those using it, much as AlphaGo surprised and mystified the Go champions. Several points can be adduced about artificial intelligence from the AlphaGo example:

◆ Artificial intelligence operates closer to how humans think and changes its own software, or way of operating, based on experience, which makes it possible to function in a changing environment.

◆ Artificial intelligence is not a typical computer

program and is more than a traditional algorithm.

♦ Artificial intelligence can be inscrutable.

♦ AI applications, for now, are skill-specific. Or, said another way, AlphaGo cannot play chess, drive a car or use medical technology, although artificial-intelligence applications have "It was impossible to make such a move. It didn't play Go as a human does."

♦ When Alpha Go defeated the Chinese world champion, the match was blocked from Chinese television, and the loss gave Beijing, as one Chinese scholar intoned, a "Sputnik moment."

To get a handle on how expansive, incredible and surprising artificial intelligence is and how it might develop in the future, we need to look at each of these areas more closely.



"I've been thinking about your bedtime story, and something doesn't make sense."

What Is Artificial Intelligence?

Silicon Valley has turned its considerable attention and resources to artificial intelligence, or AI, as it is commonly called. In fact, leaders of large corporations working in the digital field have decided that consumers will eventually make contact with the Internet not through a currently

been written or are being written to do all those things.

compiled mobile device but through some AI-enabled

device. And those market leaders prefer to say that items such as Alexa, Siri and other voice-activated software programs are AI. But are they really?

Researchers at Georgia Tech claim that true artificial intelligence has two attributes: (1) AI learns over time, adapts to changes in the environment and then alters what it does based on the data given it. That is, such software interacts with its environment, and, like a plant or an animal, changes its operating process depending on what it experiences. Said more directly, AI devices learn and change, and in doing so, they rewrite their own programs. (2) AI works in areas that are so complex or challenging that it would take humans a considerable amount of time to learn what to do (as distinguished from automated and repeated tasks). That is, AI applies its core "brain" to learn from experience and change approaches and activities to better deal with what the environment presents, much as a child might learn to alter behavior to learn how to win a game. True AI, according to Charles Isbell, one of the Georgia Tech researchers, exhibits selfgovernance, surprise and novelty. (Atlantic, 3/4/17)

AI applications identify patterns in data, and they typically do that better than humans dealing with the same data, because, if properly programmed, machines look at the data without the biases that humans have. Moreover, machines can find patterns in larger amounts of data than humans can even read through. For instance, a boardcertified dermatologist would likely see 200,000 patients in an entire career. An AI application can read through and assimilate data from 130,000 cases in a relatively short couple days. They can find patterns in those cases, program themselves to identify such patterns or anomalies in new, unseen data and thereby make a diagnosis with

a great deal of history and medical sophistication behind it. (*Adweek*, 6/12/17; *New Yorker*, 4/3/17)

An example of current AI came in 2015, when researchers sought to "teach" a deeplearning AI to diagnose cancer. It showed the device 14,000 images of skin anomalies that dermatologists and pathologists had diagnosed as cancerous. The AI application accumulated this True AI, according to Charles Isbell, one of the Georgia Tech researchers, exhibits self-governance, surprise and novelty.

data and started identifying patterns, critical images and anomalies in the examples, thereby essentially beginning to think in its own way, not in a way a programmer delineated in code. When eventually shown new images of anomalous skin images, the AI application brought to bear all of the patterns, critical imagery and learning that it had experienced and applied its own perspective to diagnosing actual cancers at a rate of accuracy reaching 72 percent. Human dermatologists identified the same cancers at a rate of 66 percent. (*New Yorker*, 4/3/17)



What Artificial Intelligence Is Not?

Traditional computer programming involves coders writing a set of standardized rules for the program to follow. For instance, if the program encounters input X, then it should do operation A, which leads to a particular outcome, Y. A set of these kinds of codes leads to an algorithm, creating a complex operation for a particular outcome. How the algorithm functions and what it does are known to the engineers writing the code. Moreover,

> the program always functions in the same way. These kinds of algorithms are everywhere and make business and technological operations much more efficient. Consider these examples of non-AI algorithms:

> DoNotPay provides an online service that can write a letter for someone wishing to challenge a parking ticket. The program asks the customer a series of questions, and based on the answers writes the letter to authorities.

The company is adding more services, such as writing letters to demand payment for delayed flights and completing forms for federal assistance.

Several newspapers, including the Washington
Post, now have software called Heliograf that can write
© 2017 Inferential Focus

news stories, based on information inputted into it, whether the input is financial news, sports outcomes or vote counts.

 \blacklozenge Lex Machina from LexisNexis can read through millions of court decisions in a rather short period of

time and provide a law firm with the chances of winning, thereby helping the firm decide which cases to pursue. (*Atlantic*, 4/17; *Wired*, 2/17)

These algorithms do not learn to get better at what they do and do not take their "own" point of view on a topic. Rather they respond and behave in accordance with the set of instructions encoded in their programs. Even though these kinds of algorithms can do some heavy lifting in terms of accessing and digesting large amounts of data, they simply follow rules

for how to do it and what to do with it. For that reason,

they are not AI but simply advanced automation.

In AI, a program can be given input X, but rather than perform а programmed operation to lead to a specific outcome, it can try out different possible operations and learn to those operations change over time. It is seeking to realize the best outcome. For instance, AlphaGo could see an opponent's move, try any number of responses

and then learn how that response works over time, and change its future responding move based on that outcome. AlphaGo learned the best outcomes by playing itself in Go several million times, more games and thus more experience than any professional Go player ever could manage to get. Because AI systems access data, teach themselves how to look at the data, change their operation procedures

> based on altered conditions or new data and make assessments based on the new perspectives that it taught itself, knowing how an AI system reached a particular conclusion is typically difficult, if not impossible, to explain, even by those who might have written the initiating software. Because it taught itself over time, its operating procedures cannot be deciphered. AlphaGo made moves that humans had never made, and the engineers who launched the software could not uncover how the AI software decided to make those particular moves.

> For a research project at New York's Mount Sinai Hospital, engineers wrote a deep-learning program called Deep Patient. The program read through the records of 700,000 patients, examining symptoms, diagnoses, treatments and outcomes. After Deep

Patient had aggregated so much data, the hospital started

feeding it data about current patients, and the software would offer a diagnosis and prescriptive suggestions. The AI application became quite good at very early diagnoses of patients with developing cancer. It also became better than human doctors at diagnosing patients who would soon have psychiatric disorders, such as schizophrenia, a diagnosis that human doctors find especially difficult to make. Yet for every diagnosis of schizophrenia, no

one could actually know how Deep Patient reached that conclusion. This new diagnostic tool, which essentially programmed its own way of operating, provided no clues as to how it might have reached its diagnosis. "We can build these models," explained Joel Dudley, team leader for the Mount Sinai experiment, "but we don't know how



number of responses and then

learn how that response works

over time, and change its

future responding move based

on that outcome.

HYPERBOLE

is Going

Artificial Intelligence Can Be Inscrutable

they work." Tommi Jaakkola, a professor at MIT who specializes in machine learning, adds this perspective: "If you had a small neural network, you might be able to understand it. But once it becomes very large, and it has thousands of units per layer [calculations per neuron] and maybe hundreds of layers, then it becomes quite ununderstandable." (*MIT Technology Review*, 5/17)

The United States military would like to use AI to drive vehicles, pilot planes, identify targets and digest large amounts of intelligence data. But if no one can ascertain how some weapon came to a particular conclusion and fired its gun, missile or bomb, the military cannot use it. The Pentagon has made "explainability" a key barrier

to the use of AI in its weapons systems, perspective а moral that other countries' military might or might not follow. At present, the Defense Advanced Research Projects Agency (DARPA) has 13 projects under way to create an explainable AI. Moreover, this

"We can build these models," explained Joel Dudley, leader for the Mount Sinai experiment, "but we don't know how they work."

inscrutable part of AI has made many technology observers, including Elon Musk and Stephen Hawking, nervous about AI and how it could lead to out-of-control robots and software. In addition, the European Union is considering invoking a rule that would require all decisions coming from AI applications be explainable (see **IF 3712**; *MIT Technology Review*, 5/17).



KUPER

"No, I'm not sure I turned the stove off. I'm not sure of anything anymore."

AlphaGo Can't Play Chess

Artificial intelligence in its current iteration can do one thing really well. Today's versions, such as AlphaGo, teach themselves to get better and better at a specific task. Al can play Go better than the world's best, or it can diagnose patients with cancer better than trained doctors or identify images from a variety of sources as accurately as humans. But this "narrow AI" or "weak AI" cannot do multiple kinds of tasks. Oren Etzioni, chief executive of the Allen Institute for Artificial Intelligence, observed,

> "AlphaGo can't even play chess. It can't talk about the game. My six-year-old is smarter than AlphaGo." (*Discover*, 4/17)

> The larger capabilities of AI – that is, capabilities that approximate those of Etzioni's six-year-old or better – come under the terms "strong AI or "general AI." These would displace humans at the top of any organizational chart, and they are, at present, the most sought after applications in the AI world. When Silicon Valley leaders speak of using AI

instead of current mobile software to access the Internet, they are hoping for general AI that would act and react through a digital device the way another human might act and react. The movie *Her* revealed something about how such interaction might take place, leading to a male



character falling in love with the female voice – that is, falling in love with an AI robot. Researchers might not know how AI applications reach decisions, and they definitely do not know how humans will react to general AI when it comes to be.

The Future of AI and China's "Sputnik Moment"

In May, when AlphaGo defeated the Chinese world champion in Go, leaders in China had what one China scholar called a "Sputnik moment," which is to say they had a sudden shock that they were far behind in an area of critical importance to the future of warfare, surveillance, intelligence and general computing capabilities. In response, the Chinese quickly put together a development plan to put the country back in the global

AI game. The new plan would have China get to the same level as leading AI countries by 2020, followed five years later by breakthroughs in AI disciplines that will become "a key impetus for economic transformation." By 2030, China plans to "become the world's premier artificial intelligence innovation center." (*New York Times*, 7/21/17)

Whether China becomes those things in the time frame officials outlined is less important than the current perception in Beijing that AI is critical for any country and economy seeking top-level status in technology. That perspective is already embedded solidly in Silicon Valley leaders, who envision AI as the future of digital communications, computation and assessments. It is also being used in space, where an AI application called Autonomous Exploration for Gathering Increased Science (AEGIS) is being used by *Curiosity*, the rover moving around the planet Mars, to select rocks that scientists back on earth would like to study. A random selection of samples would bring a 24 percent accuracy - meaning 24 percent of those samples are the kinds that scientists would be interested in. AEGIS, which does its own assessment to determine the sample's possible value, is picking samples with a 93 percent accuracy. (Space, 6/29/17)

By 2030, China plans to "become the world's premier artificial intelligence innovation center."

But why is the rush to AI happening now? After all, the concepts and even some of the software have been in existence for decades. Artificial intelligence requires an extremely large amount of computing capability, and speed is important in how that capability is applied. Why is it just now becoming possible to apply that much capability at a pace and cost that works for wide markets? For one thing, data are critical - the more data, the better the decisions - and good data, not fallacious data, have become a reality with the greater accumulation and digitization of resources, such as historical information about patients. Every new piece of information improves the learning part of an AI application. Faulty data, such as skewed racial perspectives, lead to faulty decisions, which happened in a test case in Broward County (FL), where an AI program incorrectly concluded that African-American prisoners were more likely than Caucasian prisons to repeat their crimes upon release. Better and more data have improved AI's abilities. (Discover, 4/17)

The truly short answer as to why now, however, is:

graphic processing units (GPUs). These microprocessors were created to render graphics with efficiency and speed, easily making them critical to video-gaming devices. AI breaks down processes into thousands of smaller issues, each needing its own response. GPU chips have hundreds and recently even thousands of "cores" (*e.g.*, Nvidia's recent processor has 3,584 cores), each one capable of making a calculation, and that compares very favorably to the tradition chip, the

central processing units (CPUs) that have only a few cores (*e.g.*, Intel's server CPU has 28 cores). Ingesting and digesting data for a cancer-screening AI program can require literally trillions of calculations, which done through CPUs would take weeks or even months, but with GPUs, the calculations can be completed in a few days. (*Spectrum*, 5/17; *Economist*, 2/25/17)

When the massive accumulation of data has been completed and the process for identifying, say, cancer has been established, the application can function without the massive system that essentially created it. For instance, an Israeli company, Mobile ODT (*i.e.*, optical detection technologies), has created a clip-on attachment that converts a smartphone and its camera into a device similar to a colposcope used by gynecologists to examine magnified images of a woman's cervix. The clip-on device is being used in health clinics in rural Kenya to make early diagnoses of cervical cancer. The method developed from a massive AI application, but now the process can be completed using a smartphone with an attachment. (*Spectrum*, 5/17)

So company and country interests in artificial intelligence are high, and data and processing power at a decent cost have surfaced. As we wrote in our earlier look at this topic, "Large digital tech companies, such as Google, Facebook, Baidu, Microsoft and IBM, have shifted their research focus to cognitive computing, which includes all AI approaches. Money is flowing and products are emerging" (see **IF 3712**).

Element AI, a startup that helps companies connect with AI experts, raised \$102 million in a series A round and received funds from Intel, Microsoft and Nvidia. Meanwhile, Intel gave funds to Mighty Ai, a startup that uses the general public to help train AI applications, and Microsoft Ventures contributed to a seed round for Agolo, a newly launched company that applies AI to summarize a huge number of articles and books, thereby helping companies reduce the work overload on employees. (*Venture Beat,* 6/17/17)

While AlphaGo received considerable public attention for its victory over humans, it seemed to blend in to the attention that IBM's Watson got for winning Jeopardy! But the victory over Go champions was of a different magnitude of processing. Watson retrieved answers; AlphaGo created them. Companies and countries know that, and as a result, a rush is on to claim AI as an attribute. More and more people are learning that Artificial Intelligence (Is) Not Your Average Algorithm.



Some of our previous looks at this topic:

- **IF 3813** The Cyber Silk Road: Moving Ahead Faster Than The Land And Maritime Silk Roads, 7/6/17
- **IF 3717** The Great Digital Experiment, Part I: Disruptions And Two New Realms Of Digital Experience, 12/5/16
- **IF 3712** From Mobile-First To AI-First: The Digital World's Shift To Artificial Intelligence, 9/1/16

inThought 9/22/16 The Evolutionary Brain And Digital Tech's Next Big Thing: The Human Brain And Virtual Reality – The Twain Shall Meet

inF 1106 Ubiquitous Bots: Smart Bots, Dumb Bots And The Battle For Consumer Time, 5/18/16