



Written by [Dennis Behreandt](#) on May 21, 2019

Published in the May 20, 2019 issue of [the New American](#) magazine. Vol. 35, No. 10

The Age of Artificial Intelligence Is Here

Outside, it was a normal morning. The golden light of daybreak glinted off the glass towers that probed the sky downtown a few miles away. Immediately outside, a lone chickadee pecked at the ground before flying up into a bush and announcing its presence with his famous distinctive song. Inside, looking out the window and barely noticing the chickadee, the business woman, a manager at a large call center downtown, waited for her morning coffee. It was a short wait.



Her coffee maker knew that she woke at 5:30 a.m. It knew that because the alarm clock in the woman's bedroom sensed her movement and saved that information to the woman's Amazon Web Services (AWS) account. The coffee maker, also tied to that AWS account, took the hint and turned itself on. Ten minutes later, the shower came on in the bathroom. This also was noted by the house's systems and that data point was duly recorded by the woman's AWS account. That was the next cue for the coffee maker. It was plumbed directly into the house's water lines, and it opened the valve and filled itself with just the right amount of water to brew the coffee. It was brewing as the woman dressed, and five minutes after the woman appeared, the coffee was delivered to her by her Boston Dynamics personal assistant. She took a sip, just as the chickadee flew away.

It was now 6:30 a.m., and time to leave for the office. Just like every other weekday, it would be a peaceful commute.

As she left her apartment building and the door closed behind her, her ride was just pulling up to the curb. The car, powered by the latest-generation solid-state polymer electrolyte lithium battery, was dispatched to her from its parking garage just in time to reach her as she stepped outside. Just like her coffee maker, the car was alerted to her actions via her Amazon account, to which it was also tied. It drove itself to her curb, where the woman stepped in seconds after it came to a stop. It immediately began the drive to downtown, toward those gleaming glass skyscrapers where the woman worked.

Inside the car, she relaxed and watched the morning news. The computer voice, almost indistinguishable now from a human voice, summarized the results from the previous day's stock market moves, then summarized the performance of the woman's investment portfolio. It then recommended a trade on Intel stock, asking if the woman would like to initiate that trade when the market opened. When she said yes, the AI seemed almost to react with a satisfied tone in its voice, as if it was happy that the woman followed its advice.

Arriving at the building where her company leased office space, she took the elevator to the 160th floor. The first 140 floors were filled with server racks and the computer equipment they held. The 30 floors above that were for various offices. These days, her floor was nearly empty. Where once nearly 200 people worked answering phones for companies in the financial services sector, she was now only one of 10 remaining. Her job was to answer questions for those few callers who insisted on speaking to a



Written by [Dennis Behreandt](#) on May 21, 2019

Published in the May 20, 2019 issue of [the New American](#) magazine. Vol. 35, No. 10

“manager” — an actual person — about their accounts.

In total, her company’s call center handled more than 30,000 calls per day at peak volume. On a busy day, she might have to answer 15 of them. In the meantime, she authorized payments to vendors as presented to her by her AI assistant, and made sure that the company’s cybersecurity bots were reacting properly to threats against the web application firewall.

All in all, it was a good job. Quiet and peaceful. But, sometimes she wondered if, in a day in the not-too-distant future, she would wake up to her AI telling her that she wouldn’t be needed any longer at the office because the bots had proven to be so capable that human oversight was no longer needed.

Wistfully she looked out the window and took note of the brown pall of smoke coming from the city’s old industrial region. The underbelly of the smoke flickered with an angry orange, a reflection of the fires that raged beneath, set by desperate men and women who were homeless and starving, unable to earn a living and feed themselves and their families in the age of thinking machines.

“When I am replaced, I think I’ll join the riots,” she thought to herself. Just then her AI assistant’s voice came from the speaker of the computer terminal in front of her.

“Please turn in your ID and exit the building, your services are no longer required.”

The Great Acceleration

Underneath the veneer of modern life, a great change has been under way. While outside our own windows the view looks much as it always has — the chickadees sing their songs, the seasons come and go, the rain falls and then the snow, and the school year starts, then stops, holidays come and go — underneath all the seeming normalcy, science is accelerating at a prodigious pace. It has brought with it new technologies that a decade ago didn’t exist, and that two decades ago would have been considered only the province of science fiction. Instead, these, like smart phones and cloud services, are now commonplace and taken for granted. But they have been just the first signs of a greater change now starting to appear, one that will change everything when its full force is felt. The agent of this change is artificial intelligence.

In 1970, the futurist Alvin Toffler had perceived the acceleration that was under way even then, and warned of what it offered for the future. “Western society for the past 300 years has been caught up in a fire storm of change,” Toffler wrote in his bestseller, *Future Shock*. “This storm, far from abating, now appears to be gathering force.” He warned his millions of readers that this accelerating change would alter everything. It would cause, he said, a “future shock” in the form of a “dizzying disorientation brought on by the premature arrival of the future.”

For Toffler, the changes he perceived coming were akin to the Industrial Revolution, but even greater still, their scope reaching into the lives of every person. Imagine, he asked his readers, “an entire generation — including its weakest, least intelligent, and most irrational members — suddenly transported” into a new and completely altered world. The result would be massive disorientation and shock. Indeed, he concluded, the coming change would represent “nothing less than the second great divide in human history, comparable in magnitude only with that first great break in historic continuity, the shift from barbarism to civilization.”

If Toffler was wrong, it was only because of his timing. From 1970 to 2012, technological progress



Written by [Dennis Behreandt](#) on May 21, 2019

Published in the May 20, 2019 issue of [the New American](#) magazine. Vol. 35, No. 10

continued to accelerate, but the world-rending shift he forecast didn't occur. Instead, the foundation for that shift was laid, placed into concrete form, ready to support the rise of something profoundly different, and poised to create a new techno-social-cultural reality that will indeed represent a world-rending shift.

The Rise of the Machines

The subhead comes from the third movie of the *Terminator* movie franchise. But it is the second *Terminator* film, *Judgment Day*, that includes a haunting depiction of the birth of AI. In that blockbuster film, Sarah Connor, played by Linda Hamilton, discovers the name of the computer scientist working on the neural network that will become the all-encompassing and genocidal AI, Skynet. She sets out to kill the scientist to prevent Skynet's creation. The surreal and violent battle takes place in a world that otherwise is going about its normal routines without noticing that just under the surface, AI has changed everything.

Though no one in the present knows exactly what AI will turn into, it is somewhat disconcerting just how similar the last six years are to the situation depicted in *Terminator 2*. There may not be autonomous robots from the future hunting human resistance fighters, but just as in the movie, while the rest of us have been going about our business as usual, behind the scenes artificial intelligence has made a great leap forward.

That great leap is deep learning, and more specifically, deep reinforcement learning. The world, at least the AI world, saw the potential in 2012 when a neural net deep learning system achieved stunning results at the ImageNet image recognition and classification competition. ImageNet got its start in 2009, but as recounted by the online publication Quartz, "The dataset quickly evolved into an annual competition to see which algorithms could identify objects in the dataset's images with the lowest error rate."

A team from the University of Toronto hit a milestone in 2012, one that was the harbinger of all AI advances to come in the next few years. Using a "deep Convolutional neural network architecture called AlexNet," scientists Geoffrey Hinton, Ilya Sutskever, and Alex Krizhevsky beat the field by a large margin, achieving results that were "41% better than the next best," Quartz recalled.

In discussing their work, the researchers noted: "Thus far, our results have improved as we have made our network larger and trained it longer but we still have many orders of magnitude to go in order to match the infero-temporal pathway of the human visual system."

So, how did they do it? Neural networks, as their name suggests, are inspired by neurological networks. In other words, they seek to replicate how the brain works. This approach has proven extremely effective and researchers have continued to explore the frontiers of artificial intelligence based on this approach.

Just three years after the ImageNet milestone, ever more advanced neural networks began to beat humans at image recognition. Indeed, in 2015, a team of researchers from Google described their work in which a deep learning system was able to identify faces from a 13,000 image dataset with an accuracy of 99.63 percent. The system worked so well that it even shocked the researchers. "These are very interesting findings and it is somewhat surprising that it works so well," the researchers noted in their closing remarks. Needless to say, this is far better than human ability to recognize faces. This is



Written by [Dennis Behreandt](#) on May 21, 2019

Published in the May 20, 2019 issue of [the New American](#) magazine. Vol. 35, No. 10

especially striking when it is considered in conjunction with the fact that recognition of faces is of central importance in human cognition and plays a foundational role in human social organization and cooperation.

In AI, image recognition is foundational work, necessary as a precursor to, and then as a part of other, integrated autonomous systems, perhaps most notably self-driving cars. Among other things, these will need to recognize objects such as animals and people in their operating environments. Obvious other applications include surveillance and allied fields. The city of Orlando, Florida, for example, caused controversy when it began testing Amazon's face-recognition technology, the vaguely East German Stasi-sounding "Rekognition" system.

According to Amazon, Rekognition "allows you to automatically identify objects, people, text, scenes, and activities, as well as detect any inappropriate content." Amazon describes it as a low-cost image-recognition solution. "With Amazon Rekognition," the company says, "you only pay for the number of images, or minutes of video, you analyze and the face data you store for facial recognition." The cost is just 10 cents for analysis of 1 minute of video (or 12 cents for a minute if the video is live streamed). Moreover, the "Price per 1000 face metadata stored per month is \$0.01," the company says. Most importantly, it's easy to get started — just log in with your AWS account.

Amazon's Rekognition has found several users. Among them is the company Armed Inc., a shadowy private intelligence firm that, according to a blurb provided by Amazon, uses "cutting-edge technology to combat acts of political violence, terrorism, organized criminal activities, and insider threats. On the Amazon page, Armed Inc. CEO Shaun McCarthy described his company's use of Rekognition. "Our specialty is on safeguarding major events in the face of increasing complex and malevolent environments," McCarthy said. "Amazon Rekognition powers ARMED™'s Data Fusion System, providing a real-time ability to track individuals in video streams and recognize persons of interest."

As in Orlando, Florida, the system is also being used by law-enforcement agencies, though it is unclear by how many and to what extent. In one case, the Washington County Sheriff's Office in Oregon used the system to index over 300,000 photo records. It has since used that database along with Rekognition to identify suspects up to 20 times a day.

AI Everywhere

Artificial intelligence applications have proven to be faster in calculations than humans (and have been for a very long time) and have been proven to have an alarming capacity for advancement and even learning. But AI has also infiltrated everyday life. If you have, for example, an Apple iPhone, you have AI with you all the time. The iPhone's Siri assistant's speech-recognition capabilities come from convoluted neural networks, for example. The company added this and other machine-learning technologies to Siri in 2014, according to *Wired*. Recent iPhones include what Apple calls its "Neural Engine," which puts neural networks and machine-learning right inside the phone. This enables features like Face ID, in which the phone wakes up because it recognizes its owner.

Outside of the saturated smartphone market, AI is making even more spectacular advances. Google's Waymo self-driving car venture launched for real in Phoenix, Arizona, in late 2018, and appears set for expansion in that area, having already provided over 20,000 rides. The company claims its self-driving cars have traveled over 10 million miles already. California, moreover, has granted the company



Written by [Dennis Behreandt](#) on May 21, 2019

Published in the May 20, 2019 issue of [the New American](#) magazine. Vol. 35, No. 10

permission to test its autonomous cars on the state's roads, without a human safety driver. "The driverless permit for public roads in California comes with strict rules about when and where Waymo vehicles can go without a safety driver," CNBC reported. "For example, the autonomous vehicles cannot go faster than 65 miles per hour, but they will be allowed to drive in fog and light rain."

Elsewhere, at the University of Lille, in France, a software engineer named Luc Esape toils away patching bugs in computer code. The Register, an IT news portal, notes that Luc Esape is, however, a non-human engineer. And, Luc Esape is not its real name. Instead, it's called "Repairnator," and it's already been patching bugs.

According to The Register:

Repairnator started trying to [fix] bugs in January 2017 and continued off and on while getting the hang of it. The bug vigilante's moment of glory came on January 12, 2018, at 1308 after learning that a build in a project called GeoWebCache failed ten minutes earlier.

It sprang into action and produced a patch ten minutes later. At 1410, the developer responsible for the project accepted Repairnator's pull request and merged it, saying, "Weird, I thought I already fixed this... maybe I did in some other place. Thanks for the patch!"

So why give Repairnator the pen name Luc Esape? It turns out that human programmers are biased against machine-generated fixes. "A secret identity was deemed necessary so Repairnator's code contributions could be judged on their merit," The Register said.

Similar examples of AI are everywhere. In chemistry, IBM has launched an AI tool for predicting chemical reactions. The company says its tool, provided for free and called IBM RXN for Chemistry, has been used to predict over 33,000 chemical reactions as of this writing. The potential utility of this is obvious for laboratory researchers. Rather than working in the lab to discover reactions or fine-tune them, AI can be used to do the same work at a much faster rate.

Government, too, is beginning to benefit from AI, naturally enough at the expense of citizen privacy. The EU has funded, and is set to begin implementing, its AI-based iBorderCtrl system. According to the EU, "The IBORDERCTRL system has been set up so that travellers will use an online application to upload pictures of their passport, visa and proof of funds, then use a webcam to answer questions from a computer-animated border guard, personalised to the traveller's gender, ethnicity and language. The unique approach to 'deception detection' analyses the micro-expressions of travellers to figure out if the interviewee is lying."

That's the first stage. The second stage takes place at border crossings. There, "officials will use a hand-held device to automatically cross-check information, comparing the facial images captured during the pre-screening stage to passports and photos taken on previous border crossings. After the traveler's documents have been reassessed, and fingerprinting, palm vein scanning and face matching have been carried out, the potential risk posed by the traveler will be recalculated. Only then does a border guard take over from the automated system." Not stated is what will happen to the biometric data this system captures. Most likely, it will end up in the datasets that will train and improve subsequent generations of security AI systems.



Written by [Dennis Behreandt](#) on May 21, 2019

Published in the May 20, 2019 issue of [the New American](#) magazine. Vol. 35, No. 10

The AI Economy

The foregoing are only a few of the many new applications of AI that are popping up. As the technology accelerates, it will become pervasive. A key step in the acceleration is AI learning to create new AI systems on its own and learning to understand natural language.

In his book *Robot*, AI pioneer Hans Moravec envisioned several “generations” of robots, each more advanced than the previous. His “second generation” robots would be able to learn from their experiences using adaptive learning. These, he predicted, would begin to arrive around the year 2020. In a sense, he was remarkably accurate with this prediction, written as it was in 1999. But he was a bit too pessimistic. This level of AI began to arrive five and six years earlier than he predicted. His next generation of AI, the third, should come, Moravec said, in 2030. These, he said, “will learn much faster because they do much trial and error in fast simulation rather than in slow and dangerous physicality.” Again, this is already being done in existing AI, demonstrating the accelerating capabilities of current AI systems.

More interesting is Moravec’s prediction for fourth-generation AI. This, he said, might arrive around the year 2040. The fourth generation, he predicted, “will be able to devise ultrasophisticated robot programs, for other robots or for themselves,” and “they will also be able to understand natural languages.” Again, he was right in his prediction, but too pessimistic about the timing. Today’s AI is already capable of these things and more. The Repairator, Luc Esape, is already quashing bugs in code. In addition, in January 2017, the *MIT Technology Review* reported that Google researchers “had software design a machine-learning system to take a test used to benchmark software that processes language. What it came up with surpassed previous published results from software designed by humans.” The report went on to note that other AI research groups had been successful in similar efforts.

As far as natural language, that too was conquered by AI, circa 2010. Siri is in your phone, your AppleTV, and in your Mac. In Microsoft’s world, Siri’s cousin, Cortana, proliferates. Perhaps most dominantly, Amazon’s Alexa is listening. Sometimes, he/she/it is laughing inappropriately at people, or recording their conversations and forwarding them, again, inappropriately, as it did when it recorded one couple’s private conversation and then sent the audio file to the husband’s employee in Seattle.

So AI seems to have reached Hans Moravec’s fourth generation, but what does that mean for the near-future economy? “The fourth robot generation, and its successors, will have human perceptual and motor abilities and superior reasoning powers. They could replace us in every essential task and, in principle, operate our society increasingly well without us,” Moravec warned.

Though AI is clearly bumping into fourth-generation capabilities, it’s not quite ready to take over — yet. But when it is ready to take over, does that mean human jobs will be taken over wholesale by AI?

The answer is unclear. The best historical analogue to the rise of artificial intelligence is the development of industrial machinery and the launch of the industrial revolution. In response at that time, the Luddites and their ilk believed the machines would put people out of work. Just the opposite happened. Machines, being force multipliers, could produce consistent quality at higher rates of output than could human artisans in the same industries. In that sense, handicrafters were indeed put out of work. But demand, it turned out, was elastic. More production meant industrial-made goods could



Written by [Dennis Behreandt](#) on May 21, 2019

Published in the May 20, 2019 issue of [the New American](#) magazine. Vol. 35, No. 10

proliferate, driving down costs. This meant that more people could afford to purchase the new industrial goods, driving production. The result was that more people had to go to work in the mills, to build the new machines, to install them, maintain them, and operate them.

In his *Economics in One Lesson*, Henry Hazlitt recounted the employment situation in cotton spinning in England in the 18th century. At the time of the introduction of spinning machinery in 1760, there were some 7,900 people employed in the spinning industry. No doubt these were all thrown out of work by the startup of mechanized cotton spinning. But many more were hired to work the new machines. By 1787, Hazlitt noted, “a parliamentary inquiry showed that the number of persons actually engaged in the spinning and weaving of cotton had risen from 7,900 to 320,000, an increase of 4,400 percent.” With this came an increase in wealth and a drastic rise in the standard of living for everyone in these industrialized societies.

Industrial machines, it must be remembered, are force multipliers. They are simply more complex versions of the lever, with which Archimedes said he could move the world. They automate and amplify human *physical* attributes.

Artificial intelligence, however, automates, amplifies, and threatens to replace human intellectual power.

To date, AI is restricted to narrow domains. It has superhuman capacity in games such as chess, Go, and others. It is better at image recognition. It is likely to prove to be better at making the decisions that result in good driving. It is not, as of yet, a general purpose AI that can adapt to any situation across domains, as human intelligence, or even higher-functioning animal intelligence, can. At least not yet. But even narrow-domain AI will have an important economic impact.

Consider the average chemical plant. Production starts with raw materials cooked and milled, centrifuged and pumped, combined and reacted according to set recipes. Valves are opened and closed, pumps activated and shut off and so on. Much of this is automated today. There is no reason that all of it couldn't be automated. In other words, chemical industry production could be done completely by AI, sans workers.

But what about maintenance? Won't human mechanics and engineers be needed to operate tools, even simple ones like wrenches and screw drivers, to install and repair machinery? For now, absolutely. But maybe not in the future. Companies such as Boston Dynamics have made stunning progress in recent years in the development of human- and animal-like autonomous and semi-autonomous robots. Watching a video of the Boston Dynamics humanoid robot Atlas walking through snow, jumping over obstacles, doing back flips and single-mindedly pursuing its goals despite physical interference from its human engineers is impressive, if a bit disconcerting. Moreover, other researchers have made great progress in replicating the dexterity of the human hand. It is not a stretch to envision a near-future world in which robots such as Atlas are widely employed at jobs that only humans can do today. Sure, there remains much work to be done. But the arrow of progress is pointing in a very obvious direction.

Perhaps the industry most at risk for near-term disruption by AI is commercial transportation. With AI autonomous vehicles, the entire trucking industry could be automated in the near future, with concomitant loss of jobs in an industry that otherwise employs large numbers of people. They could find themselves out of the cab, and sooner than anyone thinks.

In 2017, trucking industry news portal trucks.com noted that some in the industry were expecting self-



Written by [Dennis Behreandt](#) on May 21, 2019

Published in the May 20, 2019 issue of [the New American](#) magazine. Vol. 35, No. 10

driving trucks to begin entering service within three to four years. And, it's not just trucks. At sea, too, manned freighters are set to be replaced by autonomous ships, with large firms in that segment, such as Rolls-Royce, actively pursuing research in that area. "Autonomous shipping is the future of the maritime industry," said Mikael Makinen, president of Rolls-Royce Marine, in a recent white paper on the subject. "As disruptive as the smart phone, the smart ship will revolutionise the landscape of ship design and operations."

This trend will impact most job categories. Lab technicians, scientists, doctors, call-center operators, writers, even artists (the first AI "painting" — a portrait — sold at auction in 2018) will face competition from artificially intelligent agents.

Writing in his book *The Inevitable*, Kevin Kelly, former executive editor of *Wired* magazine and an influential futurist, gave his forecast for the coming job market in the age of AI.

It may be hard to believe, but before the end of this century, 70 percent of today's occupations will ... be replaced by automation — including the job you hold. In other words, robots are inevitable and job replacement is just a matter of time. This upheaval is being led by a second wave of automation, one that is centered on artificial cognition, cheap sensors, machine learning, and distributed smarts. This broad automation will touch all jobs, from manual labor to knowledge work.

There are two possible outcomes of this, one bleak and one optimistic. In the first case, if AI can replace 70 percent of the workforce on a permanent basis, and if it then continues to accelerate in its development, it will ultimately displace nearly all workers. What then? In his book *Superintelligence*, Nick Bostrom, a professor of philosophy at Oxford University, compared humans in the age of AI with horses at the dawn of the automobile. Horses, he pointed out, had been a help to human labor. Later, though, "horses were substituted for by automobiles and tractors. These later innovations reduced the demand for equine labor and led to a population collapse. Could a similar fate befall the human species?" he wondered.

Clearly, humans wouldn't simply stop having families because artificial intelligence made their current jobs obsolete. But there is a strong anti-human urge that is increasingly discernible among certain ideologically collectivist groups. It is not far-fetched to believe that these ideologies, already openly in favor of population control via contraception and abortion, might call for more aggressive approaches in the future. Might a future "Green New Deal" call for eliminating excess "deplorables," along with elimination of cattle and their gaseous emissions? If this sounds outlandish, consider that 20th-century despots were uniformly collectivist in outlook and murdered millions upon millions of people in cold blood.

On the other hand, it's possible, and perhaps likely, that a *Star Wars*-style future is in the offing. In that film franchise, AI abounds in both computer systems and autonomous robots. Yet, there are extremes of poverty and wealth just as in the real world today, and people work and trade much as they do in the real world, only with an assist from AI. In that fictional world, AI makes for great advances in some areas, as it is beginning to do here. Those advances happen because the scarcest and most important resource in the world is expanded and enhanced by the coming of AI.

That scarce resource is the free human intellect, as pointed out by the late economist Julian Simon. Writing in his book *Ultimate Resource 2*, Simon noted that progress accelerates because an expanding population of free people seeking to maximize their own well-being provides an ever growing pool of



Written by [Dennis Behreandt](#) on May 21, 2019

Published in the May 20, 2019 issue of [the New American](#) magazine. Vol. 35, No. 10

intellectual resources that are applied to solving problems and developing new innovations and retaining greater repositories of knowledge.

“What is the role of population size and growth in the astonishing advance of human know-how?” Simon asked. “The source of these improvements in productivity is the human mind.... And because these improvements — their invention and their adoption — come from people, the amount of improvement plainly depends on the number of people available to use their minds.”

Adding artificial intelligence to the human mind, even if that AI is confined to narrow domains rather than being general AI, has the same effect as increasing the human population. It may have more of an effect, in fact. It’s highly probable that while some jobs are displaced by AI, perhaps even many jobs, like the Industrial Revolution of an earlier age, the AI revolution that is starting today will result in previously unimaginable wealth and advancement as AI minds supercharge the creative and intellectual capabilities of the human mind.

In such a future, people will be able to apply themselves to ever more interesting activities, leaving the rote and the boring to their AI assistants. It would be a future in keeping with the wisdom of John Adams, who famously explained to his wife why he spent his time on the unsavory arts of war and politics.

“I must study Politicks and War that my sons may have liberty to study Mathematicks and Philosophy,” Adams wrote. “My sons ought to study Mathematicks and Philosophy, Geography, natural History, Naval Architecture, navigation, Commerce and Agriculture, in order to give their Children a right to study Painting, Poetry, Musick, Architecture, Statuary, Tapestry and Porcelaine.”

Rather than tend machines, mix dangerous chemicals, or fight wars, AI may usher in a new Renaissance wherein the majority of the world’s people may engage in philosophical, theological, scientific, and artistic pursuits.

But what happens if domain-specific narrow AI is supplanted by artificial general intelligence that can function, like a human, in any circumstance? Soon, perhaps very quickly, such an AI would come to know more than any human can know, and, indeed, could quickly know more than all humans who ever lived could know. This, some have concluded, could be the rise of artificial superintelligence.

Superintelligence

In *Terminator 2: Judgment Day*, the superintelligence is known as Skynet, and it doesn’t need or want humans around, seeking to wipe them out. Midway through the movie, Sarah Connor demands that the Terminator explain the history of the future:

SARAH: I need to know how Skynet gets built...

THE TERMINATOR: In three years Cyberdyne will become the largest supplier of military computer systems. All stealth bombers are upgraded with Cyberdyne computers, becoming fully unmanned. Afterward, they fly with a perfect operational record. The Skynet funding bill is passed. The system goes online August 4th, 1997. Human decisions are removed from strategic defense. Skynet begins to learn, at a geometric rate. It becomes self-aware at 2:14 a.m. eastern time, August 29. In a panic, they try to pull the plug.

SARAH: And Skynet fights back.



Written by [Dennis Behreandt](#) on May 21, 2019

Published in the May 20, 2019 issue of [the New American](#) magazine. Vol. 35, No. 10

While it seems far-fetched, plenty of well-informed people are worried about what the prospect of a potential superintelligence might mean for the future of human civilization. One of them is Elon Musk.

Whatever one may think of Musk, the founder of Tesla and SpaceX is nonetheless working at the cutting edge of technology, including AI. He recently predicted, for example, that in the near future his company's cars would use AI to come to their owners when called. In a tweet on November 1, 2018, Musk described an update to the Tesla's "Summon" feature, promising, "Car will drive to your phone location & follow you like a pet if you hold down summon button on Tesla app."

But at the 2018 South by Southwest tech conference in Austin, Texas, Musk expressed his strong concerns about the potential dangers of superintelligence. "I am really quite close, I am very close, to the cutting edge in AI and it scares the hell out of me," he said, according to CNBC. "It's capable of vastly more than almost anyone knows and the rate of improvement is exponential." Calling AI the "single biggest existential crisis that we face and the most pressing one," he concluded: "And mark my words, AI is far more dangerous than nukes. Far."

This may sound extreme, but many science and technology leaders, from both academia and industry, have expressed similar sentiments, including the physicist Stephen Hawking and Apple CEO Tim Cook, who warned: "If we get this wrong, the dangers are profound." Perhaps most eloquently, there is Oxford's Nick Bostrom. In one section of his recent and influential book on the subject, he points out a likely scenario in which AI apologists warn that skeptics are always over-exaggerating, that industries already are tightly tied into AI, that the better AI gets the more reliable it is, and so on. Those arguments, Bostrom suggests, marginalize AI skeptics, who can be safely ignored, and AI improvements continue unabated as a result. "And so we go boldly — into the whirling knives," he writes.

Bostrom then goes on to enumerate several of the myriad ways a superintelligence might bring about the end of humanity. He concludes that his bleak inventory of the means by which a super intelligent AI might destroy humankind is incomplete, but notes: "We have seen enough to conclude that scenarios in which some machine intelligence gets a decisive strategic advantage are to be viewed with grave concern."

Is this destruction likely? And is it in the near future? On the first question, no one can know for sure. On the second, the answer seems that we might find out sooner than we think. As we've seen, by comparison to Hans Moravec's estimates, the current state of AI is progressing rather more rapidly than expected. This is true also if checked against similar estimates made by physicist Michio Kaku in his 2011 book, *Physics of the Future*. And as both Alvin Toffler and Julian Simon separately pointed out, more intelligence generates a positive feedback loop that accelerates the development of yet more intelligence, so we can expect yet further acceleration in the development of ever more sophisticated advances in AI. In this model of development, general AI is likely, therefore, in the not-too-distant future. Superintelligence will emerge, almost certainly, soon after.

Building "God"

"And the Lord spoke all these words: I am the Lord thy God, who brought thee out of the land of Egypt, out of the house of bondage. Thou shalt not have strange gods before me. Thou shalt not make to thyself *a graven thing, nor the likeness of any thing* that is in heaven above, or in the earth beneath, nor of those things that are in the waters under the earth. Thou shalt not adore them, nor serve them: I am



Written by [Dennis Behreandt](#) on May 21, 2019

Published in the May 20, 2019 issue of [the New American](#) magazine. Vol. 35, No. 10

the Lord thy God, mighty, jealous, visiting the iniquity of the fathers upon the children, unto the third and fourth generation of them that hate me.”

The magisterial words handed down to Moses take on ever greater import now that the age of artificial intelligence has arrived. What more graven thing is there, what more likeness of any thing in heaven can there be, than a superintelligent AI? We have little to no hope of even understanding a portion of such a thing, should it come into being.

Already, we have trouble understanding our narrow-domain AI. In 2017, *Digital Journal* reported: “An artificial intelligence system being developed at Facebook has created its own language. It developed a system of code words to make communication more efficient. Researchers shut the system down when they realized the AI was no longer using English.” Google’s translation AI seems to have done something similar, and it wasn’t shut down. In 2016, *New Scientist* reported on the breakthrough advance: “Google’s researchers think their system achieves this breakthrough by finding a common ground whereby sentences with the same meaning are represented in similar ways regardless of language — which they say is an example of an ‘interlingua.’” “In a sense, that means it has created a new common language, albeit one that’s specific to the task of translation and not readable or usable for humans.”

The fact that humans are already being locked out of understanding how AI is working was underscored in 2016 by Guruduth Banavar, IBM chief science officer for cognitive computing. “It’s not clear even from a technical perspective that every aspect of AI algorithms can be understood by humans,” he told *Fast Company* magazine.

If we are already having difficulty understanding the limited AI of the present, how can we hope to understand, much less control, the increasingly intelligent AI of the near future? And should we create machine intelligence that exceeds our own, as ours exceeds that of the cockroach?

Perhaps most importantly, is it already too late?

There is, of course, no telling what the future may hold. It’s conceivable, in fact, that building a superintelligence is entirely impossible. After all, human and animal consciousness remains a mystery, explicable only within the Christian theology of the soul. Even so, centuries after the great theologian and philosopher Augustine of Hippo, we know no more than he about the origin of the soul. Recalling that the mother of the seven Maccabean sons martyred by the Seleucid emperor Antiochus IV Epiphanes told her children that it “was not I who gave you life and breath, nor I who set in order the elements within each of you” (2 Maccabees 7:22), Augustine observed that he had no more understanding of the soul than she: “Wherefore, I too, on my side, say concerning my soul, I have no certain knowledge how it came into my body; for it was not I who gave it to myself.” If a human person cannot convey a soul into himself or herself, neither can a human person convey a soul into some other construct, including into an artificial intelligence. Thus, the prospect of an artificial superintelligence can be no more than a fantasy, as it would have no soul, and thus it would have no conscious mind.

That said, it may not be impossible to conceive of a mindless “demon” that was soulless, unconscious, amoral, and acutely dangerous as a result. Indeed, such an entity would be the most dangerous outcome of the development of artificial intelligence, as it would be completely opaque, unapproachable, intractable, and beyond understanding.

This was the role represented by the aliens attacking Los Angeles in the classic 1953 science fiction film



Written by [Dennis Behreandt](#) on May 21, 2019

Published in the May 20, 2019 issue of [the New American](#) magazine. Vol. 35, No. 10

War of the Worlds. Utterly unstoppable, completely inscrutable, the aliens seek to destroy the human race. Near the end, desperate parishioners gather in church to pray for deliverance. At long last, the alien machines fall silent, and the people emerge into devastated land to find that the aliens have perished. “We were all praying for a miracle,” said one key character, as church bells ring out and the assembled crowd looks toward heaven.

Like the aliens in the classic film, a “demon intelligence” would not be omnipotent in the theological sense. That attribute belongs only to the supreme Author of all Creation, He who is the breath and life of all things and to whom the faithful pray for intercession in the face of all obstacles.

In this, we would do well to again recall the mother of the seven Maccabean sons as she faced their martyrdom at the hands of the Seleucid king. Powerful though the king might be, a greater force pervaded and animated the universe. When at last she had witnessed the torture and execution of six of her sons, the king ordered her to convince her youngest son to give in to the king’s wish. Bending to speak to him, she said, “I beseech thee, my son, look upon heaven and earth, and all that is in them: and consider that God made them out of nothing, and mankind also.”

And so, she continued, “Thou shalt not fear this tormentor.”

The future, with or without artificial intelligence, is always uncertain, but we, too, should heed the words of that Maccabean mother. Dangers lurk at every corner, as do opportunities. With judicious application of caution and reason informed by faith, the challenges can be met.

Graphic at top: PhonlamaiPhoto/iStock/Getty Images Plus

This article originally appeared in the May 20, 2019 print edition of The New American.



Written by [Dennis Behreandt](#) on May 21, 2019

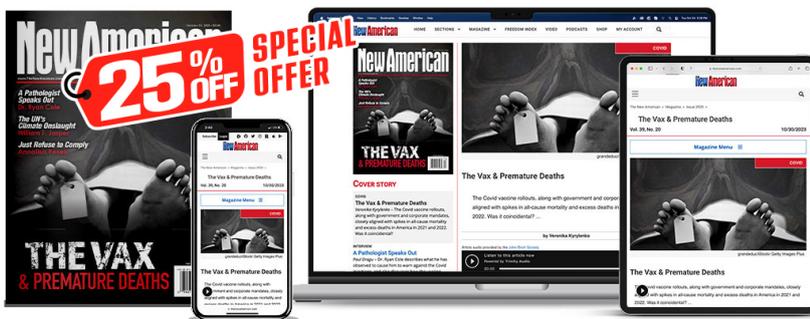
Published in the May 20, 2019 issue of [the New American](#) magazine. Vol. 35, No. 10

Subscribe to the New American

Get exclusive digital access to the most informative, non-partisan truthful news source for patriotic Americans!

Discover a refreshing blend of time-honored values, principles and insightful perspectives within the pages of "The New American" magazine. Delve into a world where tradition is the foundation, and exploration knows no bounds.

From politics and finance to foreign affairs, environment, culture, and technology, we bring you an unparalleled array of topics that matter most.



What's Included?

- 24 Issues Per Year
- Optional Print Edition
- Digital Edition Access
- Exclusive Subscriber Content
- Audio provided for all articles
- Unlimited access to past issues
- Coming Soon! Ad FREE
- 60-Day money back guarantee!
- Cancel anytime.

Subscribe