Man vs machine: The future of AI

Fear of successful cyberattacks meets fear of unintended consequences when machine learning is your first line of defense. Evan Schuman reports.

Fear can be a great motivator. If you are afraid that a human cannot make a decision fast enough to stop a cyberattack, you might opt for an artificial intelligence (AI), machine learning system. But although fear, uncertainty and doubt — the FUD factor — of not responding quickly enough might motivate you to take this action, that same FUD factor that the action your automated system takes might be wrong is an equally strong motivator *not* to employ this technology. Welcome to this year's *Catch 22*.

In the 1983 sci-fi classic *War Games*, a computer was employed to replace the soldiers who manned the intercontinental ballistic missile silos because, it was believed, the computer could launch the missiles dispassionately and not be swayed by indecision in case of a nuclear attack. A teenager hacked the system thinking it was an unreleased video game. Even someone who hasn't seen the film can imagine the plot — the machine starts running World War III scenarios and prepares a multitude of real counter-assaults, driving the military IT experts crazy.

Those are the same fears with machine learning today. Just as in *War Games*, IT can enable today's security software to not only determine if a cyberattack is occurring, but can empower a server to decide on its own to try and halt the attack, often by logging the suspected attacker off of the network or taking more aggressive actions.

The fear among "let the software do its job" opponents is that only humans should decide on an action, with the risks of autonomous software being too great. These are the experts who argued the soldiers should stay in the silos to turn the launch keys. After all, an "attack" might be false.

In fact, that very case occurred some 35 years ago. On Sept. 26, 1983, the Soviet Union's early-warning system detected an incoming missile strike from the United States. Protocol called for a retaliatory strike if such a launch is detected, but Soviet duty officer Stanisav Petrov chose to dismiss the readout as a false alarm despite electronic warnings to the contrary. Petrov was correct — there were no U.S. missiles headed at Moscow. It can be argued that Petrov personally stopped World War III. This is a classic example against using machine learning as part of a missile defense system, where the human element would not have had the opportunity to interpret the data and make a decision.

Man vs Machine

The counterargument among autonomous IT systems advocates is that there is no choice. In short, cyberattacks happen so quickly that

OUR EXPERTS: AI

Bryce Austin, IT security consultant and author
Wade Baker, professor, Virginia Tech College of Business for the MBA and Master of IT programs; member, advisory board, RSA Conference
Douglas Barbin, principal and cybersecurity practice leader, Schellman & Company
Rick Grinnell, founder and managing partner, Glasswing Ventures
James Hendler, director, Institute for Data Exploration and Applications; member, U.S. Homeland Security Science and Technology Advisory Committee
Richard Rushing, CISO, Motorola Mobility
John South, security consultant; former CSO, Heartland Payment Systems
Salvatore Stolfo, professor, Columbia University

$31B

Global revenue from AI for enterprise applications is expected to reach $31.2 billion by 2025

– Statista

only an algorithm's speed is enough to even have a shot at thwarting an attack before substantial damage is done.

"There is an unwillingness on the part of many security people to fully trust machine learning," says Wade Baker, a professor at Virginia Tech's College of Business for the MBA and Master of IT programs; he also serves on the advisory board for the RSA Conference.

"They think 'Only a human can make this decision.' Many have an emotional response," he continues. "There is a strong belief that what we do in the security industry is so hard and so nuanced. A decision needs to be made very, very quickly. There is an emotional kind of irrational thing going on there" and it is compounded by a fear of bad software decisions.

Wade Baker, professor, Virginia Tech College of Business for the MBA and Master of IT programs; member, advisory board, RSA Conference

James Hendler is director of the Institute for Data Exploration and Applications (IDEA) and the Tetherless World Professor of Computer, Web and Cognitive Sciences at Rensselaer Polytechnic Institute and a member of the U.S. Homeland Security Science and Technology Advisory Committee. Hendler agrees that speed is a concern, but if the algorithm is not reliable yet, it is still a legitimate and perhaps an unjustifiable — at this time — risk. "We do have these very fast decisions to make," Hendler says, "but technology is still not at the point where it's trustworthy to say 'let's trust it.'"

> "Technology is still not at the point where it's trustworthy to say 'let's trust it.'"
>
> *– Jim Hendler, director, Institute for Data Exploration and Applications; member, U.S. Homeland Security Science and Technology Advisory Committee*

Richard Rushing, the CISO for Motorola Mobility, says focusing on the nature of attackers — as opposed to attacks — is key to leveraging machine learning properly as a defense tool.

"Let's understand the tradecraft of the attackers. If you look at protection tools, they are set up to block based on data, seen at one time. The attackers figured this out so they change the data every time — kind of like address or ports or information and they usually hide in plain sight," Rushing says.

"What they do not change are things like time, size, process, activity, [and] steps," he continues. With artificial intelligence and machine learning, systems can look for these patterns. "This is what computers are great at doing. You just need to know what to look for but also have that specific visibility to make it happen."

Rushing adds that "layers of detection are bidirectional so you can follow the data in any direction, versus the classic outbound or inbound."

One of the almost universally accepted truths about machine learning is that it is the subject of vast amounts of hype, both from vendors trying to sell it and analysts trying to encourage its use. This buzzword status causes machine learning to be portrayed inaccurately as the ideal fix for almost any security problem, when indeed its value is limited. It is very good at dealing with massive amounts of unstructured data, but its effectiveness quickly dilutes for many other security tasks.

"A lot of folks are trying to throw something like machine learning at a problem

>95%

The best AI systems surpassed human performance accuracy of 95% in 2014

– AI Index, Stanford University

where it's not necessary," says Bryce Austin, an IT security consultant and author of the book *Secure Enough?: 20 Questions on Cybersecurity for Business Owners and Executives*. Many of these companies look to advanced efforts like machine learning when they have yet to tend to routine security matters such as multi-factor authentication, the elimination of default vendor-issued passwords and "reasonable network segmentation," he notes.

Richard Rushing, CISO, Motorola Mobility

Michael Oberlander is the former CISO for Tailored Brands (which owns Men's Warehouse, Jos. A. Bank and Moores Clothing for Men) and author of the book *CISO and Now What? How to Successfully Build Security by Design*. "Machine learning is completely overhyped. I would not spend a dime on it," Oberlander says, adding that demonstrations he saw at BlackHat 2017 — in which the tested machine learning algorithm failed to deliver — convinced him that the technology was not close to ready for the enterprise.

> "Machine learning is completely overhyped. I would not spend a dime on it"
>
> – Michael Oberlander, former CISO for Tailored Brands

But Austin says that the practical security concerns should be paramount. After all, the essence of technology exploration is trying new systems — in a secure sandbox, with no ability to do anything that would impact live systems — and see how well it does.

"We have to allow the machine to make the decision to see how many false positives we get," Austin says. "We need to let the computers try these things in real time."

Rushing's concern is that humans are not perfect. "There is this idea about some crazy bias against machines making decisions. People make mistakes on a regular basis," Rushing says. "Why do machines have to be perfect?"

Rushing argues the pragmatic security position, namely that many mass-attacks today on enterprises are so large and fast that waiting around for a person to make a decision simply can never be an effective defense. "With these attacks, a human could not stop it. They are so quick and affect so many machines so quickly. The only thing that would have saved [the enterprise] is orchestration."

In referencing the 2014 Target breach where attackers used the credentials of a heating, cooling, air-conditioning (HVAC) contractor to gain access to the internal network and ultimately the point-of-sale system network, one of Target's problems was attributed to the massive number of potential breach alerts its system generated, overloading the security staff. Ultimately, the staff overlooked the valid alerts.

"That SIEM (security information and event management) [system] shouldn't be giving me a million events," Rushing says. It should only be alerting true security attacks that merit human attention, Rushing says. "You're going to get overwhelmed because your people are missing the simple stuff."

Using machine learning to reduce the number of alerts dramatically and thereby making real threats more apparent and therefore actionable is an excellent use of the technology, Rushing says.

Machine learning in the SOC

A commonly cited security area where machine learning could do quite well is

59%

Retailers that use AI to personalize the customer experience could boost profitability rates by 59% by 2035

– Boston Consulting Group

post-login authentication. That would be where an attacker would enter the enterprise network with legitimate credentials — presumably stolen credentials — and potentially might even use the legitimate user's hijacked machine and network connection to lend even more credence to the authorization. A variation of this would involve an employee who has legitimate credentials but has chosen to exceed their authorization to engage in unauthorized and improper conduct.

In either case, machine learning could analyze the history of that user's conduct against what the suspected attacker is doing. Although this use of machine learning leverages behavioral analytics, some behaviors are effective at authenticating users before they get into the network, such as the time of day, IP address, details about their machine, number of password attempts, and typing speed. Other behaviors only work after the suspect has accessed the network. These include what files are being examined, how much downloading is happening, how many different areas are being accessed, and the number of files being accessed and viewed.

> "Moving forward, I would like to see machine learning incorporated in to web application scanning, an area untouched by AI today."
>
> *– Doug Barbin, principal and cybersecurity practice leader, Schellman & Company*

Software is the only meaningful way to review all of that data about all logged-in users and determine a potential problem before the user has had the chance to do much damage, proponents agree.

Douglas Barbin is the principal and cybersecurity practice leader at Schellman & Company, a security and privacy compliance assessor. He also considers the vast amount of data that enterprise security teams have to deal with today to be ideal for machine learning analysis.

"For SOCs (security operations centers), the data provided by traditional sources such as IDS (intrusion detection systems), firewalls, and event logs are too voluminous for any analyst team to comb through. Machine learning, in whatever form it takes today or tomorrow, is the only way to support a manageable workload of tickets — the unit of work for a SOC analyst — based on a timely and actionable event," Barbin says in an email interview.

"What is surprising is that the technology vendors, while advertising AI as a capability, still default to signature-based approaches because they can be applied across their customer bases," he continues. "This has led some of the larger enterprises and some leading MSSPs (managed security services providers) to implement their own tools to profile network traffic looking for statistical anomalies. Profiling networking traffic to identify potential deviations such as a spike in a particular port activity such as DNS (domain name servers), could indicate a potential attack before it occurs."

This raises its own practical security concerns. "By the time a signature-based detection signature fires, it is likely too late," Barbin says in an email exchange. "Sure, you can contain and respond in a reasonable manner, but the monitoring tools better support machine learning, which would be monitoring and profiling normal traffic or normal events and it could generate event tickets to instances of statistical deviation that could be profiling or a reconnaissance source that could be blocked prior to it attempting its payload.

"Moving forward," he continues, "I would like to see machine learning incorporated in to web application scanning, an area untouched by AI today. With the advanced

450%

Share of jobs requiring AI has increased 450% since 2013

– Adobe

logic of web applications and API, using machine learning to do some of the human, what-if and/or credentialed type of page traversal would increase the effectiveness of these automated vulnerability assessment technologies. As SOCs begin to take on more proactive security roles, this is a potential area of opportunity."

> "We are in a cybersecurity arms race"
>
> – Bryce Austin, IT security consultant and author

Barbin notes "that the intrusion detection and SIEM vendors focus almost entirely on signature-based approaches because they don't have to customize for a diverse customer base. Because of these limitations, leading enterprises and some MSSPs, who typically rely on these technologies, have developed their own home-grown tools to address an opportunity for incident detection that the technology vendors have not. Obviously, this capability is going to be limited to large enterprises like banks and financial tech that have the resources to be able to build in-house tools."

Bryce Austin, IT security consultant and author

There is, of course, a flipside to this argument. Defenders are not the only security experts using AI; attackers either are currently or soon will be using the technology as well. If so, why concede that advantage to the bad guys?

"The Fortune 1000 often fails to make the assumption that attackers will be using machine learning on their own," Austin says. "We are in a cybersecurity arms race."

Austin takes the argument one step further, that CISOs pushing more regular use of machine learning for security will itself force cyberthieves to do the same for attacking. "We have done very little to raise the cost to the attacker," Austin says. "If they have to use machine learning, great. At least they are having to adapt to us."

Cloud vs on-prem

Columbia University Professor Salvatore Stolfo points out how disconcerting it is that for many venture capitalists today in that any security offering today must either offer machine learning or that entrepreneur need not bother applying for money. "For VCs today, you have to present [it] as machine learning. There's no way that you'll get funding otherwise," he says.

One of those VCs is Rick Grinnell, the founder and managing partner of Glasswing Ventures, an early-stage venture capital firm. He says that he is indeed quite bullish on machine learning and especially sees its security value in dealing with hardware, particularly with the internet of things (IoT).

"Machine learning will help drive the value and usability of these products, enabling the integration and analysis of physical data from devices such as cameras, door locks, RF scanners and motion sensors, devices that today are not easily managed together today," Grinnell says.

"Integrations are typically manual and costly through systems integrators. Over time, the insights gained by combining physical data with cyber data will drive much improved defenses for situations not easily managed today. One simple example would be using camera or RF signal data to determine whether suspicious cyber activity at a computer [or] ATM is actually from a legitimate user."

~37%

The top technology that is expected to impact marketing in 2018 is AI (includes deep learning and machine learning) at 37%. Second place is Big Data at roughly 17%

– Tune

Other considerations are machine learning issues with on-premises versus the cloud, says Columbia's Stolfo. There are two very different kinds of security mass-data issues that machine learning is fine-tuned to process: The first is interactions with your company's environment; the second are all manner of activities across the internet that do not necessarily relate yet to your company, which he refers to as "internet background radiation."

> "The middle market has little to no choice but to depend on cloud machine learning"
>
> – Salvatore Stolfo, professor, Columbia University

Machine learning background radiation primarily involves watching thousands of attacks and attack attempts going on around the world, he says, looking for patterns and methodologies being used now. The intent is to learn from those attacks and to prepare defenses for when those attacks get around to going after your company.

Rick Grinnell, founder and managing partner, Glasswing Ventures

Cloud infrastructure lends itself to this kind of background radiation efforts, which is where cloud-based machine learning offerings can make a lot of sense and create a powerful argument for the technology. It combines an enterprise having its own machine learning, cloud-based or on-premises defenses that watches for attacks against that company, coupled with a vendor's cloud-based defenses that watches what is going on with everyone else.

Threat intelligence feeds could fill in some of the blanks by providing data on threats the cloud provider has yet to experience.

Granted, these are options for Fortune 1000 companies to consider, whereas small- to mid-size businesses are unlikely to have the resources for such a broad strategy. "The middle market has little to no choice but to depend on cloud machine learning," Stolfo says.

Learning to forget

One distinction between how computers normally work and how the human brain works is that, generally speaking, a computer remembers anything stored on disk unless the user deletes the data — effectively making the system forget. With AI, building in the ability to "forget" is a function that programmers are dealing with today.

Some machine learning systems today have the ability to "forget," but one of the big differences between human learning and machine learning is the human ability to forget things *selectively*. Humans can replace old knowledge with new information and make changes in our thought patterns, but machines make changes using different approaches.

Deep neural networks, for example, do not forget the same way humans do. They practice "catastrophic forgetting" — basically they delete everything and start over.

"There are techniques being worked on in the research community that fall into a category called *unlearning*, so that this 'catastrophic forgetting' is not the only form of deletion," Grinnel says.

"These include weighting prior data dynamically depending on the current context — time of day, temperature, customer or user being interacted with,

115

Number of AI start-ups acquired in 2017

– CB Insights

other parameters — so that, for example, an output behavior that would have been used for optimizing an interaction with me on a hot day in the afternoon is different than the same machine interacting with you on a cold day at night," he says in an email exchange.

"You can imagine other scenarios as well. These adaptive methods are still early in their development, but over time will be incorporated into the AI systems that protect us from cyberattacks, or recommend the next movie we should watch on Netflix," he concludes.

The process of machines forgetting data is further complicated by the European Union's new General Data Protection Regulation (GDPR) rules, which define IP addresses as personally identifiable information (PII). GDPR requires that PII cannot be preserved for any longer than is absolutely necessary.

> "These adaptive methods are still early in their development, but over time will be incorporated into the AI systems that protect us from cyberattacks, or recommend the next movie we should watch on Netflix"
>
> *– Rick Grinnell, founder and managing partner, Glasswing Ventures*

"GDPR is going to force research to be performed in this area," Austin says. "Technologies such as blockchain will further complicate the ability to selectively forget information in machine learning systems" as it is just about impossible to delete data from blockchain."

John South, a security consultant who spent seven years as the chief security officer for Heartland Payment Systems, sees the matter differently. As a practical matter, he and his team at Heartland abandoned saving perimeter-defense-captured IP addresses long before GDPR became an issue. And he did that for multiple, pragmatic reasons rather than legal ones.

First, from a security standpoint, it seemed pointless. "It's so easy for bad guys to change IP addresses, so why chase it?" he says.

Second, storing vast amounts of IP addresses started to slow down system performance. "It took longer to do searches; the SIEM took a lot longer to track things down," he says. "If somebody was doing a ping sweep or something like that, we didn't maintain" IP address history. Instead, he made sure to include plenty of trip wires beyond his perimeter to detect naughty conduct.

IP addresses "were relevant for a short period of time and then we had to age them out," he says.

Another practical concern, South says, is that once addresses were detected by tracking groups and alerts were generated, it was often too late as the cyberattackers would become aware of it and change the IP addresses they were using. "By the time they were reported by the intelligence services, [those IP addresses] were already not being used," South says. "It didn't prove to be the best use of our resources or our time."

Incidentally, for those who never saw *War Games*, here's a spoiler so stop reading now. In the end, the computer gathers enough intelligence to determine that war is not the answer. It is unclear at this point if today's AI offerings would reach that same conclusion. ■

For more information about eBooks from SC Media, *please contact Stephen Lawton, special projects editor, at stephen.lawton@haymarketmedia.com.*

If your company is interested in sponsoring an eBook, please contact David Steifman, VP, publisher, at 646-638-6008, or via email at david.steifman@haymarketmedia.com.

42%

Percentage of AI companies acquired since 2013 that have had VC backing

– CB Insights