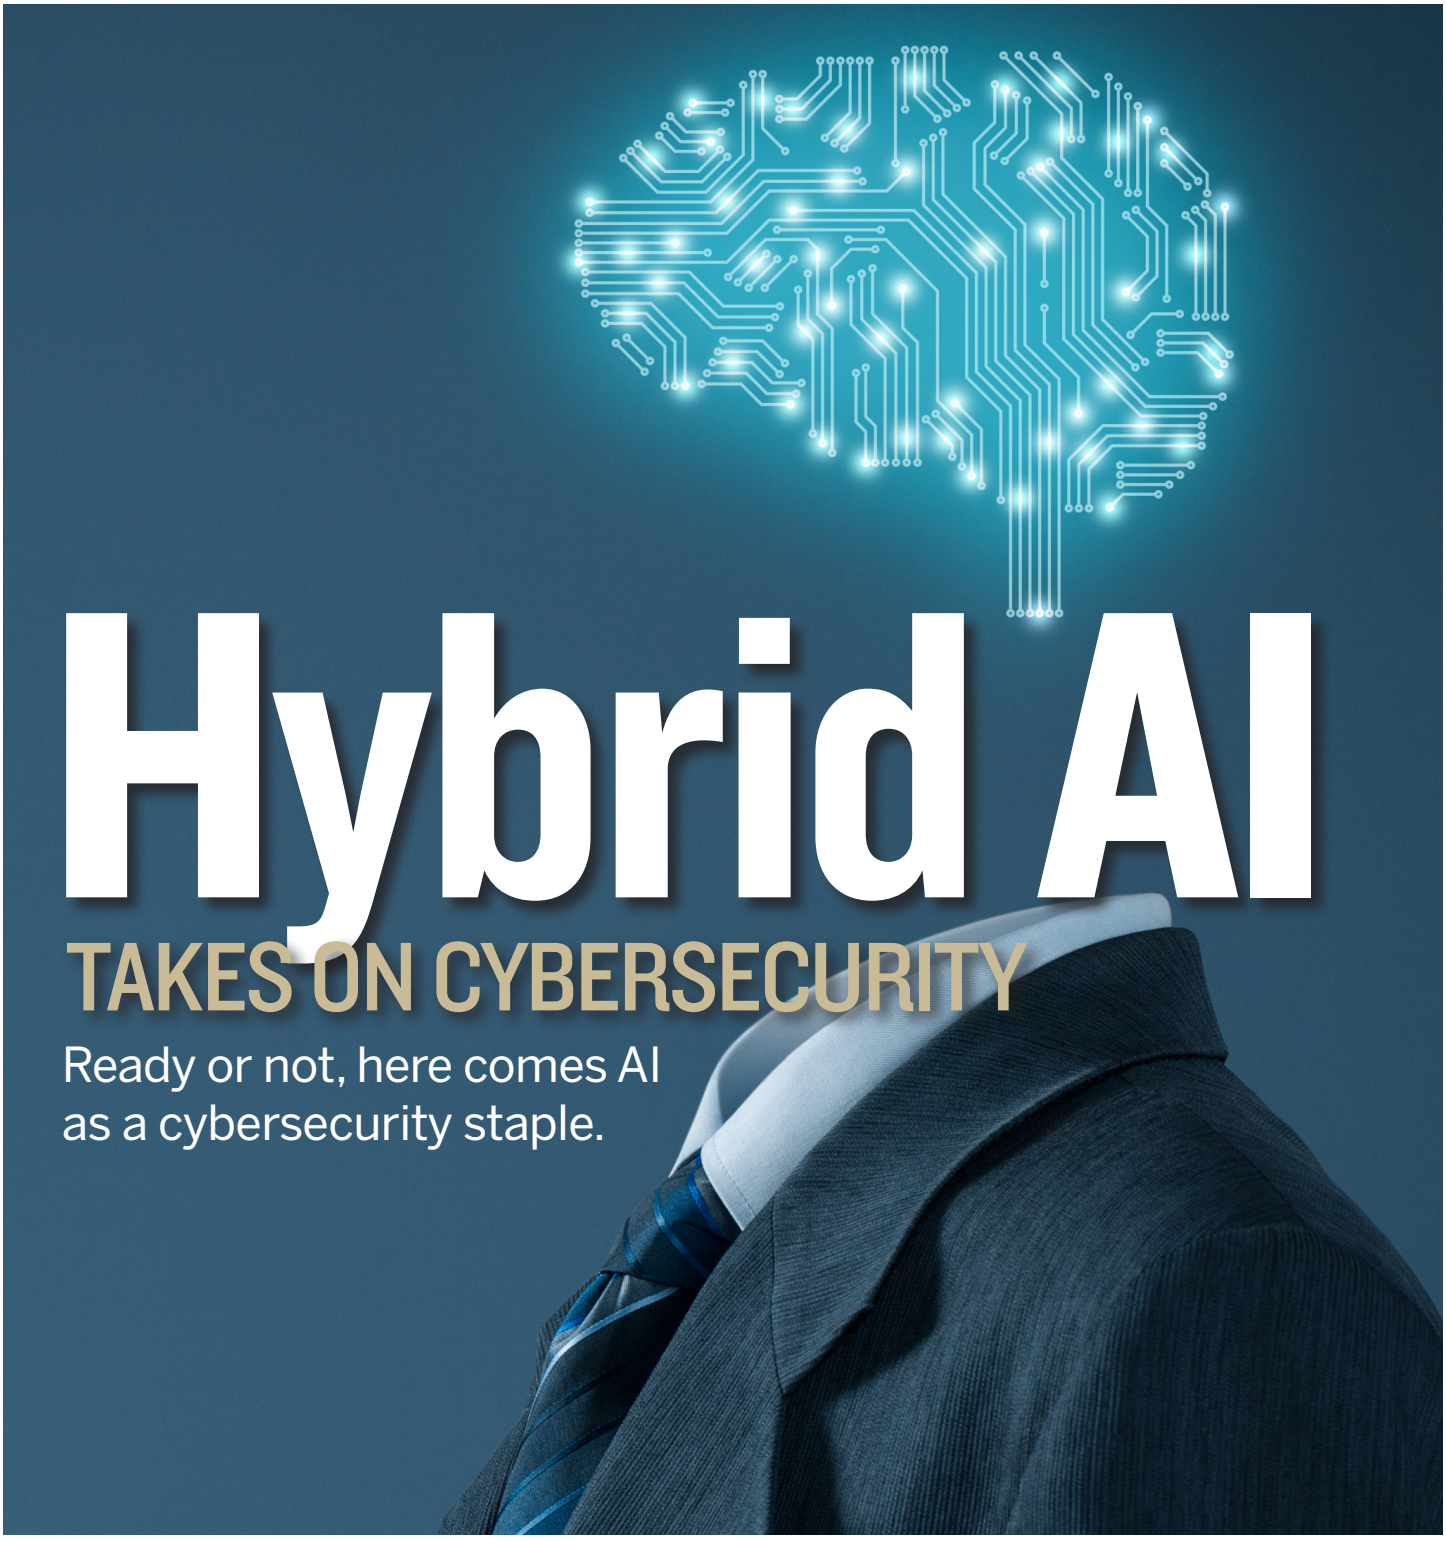

Hybrid AI

TAKES ON CYBERSECURITY

Ready or not, here comes AI as a cybersecurity staple.

special report

An SC Media publication

Sponsored by

ARCTIC WOLF

Cybersecurity analysts: Software versus mammal

Will hybrid AI disrupt your information security team? Experts agree that AI is the next step in identifying threats, but are we facing *Westworld* or just an evolutionary advancement? Evan Schuman reports.

From the perspective of the typical enterprise CFO – locked in eternal corporate combat with the company's CISO – security analysts are overly expensive.

Well, they are expensive when compared with artificial intelligence (AI) security efforts that can spot the same intrusions, note the same pattern deviations and guess which incoming emails are likely phishing scams. But these are similarities, not exact, overlapping functions. The idea of using both AI and human security analysts is well known, but the magic formula for finding the right balance between the digital and the live person, for a specific company, is rarely seen.

The facts underlying this conflict are undisputed: AI is continually improving, getting faster and more accurate; and there is an amorphous quality that humans bring to this battle. Hence, companies need to use both. But although these underlying conditions coexist without much debate, companies sometimes struggle to identify exactly where to draw the line between AI and security analysts. Put another way, how valuable is it to fight a man-in-the-middle attack by using a man-in-the-middle defense?

One can also look at this issue from the bad actor's perspective. Would they rather do digital combat with unpredictable – and creative – humans or with software programs (which is where most cyberthieves are most comfortable)?

"We're always over-thinking the whole issue," says Ed Sim, founding partner at BOLDstart Ventures, a New York-based seed stage venture capital firm. "These humans have to train the model."

He adds that if one splits security into detection and prevention, AI has the edge for detection and humans are better at prevention. "AI can be very good at endpoint security, for keeping stuff out."

Sim also argues that the "software versus mammal" debate is not primarily about saving dollars, although that certainly plays a role. More often than not, enterprises are struggling with finding experienced security analysts at (almost) any price. When that happens, it is less an "AI or person" debate and more an "AI is better than having an empty seat" effort.

Another concern is the privacy element of security: Making sure that confidential data of extreme value to thieves – payment card details, health records, financial reports and the like – is protected. To protect that data, enterprises must first know where it is. That's a critical issue in security today, with such sensitive personally identifiable information (PII) data hiding in departmental clouds, personal and mobile devices, and a 50 GB thumb drive in someone's shirt pocket.

OUR EXPERTS: Hybrid AI

Rob Enderle, principal analyst and founder, Enderle Group
Dan Faggella, CEO, TechEmergence
Kasian Franks, deep technology entrepreneur-in-residence advisor, Propel(x)
Travis Kellerman, technology consultant and serial entrepreneur, Kellerman.biz
Narayan Makaram, senior director for product marketing, Arctic Wolf Networks
Irfan Saif, principal, Deloitte & Touche's Cyber Risk Services practice
Ed Sim, founding partner, BOLDstart

Hybrid AI

85%

Percentage of cyber attacks correctly predicted by a virtual AI analyst.

– Massachusetts Institute of Technology

For hybrid AI security efforts, Sim says the holy grail is PII. "They crawl all of your systems."

Although software won't likely find data hiding in unattached places – such as the cloud, mobile or thumb drives – it will likely do a better job than humans at unearthing it within your networks. Software ignores logic, looking in places where there's no rational reason to store a customer's Social Security numbers and, often, finding it.

To state the obvious, accuracy is crucial for hybrid AI systems. If the software is too lenient, too much malicious activity gets in. If the software is too strict, too much legitimate activity gets incorrectly flagged as problematic. Either way, that just creates more work for the human security analyst and potentially forces the company to hire more human security analysts, which is the polar opposite of the argument favoring return-on-investment on an AI implementation.

Irfan Saif, principal, Deloitte & Touche's Cyber Risk Services practice

Kasian Franks, the deep technology entrepreneur-in-residence (EIR) adviser at Propel(x), a venture capital and private equity firm, says that AI systems do run the risk of messing up the signal-to-noise ratio of security analysts. "You tell it to look at normal activity and train it to assign a probability to a potential anomaly and to then send this data off to QA [quality assurance] people for an alert," Franks says. "They are going to get a lot of noise, but you set thresholds for only high probability."

Finding the right balance between software and human does not merely involve knowing the kind of attacker at issue or just the kind of attack, it also involves the nature of what is being protected.

Irfan Saif, a principal in Deloitte & Touche's Cyber Risk Services practice, says executives should consider what part of the company is being defended and the germane security implications. For example, are the areas under attack traditional IT protected resources – PII, intellectual property – or is the target the operational side of the company, with industrial control and internet of things (IoT) issues surrounding manufacturing, supply chain and facilities assets?

"This forces you to think about security and cyber risk issues a little differently," Saif says. "Are there proprietary embedded systems or smart connected devices?"

Is the attacker planning on staging a distributed denial of service (DDoS) attack, shutting down websites? Is the target PII such as payment card data or health records? Is the goal sabotage or quietly planning a trojan horse to gather data for years before it is used?

Boiled down to its most simple level, AI generally is best for dealing with known attack methods and humans are generally best for dealing with the unknown. Although true, and some would accurately say such a summary is more simplistic than simple, it belies the complexities involved.

There is an apt analogy between writing code and writing stories. It has been said many times that all English stories take the same 180,000 or so English words and simply re-arrange them. Similarly, it can be argued that all unknown attacks are little more than known attacks that have been rearranged to varying degrees. To the extent that is true, a well-written AI program has a better chance of detecting the similar patterns – or even the deviations from expected patterns – than an exhausted human.

Software doesn't become less effective at the end of a long shift, but human security analysts certainly do.

Saif sees software as having three distinct

.32%

PayPal, which uses AI for cybersecurity, saw just a tiny rise in cyber-fraud in 2015.

– MIT Technology Review, January 2016

advantages over humans: consistency (it does not get tired or cut corners), accuracy (it should detect precisely what it is told to detect – and detect every instance every time), and velocity (no cerebellum can beat the speed of a modern CPU). But, alas, humans are also the programmers of software and they suffer from the same mammal deficiencies. It all depends on "asking the right kinds of questions of the data," Saif says.

"Humans can get distracted, can misinterpret things," he adds, pointing out that

Hybrid AI in the SOC

By its very definition, hybrid AI can't possibly replace security analysts because security analysts are a co-equal component of software in a hybrid AI environment. Indeed, those analysts literally are the *hybrid* in hybrid AI.

But that doesn't mean that companies might not need fewer analysts, which is where the concerns materialize. Narayan Makaram, senior director for product marketing at Arctic Wolf Networks, argues that many SMBs, as opposed to their enterprise counterparts, are struggling to find enough security analysts to hire at anything close to what they can afford. Hence, outsourcing some of those security tasks from humans to software might not even be a debate for many companies.

So if AI is going to co-exist with security analysts, where does it make sense to draw the line? Where is AI the most effective and where is the analyst most efficient – or perhaps vice versa?

Hybrid AI is more of a "human-supervised AI," Makaram says. "In some cases, a SOC (security operations center) with hybrid AI can automatically block well-understood threats from repeat offenders, like blacklisted IPs or bad geolocations or malicious websites," Makaram says. "But whenever it comes to advanced zero-day attacks where there is malware that can really hide itself, you need a security expert to run forensics analysis to determine if it's truly a suspicious attachment. And if the suspicious attachment is truly malware, or if the suspicious URL is one where you can download malware, that's where hybrid AI comes in."

The software also does quite well when working off a lists of problematic IP addresses and other known threats. The software can, more easily than a human, "apply threat intelligence to tell if this failed login is coming from a bad IP," Makaram says. "Is this bad IP anomalous behavior coming from a bad geolocation, or is it the CEO failing a logon five times and the account being locked out?"

"Machine learning refers to the pieces that can be automated," Makaram says. "Threat intelligence [is] getting a threat intelligence subscription service where you can really block out bad IP sources or bad DNS requests based on indicators of compromise."

Malware analysis can be automated with a sandbox where one detonates suspicious binaries or attachments to see if they are good or bad, he explains. That can be automated to some extent. With behavior analytics one could automate based on standard deviations and statistical models, where the admin sees normal behavior versus abnormal behavior.

After that, human intervention becomes necessary. "After you get all of these results, you still have unknowns," Makaram says. That's where he says one needs to have leveraged customized rules that include some of this analysis in its equation to really determine what is good and what is bad.

"That's where human intuition, security expertise comes into play," he says. "Leaving all the decisions to machines, when the type of threat is a new threat, is very difficult. That's where you need to apply human talents."

35%

Increase in the annual growth rate of the US economy by 2035 if the US were to absorb AI as a new factor of production.

– Accenture and Frontier Economics

most attacks involve sifting through massive amounts of details. "In an unknown attack scenario, the reality of it is that you're looking at so much data. It's much more likely that a system will detect those kinds of things."

No matter what balance is struck with a company's hybrid AI strategy, humans are critical players. Once the software detects an issue, it will be flagged to the human security analyst. In 2017, no AI system is ready for autonomous functioning – where it is tasked with both finding and stopping all attacks. Like self-driving cars, most still require a human in the driving seat.

What Saif wants enterprise executives to do today is to be open to trying various hybrid AI approaches. "They need to be prepared to experiment," Saif says. "The technologies, the approaches, they're all very new and they are going to need some road miles to get comfortable."

Rob Enderle, principal analyst and founder, Enderle Group

Some suggest that the ultimate hybrid AI environment is one in which the software trains the human – as it detects things first – simultaneously with the human training the software. "People will see the unique cases and will train the machines at the speed of light," says Rob Enderle, principal analyst with the Enderle Group, a Bend, Ore.-based consultancy. "Going after a known exploit, that's where AI will be the strongest. But somebody has to train it."

The idea of the software training itself, known as machine learning, which remains the long-term goal for almost all AI systems, is not viable in 2017. "Machine learning, we're not there yet," Enderle says.

Dan Faggella, the CEO and founder of TechEmergence, a San Francisco-based market research firm, says the point where software can be trusted to make its own decisions and to take action on its own is maybe four to five years from now. Anomaly detection is today, he says.

One argument against AI is that its consistency – and lack of creativity or originality – could be used against it by an attacker. By studying the AI defense tactics, an attacker could guess the precise factors it's seeking and slightly tweak an attack to avoid that precise detail. That might be one scenario where a human security analyst could be superior to AI software, which by its nature is far more rigid.

"If you know what the AI is looking for, it's relatively easy to outsmart them," Enderle says. "They're relatively stupid."

Faggella agrees, saying "AI is so consistent that people can slip around it."

The AI is not going to detect really sophisticated attacks, Franks adds.

In security, deliberate variations are risky and problematic. For example, Google's search engine algorithms are deliberately changed frequently to thwart someone trying to game the system for SEO (search engine optimization) purposes. But in the SEO example, someone getting through gets good search engine placement for a few hours until the system catches up. Little harm actually is done. In security, the stakes are far higher.

Enderle argues that security professionals are likely to be watching AI combatants on both sides of these battles, with AI software powering both the attacks and the defenses. In that situation, AI defenses will likely lose – repeatedly. "We'll probably have AI delivering a more effective attack than doing an effective defense," Enderle says, pointing out that all attacks have a better success rate than defenses. It's a lot easier to attack when you can spend months preparing, he says. Defense must detect and deal with the attack – without any warnings – in milliseconds.

"There's a really good chance that a good

$250B

Estimated size of the global cybersecurity business sector by 2025.

– David Probert presentation at 34th International East/West Security Conference, 2016

AI attack will overcome a good AI defense," Enderle says. "The attacker knows how they're coming. The defender doesn't."

Faggella agrees that it's risky to trust AI systems too much. "Everybody is afraid of the machines making mistakes that fifth graders would never make," he says.

Further, he argues that today's AI system must not only flag potential problems that it spots, but it must also share with the human security analyst all of the specifics that prompted the software to make that flag. "If you're going to flag something as good or bad or whatever, I need to see all of that. I need to feel secure when the security system is making that call," Faggella says.

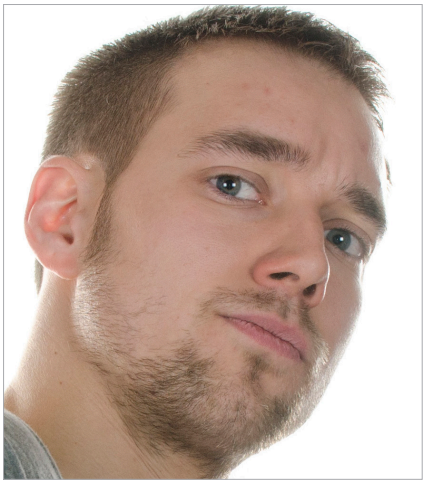

Travis Kellerman, technology consultant and serial entrepreneur, Kellerman.biz

Even letting the software be the sole entity tracking these issues is problematic. "We're sort of letting the machine take the wheel. Very few companies are going to take their hands off the wheel too much," Faggella says. "You don't let it coast. No one else is letting it coast."

Travis Kellerman, an Albuquerque-based technology consultant and serial entrepreneur, says the AI versus human security analyst debate mirrors the bigger picture security situation.

"There is this constant flow and battle," he says. "Someone creates a new machine attack and then the system response. There's this ever-present tug of war." He adds that the biggest area of vulnerability is a company's lack of understanding of the human intention of these attacks. "These AI systems are a short-term fix. Instead of just informing, the AI can handle [the full defense] more efficiently. It can respond and adapt faster."

Done properly, today's hybrid AI systems can deliver the best of both worlds, at least near term. The mammal can have the back of the software and the software can protect the back of the mammal. The inherent weaknesses and deficiencies of each can be compensated by the strengths of the other. Neither is as strong as the sum of both. The question is whether the hybrid AI combo is stronger than the attackers. Maybe and maybe not, but in 2017, some technologists believe it could be a strong contender. ■

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.

$6B

Estimated size of the software market for encrypted systems by 2020 in the cloud, mobile and database markets.

– David Probert presentation at 34th International East/West Security Conference, 2016