How artificial intelligence is changing the face of cyber security
11/11/16

By APSM on Nov 11, 2016 in Cyber Resilience, Editors Desk, Spotlight on Security

Let's inject a virus into the attacking alien spacecraft and save Earth! Let's hack into the enemy mainframe with six keystrokes and abort the torpedo launch! Cybersecurity has long been a staple of science fiction, whether it's in movies like "Independence Day" or television shows like "Star Trek." Yet in our real 21st Century world, artificial intelligence is the new face of cybersecurity, even if it doesn't sound like Hal from "2001: A Space Odyssey."

The most obvious place for added intelligence is to detect whether some pattern of network traffic is benign or hostile. Consider files being sent via email attachment. Do they contain viruses, trojans or other malware? Old-fashioned technology would examine the files for signatures – that is, bits of code that were found in previous attacks or documented from other malware-infected files.

Signature databases, distributed by antivirus/malware researchers and companies, are an imperfect method for several reasons. As the number of malware variants increase, so do the number of signatures – and this takes more time to process. Also, signatures are of limited use against zero-day attacks that nobody has seen before.

What's needed are self-learning systems that can detect malware based on the fact that it's malicious – whether or not it has a signature, or whether or not it's been encountered before. Self-learning systems – also known as machine learning – use AI techniques to foster pattern recognition that's fast and efficient.

Another area where AI is affecting cybersecurity is detecting when hackers are attempting to break through passwords and permissions. Called "authentication attacks," cybercriminals scan networks looking for vulnerabilities, such as devices or servers with no password set or a known default password. AI-based systems can monitor network traffic and detect when a malicious application is scanning the network looking for such vulnerabilities – and raise an alert or launch countermeasures automatically. How does the AI know? Because that type of network traffic doesn't fit the pattern for normal end-user or machine-to-machine usage – and rapidly recognizing patterns is what many types of AI software do best.

In these and other cases, AI gets the job done by applying mathematics to the problem. In fact, one could characterize all the types of artificial intelligence (and there are many) as advanced math. It's not a question of comparing files or examining signatures; the challenge is to solve the math problem. AI offers powerful techniques and algorithms for doing exactly that.

Documentary: How Artificial Intelligence is changing the face of Cyber Security

http://tv.netevents.org/artificial-intelligence-changing-face-cyber-security/

Is it a cat? Is it a virus?

We've all seen that software on the Internet is very, very smart about recognizing images. Facebook can often automatically identify our friends in photographs, and offer to tag them for us. Google's algorithms can identify cat videos with near-perfect accuracy. Those image-recognition algorithms aren't often called AI, but they use the same techniques of machine learning and neural networks to accomplish the same task that an AI-based malware scanner can use to detect a bad file or an authentication attack.

The question for AI-based malware and network scanners is pretty easy: Is the thing (whatever the thing is) safe or not safe? Those scanners are trained by "showing" neural networks or other self-learning systems many examples of safe and unsafe things. Once the algorithms are trained, it's ready to work in the real world, and can make a judgment call very quickly and efficiently: The file is rated as 99% probably safe and only 1% probably unsafe, so permit it to go to the end user. The request for application access is rated as 20% probably safe and 80% probably unsafe, so block it.

Why isn't it a cat photo?

Safe. Unsafe. That's an easy call to make, but sometimes administrators or end users want a little more information. Why did the AI label an Excel spreadsheet as probably unsafe? Why was that application access blocked? Why was the JavaScript on that Web page seen as almost certainly malicious? It's difficult to get answers from the machine learning systems.

Think of it as an image recognition system saying "That picture is not of a cat." Why isn't it labeled a cat? Well, it didn't look like one. That's the best you can learn. Maybe that's sufficient for a cat picture, but in a business computer network, we need to know more: Why was that file flagged as malware? In those cases, the AI system will flag the file, and then additional systems will run forensics to not only verify the initial ruling, but also analyze the malware to gain additional intelligence that can be helpful for detecting and thwarting future attacks.

Nearly every cybersecurity company is researching artificial intelligence to some extent – they can't afford not to. Three companies have taken a significant lead with artificial intelligence, and are staking a claim to be the leaders in leveraging AI for detecting and preventing attacks: Cylance, Javelin Networks and Wedge Networks.

Cylance is golden

Cylance, based in Irvine, Calif., has built what it characterizes as next-generation antivirus based on AI. The company's specialty is endpoint protection: By taking a mathematical approach to malware identification using machine learning techniques instead of reactive signatures, Cylance detects (and blocks) known and unknown malware, viruses and bots, and renders future attack methods useless.

According to Cylance, the core of its malware identification capability is a revolutionary machine learning research platform that uses algorithmic science and artificial intelligence. It analyzes and classifies hundreds of thousands of characteristics per file, breaking them down to an atomic level to discern whether an object is "good" or "bad" in real time. Those characteristics aren't the same as signatures: They are data points that must be examined for each file.

Think again about cat identification: Cats have eyes and ears and noses and a shape of the head and a texture to the fur. But so do dogs and mice and horses. It takes a myriad of data points to distinguish something that's almost certainly a cat from something that's almost certainly not a cat. The same is true with distinguishing a safe PDF from a malicious PDF: There's no one indicator. The AI has to examine lots of data in order to make a confident judgment call.

Stuart McClure, CEO, President & Founder of Cylance said: "We've cracked the code. We've figured out a way to take a very robust set of algorithms and apply them for the very first time to cyber security. These algorithms have already been proven for 30+ years in other industries such as high-frequency trading and insurance – even genome mapping and genome sequencing. These algorithms take in data then train the computer to recognise patterns and predict the future. It gave us a rare window of opportunity to take this brand-new AI application and fly with it as people were so fed up with the existing antivirus solutions."

Cylance's mathematical approach stops the execution of harmful code regardless of having prior knowledge or employing an unknown obfuscation technique. No other anti-malware product compares to the accuracy, ease of management and effectiveness of its product, says the company — especially those that rely upon comparison against traditional signature files.

Tossing the Javelin

Palo Alto, Calif.-based Javelin Networks is focused on authentication attacks, where a cybercriminal has somehow managed to penetrate a business network – and is poised to leverage its foothold to steal corporate assets. What Javelin does is use AI-based pattern recognition to instantly detect that the network has been breached, typically by recognizing activities only performed by attackers, such as scanning the network for vulnerabilities, such as easy-to-comprise user accounts, applications and servers.

When an attack is detected, Javelin leaps into action and performs several actions simultaneously.

- It raises an alert, telling corporate security professionals and IT that the attack is happening.
- It secretly isolates the compromised endpoint (such as an end-user's laptop) from accessing real corporate assets – but in a way that the attacker won't detect.
- And then it creates a fictitious but real-looking network universe for the attacker, trapping it in a maze of hundreds or thousands of appealing yet simulated resources.

While the cybercriminal attempts to further penetrate those simulated resources, looking for intellectual property, network passwords, and other goodies, Javelin plays the hacker like a salmon on a fishing line. Javelin causes the attacker to waste time while forensic tools attempt to gather intelligence on the cybercriminals themselves.

Greg Fitzgerald, COO and CMO of Javelin said "Javelin is a next generation intrusion detection technology, adopting the approach where it lets the attacker reveal themselves through their actions of what they do on a compromised machine. Javelin is changing the game on how one finds an attacker in an organisation and stops them, basically taking the assumption that an attack will happen and a machine will become compromised regardless of whether it is malware or malware less. It saves a tremendous amount of resources, both financial and human, in the collection of the data and the analysis of the information and it speeds up the time upon which an attacker is found. Common research today says it takes somewhere between 150 to 200 days to identify an attacker. Javelin has been able to prove that it can take it from months to minutes."

Javelin Networks explains that targeted attacks are all based on the attackers' knowledge of the intended victims' internal network topology. Javelin takes this knowledge away from them by masking the entire topology. The moment the attackers act on the masked topology, they are caught red-handed by the AI-based pattern recognition software – and after that, its game over for the bad guys.

Driving a Wedge into cyberattackers

Wedge Networks, based in Calgary, Alberta, focuses on preventing malware and cyberattacks from coming into a corporate network (or a customer's home) by scanning Internet and cloud traffic. The company uses many mathematical techniques to detect and prevent network users from ever seeing viruses and other malware, including ransomware. Because customers know and trust signature-based security, Wedge's cloud-based security platform employs high-speed signature scanners. However, that's only the beginning: Wedge also uses AI technology to speed the malware-recognition process, and has partnered with Cylance to apply Cylance's AI-based endpoint security system to work as part of Wedge's network-based security solution.

Wedge Advanced Malware Blocker orchestrates Cylance's AI technology and other technologies to detect and block viruses and advanced malware at the network level, to prevent them from entering enterprise networks. Orchestrating Wedge's hyper-inspection technology with Cylance's machine learning engine and threat analytics, WedgeAMB provides a critically needed break-through in malware prevention, says the company.

Explains Wedge, threats are blocked in real-time, eliminating the cost, disruption, effort and embarrassment associated with tackling threats after they've penetrated the network. It also provides real-time visibility of the network-wide threat landscape, empowering security personnel to identify and focus on the most critical risks first.

James Hamilton, CEO of Wedge Networks said: "We have a new product and it combines traditional cybersecurity policy which tends to be signature-based or heuristics-based with some of the new AI machine-learning applications. The combination of the two has proven to be very popular and very powerful. With machine learning you can understand the motivation of malware, what it looks like and what it acts like. You don't need to know exactly what it is, but you know what the attributes are. If you can learn what's going on at the endpoint, you can understand the malware's behaviour and block it."

WedgeAMB includes WedgeIQ, an automated threat intelligence engine that rolls up threat event data from enterprise-wide AMB systems to characterize, correlate, analyze and visualize the network-wide threat landscape, says the company. The comprehensive threat analytics resource provides actionable threat intelligence to further mitigate evolving threats in real-time.

Open the network bay doors, Hal

The number of threats is multiplying at a frightening rate. According to [PandaLabs](), more than 84 million new malware samples were collected over the course of 2015 – that's 230,000 new malware samples per day. That's right: A quarter of a million new malware strains detected every single day. The rate is increasing. There's no way that older technologies like signature-based scanning can keep track. Fighting malware and cybercrime requires intelligence – artificial intelligence. Every cybersecurity company is exploring AI as an essential tool in helping customers protect their networks and endpoints; Cylance, Javelin Networks and Wedge Networks are ahead of the curve. That's good news for their customers – and bad news for the cybercriminals.