Bloomberg

January 3, 2019 | bloomberg.com

Technology

Artificial Intelligence vs. the Hackers

Machine-learning algorithms watch hackers' behavior and adapt to their evolving tactics.

by Dina Bass

January 3, 2019, 5:00 AM EST *Updated on January 3, 2019, 10:22 AM EST*

Last year, Microsoft Corp.'s Azure security team detected suspicious activity in the cloud computing usage of a large retailer: One of the company's administrators, who usually logs on from New York, was trying to gain entry from Romania. And no, the admin wasn't on vacation. A hacker had broken in.

Microsoft quickly alerted its customer, and the attack was foiled before the intruder got too far.

Chalk one up to a new generation of artificially intelligent software that adapts to hackers' constantly evolving tactics. Microsoft, Alphabet Inc.'s Google, Amazon.com Inc. and various startups are moving away from solely using older "rules-based" technology designed to respond to specific kinds of intrusion and deploying machine-learning algorithms that crunch massive amounts of data on logins, behavior and previous attacks to ferret out and stop hackers.

"Machine learning is a very powerful technique for security—it's dynamic, while rules-based systems are very rigid," says Dawn Song, a professor at the University of California at Berkeley's Artificial Intelligence Research Lab. "It's a very manual intensive process to change them, whereas machine learning is automated, dynamic and you can retrain it easily."

Hackers are themselves famously adaptable, of course, so they too could harness machine learning to create fresh mischief and overwhelm the new defenses. For example, they could figure out how companies train their systems and use the data to evade or corrupt the algorithms. The big cloud services companies are painfully aware that the foe is a moving target but argue that the new technology will help tilt the balance in favor of the good guys.

"We will see an improved ability to identify threats earlier in the attack cycle and thereby reduce the total amount of damage and more quickly restore systems to a desirable state," says Amazon Chief Information Security Officer Stephen Schmidt. He acknowledges that it's impossible to stop all intrusions but says his industry will "get incrementally better at protecting systems and make it incrementally harder for attackers."

Before machine learning, security teams used blunter instruments. For example, if someone based at headquarters tried to log in from an unfamiliar locale, they were barred entry. Or spam emails featuring various misspellings of the word "Viagra" were blocked. Such systems often work.

But they also flag lots of legitimate users—as anyone prevented from using their credit card while on vacation knows. A Microsoft system designed to protect customers from fake logins had a 2.8 percent rate of false positives, according to Azure Chief Technology Officer Mark Russinovich. That might not sound like much but was deemed unacceptable since Microsoft's larger customers can generate billions of logins.

To do a better job of figuring out who is legit and who isn't, Microsoft technology learns from the data of each company using it, customizing security to that client's typical online behavior and history. Since rolling out the service, the company has managed to bring down the false positive rate to .001 percent. This is the system that outed the intruder in Romania.

Training these security algorithms falls to people like Ram Shankar Siva Kumar, a Microsoft manager who goes by the title of Data Cowboy. Siva Kumar joined Microsoft six years ago from Carnegie Mellon after accepting a second-round interview because his sister was a fan of "Grey's Anatomy," the medical drama set in Seattle. He manages a team of about 18 engineers who develop the machine learning algorithms and then make sure they're smart and fast enough to thwart hackers and work seamlessly with the software systems of companies paying big bucks for Microsoft cloud services.

Siva Kumar is one of the people who gets the call when the algorithms detect an attack. He has been woken in the middle of the night, only to discover that Microsoft's in-house "red team" of hackers were responsible. (They bought him cake to compensate for lost sleep.

The challenge is daunting. Millions of people log into Google's Gmail each day alone. "The amount of data we need to look at to make sure whether this is you or an impostor keeps growing at a rate that is too large for humans to write rules one by one," says Mark Risher, a product management director who helps prevent attacks on Google's customers.

Google now checks for security breaches even after a user has logged in, which comes in handy to nab hackers who initially look like real users. With machine learning able to analyze many different pieces of data, catching unauthorized logins is no longer a matter of a single yes or no. Rather, Google monitors various aspects of behavior throughout a user's session. Someone who looks legit initially may later exhibit signs they are not who they say they are, letting Google's software boot them out with enough time to prevent further damage.

Besides using machine learning to secure their own networks and cloud services, Amazon and Microsoft are providing the technology

Siva Kumar, a.k.a. the Data Cowboy

to customers. Amazon's Macie service uses machine learning to find sensitive data amid corporate info from customers like Netflix and then watches who is accessing it and when, alerting the company to suspicious activity. Amazon's GuardDuty monitors customers' systems for malicious or unauthorized activity. Many times the service discovers employees doing things they shouldn't–such as mining Bitcoin at work.

Dutch insurance company NN Group NV uses Microsoft's Advanced Threat Protection to manage access to its 27,000 workers and close partners, while keeping everyone else out. Earlier this year, Wilco Jansen, the company's manager of workplace services, showed employees a new feature in Microsoft's Office cloud software that blocks so-called CxO spamming, whereby spammers pose as a senior executive and instruct the receiver to transfer funds or share personal information.

Ninety minutes after the demonstration, the security operations center called to report that someone had tried that exact attack on NN Group's CEO. "We were like 'oh, this feature could already have prevented this from happening,'" Jansen says. "We need to be on constant alert, and these tools help us see things that we cannot manually follow."

Machine learning security systems don't work in all instances, particularly when there is insufficient data to train them. And researchers and companies worry constantly that they can be exploited by hackers.

For example, they could mimic users' activity to foil algorithms that screen for typical behavior. Or hackers could tamper with the data used to train the algorithms and warp it for their own ends–so-called poisoning. That's why it's so important for companies to keep their algorithmic criteria secret and change the formulas regularly, says Battista Biggio, a professor at the University of Cagliari's Pattern Recognition and Applications Lab in Sardinia, Italy.

So far, these threats feature more in research papers than real life. But that's likely to change As Biggio wrote in a paper last year: "Security is an arms race, and the security of machine learning and pattern recognition systems is not an exception."

(Updates with Amazon's Macie service. A previous version of this story corrected the false positive rate in ninth paragraph.)

Posted from *Bloomberg.com*, January 3, 2019, copyright by Bloomberg L.P. with all rights reserved.
This reprint implies no endorsement, either tacit or expressed, of any company, product, service or investment opportunity.
#C96199 Managed by The YGS Group, 800.290.5460. For more information visit www.theYGSgroup.com/content.