Technology | DOI:10.1145/2898969 Samuel Greengard

Cybersecurity Gets Smart

Researchers aim to apply artificial intelligence and machine-learning methods to take cybersecurity to a new, higher, and better level.

OVER THE LAST decade, cybersecurity has careened into the danger zone. Corporations, governments and individuals increasingly have found themselves staring down the barrel of hacks, attacks, and breaches. What is more, as the intensity and frequency of these incidents grows—and far more sophisticated malware, social engineering, and state-sponsored cyberterrorism takes shape—researchers and other security experts are desperately searching for new and better ways to address threats. The traditional approach of using signature-based malware detection, heuristics, and tools such as firewalls and data loss prevention (DLP) simply is not getting the job done.

"We have entered a different era. As society digitizes everything of value, we create irresistible targets for people who want to engage in criminal activities," observes Bruce Daley, principal analyst at technology research and consulting firm Tractica and author of *Where Data Is Wealth* (Play Technologies, 2015). Today, Daley says, "We're seeing a level of ingenuity and sophistication from criminals that places even the most modern and sophisticated IT systems at risk."

Adds Duen Horng Chau, an assistant professor in the College of Computing at the Georgia Institute of Technology, "Traditional security methods aren't keeping up with cyberthieves. Organizations are constantly forced to react and respond to attacks."

As a result, researchers are now developing systems that tap artificial intelligence (AI) and machine learning to create radically different, far more sophisticated defense models. These techniques revolve around technologies such as big data, pattern mapping and matching, cognitive computing, and deep learning methods that simulate the way the human mind works. The goal, quite simply, is to better identify suspicious patterns and behavior, and build security frameworks that are more resilient and adaptable—often on their own. Says Chau, "Malware and threats have been around since the invention of computers, and they aren't going away. But a variety of new techniques offer hope that we can take a more effective approach."

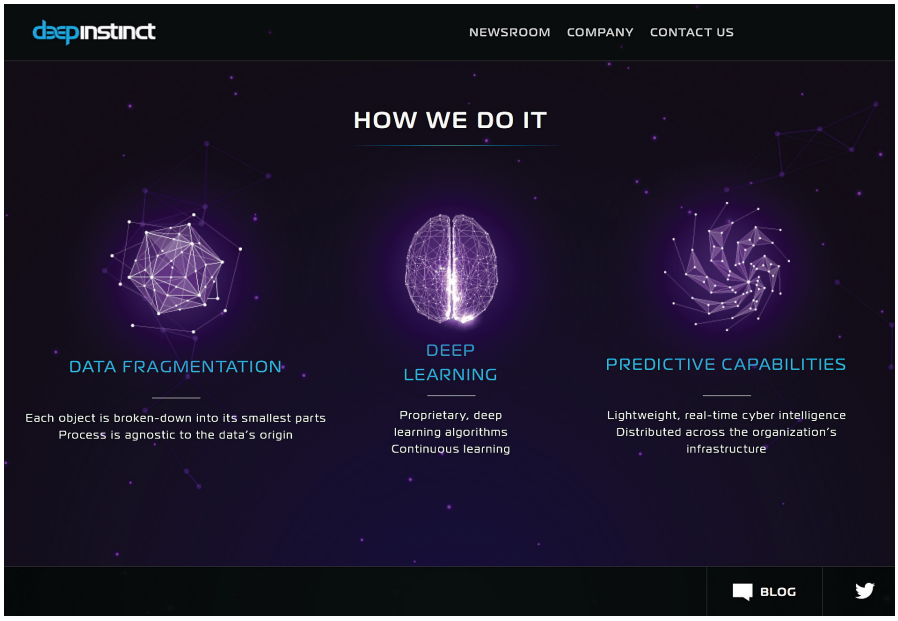

SCREENSHOT FROM DEEP INSTINCT WEBSITE

Israel's Deep Instinct describes how its security software utilizes an artificial neural network.

Protection Schemes

The scope of today's cybersecurity problem is somewhat mind-boggling. Security breaches due to malware cost victims more than $500 billion each year worldwide, according to a 2014 study by market research firm IDC and the University of Singapore.

The problem is growing worse by the day. Ponemon Institute's 2015 Cyber-Crime Study found that the cost of digital crime rose by 19% in the last year, and the average annual loss to companies worldwide is $7.7 million. "The number of files that a security company must deal with now exceeds hundreds of thousands every day and it's continuing to grow," Chau says. "Manual approaches and signature-based approaches are no longer effective."

Adding to the problem is the growing prevalence of zero-day attacks (those that strike and spread immediately), polymorphous malware (which can alter its code to escape detection), viruses, Trojan horses that can hide in systems and graphics processing units (GPUs) for months or years, and malware code that can detect when it has been captured in a sandbox. Meanwhile, firewalls have become less effective as cloud computing and APIs string together data across enterprise boundaries.

Says David Garlan, a professor at Carnegie Mellon University (CMU), "It has become very difficult to pick out malicious events from all the non-ma-

licious events." This has contributed to an uptick in missed detections, as well as false positives. Likewise, other cybersecurity threats ranging from social engineering techniques like phishing and spear phishing (which target an individual with a realistic-looking message or file) to botnet attacks have become more difficult to pinpoint and block from the network because they use cloaking techniques and alias IP addresses.

Researchers are now taking aim at the problem and attempting to move beyond the traditional cat-and-mouse game that includes whitelists, blacklists, and basic code matching. Enormous advances in computing power, far more sophisticated analytics capabilities, and the emergence of cognitive computing and deep learning, are introducing opportunities to spot malware and attack patterns instantly and adapt systems dynamically.

Cognitive computing scans files and data using techniques such as natural language processing (NLP) to analyze code and data on a continuous basis. As a result, it is better able to build, maintain, and update algorithms that better detect cyberattacks, including Advanced Persistent Threats (APTs) that rely on long, slow, continuous probing at an almost-imperceptible level in order to carry out a cyberattack.

Cybersecurity threats have become more difficult to pinpoint and block from the network because they use cloaking techniques and alias IP addresses.

Deep learning attempts to mimic the human brain and use neural networks to take data and pattern recognition to a level that is practically beyond human grasp. It, too, taps machine learning to update algorithms in an unsupervised or semi-supervised way. The technique is now used for speech and pattern recognition tools that mimic human sight and hearing.

One company at the vanguard of AI is Tel Aviv, Israel-based Deep Instinct, which has introduced security software that uses an artificial neural network (ANN) to digest huge volumes of data and put it to use quickly and effectively. "The vast majority of new malware consists of very small mutations in comparison to known malware; by some estimates, there's less than a 2% difference in code," says Eli David, Deep Instinct's chief technology officer and a lecturer in machine learning at Bar-Ilan University in Tel Aviv. However, that is enough to throw off most conventional malware detection tools, including those that take aim at zero-day attacks and APTs. For example, the 2014 attack on Sony was a simple and relatively unsophisticated attack, David notes. "It's just that the malware was new and, as a result, the existing security solutions did not detect it."

Deep learning strives to reduce, if not eliminate, the endless cycle of manually updating signatures in response to the latest permutation. It examines patterns by training on millions of critical parameters and elements by clusters of graphical processing units (GPUs). These processors operate faster than central processing units (CPUs) and, thus, are critical for training deep neural networks. The firm then distributes the trained neural network to end users, who can use it to identify threats with a lightweight endpoint software agent that installs on servers, laptops, and mobile devices.

Deep Instinct has tested 61 other

Milestones

Computer Scientist Receives MacArthur "Genius" Fellowship

A computer scientist was among the recipients of the 2015 fellowships from the John D. and Catherine T. MacArthur Foundation. The no-strings fellowships are awarded to "individuals who show exceptional creativity in their work and the prospect for still more in the future."

The winners receive $625,000 over five years, and people cannot apply for the fellowship; the Foundation simply makes its selections.

Among the 24 recipients of MacArthur Fellowships was Christopher Ré, assistant professor of computer science at Stanford University, who is "democratizing big data analytics through theoretical advances in statistics and logic and groundbreaking data-processing applications for solving practical problems."

Ré leveraged his training in databases and knowledge of machine learning to create the DeepDive inference engine, which can analyze data of a kind and at a scale that is beyond the current capabilities of traditional databases. DeepDive analyzes "dark data"—unprocessable data buried in texts, illustrations, images, etc.; it then extracts relationships among entities (such as real-world objects) in the data and infers facts involving those entities. These facts, or assertions, form a knowledge base, which can then be integrated into an existing database. DeepDive has proved to be more accurate than human annotation, and it can be "trained," even by users without computer science expertise, to improve the quality of results through simple domain-specific rules and low-level feedback about correct or erroneous predictions.

Ré is also working to overcome the technical challenge of computing the many possible interpretations of each data item with speed and efficiency. He has made significant advances in Incremental (Stochastic) Gradient Descent (IGD) and has prototyped Hogwild!, which enables data analysis on a multicore machine, and Bismarck, which integrates various analytical tasks into a traditional database system without the need for separate code paths. These improvements have led to the application of DeepDive in a wide range of settings, from scientific laboratories to law enforcement.

After receiving a B.S. from Cornell University and a Ph.D. from the University of Washington at Seattle, Ré spent four years as an assistant professor at the University of Wisconsin at Madison, before joining the faculty of Stanford University.

solutions on datasets of new malware (zero-day) and, according to its internal analysis, Deep Instinct achieved a 98.8% detection rate, compared to 79% for the next-best solution (many of the most famous solutions on the market only identified 20% to 40% of the malware, David says). In addition, the software—which now trains using millions of malware files and data—is more difficult to confuse, once the training process has taken place.

The advantage is clear: "The underlying principle is that you don't have to understand the reason the software works or why it is detecting malware, it just accomplishes the task." Consequently, retraining the ANN can take place in near-real time, as conditions and requirements change. David likens it to someone tossing a ball and another person catching it reflexively; "your brain operates in prediction mode and instinct takes over."

A Healthy Dose of AI

Deep learning is not the only approach that relies on AI and machine-learning methods. Researchers are also exploring a number of other technologies. For example, at Georgia Tech, one initiative revolves around the use of an algorithm that analyzes relationships with peer files using locality-sensitive hashing and graph mining, which clusters risks by probability. The method, now patented and used by Symantec, determines whether a specific file is good or bad. Tests show the approach identifies 99% of benign files and 79% of malicious files a week earlier than other technologies.

Another project relies on touch signature to authenticate a mobile device, but then takes things a step further by constantly comparing the signature with the behavior of the person using the device at any given moment.

At CMU, Garlan and other researchers are pushing the boundaries of AI planning models. "A basic problem with security systems is that you can't engineer every security feature into a system from the start. You simply cannot anticipate everything that's going to go wrong," he says. Consequently, he is studying how to move beyond using AI merely on the detection side and harness it on the repair side.

"Today, repair takes place either at a very low level or a very high level. The low-level mechanisms tend to be myopic because they are local; they often aren't able to diagnose what's really happening because they lack a global perspective. By contrast, human operators are very good at achieving a global perspective on the system, but tend to work slowly and they often aren't particularly reliable for routine tasks."

"We will never move away from the need for human participation. The problem will never completely disappear."

Garlan and other researchers hope to create a middle ground between machine and human that would ultimately lead to autonomous and self-adaptive systems. This means that a security system, when it detects a suspicious event based on code patterns, the time of day, IP address, or numerous other factors, would generate a Captcha or request a two-factor authentication code without involving a human operator.

The ability to view events in a more granular and contextual way could also allow the system to adapt to events and risks in real time. For instance, if a system detects a flurry of attacks or transaction volumes are high, the system might dial up or dial down its level of aggressiveness. This approach could be used for everything from malware detection to overseeing fired and disgruntled employees, and even the inadvertent sharing of sensitive data.

Garlan says the key to moving further into the realm of AI is the ability to process large volumes of data along with continued research on algorithms, machine learning, and neural networks. Many of the answers already exist, he says; it is simply a matter of combining data points, crunching huge volumes of data, and rethinking interfaces to introduce more streamlined and functional machine-human interaction.

"We are going to see systems become smarter and better at analyzing situations through machine learning and AI, including pattern and anomaly detection on the prevention side of the equation, and greater resilience and adaptability and flexibility on the repair side," Garlan says.

In the end, Daley likens the process of advancing security to castle-building during the Middle Ages: basic moats led to wooden stockades and more advanced stone walls with turrets, which led to flaming weapons. In today's world, there is a need for data scientists to frame problems and opportunities in new and different ways and incorporate machine-learning methods that can better adapt to the environment while adding new capabilities on the fly.

Regardless of advances in AI and cybersecurity, however, "We will never completely move away from the need for human participation," says Daley. "The problem will never completely disappear. As long as there are people, there will be cybersecurity risks and threats."

Further Reading

Tamersoy, A., Roundy, K., and Chau, D.H.
Guilt by Association: Large Scale Malware Detection by Mining File-relation Graphs. College of Computing, Georgia Institute of Technology. KDD'14, August 24–27, 2014, New York, NY, USA. http://www.cc.gatech.edu/~dchau/papers/14_kdd_aesop.pdf

Saravanan, P., Clarke, S., Chau, DH., and Zha, H.
LatentGesture: Active User Authentication through Background Touch Analysis. College of Computing, Georgia Institute of Technology. Chinese CHI '14, April 26-27, 2014, Toronto, ON, Canada http://www.cc.gatech.edu/~dchau/papers/LatentGesture.pdf

Morel, B.
Artificial Intelligence and the Future of Cybersecurity. In *Proceedings of the 4th ACM workshop on Security and artificial intelligence* (AISec '11). 2011. ACM, New York, NY, USA, 93-98. DOI. http://dx.doi.org/10.1145/2046684.2046699

Xing Fang; Koceja, N.; Zhan, J.; Dozier, G.; and Dipankar, D.
An Artificial Immune System for Phishing Detection, *Evolutionary Computation* (CEC), 2012 IEEE Congress, vol., no., pp.1-7, 10-15 June 2012. http://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=6256518&isnumber=6252855

Samuel Greengard is an author and journalist based in West Linn, OR.