Demystifying AI in Cybersecurity

A look at the technology, myth vs. reality, and how it benefits the cybersecurity industry

Table of Contents

Introduction

Any time a new technology emerges, it's likely to follow a predictable path in terms of public opinion. It starts with initial interest, moves to intrigue, then on to (sometimes wildly) unreasonable expectations, which are followed by an uproar around early failures, and then disillusionment—all before the technology has had a chance to prove itself.

Although artificial intelligence (AI) has been in use for some time, its entrance into the mainstream cybersecurity discussion followed the same pattern. Depending on whom you ask, artificial intelligence (and two closely related subsets of AI: machine learning and deep learning) could be the greatest thing ever or yesterday's news. Some of us have been offered visions of autonomous cybersecurity where you can simply point AI at the problem and "security happens". Then again, we've read that AI is pure hype—everyone claims they have it, but its real role in security has yet to be determined.

The reality is somewhere in between. If we dig deeper into AI, we can define what it is and, more importantly, what it isn't; discover what it needs to be successful; and determine concrete areas where it can dramatically enhance security.

> ***"As organizations continue to more frequently incorporate AI into their business practices, they need to ensure they know all the AI offerings available and choose based on what would best serve their needs. Not all AI is created equal."***
>
> – Krishna Roy, Senior Analyst, 451 Research

AI: What it is and What it isn't

Artificial intelligence, as a term, isn't especially descriptive, nor even all that accurate. AI isn't artificial in the sense of a cheap imitation of the real thing, and it's also not equivalent to the way human intelligence works. Back in 1950, Alan Turing addressed the question of whether machines could think[i]; six years later during the first academic conference on the topic, John McCarthy (known as the godfather of AI) coined the term artificial intelligence, calling it "the science and engineering of making intelligent machines.[ii]" In other words, he defined it as machines that can perform tasks that exhibit characteristics of human intelligence.

AI means different things to different people. The term is used to talk about everything from basic machine learning all the way up to the Turing test (which is capable of fooling a human into thinking they are interacting with another human) and beyond. But, at its most basic level, AI can be described as a simple compilation of if-then statements. As Google's Chief Decision Scientist Cassie Kozyrkov explains, it's like a smart labeling machine.[iii] Now, labeling itself is fine for situations where the rules are reasonably cut-and-dried, and the machine can be taught to make simple, binary decisions. But, as use cases like cybersecurity become more complex, binary labeling doesn't cut it. You need more depth, which is where more advanced AI concepts like machine learning and deep learning come in.

Machine Learning: the Next Step

Although many people believe machine learning (ML) and AI are synonymous[iv], the reality is that machine learning is a subset of AI based on statistical analysis. Machines take in data and "learn" to recognize patterns on their own, becoming smarter over time through continuous retraining. Dr. Arthur Samuel, a pioneer in ML, defines it as a "field of study that gives computers the ability to learn without being explicitly programmed.[v]" When implemented properly, it can act like a team of threat researchers, who can filter through data and identify interesting patterns; which is a resource that many companies can't afford to hire or retain.

Machine learning has two key requirements: lots of data and lots of expertise. To be effective, the data must be collected over an extended period of time, and the machines processing it must be guided by humans who can continually fine-tune the algorithms and validate the machine's output.

Done right, machine learning algorithms can become increasingly accurate over time through consistent refinement and training. Here, the problem-solving is not done via simple if-then statements; instead, it consists of more advanced techniques that enable the system to derive conclusions from the data itself, once it has been trained on the basics. ML is used in more complex cases such as fraud detection, where it analyzes transaction details in real time and classifies a given transaction as either legitimate or fraudulent. Machine learning can also perform anomaly detection in which it learns normal behavior patterns, then detects anomalous instances. Again, continuous training by experts who can impart their knowledge is key to success.

> ***"To me, machine learning means speed and scale. It's the ability to mine huge data sets quickly to identify problems and patterns, knowledge sequences, and unanticipated behaviors. The power to sift through so much data so quickly can be a real game-changer in security, if you know what to do with it."***
>
> – Hal Lonas, CTO, Webroot

Deep Learning: Putting a Finer Point on it

From simple if-then statements to more complex models in which machines begin to draw their own conclusions, recent advances take AI a step further. The next evolution, deep learning (DL), allows machines to use techniques that more closely simulate human decision-making. The mechanism in use here is multi-layer artificial neural networks along with sophisticated algorithms. The term "deep" refers to the depth or number of layers in a neural network.

DL results in higher accuracy and performs very well on tasks involving unstructured data. Of course, it requires more training time and expertise (and hardware resources) to be able to train, test and rely on DL.

> ***"Deep learning can achieve state-of-the-art performance on sequence-based classification tasks, but it presents different computational challenges. Training a multi-layer artificial neural networks on large data sets that contain huge sequences takes a lot of computing power."***
>
> – Maurice Schmidtler, Lead Scientist, Webroot

Plenty of Promise, for Better or Worse

Interest in AI peaked early, but, as with any new technology, the name became widely used before it was sufficiently understood. "AI" started being used to refer to just about anything smarter than a calculator. In fact, UC Berkeley Professor Michael Jordan calls the current public dialog about AI "an …intellectual wildcard".[vi] Lack of specificity allows anyone to say anything they want about AI. (For example, you may have heard of an electric toothbrush "powered by AI" that simply checks to make sure you are brushing for the proper amount of time, or a supposed AI program that comes up with new recipes for pizza[vii].)

The security industry certainly hasn't been immune to hype around AI. In fact, it's hard to find a security vendor that does not claim to use AI and ML in their solutions. With so much hype, expectations have run high.

Many people assumed AI would quickly lead to "set and forget" security solutions, which would detect and deflect all threats at the onset of deployment, leaving organizations immediately better protected. Obviously that expectation was exaggerated; but AI has, in fact, massively improved cybersecurity in a number of ways.

The Difference AI Makes in Cybersecurity

AI and its associated technologies are a boon when it comes to solving security problems. With the incredibly fast pace of threats, attacks that come from all over the world and morph constantly, and the alert overload that beleaguers even the most seasoned security analysts, it's clear traditional approaches aren't working. We need to harness the power of AI to make really fast, really smart decisions.

> ***"A key benefit of AI is making your workforce more efficient. Machines never get tired or need rest; they work 24x7. They can identify patterns in massive amounts of data or alert streams to help security analysts and IT security professionals predict new emergent behaviors, which helps security itself become more proactive."***
>
> – Cathy Yang, Product Manager, Threat Intelligence Partnerships, Webroot

Recently there has been tremendous progress in applying AI to three security use cases: malware detection, classification and scoring of URLs and IPs, and phishing detection. In each case, ML and DL are working to bring measurable results.

Malware detection

In years past, bad actors would place the same malware on any number of endpoints; it was relatively easy to detect the malware simply by looking for the correct "signature". Today, attackers generally use polymorphic malware: they make small changes to a single instance of malware, so it appears to be different everywhere it is running.

95% of malware is unique to a single PC.[viii]

By making changes to names, encryption keys, hashes, function instructions, or the order of execution, attackers evade detection. To effectively detect and mitigate the impact of polymorphic malware, ML solutions look at data from millions of nodes around the world, performing static and dynamic analysis and monitoring behavior—in real time. Without the power of AI, this would be an impossible feat for human analysts.

Classification and Scoring

High-risk URLs use many techniques to avoid detection, such as living under otherwise benign domains or only offering malicious content to site visitors in certain geographic locations.

1 in 4 malicious URLs are hosted on trusted domains.[ix]

Similarly, IP addresses cycle from malicious to good and back again to bypass security measures. To protect users, organizations need to quickly determine URLs and IPs are good versus the bad right now. While some businesses rely on static lists of potentially malicious IPs and URLs, that approach is not fast enough nor sufficiently accurate. Enter AI-based solutions that perform deep analysis of seemingly benign URLs and IPs to determine their most current and accurate threat level; AI uses that knowledge to inform and improve detection and mitigation solutions.

Phishing detection

Although it's common knowledge in the security industry that phishing is a prevalent attack vector, this threat isn't showing any signs of slowing down. The Anti-Phishing Working Group reports the number of phishing attacks in Q2 2019 was much higher than that seen in the first quarter, and far above the number of attacks reported in the second half of 2018.[x] Millions of phishing and other fraudulent websites trick users into giving up their credentials and personal information. The level of sophistication is impressive, with phishing websites appearing virtually identical to the real thing.

29% of phishing sites use HTTPS to give internet users a false sense of authenticity.[xi]

The short lifespan of these sites also contributes to the difficulty of detection, as most phishing sites are only active for a few hours.[xii] Here, again, the traditional approach of using static lists of IP and web addresses can't possibly be fast or accurate enough to provide real-time detection and protection.

Key Requirements for AI/ML-based Security

The goal of much better cybersecurity through AI and ML is 100% achievable, but only when the following key requirements are met: lots of data, lots of speed, and lots of expertise.

Data

It's easy to say a security product uses AI or ML. But it's not so easy to gather and analyze the huge amounts of data required to detect patterns, spot anomalies, and stay on top of the constantly changing threat environment. Without years' worth of collected data and the expertise of analysts who have been in the trenches for a long time, accuracy can suffer; and false positives and false negatives can render the results of AI or ML less than useful. Security vendors who are just beginning to implement machine learning typically don't have such data at their fingertips unless they're sourcing it from someone else. In many cases, these vendors have not had the benefit of lessons learned over time, so their ability to draw meaningful conclusions may be limited.

Speed

Effectively analyzing huge amounts of data calls for speed and massive computing power. Today, this means doing a lot of the computing in the cloud. Some vendors supplement cloud-based analysis by using a supercomputer center, such as the one at the University of California, San Diego. Given an amount of data that would normally take days to analyze via a standard cloud infrastructure, a supercomputer can perform analysis, make decisions, and provide results in less than an hour. Not every organization has access to supercomputers, but those who do can provide even speedier decision-making.

Expertise

Many people believe AI-based security solutions will make security experts obsolete, but nothing could be further from the truth. Security analysts' expertise is crucial for training systems, overseeing algorithms and models for ML and DL, and validating the output. In fact, while there are approximately 2.8 million professionals working in the cybersecurity field, one study determined an additional 4 million trained workers would be necessary to close the skills gap.[xiii] AI may be the perfect way to overcome the scarcity of security experts, but it by no means negates the need for the human expert in the first place. What AI actually does is amplify the existing security team, while simultaneously freeing trained analysts to focus their attention on outliers and root cause investigation.

The Webroot Perspective: AI Today and Tomorrow

At Webroot, we're harnessing the power of AI, ML, and DL to protect our customers around the world in a matter of minutes[xiv], as we uncover new and emerging threats. Every Webroot deployment—that is: every Webroot-protected consumer or business endpoint, every network using Webroot DNS Protection, and each of the tens of millions of users leveraging Webroot threat intelligence through our security and technology partner integrations—is a sensor sending back threat telemetry to add to the data we have been amassing for more than a decade.

By leveraging the power and speed of the San Diego Supercomputer Center, we can very quickly analyze billions of incoming links, URLs, HTML artifacts, sequences of actions—in short, all the context around any online activity—and discover phishing attacks, uncover malware in real time, differentiate malicious URLs and IP addresses from safe ones, and generate predictive risk scores.

> ***"We use AI to evaluate URLs and IPs and produce a reputation score, derived over time, not just a binary good or bad value. This score incorporates the historical activity of a URL or IP. For example, a URL may be safe right now, but it could have a risky history that tells us the likelihood it will become bad again."***
>
> – Cathy Yang, Product Manager, Threat Intelligence Partnerships, Webroot

More data means more learning for the ML and DL systems. And our data science team continuously updates models and algorithms to ensure accurate real-time results. Each Webroot deployment is also the beneficiary of this analysis, with threat definitions updated in a matter of minutes. That means with each passing moment, Webroot protection technology is smarter than it was before.

A Few Predictions

As we move forward, we expect to continue to see the need for security experts to help train ML and DL systems. In fact, a study by MIT[xv] showed how the "human-in-the-loop" approach to AI outperformed AI alone, and humans alone, in attack detection.

> ***"The key is to leverage the expertise you have (or your vendor has) to work smarter and faster. ML gives us the ability to magnify the ability of humans, merging the abilities of multiple experts into something that is much greater than the sum of the parts. At the same time, by taking work off the table, we can free up seasoned analysts to do more meaningful work."***
>
> – Hal Lonas, CTO, Webroot

We also expect to see an alarming escalation in fakes. Reports of fake websites and fake news have been in the public arena for several years, but fakes are getting harder to spot and the types of content being faked are growing broader. According to Grayson Milbourne, security intelligence director at Webroot, "deep fakes" (i.e. videos and photos that look legitimate but are, in fact, sophisticated graphics and voice and video technology made to look like real content) are likely to grow in prevalence. Whether for political gain, to sway public opinion, or merely for the thrill of it, he warns that the scenario is one to be monitored.

> ***"As we see more deep fakes, ML and DL may be what it takes to differentiate between real, authenticated communication and identities, and fraudulent ones."***
>
> – Grayson Milbourne, Security Intelligence Director, Webroot

Third, because malicious actors are constantly on the lookout for technology that will help their attacks achieve success, we expect them to be big adopters of AI. We have already seen AI being used against security products via automated penetration testing using ML, or injecting corrupt data into algorithms and statistical models to throw off their learning.

Conclusion

AI, like all new technologies, has gone through a cycle of initial excitement, followed by overhyped expectations and finally disappointment and distrust. However, it's worthwhile taking a deeper look at how AI in practice—via machine learning and deep learning—can bring real benefits to security.

When it is done right, based on years of data, the expertise of threat researchers and massive compute power, AI can solve thorny security problems in real time. As we face even more serious issues, such as sophisticated fake videos and increasing use of AI by attackers, ML and DL will become even more necessary in winning the battle.

[i] Turing, Alan M. Computing Machinery and Intelligence, 49 Mind 433-460 (1950). Retrieved from www.csee.umbc.edu/courses/471/papers/turing.pdf.
[ii] Gil Press, Artificial Intelligence Defined as a New Research Discipline: This Week in Tech History, FORBES. Aug. 28, 2016. Retrieved from www.forbes.com/sites/gilpress/2016/08/28/artificial-intelligence-defined-as-a-newresearch-discipline-this-week-in-tech-history/#6913216e6dd1.
[iii] Kozyrkov, Cassie, Decision Intelligence. TNW Conference 2018, June 2018. Retrieved from www.youtube.com/watch?v=iLu9XyZ55ol.
[iv] Kozyrkov, op. cit
[v] www.ibm.com/developerworks/community/blogs/jfp/entry/What_Is_Machine_Learning?lang=en.
[vi] Jordan, Michael. Artificial Intelligence – The Revolution Hasn't Happened Yet. Medium. Apr 2018. Retrieved from medium.com/@mijordan3/artificial-intelligence-the-revolution-hasnt-happened-yet-5e1d5812e1e7.
[vii] www.techworld.com/picture-gallery/tech-innovation/weirdest-uses-of-ai-strange-uses-of-ai-3677707/
[viii] Webroot. 2019 Webroot Threat Report: Mid-year Update. Sept. 2019. Retrieved from www.webroot.com.
[ix] Webroot, op. cit.
[x] Summary – 2nd Quarter 2019. Anti-Phishing Working Group. Sept. 2019. Retrieved from apwg.org/trendsreports.
[xi] Webroot, op. cit.
[xii] Webroot. 2018 Webroot Threat Report. March 2018. Retrieved from www.webroot.com.
[xiii] www.cnbc.com/2019/11/01/jobs-companies-need-cybersecurity-workers-as-attacks-intensify.html
[xiv] Based on internal Webroot data, processing power, and testing.
[xv] Ransbotham, Sam. Justifying Human Involvement in the AI Decision-Making Loop. MIT Sloan Management Review, Oct. 2017. Retrieved from sloanreview.mit.edu/article/justifying-human-involvement-in-the-ai-decision-making-loop.

About Webroot

Webroot, a Carbonite company, harnesses the cloud and artificial intelligence to protect businesses and individuals against cyber threats. We provide endpoint protection, network protection, and security awareness training solutions purpose built for managed service providers and small businesses. Webroot BrightCloud® Threat Intelligence Services are used by market leading companies like Cisco, F5 Networks, Citrix, Aruba, A10 Networks, and more. Leveraging the power of machine learning to protect millions of businesses and individuals, Webroot secures the connected world. Webroot operates globally across North America, Europe, Australia and Asia. Discover Smarter Cybersecurity® solutions atwebroot.com.

World Headquarters
385 Interlocken Crescent
Suite 800
Broomfield, Colorado 80021 USA
+1 800 772 9383

Webroot EMEA
6th floor, Block A
1 George's Quay Plaza
George's Quay, Dublin 2, Ireland
+44 (0) 870 1417 070

Webroot APAC
Suite 1402, Level 14, Tower A
821 Pacific Highway
Chatswood, NSW 2067, Australia
+61 (0) 2 8071 1900