Engineering Original AI for Cyber Physical Security

By Roland Cozzolino, CTO InsightCyber

Visit the website of almost any cybersecurity vendor and you will probably see Artificial Intelligence (AI) cited as a kind of secret sauce in the solution they offer. This makes perfect sense, not only because AI can be so powerful but also because companies need to say they're using AI if they want to appear innovative.

The paradox is that the more people talk about "AI," the less it practically means. The term became fashionable long before most people had a clue what it meant. Many still don't.

The fact is that AI is a sprawling arena of ideas, aspirations, and computing approaches that date back more 60 years as an academic and technological pursuit. Its application to modern life has accelerated dramatically in the last decade or two. AI now underlies every search you perform, every online ad you see, every conversation you have with Siri or Alexa.

It's only natural, then, that it is being employed in cybersecurity. One should not wonder *if* AI is being used, but rather *how*

Dig a little bit, and you'll discover that most cybersecurity vendors aren't using AI for much more than optimizing systems and processes which, at the core, still require humans to do the heavy lifting of analysis and response. For many companies, AI is a supporting cast member, an efficiency-driving supplement to the labor-intensive things they originally designed before AI became widely accessible.

At InsightCyber, AI is not a complement to our approach. AI is the heart of our approach. Here's a look at how we employ AI and why it's unlike anything any other cybersecurity provider offers today.

The impossible challenge

When I struck up a conversation with my old friend Francis Cianfrocca (InsightCyber's CEO) in the parking lot of a Long Island Home Depot in 2017, I had little expertise in cybersecurity. My domain was digital advertising, specifically building large-scale, secure solutions for real time ad buying, which gave me a deep understanding and respect for AI. I'd had a lot of success, I'd recently sold my company, and I was up for another ridiculously hard challenge.

Francis had a vision for protecting the world's critical infrastructure and operational technology (OT), a mission he'd spent several decades obsessing over. He'd devised an incredibly original idea: to continuously monitor connected devices in utilities, factories, banks, hospitals, and transportation systems—pretty much every bit of hardware running the world—and to use AI to make sense of the data, spot the tiny operational anomalies that signal that a cyberattack is either imminent or happening in the early stages, and deliver guidance to remediate issues immediately, with high efficacy and minimal organizational impact.

The challenge instantly appealed to me in multiple ways. First, because I share with Francis a deep concern that the nations and economies of the West are dangerously vulnerable to cyberattacks from well-funded adversaries. You need only look at Stuxnet, NotPetya, and Colonial Pipeline to get an idea of the potential damage that nation-state sponsored cyberattacks can do. Until we build radically better new cybersecurity defenses, it's just a matter of time before we find out.

Beyond this, I was inspired by what Francis had devised—a method for recognizing tiny but significant anomalies in massive environments and responding rapidly and effectively. His technical approach to this was radically different from anything I'd seen. I was soon on board.

I set out to address Francis's vision by taking the algorithms I'd written in the advertising and finance worlds, and literally almost reversing them. In those scenarios, the idea is to find something—for example, the guy in Sheboygan, Wisconsin who wants to buy a video game at 2:00 in the morning—and put the exact right offer in front of him. But here, the task is inverted. Now we're saying: Please, God, tell me nothing is wrong, and nothing needs to be done! And when something is wrong, what's the best way to figure out why? The difference might seem subtle, but as a computing challenge, it is monumental.

We spent nearly three years building and refining the AI that powers our service. From the outset, we had to address the problem that rendered any conventional approach impractical: the sheer, overwhelming volume of data before us. Many AI solutions are built on "knowledge representation," which essentially means understanding everything that exists in an environment or situation through mathematical modeling. When you are dealing with a complex global network that comprises millions of devices that are all generating data, that representation becomes rather substantial. If not impossible.

A new breed of AI

What we ultimately employed was a mix of flavors of what I might refer to as "Traditional AI," which we use in conjunction with one another. Our approach involves Neural Networks and Deep Learning, and includes game theory, video game technology, semantic engines, real-time investigation, and other subtle components. We also employ Self-Learning AI—a brand of Machine Learning—which basically says that as more and more data comes in, the system teaches itself the best way to identify, understand, and respond to a problem.

As we innovated our AI, we spent nearly three years honing and fine-tuning the service. We refined our models so that they can diagnose issues themselves, versus requiring humans to do it. And we had the critical advantage of working with a huge amount of new, real-world data we generated through companies and partners we worked closely with, and that enabled us to train the system at a deep level.

So, even though ours is a new offering, it is remarkably well educated. That's something you can't fake or purchase; it simply takes a lot of time. Now, we can understand exactly what exists in a customer's myriad environments, and we have the best mathematical approaches for things like determining significant anomalies and generating remediation playbooks.

The solution is powerful today, and it will only get better over time—not just for a single customer's discrete requirements, but across industries and geographies. If, for example, we see an anomaly at Site A, then a very similar anomaly with the same heuristics at Site Z, we can assume something may be propagating and then investigate what's happening. When every company is using a solution like this, the odds are significantly enhanced that a broad threat unfolding simultaneously across companies, industries, or geographies is much more likely to be promptly recognized and stopped.

For me, the value of our innovation is in helping to keep the world safe. I hope to look back and say that we played a significant role in devising a fundamentally new way to protect things. If we can prevent the kinds of devastating attacks we're just now starting to see, and if we can put the bad guys forever on the defensive, this will be a quiet success—and may prove to be the most important work I ever do.

About InsightCyber

InsightCyber is on a mission to keep the world's critical infrastructure, supply chains, and manufacturing operations cyber-safe, preventing attacks that can have catastrophic human and economic impact. The company's AI-powered security service continuously monitors an industrial enterprise's environment, providing insight and protection against a wide range of cyberthreats. To learn more, please visit https://insightcyber.com.