paloalto NETWORKS | CYBERSECURITY PARTNER OF CHOICE

A New Era of Cybersecurity with AI: Predictions for 2024

By Dr. May Wang, CTO of IoT Security
Palo Alto Networks

Artificial intelligence (AI) has been table stakes in cybersecurity for several years now, but the broad adoption of Large Language Models (LLMs) made 2023 an especially exciting year. In fact, LLMs have already started transforming the entire landscape of cybersecurity. However, it is also generating unprecedented challenges.

On one hand, LLMs make it easy to process large amounts of information and for everybody to leverage AI. They can provide tremendous efficiency, intelligence, and scalability for managing vulnerabilities, preventing attacks, handling alerts, and responding to incidents.

On the other hand, adversaries can also leverage LLMs to make attacks more efficient, exploit additional vulnerabilities introduced by LLMs, and misuse of LLMs can create more cybersecurity issues such as unintentional data leakage due to the ubiquitous use of AI.

Deployment of LLMs requires a new way of thinking about cybersecurity. It is a lot more dynamic, interactive, and customized. During the days of hardware products, hardware was only changed when it was replaced by the next new version of hardware. In the era of cloud, software could be updated and customer data were collected and analyzed to improve the next version of software, but only when a new version or patch was released.

Now, in the new era of AI, the model used by customers has its own intelligence, can keep learning, and change based on customer usage — to either better serve customers or skew in the wrong direction. Therefore, not only do we need to build security in design – make sure we build secure models and prevent training data from being poisoned — but also continue evaluating and monitoring LLM systems after deployment for their safety, security, and ethics.

Most importantly, we need to have built-in intelligence in our security systems (like instilling the right moral standards in children instead of just regulating their behaviors) so that they can be adaptive to make right and robust judgment calls without drifting away easily by bad inputs.

What have LLMs brought for cybersecurity, good or bad? I will share what we have learned in the past year and my predictions for 2024.

Looking Back on 2023

When I wrote The Future of Machine Learning in Cybersecurity a year ago (before the LLM era), I pointed out three unique challenges for AI in cybersecurity: **accuracy**, **data shortage**, and **lack of ground truth**, as well as three common AI challenges but more severe in cybersecurity: **explainability**, **talent scarcity**, and **AI security**.

Now, a year later after lots of explorations, we identify LLMs' big help in four out of these six areas: data shortage, lack of ground truth, explainability, and talent scarcity. The other two areas, accuracy and AI security, are extremely critical yet still very challenging.

I summarize the biggest advantages of using LLMs in cybersecurity in two areas:

1. Data

Labeled data

Using LLMs has helped us overcome the challenge of not having enough "labeled data".

High-quality labeled data are necessary to make AI models and predictions more accurate and appropriate for cybersecurity use cases. Yet, these data are hard to come by. For example, it is hard to uncover malware samples that allow us to learn about attack data. Organizations that have been breached aren't exactly excited about sharing that information.

LLMs are helpful at gathering initial data and synthesizing data based on existing real data, expanding upon it to generate new data about attack sources, vectors, methods, and intentions, This information is then used to build for new detections without limiting us to field data.

Ground truth

As mentioned in my article a year ago, we don't always have the ground truth in cybersecurity. We can use LLMs to improve ground truth dramatically by finding gaps in our detection and multiple malware databases, reducing False Negative rates, and retraining models frequently.

2. Tools

LLMs are great at making cybersecurity operations easier, more user-friendly, and more actionable. The biggest impact of LLMs on cybersecurity so far is for the Security Operations Center (SOC).

For example, the key capability behind SOC automation with LLM is function calling, which helps translate natural language instructions to API calls that can directly operate SOC. LLMs can also assist security analysts in handling alerts and incident responses much more intelligently and faster. LLMs allow us to integrate sophisticated cybersecurity tools by taking natural language commands directly from the user.

Explainability

Previous Machine Learning models performed well, but could not answer the question of "why?" LLMs have the potential to change the game by explaining the reason with accuracy and confidence, which will fundamentally change threat detection and risk assessment.

LLMs' capability to quickly analyze large amounts of information is helpful in correlating data from different tools: events, logs, malware family names, information from Common Vulnerabilities and Exposures (CVE), and internal and external databases. This will not only help find the root cause of an alert or an incident but also immensely reduce the Mean Time to Resolve (MTTR) for incident management.

Talent scarcity

The cybersecurity industry has a negative unemployment rate. We don't have enough experts, and humans cannot keep up with the massive number of alerts. LLMs reduce the workload of security analysts enormously thanks to LLMs' advantages: assembling and digesting large amounts of information quickly, understanding commands in natural language, breaking them down into necessary steps, and finding the right tools to execute tasks.

From acquiring domain knowledge and data to dissecting new samples and malware, LLMs can help expedite building new detection tools faster and more effectively that allow us to do things automatically from identifying and analyzing new malware to pinpointing bad actors.

We also need to build the right tools for the AI infrastructure so that not everybody has to be a cybersecurity expert or an AI expert to benefit from leveraging AI in cybersecurity.

Three Predictions for 2024

When it comes to the growing use of AI in cybersecurity, it's very clear that we are at the beginning of a new era – the early stage of what's often called "hockey stick" growth. The more we learn about LLMs that allow us to improve our security posture, the better the likelihood we will be ahead of the curve (and our adversaries) in getting the most out of AI.

While I think there are a lot of areas in cybersecurity ripe for discussion about the growing use of AI as a force multiplier to fight complexity and widening attack vectors, three things stand out:

1. Models

AI models will make huge steps forward in the creation of in-depth domain knowledge that is rooted in cybersecurity's needs.

Last year, there was a lot of attention devoted to improving general LLM models. Researchers worked hard to make models more intelligent, faster, and cheaper. However, there exists a huge gap between what these general-purpose models can deliver and what cybersecurity needs.

Specifically, our industry doesn't necessarily need a huge model that can answer questions as diverse as "How to make Eggs Florentine" or "Who discovered America". Instead, cybersecurity needs hyper-accurate models with in-depth domain knowledge of cybersecurity threats, processes, and more.

In cybersecurity, accuracy is mission-critical. For example, we process 75TB+ amount of data every day at Palo Alto Networks from SOCs around the world. Even 0.01% of wrong detection verdicts can be catastrophic. We need high-accuracy AI with a rich security background and knowledge to deliver tailored services focused on customers' security requirements. In other words, these models need to conduct fewer specific tasks but with much higher precision.

Engineers are making great progress in creating models with more vertical-industry and domain-specific knowledge, and I'm confident that a cybersecurity-centric LLM will emerge in 2024.

2. Use Cases

Transformative use cases for LLMs in cybersecurity will emerge. This will make LLMs indispensable for cybersecurity.

In 2023, everybody was super excited about the amazing capabilities of LLMs. People were using that "hammer" to try every single "nail".

In 2024, we will understand that not every use case is the best fit for LLMs. We will have real LLM-enabled cybersecurity products targeted at specific tasks that match well with LLMs' strengths. This will truly increase efficiency, improve productivity, enhance usability, solve real-world issues, and reduce costs for customers.

Imagine being able to read thousands of playbooks for security issues such as configuring endpoint security appliances, troubleshooting performance problems, onboarding new users with proper security credentials and privileges, and breaking down security architectural design on a vendor-by-vendor basis.

LLMs' ability to consume, summarize, analyze, and produce the right information in a scalable and fast way will transform Security Operations Centers and revolutionize how, where, and when to deploy security professionals.

3. AI Security and Safety

In addition to using AI for cybersecurity, how to build secure AI and secure AI usage, without jeopardizing AI models' intelligence, are big topics. There have already been many discussions and great work in this direction. In 2024, real solutions will be deployed, and even though they might be preliminary, they will be steps in the right direction. Also, an intelligent evaluation framework needs to be established to dynamically assess the security and safety of an AI system.

Remember, LLMs are also accessible to bad actors. For example, hackers can easily generate significantly larger numbers of phishing emails at much higher quality using LLMs. They can also leverage LLMs to create brand-new malware. But the industry is acting more collaboratively and strategically in the usage of LLMs, helping us get ahead and stay ahead of the bad guys.

On October 30, 2023, U.S. President Joseph Biden issued an executive order covering the responsible and appropriate use of AI technologies, products, and tools. The purpose of this order touched upon the need for AI vendors to take all necessary steps to ensure their solutions are used for proper applications rather than malicious purposes.

AI security and safety represent a real threat — one that we must take seriously and assume hackers are already engineering to deploy against our defenses. The simple fact that AI models are already in wide use has resulted in a major expansion of attack surfaces and threat vectors.

This is a very dynamic field. AI models are progressing on a daily basis. Even after AI solutions are deployed, the models are constantly evolving and never stay static. Continuous evaluation, monitoring, protection, and improvement are very much needed.

More and more attacks will use AI. As an industry, we must make it a top priority to develop secure AI frameworks. This will require a present-day moonshot involving the collaboration of vendors, corporations, academic institutions, policymakers, regulators — the entire technology ecosystem. This will be a tough one, without question, but I think we all realize how critical a task this is.

Conclusion: The Best Is Yet to Come

In a way, the success of general-purpose AI models like ChatGPT and others have spoiled us in cybersecurity. We all hoped we could build, test, deploy, and continuously improve our LLMs in making them more cybersecurity-centric, only to be reminded that cybersecurity is a very unique, specialized, and tricky area to apply AI. We need to get all four critical aspects right to make it work: data, tools, models, and use cases.

The good news is that we have access to many smart, determined people who have the vision to understand why we must press forward on more precise systems that combine power, intelligence, ease of use, and, perhaps above all else, cybersecurity relevance.

I've been fortunate to work in this space for quite some time, and I never fail to be excited and gratified by the progress my colleagues inside Palo Alto Networks and in the industry around us make every day.

Getting back to the tricky part of being a prognosticator, it's hard to know much about the future with absolute certainty. But I do know these two things:

- 2024 will be a phenomenal year in the utilization of AI in cybersecurity.
- 2024 will pale by comparison to what is yet to come.

Dr. May Wang is the CTO for IoT Security at Palo Alto Networks.

paloalto® NETWORKS

3000 Tannery Way
Santa Clara, CA 95054

Main: +1.408.753.4000
Sales: +1.866.320.4788
Support: +1.866.898.9087

www.paloaltonetworks.com