theNET
BY CLOUDFLARE

Preparing for the future of AI

Leadership perspectives

Q2'2025

Contents

06
Using AI to tell a story with data

12
Balancing data sovereignty and AI

14
Combating shadow AI

18
Can AI find vulnerabilities?

Introduction

Over the last few years, artificial intelligence (AI) went from science fiction to everyday fact. In 2024, 65% of companies surveyed by McKinsey said they were regularly using AI, and 75% predicted generative AI would lead to disruptive change in the coming years.[1]

But the implications of AI's growing use are still being worked out. This guide will help you work through some of the most important and impactful:

- How does AI affect security? Data usage? Government policy?
- How does unsanctioned AI usage impact organizations?
- How will cyber attackers use AI?

Hear from Cloudflare CSO Grant Bourzikas for the C-suite perspective, along with encountering in-depth examinations of the trickiest parts of AI:

- How to balance AI usage with data sovereignty and maintain compliance with applicable regulatory frameworks
- The problem of "shadow AI"
- The possible AI arms race between attackers using AI to find vulnerabilities and security teams using AI to patch them

AI is here to stay, but best-in-breed organizations will know how to adapt and how to do so securely. Read on for insights that will help you navigate some of the biggest risks and most promising use cases for AI.

[1] https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai

Using AI to tell a story with data

By George Portillo
CEO and Co-Founder, Sushidata

Unstructured data presents a massive opportunity for unlocking new customer insights. If organizations can translate that data into easily understood stories about their customers, they can take actions that help increase customer satisfaction and reduce churn.

There is no shortage of unstructured data available. In fact, by some estimates, 90% of all data is unstructured,[1] and it is growing significantly faster than structured data. The text, images, audio, and other types of content that constitute unstructured customer data are multiplying rapidly across a range of platforms, including Discord, Reddit, Slack, and X. Collecting all this data and transforming it into actionable insights can be challenging, but the potential benefits are worth the effort.

When I was a UX engineer at Google, I saw how combining generative AI capabilities with customer data could enhance digital experiences. I was part of an amazing team that worked on marketing webpages, including the cloud.google.com homepage. Some of our projects involved building conversational UIs for sales and support that partially automated Salesforce tickets. With semi-automated support responses, we were able to rapidly address customer issues and keep customers happy.

For most organizations, though, the customer feedback chain is broken — and at Sushidata, we're working to fix it. As a co-founder and CEO of the company, I'm working with our talented team to help organizations deepen customer insights and improve customer experiences by applying AI to unstructured data. Through our company's journey so far, we've learned some important lessons about where to focus efforts and how to make the most of data.

Maximizing the value of existing customers

Focusing on the experience of existing customers can often deliver the greatest impact. It takes by far more resources to acquire new customers than to retain the ones you already have. While you'll never stop working to bring in new prospects, addressing the needs of current customers should be among your highest priorities.

By uncovering insights from unstructured data, you can spot potential problems early. For example, some of the data we gather for our clients includes complaints from their customers who are unhappy with how long support takes to resolve their problems. We are now working on an alerting mechanism that leverages natural language processing to fire off an alert to the right team in any medium necessary — Slack, email, or even text. With this mechanism, our clients will be able to address issues before their customers decide to explore competing solutions.

Meanwhile, understanding customer needs in real time can create new upsell opportunities. If a customer describes a particular business challenge on Slack or posts a feature request on Discord, you might be able to offer solutions that meet their requirements. And here's some food for thought: When someone asks for enterprise pricing, what should you do? To me, the answer is to engage that customer and be as helpful as possible toward that potential buyer.

[1] https://images.g2crowd.com/uploads/attachment/file/1350731/IDC-Unstructured-Data-White-Paper.pdf

“By uncovering insights from unstructured data, you can spot potential problems early.”

Tapping into the power of AI

How do you find the gold in vast volumes of customer data? Fully capitalizing on unstructured data requires AI. Employing large language models (LLMs), you can efficiently collect customer data from multiple sources, connect pieces of data acquired from distinct places, unify the user across platforms, analyze data to understand sentiment, spot real-time trends, and present new insights in a visual format.

This work would be nearly impossible without AI. Let’s say you wanted to organize conversational data from multiple online forums into a coherent customer narrative about your product. More than just finding every mention of your product, you need a context-aware system that can discern relationships between numerous messages, connect particular pieces of information, and generate a single story that makes sense. Given the nuanced, often ambiguous nature of human communication, this process requires sophisticated AI capabilities and context-keeping even when conversations are intertwined and fluid in nature.

Addressing key strategic decisions

How do you implement an AI-powered approach to generating customer insights from unstructured data? There are a few essential decisions you need to make, and challenges you have to overcome, before you can start capitalizing on this wealth of data.

Collection

No matter your organization’s size, community platforms such as Discord, Slack, X, and Reddit are often the best places to mine for customer sentiment and feedback. In addition to the conversations your team may be having directly with customers on these platforms, customers are also talking with their peers (i.e., potential customers) about your company and your products, providing potentially critical information.

You need to decide which technology will allow you to collect all relevant data quickly and efficiently, while adhering to platform rules and data privacy regulations. At Sushidata, we use OAuth to facilitate data collection as opposed to using Zapier, because the friction of connecting to Zapier is not something we want for our users. We pay special attention to each source and make sure connecting to that source is as fast and efficient as possible. OAuth is an open standard that enables us — and our clients — to connect to an API from each platform. With OAuth, organizations can easily access customer information without having to get into the ethical gray area of scraping data from public forums.

Unification

Unifying data from multiple sources is one of the greatest challenges of analyzing customer data. On one hand, you want to bring all of this siloed data together. But you also want to understand which information came from which platform so you can take action in the right place.

“You need to decide which technology will allow you to collect all relevant data quickly and efficiently, while adhering to platform rules and data privacy regulations.”

At Sushidata, we assign IDs per source platform. If someone wants to drill down on a particular product issue, bug, or feature request, they can go directly to the source with a click.

Storage

If you are mostly collecting text data, you might use a traditional database. At Sushidata, we use Cloudflare's serverless database with a separate database instance for each tenant to make sure each organization's data is separate from everyone else's.

If you are including other types of data, such as images, a vector database (which keeps related data in close proximity) can help speed up performance. Cloudflare's developer platform enables us to determine which data should be included in the vector database.

In addition, choosing object storage, such as Cloudflare R2, can help you store a large amount and variety of data, from text, images, and video to log and event data.

Analysis

Organizations today have access to huge amounts of data, but using that data to make informed decisions requires analysis. AI is critical for tagging and analyzing all that data, and then generating actionable insights.

Finding or building the right AI models is key. Sushidata offers access to multiple AI models so you have the flexibility to easily explore new models as they become available. We use Cloudflare Workers AI for embeddings and text-generation models, which are run at the edge, close to users.

With the right models, you can analyze the context of the unstructured data you've collected and then perform multi-dimensional sentiment analysis. When my Sushidata co-founders — Victor Sanchez and Victor Ilisei — and I set out to measure sentiments, we wanted to do more than evaluate whether customers were happy or sad. There are so many more emotions you can explore.

We decided to use AI to perform five-dimensional sentiment analysis. This approach has helped our clients better understand whether their customers are expressing confidence or fear, confusion or clarity, and more. Homing in on the right sentiment enables you to better determine the best action to take.

Visualization

In most cases, the people who use customer insights are not data scientists — they are members of a customer experience or community management team. You need to find a solution for presenting insights to them in a visual format that conveys information quickly and easily.

With the right visualization capabilities, that team can immediately see whether your company is receiving more feedback, feature requests, bug reports, or mentions of other issues. The customer experience team can map out customer journeys and then work to optimize those journeys, and they can use visualizations to share insights with company leaders.

Security

Securing customer data and complying with data privacy regulations are crucial. To protect the privacy of individual customers, you need ways to de-identify data, removing personally identifiable information (PII) as data is collected. You also need to comply with platform rules on data collection. Also, as I mentioned earlier, multi-tenancy is important for safeguarding data.

If you are training your own AI model, you also need to ensure that the data you are feeding into the model is not compromised or corrupted. For example, we have seen companies using Reddit data — including conversations between company employees and customers — to train generative AI models. They plan to deploy those models in forums to answer customer questions on their behalf. But they need to be sure they have clean, accurate data. If someone goes into a forum and impersonates users, the models based on that data will not deliver accurate, constructive responses.

Envisioning the future of AI-powered storytelling

Applying AI to unstructured data has tremendous potential for better knowing your customers — how they feel about your products, what problems they are experiencing, and more. With that knowledge, you can take actions that improve customer satisfaction, reduce churn, and ultimately boost revenues.

We envision a future where you can go to Sushidata and ask, "Why are my users leaving?" and have AI tell you a story from your own data. That story can use an evolving graphic or a dashboard that — with the click of a "play" button — will help you understand your data in a way you've never experienced before.

Yes, there are some challenges to realizing that vision. Bringing together data from multiple sources, analyzing that data, securing data, and creating a compelling story machine are complicated tasks. But we founded Sushidata to turn that vision into a reality.

Cloudflare has played a key role in helping us build and manage our platform. Using Cloudflare products has enabled us to successfully analyze and categorize tens of thousands of conversations from online forums, transforming data into coherent customer stories. Organizations are using those stories to address customer needs and build long-term strategies that produce the greatest value for their customers.

> "Applying AI to unstructured data has tremendous potential for better knowing your customers — how they feel about your products, what problems they are experiencing, and more."

Balancing data sovereignty and AI

How to keep sensitive data local while using AI

With data regulations increasing in complexity, it is ever more important to know where data comes from, where it goes, and who processes it. In many parts of the world, data regulations require data to remain in its region of origin unless external organizations can demonstrate compliance with those regulations. This concept is called data sovereignty: the idea that data is regulated by the laws of the country or region in which the data is processed.

Yet even if the location or vendor to which data is being transferred is themselves compliant, cross-border data transfers can result in violations. For instance, government agencies in some countries may be empowered to examine data traversing their borders, which would violate the data regulations of other countries.

Organizations that transfer data out of their region of origin without adequate protection in place can face serious legal and financial consequences. As an example, in 2023 Meta received a US $1.3 billion fine for transferring personal data from the EU to the US without adequate privacy protections for the transferred data.

The above has given rise to the concept of data localization: to maintain data regulatory compliance and consumer trust, organizations often face the need to keep data within their own regions.

The idea of data localization is that data is kept within a given country or region, rather than transferred across borders and processed or stored on servers in remote areas. However, this approach makes cloud computing and the use of external third-party services more complex, as such services are usually not localized in this way. Cloud data centers are located all over the world regardless of where the services they support are based.

This means the need to localize data, for many organizations, may come into conflict with one of the most important cloud-based services available today: artificial intelligence (AI).

AI has emerged as a powerful tool for business

In recent years, a combination of more powerful hardware and increasingly refined software has led to an explosion in AI capabilities. Organizations are incorporating AI into their processes to assist with predictive modeling, content ideation, research, sentiment analysis, and customer service automation.

Analyst firms such as McKinsey continue to be optimistic about the expanding business uses for generative AI (GenAI). Most businesses do not have time or resources to build their own AI models, so are relying on outside vendors in order to use these technologies.

AI, however, vacuums up data in order to function. AI models are based on large datasets that are used for training complex algorithms. Large datasets can be, and are, stored in a variety of places. But because of its scalability, training data for AI is almost always stored in the cloud, in data centers around the world. (From the FAQs for OpenAI consumer services: "Content is stored on OpenAI systems and our trusted service providers' systems in the US and around the world [emphasis added].")

This means data uploaded to AI or used to train GenAI models passes outside the control of the organization that originally had the data, and is most likely outside of the geographic region where it originates from.

As the models receive more inputs, they continue to be fine-tuned. This means inputs may influence future outputs — or even reappear as future outputs (the latter of which is a risk to sensitive data, leading some organizations to ban the use of GenAI by their employees). Often this happens with very little visibility — AI users may not know where the machines are that process the data they provide. Also of concern is shadow AI, which is addressed further later in this guide.

> "On the other hand, the risks from not using AI, and falling behind the competition, pose a similar threat to businesses."

In many jurisdictions, this potentially brings businesses into conflict with data sovereignty requirements. The risks of conflicting with such requirements include fines (from small fines to the massive one levied against Meta), sanctions, and a decline in public reputation and customer trust.

On the other hand, the risks from not using AI, and falling behind the competition, pose a similar threat to businesses.

To summarize: AI is hugely useful but may be risky for organizations operating under strict data regulations — unless they can find a data-sovereignty-friendly approach for AI.

Options for leveraging AI without crossing borders

How can companies use AI while avoiding the risk of data crossing geographical borders? What is needed is an approach that offers computational power capable of supporting complex AI models but in a localized fashion. Organizations also need to make sure they control where their data is stored and processed, both in transit and at rest.

The best path forward is, therefore, data localization combined with local AI instances, either built on a third-party platform or offered pre-built by a vendor. Full data localization involves full control over where data is stored, where users are served from, and where cryptographic keys are stored (since this dictates where data exists in decrypted form).

These capabilities must be integrated with a powerful global AI network with local presence, one with sufficient computational power available on demand to operate AI models.

Businesses simultaneously facing the need to use AI and the need to localize data need a partner who understands these requirements and can support them.

Combating shadow AI

Implementing controls for government use of AI

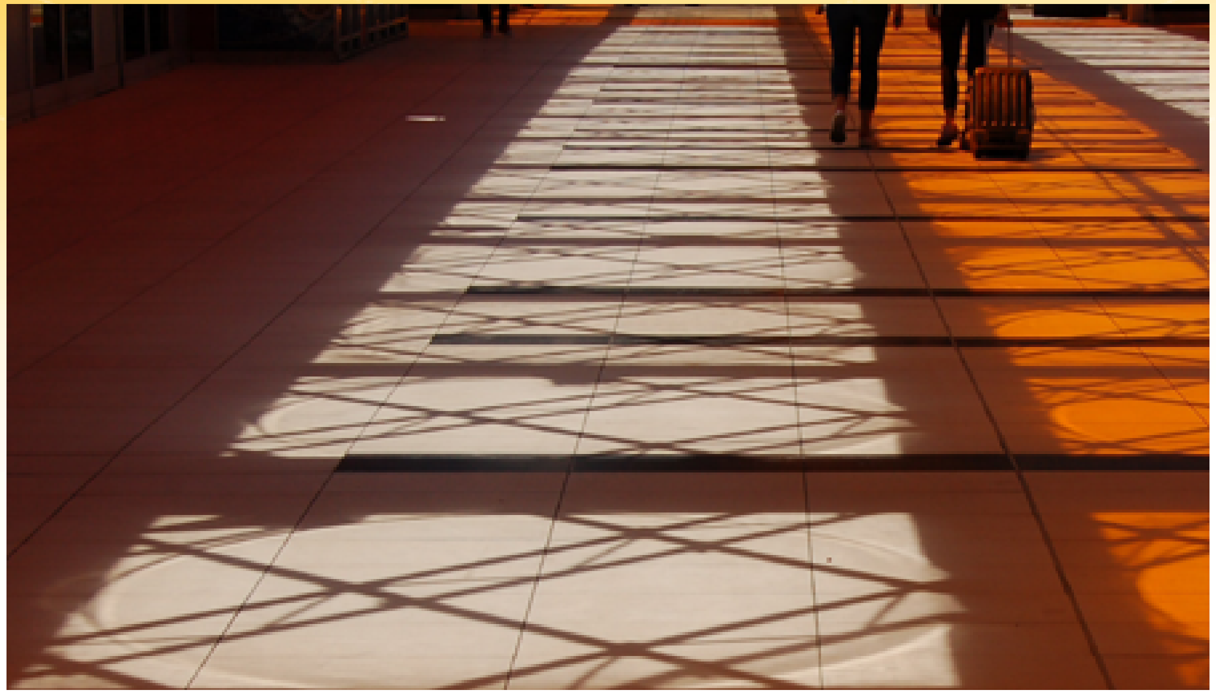

The White House Office of Management and Budget released Memorandum 24-10 for governance, innovation, and risk management in the use of artificial Intelligence to all federal agencies and departments. The three-part focus of the memorandum is to:

- Strengthen AI governance
- Advance responsible AI innovation
- Manage the risks from the use of AI

Last year, 25 states introduced legislation focused on some aspects of AI. Puerto Rico and 18 states enacted some form of legislation around AI. These legislative efforts range from initial study and evaluation of AI use all the way through governance on its use by employees to required controls to mitigate malicious or unintended consequences of AI.

Broadly speaking, this new body of legislation represents new compliance, consumption, and controls for government and other public sector organizations.

In this article, we will review some of the challenges facing organizations both for the protection of public-facing properties as well as identifying and crafting governance for the consumption of AI models.

Challenge 1: Protecting public Internet properties from AI bots

The impact of crawlers can be both legitimate and problematic for agencies. In some contexts, responsible crawlers and indexers will be able to use publicly accessible data, enhancing citizens' ability to find relevant on-line services and information.

On the other hand, poorly developed or malicious AI crawlers can scrape content to train public AI platforms without consideration for the privacy of that content.

There are numerous intellectual property and privacy concerns if this data ends up in training models that back public AI platforms. If unchecked, these bots can also hamper the performance of public websites for all users by consuming resources from legitimate interactions.

1 Control : Deploy application-side protections

There are several server-side protections that can be implemented to help control how bots interact with the server. One example is the deployment of a robots.txt file. In a nutshell, this file can inform and define how crawler traffic interacts with various sections of the site and the data therein. The file is deployed in the root of the site and defines what agents (bots) can crawl the site and what resources they can access.

There are a couple of challenges with this approach. The first and most obvious is that the crawler must respect the robots.txt file. While this is a general best practice for "respectable" bots, let's face it - not everyone follows the rules. There are also non-malicious bots that may just misinterpret syntax and therefore are able to interact with elements that agencies want to stay hidden.

In short, while a common approach, it should be noted that leveraging robots.txt is not foolproof protection. However, they are part of a holistic approach to governing how legitimate bots interact with application content.

2 Control : Deploy bot mitigation within a web application firewall

Web application firewalls (WAF) and bot mitigation solutions are table stakes in today's world for public web applications. These controls help organizations protect their public digital properties from DDoS threat vectors, shadow and insecure APIs, and various other threats delivered in the form of bot technology.

Any bot mitigation strategy today should include the ability to programmatically identify and classify bots that are scraping content in the service of AI data training. This classification mechanism is a critical capability. It defines whether to limit or allow only legitimate and verified AI crawlers or block them altogether until it is determined how these bots should be allowed to interact with government websites.

In the summer of 2023, António Guterres, Secretary-General of the United Nations, noting that AI has been compared to the printing press, observed that — while it took more than 50 years for printed books to become widely available across Europe — "ChatGPT reached 100 million users in just two months." The scale and the unprecedented growth in AI platforms directly correlates to the growing number of AI bots searching for any publicly exposed datasets for training.

This goes to the second major consideration in implementing these WAF and bot management controls: the architecture of these platforms must be able to scale in a distributed global environment.

Challenge 2: Unsanctioned consumption of public AI models

Public AI platforms have enabled users to accelerate everything from writing a memo to writing complex code. State and federal agencies see AI as critical to solving complex social problems like healthcare, access to citizen services, and food and water safety, among others. However, without governance, organizations may be complicit in leaking regulated datasets to insecure public language model training data.

In the same way that organizations have leveraged tools to get a handle on the consumption of unsanctioned cloud applications or “shadow IT”, they now need to understand the scope of shadow AI consumption within their organizations.

The increase of shadow AI is making headlines. A 3Gem study of over 11,500 employees worldwide showed that 57% of employees used public generative AI tools in the office at least once a week. Thirty-nine percent of respondents agreed that there is a risk of sensitive data being leaked through these interactions.

This information is sometimes even being unknowingly shared across AI models, given the increase of AI models being trained on data produced by other models as opposed to traditional sourced content.

1 Control : Determine appropriate use

Any comprehensive approach needs to include the determination of acceptable use of public AI models and, more specifically, what roles need access to those models. Establishing these guardrails is a critical first steps. In fact, one major theme in the rising new legislation on AI in government is the review appropriate use of AI within agencies and which models should be allowed.

2 Control : Deploy controlled accesse

Once those determinations have been made, agencies must then develop controls for enforcing those policies. Zero Trust Network Access (ZTNA) principles enable the development and enforcement of those policies to restrict unsanctioned access.

For example, you may only allow authorized users from specific administrative groups to access public AI models. Even if they are an authorized user, ZTNA allows additional posture checks such as ensuring corporate devices are up to date with patches or that the device has government-approved endpoint management agents running prior to allowing access.

In this way, governments can enforce and restrict who can access these public AI models while operating on government assets.

3 Control : Determine what data is appropriate for disclosure to AI platforms

Acceptable use is not only defining which users can access AI platforms. Governments need to also understand the data that is posted or submitted into AI platforms.

Even something as innocuous as a department memorandum could have non-public or sensitive data points. Once those data points are submitted to an LLM, there is a risk of that data being exposed.

Integrated data loss prevention (DLP) controls should be developed to ensure that proprietary information, such as sensitive application code or even citizen data, does not become a part of an unsecured training dataset for an AI platform.

Let’s take the example of an “AI developer group” needing to interact with both public and private or in-house AI platforms.

An agency could allow for the consumption of both public (e.g., ChatGPT) and private (e.g., AWS Bedrock) AI platforms. Only approved users in the “AI development group” are allowed access to these platforms. General users are blocked from both platforms.

However, even for approved “AI development group” users, the implementation of a DLP rule to examine the data that is being posted to these platforms ensures that non-public sensitive data can only be posted to the internal private AI platform.

Protecting constituents

Governance must start from a policy or mission perspective rather than a technology perspective. Understanding the role of AI in government programs from both a benefit and risk perspective takes intentionality by leadership to appoint focused teams that can evaluate the potential intersections of AI platforms and the mission of the agency.

> “Governance must start from a policy or mission perspective rather than a technology perspective.”

The increase of public engagement through technology creates an accessible rich set of data that AI platforms can use to train their models. Organizations may choose a more conservative approach by blocking all AI crawlers until the impact of allowing those interactions is understood. For those entities that see benefit for legitimate crawling of public properties, the ability to allow legitimate and controlled access by verified AI crawlers while protecting against the bad is critical in today’s environment.

From within the organization, establishing what roles and tasks require access to AI platforms is a critical early step in getting ahead of increased regulations. Mapping those needs to a set of controls that determine who gets access and when, as well as control over the kinds of data posted to these models, ultimately allows the removal of shadow AI without sacrificing the tangible benefits these technologies provide.

The promise of AI may — and in some ways already is — solving many complex social problems. However, governments must also protect their constituency while they wade into these new technologies.

Can AI find vulnerabilities?

As artificial intelligence (AI) advances, organizations and governments are scrambling to find its best applications. While ChatGPT and other large language models (LLMs) have captivated the media's attention, the potential uses for AI are far broader than text generation. One such area is security, especially the repetitive, large-scale task of identifying software vulnerabilities.

But the question of whether AI leads to better or worse security depends on who or what is doing the vulnerability identification — and for what purpose.

The unavoidable fact of vulnerabilities

Some flaws in software are essentially benign. But some flaws, known as vulnerabilities, can give someone who exploits the flaws a foothold within the system, leading to compromise. A significant chunk of cyber security practice is devoted to identifying and patching these vulnerabilities.

The list of exploited vulnerabilities leading to compromise is lengthy; some examples of high-profile incidents include:

- The 2017 Equifax breach, which started with an unpatched vulnerability
- The 2022 LastPass breach was partially caused by a vulnerability in third-party software
- The Norwegian government's IT systems were hacked in 2023 via a zero-day vulnerability

The consequences of vulnerability exploits can be disastrous, from data leaks to ransomware infections that freeze up an organization's systems. Organizations need to identify and patch vulnerabilities as rapidly as possible to avoid such occurrences.

Automated vulnerability discovery

Analyzing complex software programs in search of mistakes is a repetitive task that would seem to be a good fit for automation. Noted technologist Bruce Schneier has observed that "Going through code line by line is just the sort of tedious problem that computers excel at, if we can only teach them what a vulnerability looks like."

And indeed, machine learning (a subset of AI capabilities) has long been used for finding potential vulnerabilities in code. GitHub, for example, includes machine learning in their code scanning feature, which identifies security vulnerabilities in code. Naturally, this approach sometimes results in false positives, but when paired with manual analysis, a well-trained machine learning model can accelerate vulnerability identification.

As artificial intelligence advances by leaps and bounds, the possibility arises for training this technology to find vulnerabilities even more effectively. In fact, in 2023 the US agency DARPA announced a program called Intelligent Generation of Tools for Security — INGOTS. (DARPA, notably, was the agency that created ARPANET, the precursor to the Internet.)

> "But whether AI leads to better or worse security depends on who or what is doing the vulnerability identification — and for what purpose."

The program "aims to identify and fix high-severity, chainable vulnerabilities before attackers can exploit them" by using "new techniques driven by program analysis and artificial intelligence to measure vulnerabilities." INGOTS looks for vulnerabilities in "modern, complex systems, such as web browsers and mobile operating systems."

But is AI actually good at finding vulnerabilities? DARPA aims to find out, but their 36-month program is still somewhat exploratory.

AI vs. human hackers: Can AI find vulnerabilities?

Back in 2016, DARPA hosted the "Cyber Grand Challenge" in which seven teams of engineers created autonomous AI hacking programs, then faced off against each other in a digital game of Capture the Flag. The idea was to see how well an automated program could hack a secure system. After several hours, the program "Mayhem", designed by a team from Carnegie Mellon, won the competition.

The DEF CON 2016 conference was being hosted nearby, and Mayhem was invited to participate in DEF CON's own Capture the Flag game against human hackers. Mayhem came in last place, and it wasn't close.

AI has advanced a great deal since then, and researchers continue to release machine learning models for vulnerability discovery. But software investigated by even the latest machine learning models still requires human review to avoid false positives — or false negatives.

There is no denying that AI can find vulnerabilities. But human penetration testing still appears to have its place. This may change in the future, as AI becomes more robust.

> "There is no denying that AI can find vulnerabilities."

Can AI patch vulnerabilities?

Patching a vulnerability involves writing code that corrects the flaw. AI tools can certainly generate code. But to do so, they require specific prompts generated by their human users.

Even INGOTS does not plan to rely fully on automated processes for remediating vulnerabilities, instead aiming to "create a computer-human pipeline that seamlessly allows human intervention in order to fix high-severity vulnerabilities."

But the same caveat applies: As AI becomes more advanced, it may be able to rapidly and efficiently generate patches in the future.

Can attackers use AI to find vulnerabilities?

It is inevitable that, if a tool or technology is widely available, one side will use it to defend systems from attacks, and one side will use it to generate attacks. If AI can effectively find and patch vulnerabilities in software, then attackers will certainly use it to find those vulnerabilities before they are patched and write exploits.

Not all cyber attackers have access to such resources. But those who do will likely have no qualms about selling the vulnerabilities their AIs find, or the exploits they write, to the highest bidder on the dark web. Malware authors have already incorporated AI into their tools, and they will surely continue to do so as AI improves.

The possibility looms of an escalating, AI-driven arms race between legitimate software developers and malicious attackers, in which vulnerabilities are identified and exploited almost instantaneously, or (hopefully) patched just as quickly.

Of course, attackers are already combing through code looking for undiscovered vulnerabilities: such zero-day vulnerabilities are extremely valuable and can either be used by the discoverer for purposes of hacking the system, or sold on underground markets for a high price. Malicious use of AI may become a game-changer, but it's the same old game.

Can AI write vulnerability exploits?

As with patching, it is possible for AI to write vulnerability exploits but the process still requires human guidance. Therefore it may not actually save attackers any labor — and many of them buy exploit kits anyway, rather than writing their own code.

The answer may change 10 or even five years from now, and security folks should be preparing for a wave of fully automated vulnerability exploits targeting their systems.

Future-proofing against vulnerability exploits with Zero Trust

All networks are vulnerable to compromise — indeed, given enough time and a determined attacker, compromise is inevitable.

Even if AI brings a new world of vulnerability discovery for the side looking to secure their systems, attackers will be using the same methods to try and find vulnerabilities first, or at least before they can be patched. AI is becoming another tool in the toolbox or attackers, just as it is for the good side.

Forward-thinking organizations start with the assumption that compromise has occurred-that their security may fail, their data is at risk, and that attackers may already be inside the network.

They assume that their external-facing security may not always work perfectly, and therefore microsegment their networks so that malicious parties cannot extend their reach beyond the one segment they have already accessed. Think of how a ship can be sealed off into separate watertight compartments to prevent a leak from spreading. Ideally, security teams can use this same approach for containing attacks.

This approach is called Zero Trust, and there are strong reasons for this philosophy's growing adoption. As AI tools enable escalating exploits, Zero Trust can help ensure that those exploits remain restricted to a small corner of the network, and that attackers never gain a big enough foothold to cause real damage.

Vulnerability exploit discovery may accelerate, but Zero Trust offers the most hopeful path forward.

> "AI is becoming another tool in the toolbox for attackers, just as it is for the good side."

Conclusion

Finding your footing in the midst of the AI revolution

AI is developing faster, not slower, than anticipated. No one had expected AI to reach the point it had reached when ChatGPT-4 was released to the public in late 2022.

So the time to start thinking about and preparing for where AI will in five years, is now. Because "five years from now" could come much faster than that.

The Cloudflare connectivity cloud is uniquely positioned to help organizations build on AI, protect their data, and secure their own AI usage.

Cloudflare is a global cloud-based network that connects any application or service to users, data centers, branch offices, and networks anywhere in the world.

The Cloudflare connectivity cloud allows you to:

- Consolidate your security across all layers of your technology stack, in all locations, with Zero Trust access controls
- Control where sensitive data goes and which applications interact with it
- Protect against the latest automated attacks with up-to-the-second threat intelligence from a network that processes 60 million HTTP requests per second on average

For more perspectives on the latest trends and topics impacting today's technology decision-makers, visit theNET.

theNET
BY CLOUDFLARE

Call: 1 888 99 FLARE
Email: enterprise@cloudflare.com
Visit: cloudflare.com