THREAT WARRIOR

Unsupervised Neural Networks in Cybersecurity

Moving Well Beyond Traditional AI

Overview

Global cybercrime will be a $10.5 trillion USD business by 2025, up from $3 trillion USD in 2015, according to Cybersecurity Ventures. Let's put that into perspective. This is a larger market than the global trade of all major illegal drugs combined. Stolen money (personal and financial data theft, embezzlement, fraud), and intellectual property theft represent the greatest transfer of economic wealth in history. And its impact does not stop there. The costs associated with damage and destruction of data, lost productivity, post-attack business disruption, forensic investigations, restoration and deletion of hacked data and systems, and reputational harm are enormous.

It should come as no surprise as to why we have more than 3,500 US cybersecurity vendors. At the same time, everyone in the cybersecurity business is well familiar with the fact that we are abysmally behind in the required human talent to fight the good fight. While the global cybersecurity workforce grew to 4.7 million people in 2022, we remain 3.4 million security professionals short.

At first blush, this feels astonishingly hopeless. But there is a bright spot on the horizon: the application of true artificial intelligence to the global problem of cybersecurity.

We'd like to accomplish two things in this paper. First, the term 'artificial intelligence' is relatively overloaded these days. We'll provide a primer - so readers will have a foundational perspective on the technology at large. Second, we'll discuss how AI specifically changes the game of faster detection of company network breaches and incidents that underpin top of mind security concerns - like ransomware, phishing, data leaks and more - that keep every CISO on the planet awake at night.

AI Primer

Certainly much has been written on artificial intelligence, machine learning, and even neural networks. And yet, there remains confusion on what each term means - which makes it difficult to distinguish vendors' use of the terms. Let's start with a more comprehensive construct and work our way down to those three.
What is Artificial Intelligence?

Artificial Intelligence (AI) refers to the ability of machines to simulate human intelligence and perform tasks that typically require human cognition, such as learning, problem-solving, decision-making, and pattern recognition. It leverages algorithms, models, and systems to analyze data, extract meaningful insights, and make predictions or recommendations based on that analysis.

What are some example applications of AI?

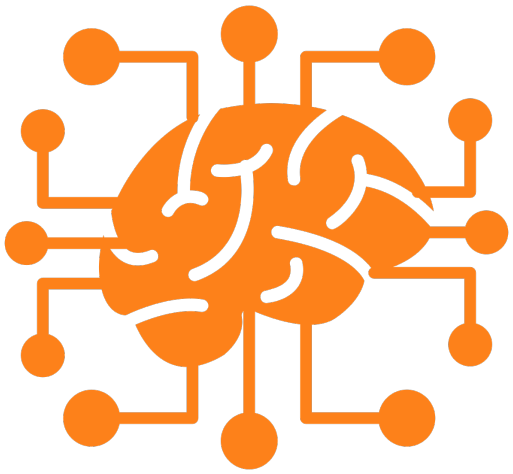

AI is used in a wide range of applications, including:

Virtual assistants: Siri, Alexa, and Google Assistant can answer questions, perform tasks, and provide information to users via voice or text-based interactions.

Image and speech recognition: AI can analyze images and identify objects, people, and other features with high accuracy. Speech recognition technology can transcribe spoken words into text for applications like dictation, language translation, and speech-to-text interfaces.

Predictive analytics: AI can analyze data from various sources and predict future outcomes with high accuracy. Examples include fraud detection, credit scoring, and predictive modeling for disease outcomes.

Autonomous vehicles: AI-powered autonomous vehicles use sensors and algorithms to navigate roads and make driving decisions without human intervention. Obviously companies like Tesla, Waymo, and Uber are advancing this space rapidly.

How do we even think about the word 'Intelligence'?

All intelligence - whether natural or artificial - requires three things:

Data: Humans learn from experiences and information. AI learns from information (for the foreseeable future). Information can be in the form of images, text, numbers, etc.

Algorithms: These are the instructions that tell intelligence what to do with data.

Compute power: Of course compute power is needed (and lots of it) to analyze large amounts of data quickly in order to learn from it and make predictions or decisions.

This is where it starts to get interesting. Many 'learning models' can be used to build intelligence. They include (but are not limited to):

Rule-based systems: These systems use a set of predefined rules to make decisions or draw conclusions. Rules are typically expressed in the form of "if-then" statements.

Symbolic reasoning: These systems use logical rules to manipulate symbols and infer new information. Symbols can represent any kind of concept, such as objects, properties, or relations, and the rules can specify how the symbols can be combined or transformed. Symbolic reasoning systems are often used in natural language processing (NLP), which is the field of AI that focuses on understanding and generating human language.

Evolutionary algorithms: These are inspired by the process of natural selection, where the fittest individuals are more likely to survive and reproduce. In evolutionary algorithms, a population of candidate solutions is evolved over many generations through the application of mutation, crossover, and selection operators.

Swarm intelligence: This is a collective problem-solving behavior that is observed in some animal societies, such as ants or bees. In swarm intelligence, a group of individuals work together to solve a problem by sharing information and coordinating their actions.

Cognitive architectures: These are frameworks for building intelligent agents that are inspired by the structure and function of the human brain. Cognitive architectures typically include modules for perception, attention, memory, reasoning, and decision-making.

Bayesian networks: These are probabilistic graphical models that represent uncertain relationships between variables. Bayesian networks can be used for tasks like diagnosis, prediction, and decision-making.

Machine learning: Machine learning (ML) is a subset of AI that involves training algorithms to learn patterns and insights from data without being explicitly programmed. ML systems can automatically improve their performance by learning from data and making predictions or decisions based on what they have learned.

Delving into each learning model is beyond the scope of this paper. But we'll pay particular attention to three of them: rule-based systems, Bayesian reasoning, and ML - as that serves our purpose with respect to the application of AI to cybersecurity.

Rule-based systems and Bayesian reasoning have been around for quite some time, and are probably familiar territory for anyone reading this paper. So, we'll assume that as a baseline. Let's turn our attention to a closer look at ML.

ML: a broader space than you may have thought

Not surprisingly, there are multiple types of ML. They include:

> "Machine learning is all about extracting valuable information from data, enabling machines to learn by experience. It helps address cases where it's infeasible to develop specific instructions for performing a task.
>
> **PETE SLADE**
> CTO, ThreatWarrior

Supervised learning: An algorithm is trained on labeled data, where each example in the training data is associated with a specific output or label. The algorithm learns to map inputs to outputs by generalizing from the training examples, and then makes predictions on new, unseen data. Everyday examples include email spam filtering and voice assistants like Siri, Alexa, and Google Assistant.

Unsupervised learning: An algorithm is trained on unlabeled data, where there are no specified output labels. The algorithm learns to find patterns or structures in the data, such as clusters or associations, and can be used for tasks like data clustering and anomaly detection. Common examples include customer segmentation where online retailers use unsupervised learning to group customers into segments based on their purchasing behavior, preferences, and demographic information; and credit card company detection of fraudulent transactions.

Semi-supervised learning: Supervised and unsupervised learning are combined such that the algorithm is trained on both labeled and unlabeled data. This can be useful in cases where labeled data is scarce or expensive to obtain.

Reinforcement learning: An algorithm learns through trial and error by receiving feedback in the form of rewards or penalties. The algorithm learns to maximize its reward by taking actions that lead to positive outcomes and avoiding actions that lead to negative outcomes.

Transfer learning: A learning technique where a model trained on one task is used as a starting point for training on a different, but related task. This can be useful when there is limited labeled data available for the second task.

Let's telescope down into Neural Networks

Neural networks consist of layers of interconnected nodes, or artificial neurons, that process information and learn patterns from data. Each node takes input data, applies weights and biases to that data, and then passes the transformed data on to the next layer. By adjusting the weights and biases between layers, neural networks can learn to identify patterns, recognize images or speech, make predictions, and perform other tasks. In other words, neural networks are a specific type of machine learning algorithm that are designed to mimic the behavior of the human brain.

It is important to remember that neural networks are a subset of ML algorithms. ML encompasses a broader range of techniques and algorithms that enable computers to learn from data and improve their performance over time. Neural networks use a specific architecture to learn patterns from data.

One last drill-down: Deep Learning

At times we see neural networking and deep learning being used interchangeably. They are not synonymous.

As discussed above, neural networks are a type of machine learning algorithm that are inspired by the structure and function of the human brain. Deep learning is different. It is a subfield of ML that uses neural networks with many layers, hence the term 'deep'. Deep learning algorithms can learn complex patterns and representations in data, which makes them particularly effective for tasks like image recognition, speech recognition and complex network traffic analysis.

A major advantage of deep learning, over say older techniques like Bayesian estimation, is that it excels at feature extraction - which enables it to build complex hierarchies of meaning to express information from raw data.

Summary

Let's summarize the key points of our primer as follows:

1. AI is a broad concept that enables machines to simulate human intelligence and perform human tasks.
2. Machine learning (ML) is a subset of AI focused on how machines learn patterns and insights from data without being explicitly programmed.
3. Neural networking is a subset of ML that is modeled after the structure and function of the human brain.
4. Deep learning uses neural networks with many layers to learn complex patterns and representations in data.

Applying AI to Cybersecurity

Traditional AI Approaches

Without question, AI has been employed within security products for decades. However, until fairly recently the vast majority of AI usage would have been predominantly underpinned by rule-based algorithms and Bayesian reasoning.

Supervised and Unsupervised ML

In the last few years many security solution providers have expanded detection, analysis, and prevention capabilities through the use of supervised and unsupervised ML applied to system logs, alerts, NetFlow data, and even full packet capture.

These advancements are quite worthy. At the same time, we remain short of where we need to be as an industry. Trace back to the beginning of this paper. The root problem in cybersecurity here in 2023 is that we still do not have enough human capital to stave off attackers, nor will we be able to close the gap any time soon.

Unsupervised Neural Networks

But there is a new frontier that can tip the scales faster. It is called unsupervised neural networks. Unsupervised neural networks provide a supercharged approach for discovering hidden patterns and relationships in data. As the name suggests, unsupervised neural networks are trained on unlabeled data - enabling them to discover patterns and structure in data without explicit feedback or guidance from labeled examples.

Within cybersecurity, most AI-based solutions on the market today are centered on supervised ML. Traffic is classified as 'known good' or 'known bad' and enforcement policies are then invoked. This is not only a time-consuming and costly endeavor, it is error prone even for networks of modest complexity.

Benefits of Unsupervised Neural Networks in Cybersecurity

Deep learning systems that leverage unsupervised neural networks provide a major step forward in helping security personnel find and stop never-before seen malware, novel attacks, ransomware, unknown threat variants, insider attacks, and more. Benefits include:

Zero-day discovery: One of the primary benefits of unsupervised neural networks in cybersecurity is their ability to detect anomalies in network traffic. By analyzing patterns in network traffic over time, unsupervised neural networks can learn what "normal" traffic looks like and quickly identify any unusual activity. This can be particularly useful for detecting zero-day attacks, which are previously unknown attacks that do not have a signature that can be detected by traditional security tools.

Malware mutation discovery: Unsupervised neural networks can be used for clustering, which involves grouping similar data points together. In cybersecurity, clustering can be used to group together similar malware samples, which can help identify new variants of known malware families.

Faster threat analysis: Unsupervised neural networks can be used for dimensionality reduction, which involves reducing the number of variables in a dataset. This is useful in cybersecurity when dealing with large datasets that may have many features that are not relevant to the problem at hand. By reducing the number of features, unsupervised neural networks improve the efficiency of other machine learning algorithms that may be used for classification or prediction.

Faster suspicious activity discovery: Unsupervised neural networks excel at analyzing network traffic in real-time to help identify and prevent cyberattacks. By analyzing the patterns and behaviors of network traffic, unsupervised neural networks can quickly flag any suspicious activity and alert security teams.

Focused threat hunting: Unsupervised neural networks can be used for threat hunting, which involves proactively searching for potential threats on a network. By analyzing large volumes of data, unsupervised neural networks can identify potential indicators of compromise (IoCs) that may have been missed by other security tools. This helps organizations identify and remediate threats more quickly and efficiently.

Dwell time reduction: Unsupervised neural networks can be used to analyze large amounts of network traffic data in real-time to identify unusual patterns that may indicate an attack. This can include anomalies in network traffic, unusual user behavior, or even new types of malware that have not been seen before. By leveraging the power of neural networks, cybersecurity professionals can quickly identify potential threats and respond to them before they can cause damage.

SOC/SecOps efficiency: Unsupervised neural networks are capable of recognizing complex patterns in data that may be difficult for humans or traditional machine learning algorithms to identify. This results in more accurate and reliable threat detection and response, reducing the risk of false positives or false negatives that can lead to security breaches.

Evergreen threat detection: Unsupervised neural networks are designed to be adaptive, meaning they can learn and improve over time as they are exposed to new data. This is particularly useful in cybersecurity, where new threats are constantly emerging and evolving. By continually training neural networks on the latest threat data, cybersecurity professionals can stay ahead of the curve and respond quickly to new attacks.

ThreatWarrior Unsupervised Neural Networks

ThreatWarrior leverages unsupervised neural networks to power our network detection and response (NDR) platform. Our implementation of neural networking is made unique and valuable through a set of core design principles:

1. Unsupervised ML
2. Meta representations
3. Custom feature hierarchies
4. Automated feature engineering
5. Training granularity

Specifically with respect to unsupervised ML, we'd like to dive a little deeper in a way that can equip buyers to ask informed questions when evaluating which NDR solution would be best for their needs.

As AI is a raging topic at present, no vendor wants to be perceived as missing in action or behind the curve. Not surprisingly, there is a tendency for vendors to conflate terms - and descriptions of underlying technologies. This can leave buyers perplexed as to whether there are material differences from one vendor's approach to the next. At Threat Warrior, we believe the use of specific machine learning approaches - as well as their actual implementations - create substantial deltas in the scope, speed, and fidelity of suspicious and malicious activity identification.

There are vendors who lean almost entirely on supervised ML approaches. As has been noted earlier in this paper, supervised ML alone has some important limitations. Next, there are vendor claims of unsupervised ML utilization. But often, the unsupervised learning is confined to the boundaries of their legacy-designed Bayesian detection models. Solutions built on 'unsupervised ML' that discretely observes field incidents within specific customer deployments are really just extending a Bayesian approach - where labeled data sets from one customer deployment determines what 'normal' looks like on your network.

This is quite different from an unsupervised, layered neural network approach - where AI requires no guidance to begin learning based on what it is observing. Buyers should be clear that Bayesian theory of probability is not true learning, and can introduce error-prone biases that send analysts and threat hunters down the proverbial rabbit hole - wasting valuable time, energy, and money.

In contrast, ThreatWarrior leverages a hybrid learning approach - driven primarily by unsupervised neural networking, but also leveraging supervised ML. Learning is unsupervised, and takes place within specific environments, actually down to a specific network segment (which has its own unique behaviors). Global threat signature rule sets (updated hourly, parsed in real-time via NLP since they enter the engine as raw data) are applied using supervised techniques that can be quickly tuned by a SOC or ThreatWarrior as needed. For any observed signature - malicious or benign - there is an option to train the system around that signature right from the alert.

The difference in the two approaches is substantial - with the latter leading to far greater scope, speed, and fidelity of suspicious and malicious activity identification.

As an example, zero-day exploits have no threat signature, and some are so cleverly designed that they will not shift the pattern of traffic in a way that is significantly different from 'normal' on any network but your own. Given those nuances, a layered-neural network approach is far better positioned to trigger on zero-day activity than more simplistic learning models.

We believe our design principles translate directly to a more effective NDR solution:

Faster, more accurate traffic understanding. Through unsupervised ML, ThreatWarrior's platform learns faster and scales further than more commonly employed supervised ML approaches. Rather than feeding our engines with network data that has been labeled good or bad, ThreatWarrior explores inputs, analyzes each, and outputs a richer set of pattern findings.

No 'normal behavior' bias. ThreatWarrior distills traffic patterns down to unique 'meta' representations. Meta representations prevent the introduction of human bias. Biases can - without intention - slant or steer a neural network into a 'line of thinking'. That can eventually create blind spots in a machine's 'thinking'. With ThreatWarrior, deep learning trains machines with no preconceived notions of 'normal behavior'.

No traffic misclassification. Our deep neural networks develop custom feature hierarchies that capture the essence of a specific network as opposed to the features a product engineer believes should matter. By allowing a system to form its own thought around data classification and segmentation, anomaly discovery becomes more comprehensive and accurate.

Faster data training. ThreatWarrior's deep learning approach allows us to circumvent the limitations of feature engineering. In the past it was not feasible for a computer to understand raw data with zero guidance. One had to guide it through supervision to help it understand what it was seeing. For example, a request like 'learn to recognize faces by looking at pictures of faces' was a cognitive inability. To guide machines, computer scientists hand-crafted algorithms to detect different kinds of lines and edges. Researchers made educated guesses as to what kind of features would even prove useful. Eventually, layer upon layer of hand-crafted features built up from simple lines and shapes could crudely describe facial features like eyes, ears, noses and lips. Those high-level, engineered features would then be the input to the engine instead of raw data. With neural networks we can skip the tedious and expensive human-based 'pre-training' and allow computers to learn directly from raw data - freeing researchers and programmers to focus on higher order work.

Unique customer and network segmentation. ThreatWarrior trains custom models for each customer and each network segment within a customer's IT environment. For each monitored network segment, a deep neural network is built and trained on samples of traffic metadata from only that segment. Further, datasets are sampled across different time strata. As a result, ThreatWarrior baselines can be tightly bound to each observed network segment. For example, the traffic patterns for a network segment composed of office workers will be quite different from a segment that contains outward facing e-commerce traffic. Once each segment is 'mapped' it becomes easier to discern abnormal east-west and north-south traffic. This is a fundamentally different approach than simply observing network traffic en masse and aggregating it into a cloud instance that then drives a more generic set of anomaly patterns to a set of customers.

The speed, efficiency, and accuracy with which ThreatWarrior is able to detect and stop cyber threats is not possible with antiquated, less-powerful applications of artificial intelligence.

Summary

In 2021, organizations around the world spent around $150 billion on cybersecurity. And yet the cost of cybercrime is an order of magnitude larger. There is simply no way that traditional defense-in-depth approaches and the hope for a massive human talent infusion will stem the tide. We have no choice but to leverage artificial intelligence to provide a human capital multiplier effect to the work of monitoring network traffic, recognizing the root of attacker activity faster and with greater accuracy, and converting those findings into actionable intelligence. Many security solutions are currently moving past rule-based and Bayesian reasoning algorithms into supervised machine learning. While this is a solid step forward, it lacks the ability to scale cost-effectively. Unsupervised neural networks are the new frontier for making sense of network traffic and enabling humans to fend off determined adversaries. ThreatWarrior's NDR solution has been designed around unsupervised neural networks and deep learning technology from inception. Our NDR solution can help organizations close the cybersecurity gap faster and more cost-effectively than less advanced approaches.

About ThreatWarrior

ThreatWarrior is a leader in network detection and response (NDR) for on-premises, hybrid and multi-cloud enterprises. Our AI-powered platform helps organizations identify threats in real time and stop them before they cause breaches. By combining complete visibility, deep packet inspection, behavioral anomaly detection, forensics and threat hunting, ThreatWarrior delivers the context security teams need to detect, investigate, and respond to threats. Leading organizations use ThreatWarrior to defend against APTs, zero-day exploits, ransomware, and more.

THREAT WARRIOR

info@threatwarrior.com

threatwarrior.com

844.463.9440