

DETECTING MALICIOUS DOMAINS USING ARTIFICIAL INTELLIGENCE AND MACHINE LEARNING

THE CHALLENGE

Cyberattacks are a top priority in most IT organizations; the risk associated with ransomware attacks, data breaches, business email compromise, and supply chain attacks has garnered that significant attention be made to IT resources and budget to address these threats. The challenge in building a comprehensive security strategy designed to prevent attacks is the ever-changing threat landscape.

New cybercriminal organizations and threat actors pop up daily with the intent to design attacks that evade detection. These elusive techniques obfuscate attack behaviors and make nefarious actions appear benign. One technique is to constantly change domain names used as part of an attack. If a threat actor used the same domain throughout an attack, eventually they'd be blocked. These simple acts inevitably become the norm and make a defender's job that much more difficult.

Just like a regular business that relies heavily on technology, threat actors build out their own malicious infrastructure from which to run and support their campaigns. So, the use of easily discard-able domains helps cybercriminals to avoid detection by security solutions.

Assuming an attack relies on one of these domains, you could block malicious domains and effectively prevent attacks, right? Sounds simple enough, but in practical application, this isn't an easy task.

DOMAINS ARE TYPICALLY USED IN A CYBERATTACK FOR:

MALWARE

If a threat actor's goal is to download malware and infect an endpoint, domains can be used as part of drive-by downloads, communication with a command and control (C2) server, downloads of nefarious payloads, and exfiltration of data to a specific server.

PHISHING

Emails impersonating legitimate brands and domains (e.g., cha5e.com) can be used as part of business email compromise (BEC) scams. Credential harvesting attacks that seek to fool users into providing their cloud credentials to, say, Office 365 may point the user back to a specific spoofed domain.

SPAM

While spam is considered benign, some spam is borderline phishing, focused on obtaining credit card information, getting users to click links to compromised websites, etc. which can still present a threat to users and their organizations.

SO, HOW MANY DOMAINS ARE WE TALKING ABOUT AND HOW QUICKLY ARE THESE DOMAINS BEING CREATED?

In 2019, there were over 275,000 domains registered on a daily basis! The continual creation of so many domains further challenges IT organizations; it's impossible for IT pros to even attempt to manually keep up with which domains are legitimate and which are malicious.

In recent years, we've heard more and more about the use of Artificial Intelligence (AI) and Machine Learning (ML) to help make security efforts more current, effective, and responsive. But every vendor is claiming to use AI/ML today. It can create a ton of buzz and hype, reflecting the "state of the art," but the question remains **does it bring any practical value?**

To help clear up the confusion, in this paper, we'll provide a high-level definition of both AI and ML from a security perspective, as well as practically apply the principles of each to domains in an effort to demonstrate how they are used to spot malicious traffic before it becomes a problem.

TURNING DOMAIN DETAILS INTO PREDICTIVE SECURITY: DOMAINTOOLS

Most organizations need to wait until some kind of malicious behavior occurs before any response can be taken. This puts organizations at risk, as threat actors take steps to avoid detection. But the use of domains for malicious communications can serve as a proactive means to identify potentially threatening actions **before** they happen.

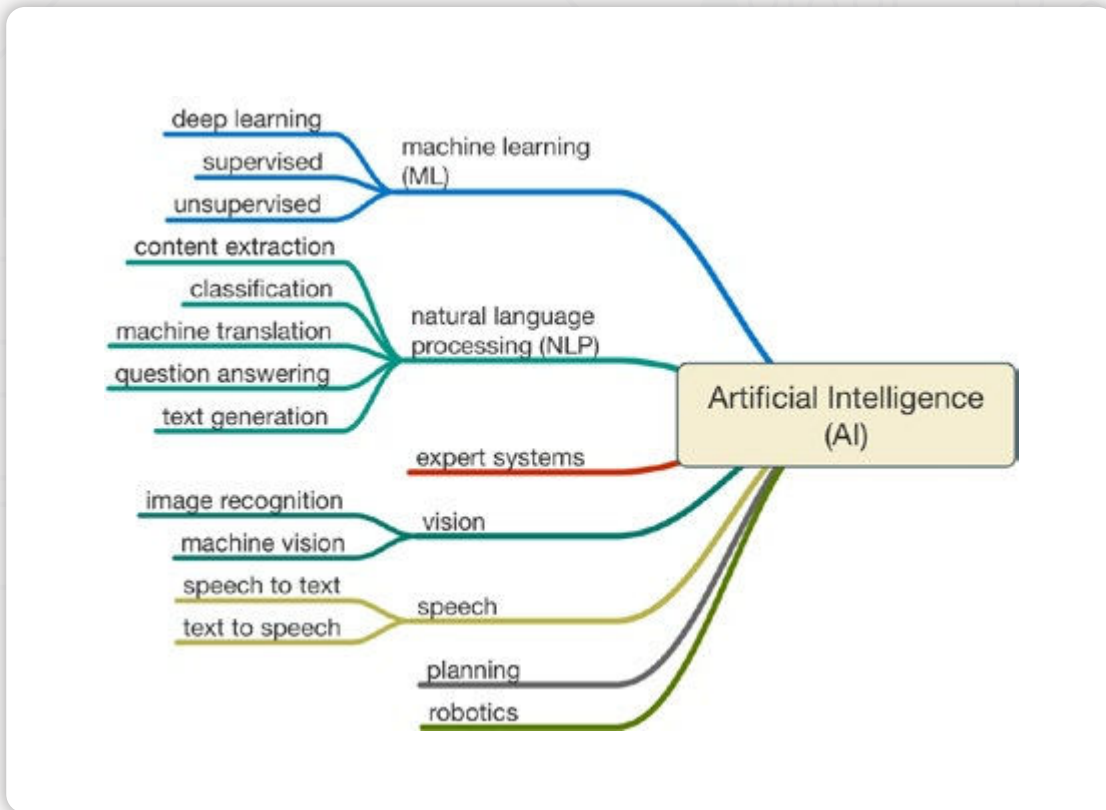
DomainTools Risk Score helps organizations identify suspect domains based on a variety of factors, assessing reputation and assigning risk on over 300M domains to help organizations mitigate threats before they start. Look for insights from DomainTools throughout this paper.

AI & ML: A PRIMER

Every IT professional has heard a lot about Artificial Intelligence and Machine Learning over the last few years. So much so, that it's created both a ton of awareness for the need to have advanced methods of learning to address security problems, as well as confusion around what exactly it means to apply these technologies and their benefits. There are plenty of security vendors saying "we do AI", but what exactly does that mean? Let's start by defining each technology and then look at how it can be applied to the problem of detecting malicious domains.

Artificial Intelligence

Artificial Intelligence (AI) is a branch of computer science focused on helping computers learn on their own, adjust to new inputs, and perform tasks - all without human intervention. Made up of a number of different types of learning (shown below), AI is useful in everything from IT security, to robotics, to DeepFake videos, to real-time translation of conversations.



Machine Learning

Machine Learning (ML) is just one family within AI that is most applicable to cybersecurity. ML focuses on devising algorithms that allow a computer to improve its understanding of the data presented automatically through continual exposure to data. In essence, you give a ML process some data and it attempts to learn very specific things about the data - whether guided by human input or on its own.

Learning Phase

As the name implies, there is a learning process that needs to happen before the ML algorithm can be useful. The learning process allows for a model to be created – in the case of an application to cybersecurity, it could be a model that identifies which emails, attachments, or domains are potentially malicious or benign. There are two types of learning:

Supervised

Also referred to as **Classification**, this method of learning uses externally-defined “buckets” of labeled data to help the ML system, in essence, learn from a human teacher. For a given data set, there are human-defined characteristics (also called features) that describe the data. The goal is to teach the machine to use these features to build a model that will serve as the basis for processing future unlabeled data to decide which “bucket” a piece of new data belongs in.

Unsupervised

Also referred to as **Clustering**, this method of learning looks at unlabeled data and learns “on its own”, self-organizing the data and determining its’ own patterns and structures. Unsupervised learning can produce patterns not previously seen by humans, but can generate more unpredictable models.

Testing Phase

Once the model is built, labeled test data is put through the model to determine if the machine can correctly identify data that are known to have specific features. The goal is to compare the testing results (which establishes what the ML model “thinks” of the test data) to what a human thinks the ML algorithm results should have been. If the accuracy (as determined by the human) is low, the model needs to be readdressed and both the learning and testing must be redone.

How then can AI/ML be used to identify potentially or actively malicious domains?

APPLYING AI/ML TO MALICIOUS DOMAIN NAMES

Rarely does a malicious attack today not involve some form of communication over the Internet. Malware needs to use command and control (C2) servers to move laterally inside a network, business email compromise attacks often start with a spoofed domain made to look like a credible business, and malware-less phishing attacks still link to malware-laden websites. In all these cases, a domain name of some kind is used. So, leveraging AI/ML as a means to identify a suspicious domain based on known features of previously identified malicious domains makes sense.

The benefit for organizations is block traffic to a domain before its operationalized and is used to perform some malicious action on your network.

DOMAINTOOLS INSIGHTS - ACCURATELY PREDICTING MALICIOUS DOMAINS

If a domain hasn't shown itself to be malicious, how is it possible for ML to learn what is an isn't suspect when it comes to domains? DomainTools uses data from domain blacklists – lists of domains that have already exhibited malicious behavior – in both the Learning and Testing phases to develop algorithms that can predict a domain does or will have malicious use with a high degree of accuracy.

USING DOMAIN FEATURES TO SPOT SUSPECT DOMAINS

Domain names are more than just an identifier and a top-level domain (.com, .net, etc.). In fact, there's a lot of supporting detail within the DNS forensic data ML can use to understand the intent and potential use of a given domain. As part of the training process, ML can consider a large number of domain features to eventually spot malicious domains, including:

- **Name Composition or Entropy** – The makeup of a domain name matters. Examples include domains made up of random characters from the keyboard rather than those spelling out one or more words are suspect, “look-alike” domains with one character changed, or domains with an unusual number of vowels or consonants inconsistent with normal domain naming. Even the presence of a hyphen in the domain can be suspect; a domain like logon-chase.com could be easily used to attempt to trick victims into thinking they are at a chase.com login page.
- **Domain Age** – When a domain is registered has an impact on its’ reputation. Malicious domains don’t typically have a very long lifespan, so one that’s a week old is considered far more suspicious than one that’s a year old.
- **Domain Name Length** – The number of characters in a non-malicious domain is usually limited to something palatable to humans. Domains with extended lengths are suspect.
- **Name Server Details** – The name server hosting the domain can have an impact; IP address of the DNS server, which other domains are on that same server, and more can all indicate a potential issue with a domain.

CREATING A THREAT PROFILE

A threat profile is set of supervised machine learning classifiers that can be used to predict whether a domain is likely to be malicious today, may have malicious intent at some point in the future, or is benign in nature. Keep in mind that a profile is likely only good for a very short period of time – in some cases measured in days to even just hours. So, to have an effective ability to spot malicious domains, the models need to be rebuilt on a regular basis, requiring new data to understand how threat actors are changing their tactics around domains and their use.

ASSIGNING A RISK SCORE

The desired output of ML is a domain risk score – in essence, a reputational score that defines how likely a domain is to be harmful. As previously mentioned, with over 275,000 domains created daily, the question becomes which ones are created with malicious intent and which are simply someone with a great idea to start a new website? The features above – and many more – are used to help ML algorithms learn about the composition of a domain and eventually allow it to predict whether a domain is suspicious or not.

TAKING ACTION ON SUSPECT MALICIOUS DOMAINS

The most natural response to a risky domain would be to simply block it using, for example, a web gateway block list. But that may cause more problems than it solves – especially in cases where domains that meet high risk score criteria haven’t even shown themselves to actually be malicious in nature. So, while blocking traffic is one of the initial actions that can be taken based on a given risk score, the other two most prevalent responses include alerting the appropriate groups, and launching an investigation. But, depending on whether an organization has a dedicated security team and established security response procedures based on domain risk scores, other actions can include passing information over to a SIEM, threat intelligence platform, threat aggregators, or security orchestration & automation solutions for further processing and investigation.

DOMAINTOOLS INSIGHTS

THE BONEYARD

Some features don’t help the model. When a feature doesn’t help the model, DomainTools places the feature in the “boneyard”. As threat methods change, the models are re-evaluated against features found in the boneyard to make certain bad guys aren’t changing tactics with domains where legacy features may provide insight into detecting malicious intent.

BREAKING OUT DOMAINS BY THREAT

Initially phishing, malware, and spam domains were grouped together with a single classification process, but the output wasn’t terribly accurate. So instead, each attack vector or type is given its own classification to improve the accuracy, building different learning sets, and customizing the features of the domains used in each attack vector. The resulting output is three different independent reputation scores – one per vector or type.

DOMAINTOOLS INSIGHTS

RISK SCORE

Waiting until a domain is used as part of some form of malicious activity creates a window of opportunity for an attack to be successful. DomainTools Iris proactively assesses domains using over 100 features as soon as a domain is registered to predict the likelihood that it may be used as part of an attack.



As shown above, separate scores are assigned for three primary domain name-based attack vectors – **phishing**, **malware**, and **spam**, allowing organizations to respond as desired to each potential threat.

MAKING RISK SCORE DATA ACCESSIBLE

While DomainTools Iris provides organizations with the toolset necessary to comprehensively investigate domains, their history, and SSL profile, there's also the need to leverage this data within other security tools. DomainTools offers a number of APIs as integration points with other security solutions such as Splunk, Splunk Phantom, IBM QRadar, IBM Resilient, Anomali, Maltego, ThreatConnect, Demisto and more, to provide security teams with the ability to leverage DomainTools domain insights within the context of established IT security workflows and processes.

PROACTIVE THREAT PREDICTION AND PROTECTION WITH MACHINE LEARNING

With the speed at which new threat actors surface, phishing campaigns are launched, malware is created, and domains are utilized, it's impossible for an individual security professional to stay current enough to have a positive impact on their organization's security. The use of ML is an absolute necessity in today's ever-changing threat landscape to ensure that defenses are ready and effective.

ML empowers security solutions to continually evolve with changes in threat tactics – something no human can do with an appropriate level of speed and accuracy. Through its ability to identify malicious domain patterns, ML gives the advantage to security teams; the use of Domain Risk Score provides organizations with the ability to proactively prevent attacks from ever coming to fruition. By leveraging ML to identify patterns in the domain data, active and future malicious domains can be identified, giving organizations a leg up on attackers before their organization is even a target.

Whether used as part of spotting malicious domains, or some other facet of your organization's cybersecurity strategy, leveraging solutions that rely on ML to identify potential risks at all levels of the security stack will provide stronger levels of security than those not using ML. As you consider such solutions, don't simply take the vendor's word that "they have ML" and make assumptions as to what that means; ask questions, attempt to understand how their ML specifically supports your security efforts.

By using predictive ML-based security solutions – especially in the case of identifying and blocking malicious domains – your organization will be ready for the next attack before it even occurs.

ABOUT THE AUTHORS



Randy Franklin Smith is an internationally recognized expert on the security and control of Windows and AD security. Randy publishes www.UltimateWindowsSecurity.com and wrote *The Windows Server 2008 Security Log Revealed*—the only book devoted to the Windows security log. Randy is the creator of LOGbinder software, which makes cryptic application logs understandable and available to log-management and SIEM solutions. As a Certified Information Systems Auditor, Randy performs security reviews for clients ranging from small, privately held firms to Fortune 500 companies, national, and international organizations. Randy is also a Microsoft Security Most Valuable Professional.



John “Turbo” Conwell is a Senior Data Scientist at DomainTools. He brings 10 years experience in data science and machine learning to bear on cybersecurity. He is currently focusing on building models to identify domains created for malicious intent as soon as they are created.

ABOUT DOMAINTOOLS

DomainTools helps security analysts turn threat data into threat intelligence. We take indicators from your network, including domains and IPs, and connect them with nearly every active domain on the Internet. Those connections inform risk assessments, help profile attackers, guide online fraud investigations, and map cyber activity to attacker infrastructure. Fortune 1000 companies, global government agencies, and leading security solution vendors use the DomainTools platform as a critical ingredient in their threat investigation and mitigation work. Learn more about how to connect the dots on malicious activity at <http://www.domaintools.com> or follow us on Twitter: @domaintools.