Veracity Trust Network

VERACITY CYBER BRIEF

AI in the Battle Against Malicious Bots

This Cyber Brief highlights the critical role of AI in countering malicious bot attacks on your web assets. It covers AI's significance in cybersecurity and Veracity's approach to industry leading proactive bot detection, and the roadmap for continued innovation in proactive bot detection.

Veracity Trust Network

01

Introduction to AI, Cybersecurity, and the Danger of Malicious Bots

The Advancement of Artifical Intelligence

We are currently living in a period where Artificial Intelligence (AI) and the subset of AI, Machine Learning (ML), are at the forefront of technological and societal progress. Solutions to problems previously deemed to be incredibly complex and challenging, or even impossible, are being unlocked by developments in the field of ML.

ML algorithms attempt to replicate human problem solving by analysing vast quantities of data that describe aspects of the problem at hand and attempts to learn the behaviours and nuances of the problem. Once an algorithm has "learnt" the correct behaviours it can apply these behaviours to new data it has not come across before. Once trained, such ML algorithms can make predictions on data at a phenomenal speed, vastly quicker than the rate at which humans work. They can, therefore, be used to automate complex and time-consuming tasks and to solve problems that are beyond the capability of humans to understand.

As a result of these rapid advancements, AI is constantly being applied to new fields and problems, across academia and commercial industries, such as healthcare, finance, autonomous vehicles and even the creative arts.

Why AI Matters in Cybersecurity

Cybersecurity is about safeguarding an organisation's systems, data, as well as the information of employees, clients, and affiliated parties, to prevent unauthorised access.

Modern businesses are built around the data they gather, store, and analyse for competitive advantage, improved service delivery and profit/ROI optimisation. Such data usually includes information that is confidential, protected under GDPR, CCPA, or similar data protection regulation, or which may be subject to regulatory control, or which is otherwise critical to the organisation, and which must not be divulged apart from to authorised parties.

Such data is valuable to bad actors for a whole number of reasons, ranging from non-commercial malicious breaches, for bragging rights or "fun", through general phishing or information gathering to prepare for a future attack, through to highly targeted activity to steal specific data or damage a specific target. Much of this activity has a financial purpose at the heart of it, whether stealing data to sell-back or sell-on, stealing money, content theft, or replacement, or more. Figure 1 indicates some of the problems different sectors face.

Damaging Bot Activity by Market Sector

Market Sector	Damaging Bot Activity
Automotive	Price scraping, data scraping, inventory checking
Business Services	Attacks on the API layer, data scraping, account takeover
Education/Online Learning Platforms	Account takeover for students & course availability, scraping proprietary research papers and data
Entertainment & Arts	Account takeover, price scraping, inventory checking, scalping
Financial Services, Banking, Insurance	Account takeover, card cracking, content scraping
Food & Beverage	Credit card fraud, gift card fraud, account takeover
Gaming & Gambling	Account takeover, odds scraping, account creation for promotion abuse
Government & Public Services	Account takeover, data scraping of business & voter information
Healthcare Services	Account takeover, content scraping, inventory checking, vaccine appointments/availability
Information Tech & IT Services	Account takeover, scraping
Marketing & Agencies	Custom content scraping, ad fraud, denial of service
News & Media	Custom content scraping, ad fraud, comment spam, fake accounts
Retail & eCommerce	Denial of inventory, credit card fraud, gift card fraud, account takeover, data and price scraping
Sports	Sports updates, news, live score services data scraping (live scores, odds, etc)
Telco	Account takeover, competitive price scraping
Travel/Airlines	Price & data scraping, skewing of look-to-book ratio, denial of service, price scraping, account takeover

Figure 1

Why Malicious Web Bots are Difficult to Prevent

Bots are the foot soldiers in this war, used for finding vulnerabilities, or for scraping information and content that may be used directly, or for future attacks.

Bots are nothing more than software tools programmed to behave in a certain way; the term “bot” is a shortened version of the term “software robot”.

Bots are used for good, legitimate purposes as well as bad. Our interest is in detecting and stopping malicious bots from being able to carry out their activities.

The impact on Cybersecurity and bot detection of the recent rapid developments in AI are profound and long-lasting. Generative AI has allowed bad actors to create authentic looking data more easily to support phishing and vishing attacks; ML is already in use for the generation of more human looking bots in bot

networks; and the big data systems that underly ML are being used to store and analyse data on potential attack vectors, using data gathered daily by malicious bots from hundreds of thousands of websites.

The simplest, although not foolproof, way to protect a system from intrusion is not to connect it to the Internet. However, websites, web applications and APIs are designed to be used on the Internet, and often to be open to all visitors – because it what they exist for. There is very little built-in protection on the Internet to prevent a visitor to a website to fake their origin, IP address, browser data and actions and to pretend to be something that they are not. This intrinsic openness is what malicious bots exploit.

AI powered malicious bots can bypass traditional bot detection systems with ease. Less sophisticated bot detection systems will result in a high false positive rate (where human visitors are incorrectly identified as bots) and reduced detection of malicious bots. Sophisticated AI bot detection as provided by Veracity Trust Network is the only way to combat this danger.

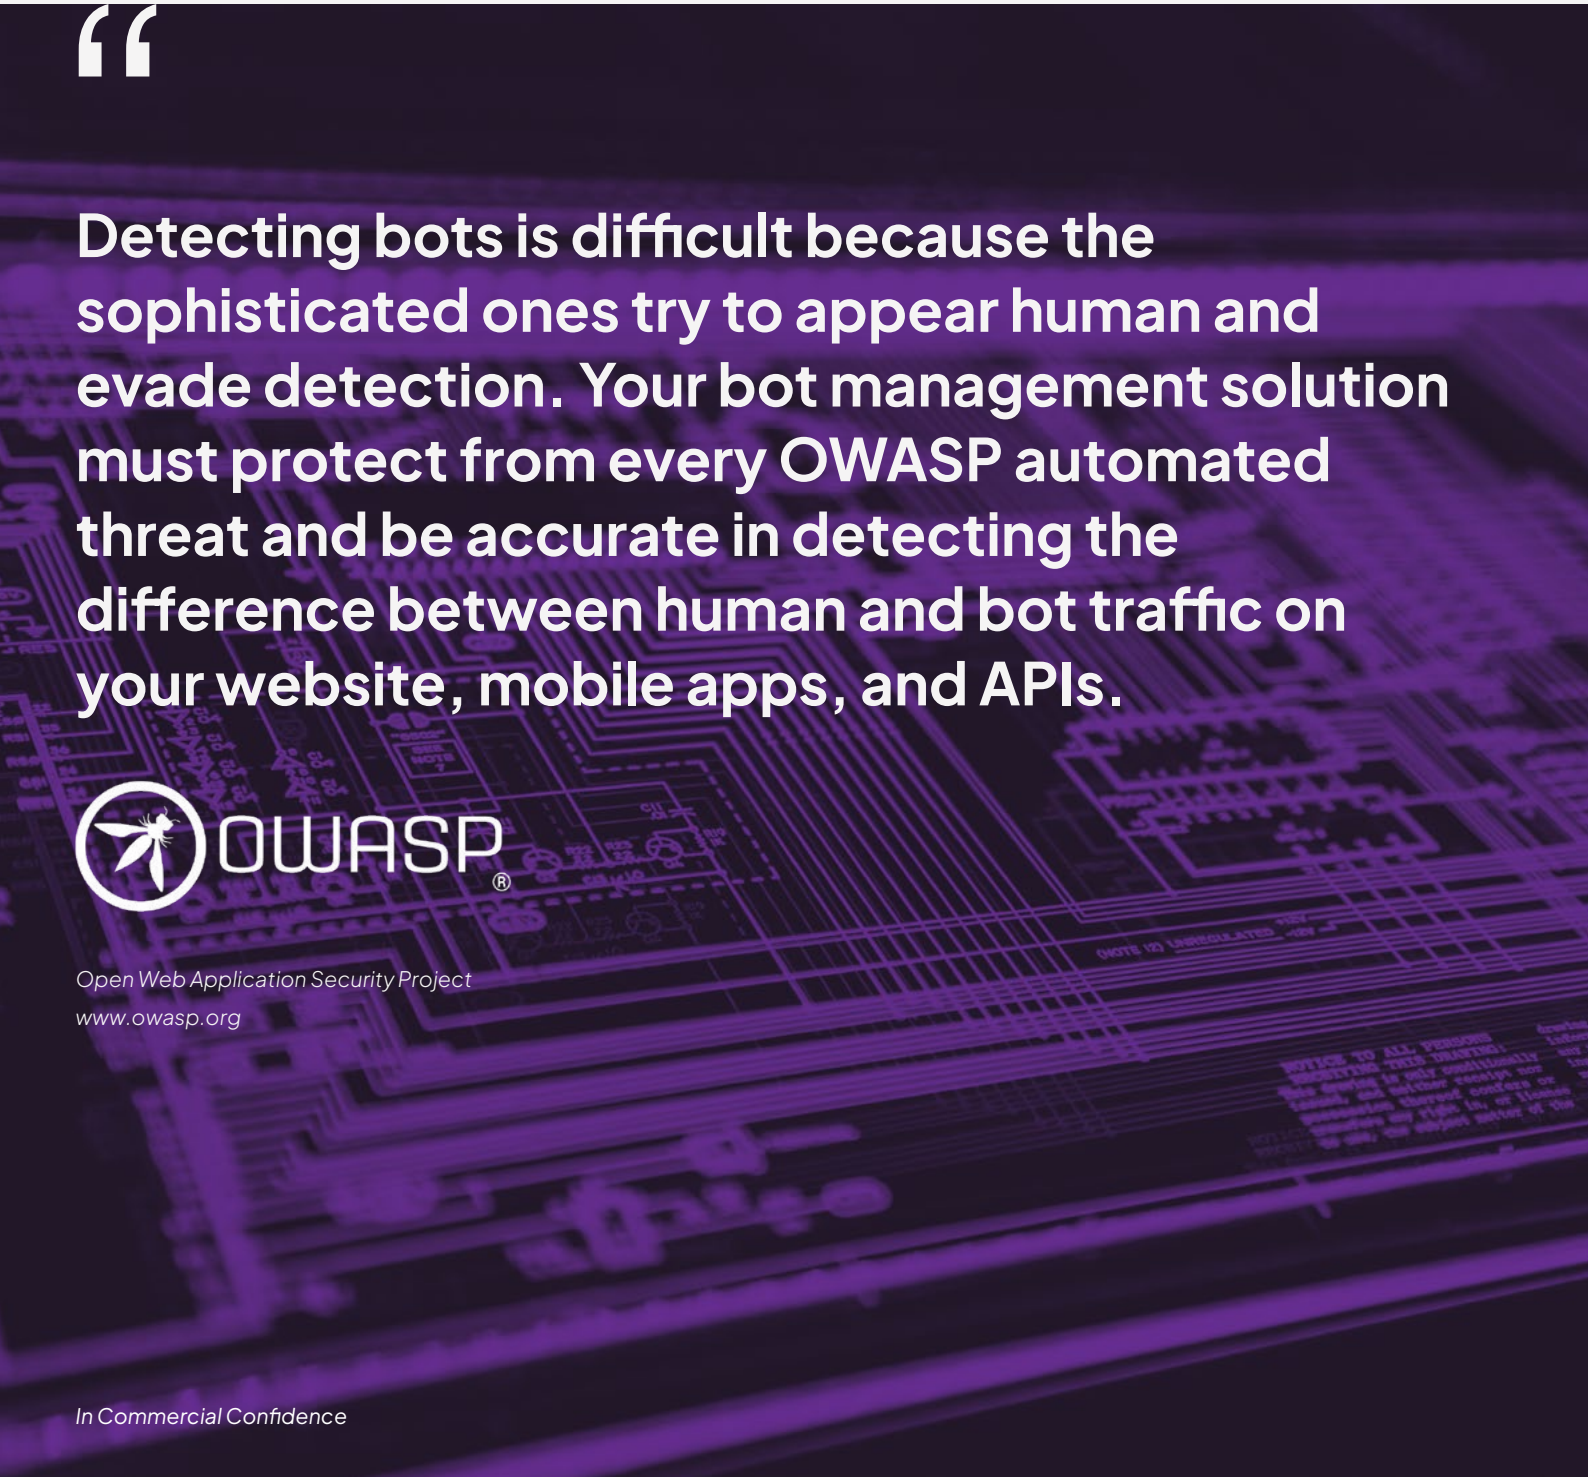

Detecting bots is difficult because the sophisticated ones try to appear human and evade detection. Your bot management solution must protect from every OWASP automated threat and be accurate in detecting the difference between human and bot traffic on your website, mobile apps, and APIs.

OWASP®

Open Web Application Security Project
www.owasp.org

Veracity Trust Network

02

How Veracity Uses AI to Detect Bots

The field of bot detection and protection is both sophisticated and rapidly changing.

Traditional bot detection systems are "rules-based". Algorithms of this kind are strictly defined by a set of rules that define the parameters of the problem and tend to be very poor at handling sophisticated and/or rapidly changing environments.

The problem with rules-based solutions is that they rely on explicitly defined rules. If a new type of behaviour emerges within the problem, the existing rule set will not account for this, and the rule set must be updated. This is a time-consuming and, often complex procedure.

As bot sophistication is constantly improving at an exponential rate thanks to AI advancements, it is becoming impossible to update rules-based algorithms quick enough to keep up with new types of bots. A more dynamic solution that does not require constant updating is, therefore, required to prevent the irreparable harm that sophisticated, malicious bots can reap.

At Veracity Trust Network, we have developed a Machine Learning (ML) based solution based upon an algorithm called "Random Forest", which has proven highly effective at detecting malicious bots.

What is a Random Forest?

A Random Forest is a network of decision trees, each of which can make a prediction from a data point as to whether it is a human or a malicious bot.

The algorithm learns how to correctly classify data points by finding an optimal weighting for all the trees based on how important it believes each of the trees to be. Trees determined to be more important are given a higher weighting and have more say over the final output.

When making predictions on data, it uses a voting system across all trees while taking their weighting into account.

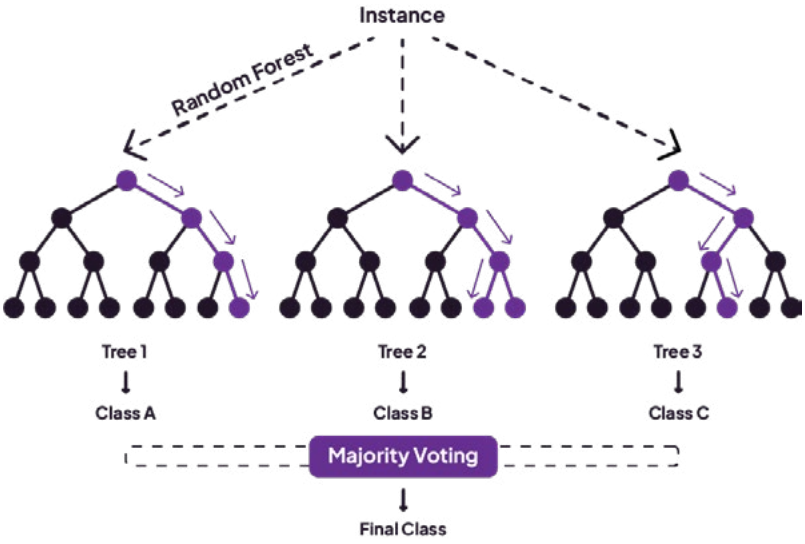

Veracity Trust Network

03

Futureproofing Industry Leading Bot Protection

Our current systems can detect and block 98.5% of malicious bots within 0.3 seconds of first page load, with the remaining 1.5% blocked on second page load.

These are excellent, industry leading detection rates, but we know that bot detection is an arms race and now that the bad actors have access to sophisticated AI at a fraction of the cost of 5–years ago, that we need to continue to innovate to maintain and even improve on the precision and speed of our bot protection capabilities.

At Veracity we place great emphasis on the importance of continuous research-backed development. The difficulty of the problem at hand requires the most cutting-edge academic research to be reviewed and implemented; we work with academic institutions across the UK and further afield to commission specific research work in areas of particular interest.

The goal of our Research & Development (R&D) is to future proof our bot detection solution so that it can adapt to the increase in bot sophistication for the foreseeable future. There are several research projects we are currently working actively, or which we have partially researched and plan to implement in the near future.

To properly outline the progression of our research and present a story detailing how the Veracity system is progressing, this section will be split into two sections:

1. Current R&D Projects
2. Planned R&D Projects

Current R&D Projects

There are two main projects that Veracity Trust Network are actively developing:

Computer Vision Model

Veracity's current ML model performs with high precision and can accurately identify the vast majority of malicious bots. The bots that it can accurately identify range from simple bots to competent and sophisticated bots that can simulate many human attributes, such as how long a human might spend on a web page or the number of mouse movements a human makes.

Due to the rapid increase in bot sophistication, there are now bots which can simulate more complex aspects of human behaviour such as superficially legitimate mouse movements or touch screen activity. However, human behaviour at a low level of granularity is almost impossible to properly replicate as there is an inherent level of randomness in human behaviour (in this case the way we interact with a webpage).

We have developed a solution, currently in beta testing, that focuses solely on differentiating the plot of a human's mouse movements over time and a bot's movements. Our solution is a "Computer Vision" algorithm, which is a subset of ML models that take in images as input rather than numeric input.

By taking historic user journey data from our many millions of anonymised journeys drawn from customer data over the past five years, and by plotting the mouse movements we can build up a training set consisting of images and use these to train the model. We have

developed a proof of concept and have already recorded very promising results on a small set of training data.

MLOps Framework

To keep up with the most recent developments in bot behaviour, our ML algorithm is regularly retrained on the most recent set of data to ensure that it can reflect new bot behaviours as they arise. However, retraining the model can be a manual, time-consuming process.

To remove the need for a manual process, Veracity is currently developing an "MLOps Framework", which is infrastructure built around a ML model that constantly monitors the performance of the algorithm and if the performance falls below a pre-defined threshold, a retraining cycle is automatically triggered. The final performance is then reported back to the Veracity team, who can decide whether further optimisation of the model is required. This saves the time manually retraining the algorithm, but more importantly it is a much more robust and reliable way to ensure that the model is maintaining high performance and automatically keeping up with the most recent developments in bot behaviours.

Veracity Trust Network have already implemented aspects of this framework, but in the future, Veracity plan to create a highly sophisticated MLOps framework that can support multiple Machine Learning models and handle a wide range of scenarios.

Veracity Trust Network

Planned R&D Projects

There are several research directions for which preliminary research has been started and where we plan to create proof of concepts in the near future.

An Unsupervised Learning Model

Veracity has investigated several ML models that operate by taking in training data along with a series of correct responses to each of these data points. This set of ML algorithms are known as "Supervised Learning" methods.

There are also a set of algorithms known as "Unsupervised Learning" algorithms which do not require labelled data at all. Instead, they learn by attempting, unsupervised, to find patterns or structure within the data to differentiate between different types of data. This is advantageous as it means that the data does not have to be labelled with correct responses, an activity which inevitably introduces a degree of bias into the model, although it is attempted to remove this bias as much as possible.

Veracity is currently exploring two options for Unsupervised Learning. The first of these is a "K-Means Clustering" algorithm which works by trying to group together data points based on how numerically similar they are. The second, and by far the most interesting and innovative, is a "Reinforcement Learning" algorithm. This model learns by trial and error to reach an optimal set of behaviours and is, therefore, a self-sustaining model that is constantly learning and improving. This model is at the forefront of innovation in the field of Machine Learning and the application of this to bot behaviour is a concept that is still in development even in the academic field.

Polynomial Curve Fitting

Returning to the concept of malicious bots increasingly being able to reproduce human characteristics, one behaviour we have observed is bots trying to simulate human mouse movements by tracing "Polynomial Curves". Put simply, a polynomial curve describes a line that is not linear (straight) and is a method used by malicious bots to avoid detection based on their mouse tracking.

Our research has indicated that this can be used to differentiate bots from humans as it is incredibly unlikely that a human would produce a mouse plot that exactly matches a polynomial curve. As this is a clearly observed and well-defined behaviour it makes sense to add functionality to the system that solely identifies this type of behaviour. We have performed some preliminary research into using a mathematical method known as "Regression" which attempts to fit a given curve to known polynomial curves and returns a similarity metric for each of the known curves. The theory is that if a user's mouse movements closely match a known polynomial curve, we can say with high certainty that this user is a malicious bot.

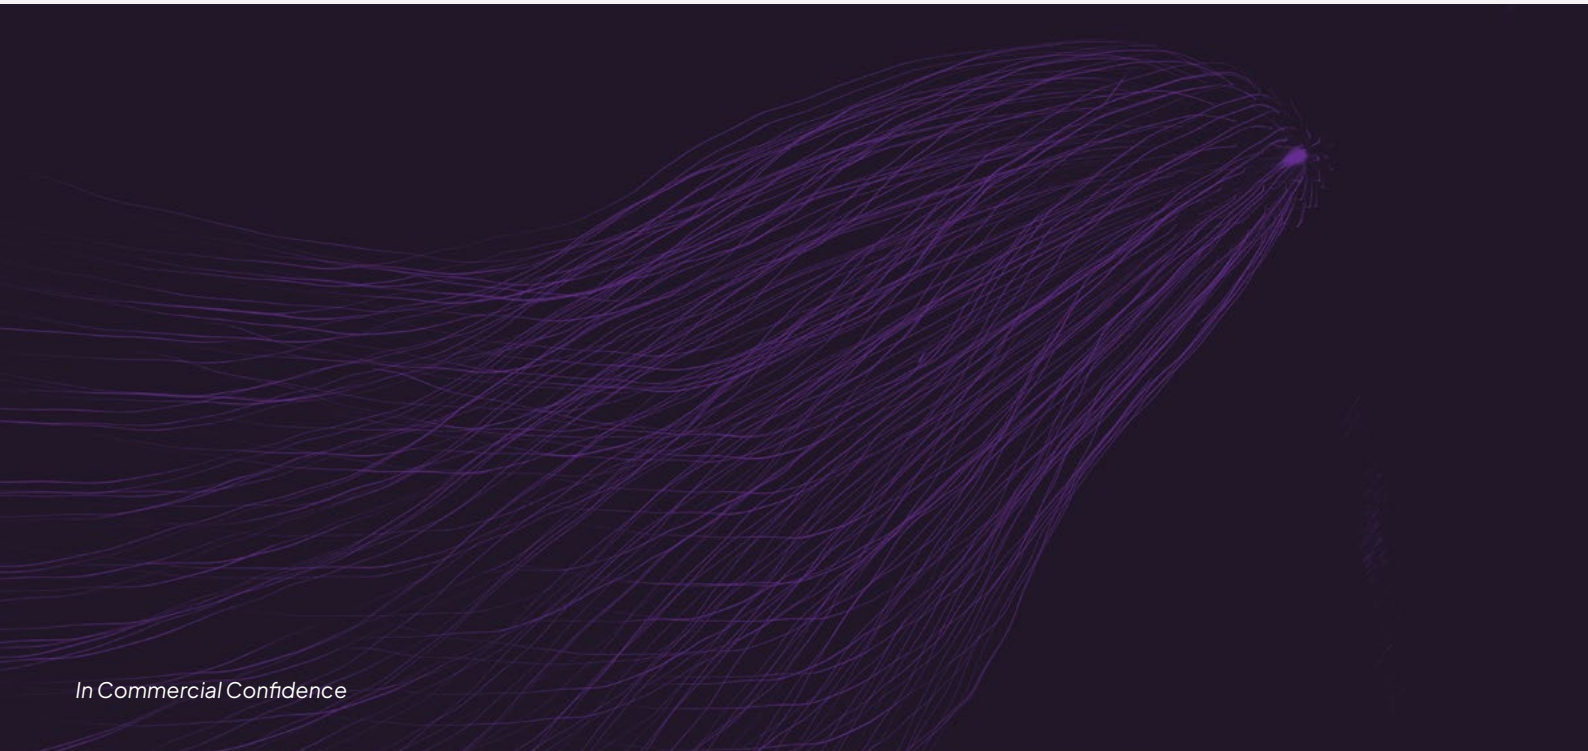

Veracity Trust Network

04

Conclusion by Stewart Boutcher, Veracity CTO

Artificial Intelligence is a powerful tool and there is no doubt that it is currently being adopted by many organisations for good. The recent McKinsey report on "The state of AI in 2023" indicated that two thirds of organisations are using AI or plan to use AI, although in many cases this is generative AI – such as ChatGPT – rather than Machine Learning algorithms, which are more limited in their scopes.

Using the Machine Learning tools we discuss in this cyber brief to detect malicious bot activity more quickly and accurately is, I am certain, critical to safeguarding the Internet and web as we know them.

AI has reached a stage where adoption is substantially easier for everyone, including to those who wish to use them for malicious purposes. We are already seeing the tell-tale traces of generative AI in more sophisticated bot types to mask malicious behaviours; this will only accelerate over time.

At Veracity Trust Network, our focus is continuing to innovate in malicious bot detection and prevention using whatever technology we can, keeping our customers and their legitimate visitors safe online.

Veracity Trust Network

05

About Veracity Trust Network

Veracity Trust Network is an award-winning cybersecurity company, specialising in the detection and prevention of malicious bot activity in the web sphere: websites, web and mobile-web applications and APIs, with native mobile protection to follow.

Veracity Trust Network are expanding the boundaries of the application of ML to enhance the speed and accuracy of malicious bot detection on the web. We are award winning, tech-led, and solution-focused.

Authored by:
Stewart Boutcher, CTO Veracity Trust Network
Reuben Sodhi, Lead Data Scientist, Veracity Trust Network
Shivang Chaudhary, Data Researcher, Veracity Trust Network
Muheeb Ahmed, Data Researcher, Veracity Trust Network

veracitytrustnetwork.com
UK +44 (0) 5603 861 037
US +1 (833) 286 6284

hello@vtn.live

VeracityTrust
veracitytrustnetwork