NETACEA

Technical Showcase: How Netacea developed Intent Clustering technology

IN PARTNERSHIP WITH THE UNIVERSITY OF NOTTINGHAM AND KTP

EXECUTIVE SUMMARY

Malicious bots attack at speed, so it is vital that defenses react quickly to stop them. With reams of requests being processed on servers every minute, distinguishing between a full spectrum of malicious and legitimate traffic quickly and accurately requires advanced technology developed with specific expertise.

Most traditional bot management solutions make compromises (such as speed, accuracy or adaptability) when attempting to mitigate malicious traffic due to limitations in technology, infrastructure or knowledge.

Netacea wanted to add adaptive real-time anomaly detection to its bot management offering. To augment the project with the highest level of expertise, Netacea embarked on a Knowledge Transfer Partnership (KTP) with the University of Nottingham.

The KTP allowed academics from the University of Nottingham to apply their “clean room” research into machine learning to the real-world constraints of a business environment, bringing fresh perspectives and unique insights to the project.

The resulting technology, Intent Clustering, overcomes the constraints of other solutions to provide a solution that is scalable, flexible to new threat types, and can mitigate malicious bot traffic in real time with a high level of accuracy.

ABOUT THE KNOWLEDGE TRANSFER PARTNERSHIP

The KTP scheme is funded by the UK Government's InnovateUK initiative and encourages collaboration between industry and academia, fostering business growth and innovation. Academic institutions get access to real world data and scenarios to develop their research, whilst businesses benefit from access to the latest cutting-edge research to help develop new solutions or augment their existing offerings to customers.

This KTP had several goals, with one specific objective to support Netacea in evolving our anomaly detection solution, which identifies unusual patterns of behavior within website traffic to help better detect bot attacks. Netacea wanted to categorize traffic into "anomalous" or "normal", understand the users' profile, and uncover their intent.

To achieve this, the University of Nottingham worked closely with Netacea's data science team in developing new artificial intelligence systems. As well as producing a new piece of technology, David Fricker, who has a Master's degree in computer science from the University of Nottingham, advanced his career in cybersecurity by joining the Netacea team as a data scientist. He is now Senior Machine Learning Engineer at Netacea.

Meet the team

The KTP team for this project brought together experts from Netacea and the University of Nottingham:

Netacea

- Andy Still, CTO
- Dr Mark Greenwood, Chief Architect
- Dr Matthew Jackson, Head of Data Science
- David Fricker, Senior Machine Learning Engineer

University of Nottingham

- Prof Bob John, Professor in Fuzzy Logic
- Prof Dario Landa-Silva, Professor in Optimization
- Dr Mercedes Torres Torres, Assistant Professor in Machine Learning

THE PROBLEM

Fraudsters are deploying business logic attacks against their targets at an increasing rate, utilizing automation and sophisticated bots to do so. As businesses implement defenses against such threats, bad actors adapt and re-tool, evolving the threat landscape in a constant game of cat and mouse.

As they use automation, bot attacks are typically very aggressive and can fire hundreds of requests at servers within a short space of time. Yet, on large websites these malicious requests can get lost in the noise of busy web logs. A recent Netacea survey of 440 enterprise businesses revealed that bot attacks go unnoticed for an average of 14 weeks – usually only being discovered after their damage has been done.

Netacea wanted to gain a deeper understanding of user intent by investigating methods to detect this anomalous behavior as requests are logged on the server in real time.

Target attack types

Two common attack types that this kind of real-time solution would be most effective against, due to the volumetric nature of most attacks, are account takeover (ATO) and credential stuffing. Both attacks are high risk and can have immediate tangible impact on their targets, making timely detection imperative.

Account takeover (ATO)

ATO occurs when an attacker illegally logs in to a user's account. The threat puts both customer and business at significant risk, exposing personally identifiable information (PII), preventing access to accounts and enabling the attacker to make fraudulent transactions.

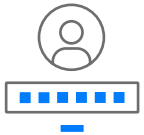

Credential stuffing

Credential stuffing is a common account takeover technique used to gain brute force access to an account by continually and automatically injecting usernames and passwords into website login forms until a match is found.

What are Bots Costing Your Business?
Read the full report from our recent survey

Robust to concept drift

While ATO and credential stuffing are well-known attack types, new types of attacks are constantly being developed and could strike businesses at any time. The methods and tools by which bad actors attack their victims are also constantly evolving. Websites themselves also change over time, as do their traffic patterns, for example during peak periods like Black Friday. These constant changes have previously made handling the defense of such sites a manual and time-consuming process.

To manage this, we need a solution that is flexible to different types of attacks and robust against concept drift.

Bringing academic research into real-world situations

The advantage of academic research is that solutions are often unconstrained by issues such as budgets, time limitations or client requirements. Researchers can think “outside the box” and explore cutting-edge ideas in a “clean room” environment, even in areas like machine learning.

Calling on the expertise of professors in machine learning from the University of Nottingham presented new perspectives for the Netacea data science team. However, these concepts had to fit into the requirements and restraints of the business objectives, which posed their own challenges.

Real-world business constraints

Academic researchers in the field of machine learning often have the luxury of working with batches of data, reiterating their algorithms to explore different outputs. Netacea required the algorithms to work in real time rather than batches, which is much more challenging to achieve.

Infrastructure constraints

In academia, there is freedom to explore solutions without needing to be overly concerned with cost effectiveness. This is not the case for businesses, where financial resources must be accounted for. This presents challenges when scaling solutions, meaning the view of the world must be shortened and optimized.

COMPROMISES OF CURRENT SOLUTIONS

In most bot detection solutions, corners are cut to make detection easier or faster. Taking these shortcuts results in less accurate, slower solutions that fail to spot new attack types or threat actors.

Common shortcuts include:

Processing in batches, not real-time

Most solutions analyze data in batches, sacrificing the speed of their mitigations. By sorting through data in “windows” of time, most solutions either overlap their analysis, processing the same data more than once, or leave gaps where data is missed, resulting in lower quality outputs. Real-time mitigation is crucial, as every second counts during an attack.

Reliance on static lists

Some solutions have relied on detecting known bad actors to mitigate their attacks using static reputation lists; this has allowed attackers to bypass defenses by distributing their attacks through botnets, residential proxies and rotating IP addresses.

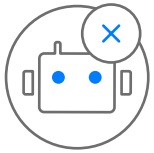

Dependence on client-side bot detection

The overwhelming amount of data logged by user requests into web systems has also discouraged solutions from creating actionable information from this source. But client-side detection such as JavaScript is visible to attackers and can be reverse-engineered using easily accessible tools, making it straightforward to bypass without much technical skill needed.

THE SOLUTION

While client-side detection can spot less sophisticated threats, and reputational lists are a useful resource to speed up detection of known bad actors, in isolation both leave large gaps that previously unseen or more sophisticated attacks can easily slip through. Therefore, Netacea's solution required a different approach that did not rely on either method.

Instead, the teams focused on something no attacker can ultimately mask – intent. No matter where malicious bots originate or how sophisticated they are, they will always behave in particular ways to achieve their core objective.

By avoiding the previously mentioned shortcuts, Netacea and the University of Nottingham aimed to produce a solution that detects bots in real time, is adaptive to new threats and environments, and is comprehensive of all anomaly types – both known and previously unseen.

INTENT CLUSTERING – A TECHNICAL OVERVIEW

The result of this part of the KTP between Netacea and the University of Nottingham is a flexible, robust end-to-end streaming detection engine of novel attacks. Intent Clustering works in three main steps.

Log data pre-processing

The first step is to gather log data wherever they are made. At Netacea we collect logs at the server level to capture all traffic across websites, apps and APIs.

```
127.0.0.1 user_agent_1 [17/Sept/2020:13:55:36 -0700] "GET /robot.txt HTTP/1.0" 200 2326
127.0.0.1 user_agent_1 [17/Sept/2020:13:55:38 -0700] "GET /products HTTP/1.0" 200 4684
127.0.0.1 user_agent_1 [17/Sept/2020:13:55:40 -0700] "GET /products/list_1 HTTP/1.0" 200 5117
127.0.0.1 user_agent_1 [17/Sept/2020:13:55:42 -0700] "GET /products_list_2 HTTP/1.0" 404 5171
127.0.0.1 user_agent_2 [17/Sept/2020:13:55:44 -0700] "GET /products_list_3 HTTP/1.0" 200 7852
127.0.0.1 user_agent_2 [17/Sept/2020:13:55:44 -0700] "GET /products_list_4 HTTP/1.0" 200 7852
```

With our data collected, we then validate and normalize the logs. This means that wherever the logs originated, our system can work with them in the same way and transfer findings across sources and clients.

This normalized data is enriched with multiple context and intelligence streams, outputting enhanced logs with much more useful information associated with them for the next stage.

Real-time clustering

We apply real-time dynamic clustering to each enriched log line by implementing a custom algorithm. The goal of the algorithm is to create clusters of data that share similar intent.

This algorithm was developed in collaboration with the machine learning team at the University of Nottingham, using the latest research in the field and customized to the exact needs of the project.

Because the clustering is dynamic, it not only acts in real-time but is also robust to concept drift. Over time and in response to new data, clusters will move, merge or even die off if no longer relevant.

Real-time clustering using this algorithm is also highly resource efficient, with guaranteed caps on memory usage. Bayesian optimization is applied to the parameters of clusters for each customer to improve performance and reduce the time taken to onboard new customers, as the solution is very flexible to nonstandard log types and does not need manual intervention to adapt to new environments.

Machine learning classification

Clustering creates an abstraction on the data that helps build robust classification systems, sorting requests into “benign” and “malicious”, for example. Automatically generated models are deployed to investigate the intent of each cluster in real time.

An investigation is triggered if anomalous clusters are found, which can be reviewed by hand to verify the information. These classifications are also easily interpretable, providing clients with a safe level of service and evidence of logical outputs.

The output: Recommendations

Once malicious bots have been classified and identified, we are able to make recommendations to the customer. The recommendation to block can be passed on via a threat intelligence feed, which allows customers to react to attacks according to their own appetite for risk. This typically depends on factors such as their vulnerability to certain types of attack (for example, any site with a login page is susceptible to ATO and credential stuffing attacks) and the type of information exposed to bots.

Netacea can also instantly block bad bots upon detection. The type and level of blocking (from single IP or ranges of IP escalating up to specific user agents) is applied intelligently based upon learned heuristics.

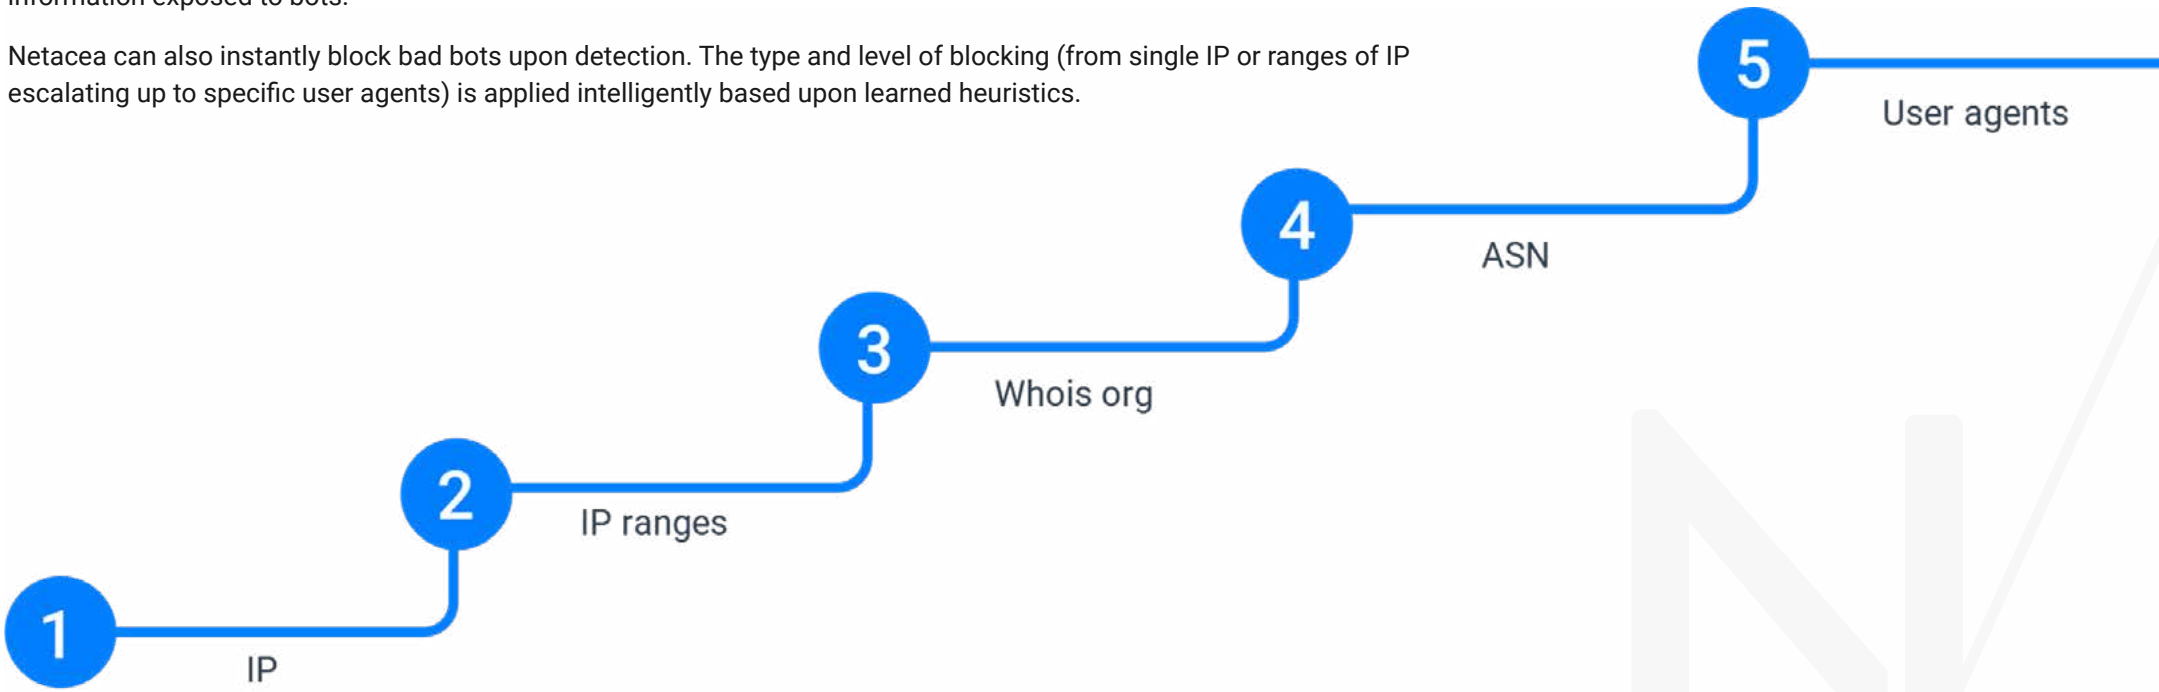

INTENT CLUSTERING IN ACTION

The charts below show the results of Intent Clustering in a live environment protecting two different client sites.

In Fig.1, we can see a short but high-volume spike of mitigated (bad) traffic in red, with the remaining (good) traffic in green. Fig.2 shows the traffic levels once the malicious traffic is removed, which highlights the impact of Intent Clustering in removing anomalous traffic.

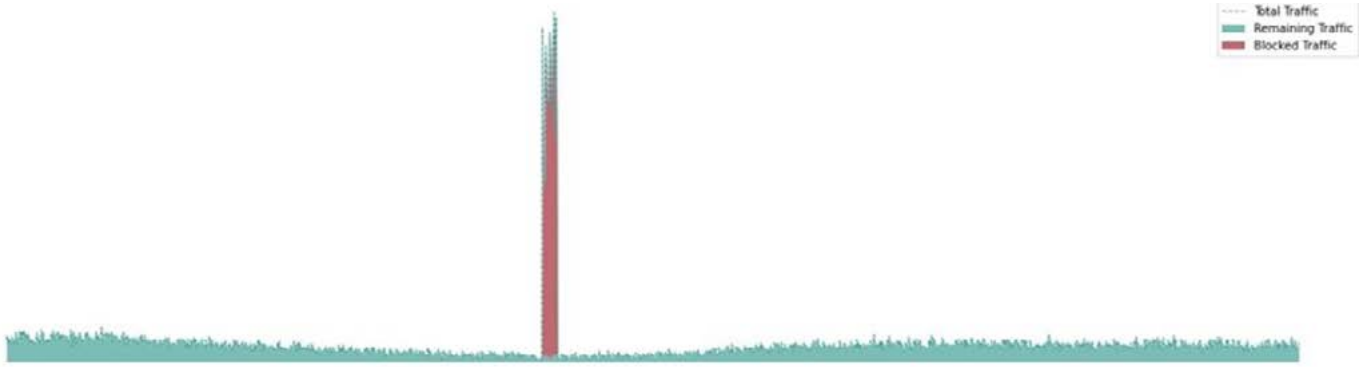

Fig.1

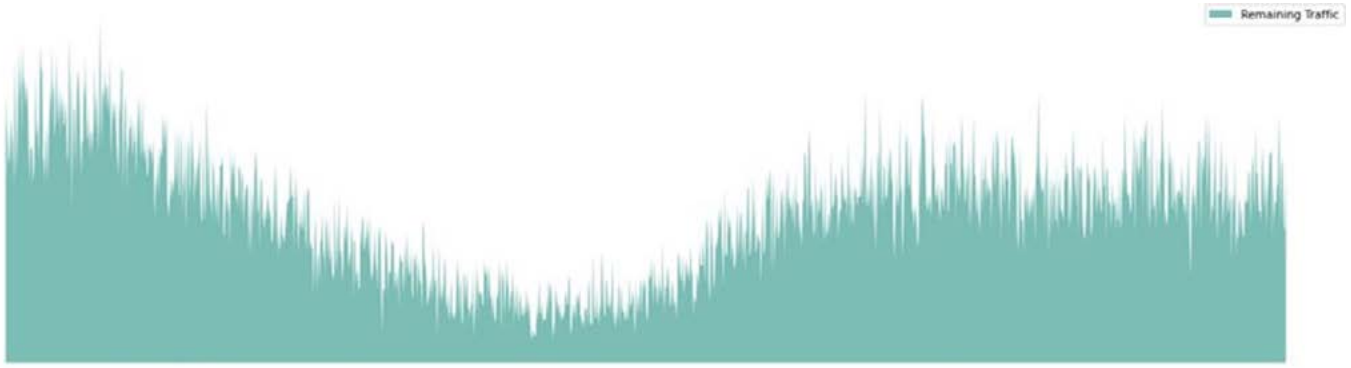

Fig.2

NETACEA

In a second example from another client, Fig.3 shows a huge number of malicious attempts to access a website in the red spikes, sustained throughout a full day. Once this is removed, it is clear to see traffic levels are much more normal and expected in Fig.4, with low traffic early in the morning, picking up during the day and then slowly lowering into the evening and night.

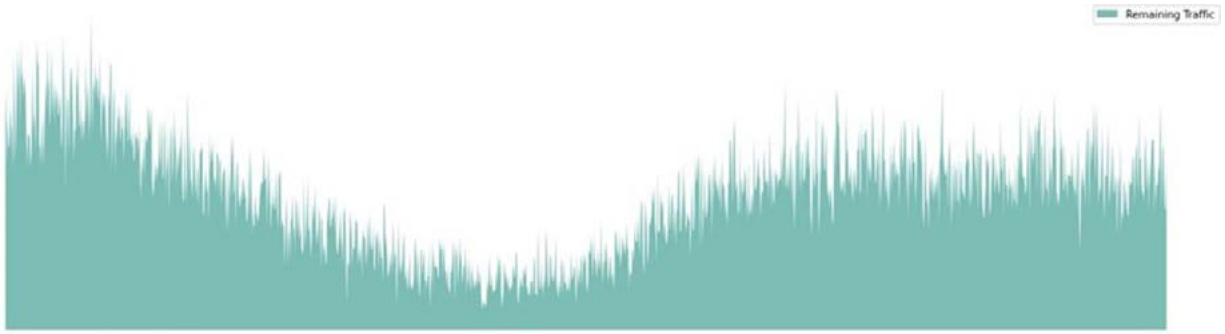

Fig.2

Fig.3

CONCLUSIONS

Intent Clustering is simple in concept. It analyzes every incoming request, groups them into clusters, classifies their intent based on behavior signals and then makes recommendations for mitigation.

In reality, the technology is complex and required cutting-edge research to actualize. Classifying such large amounts of data in real time to output accurate recommendations is extremely complicated. Collaboration between academics at the University of Nottingham and machine learning specialists at Netacea created a unique environment for the data science models to be developed.

The result was a solution that scales well and not only detects known threats, but can also detect novel attack types and threat actors. Additionally, Intent Clustering is fully adaptable to new log types, allowing new customers with unique logging formats to onboard and see benefits much faster than before.

The successful KTP project has produced a system that has since been proven highly effective on thousands of attacks in production. Intent Clustering is now actively detecting anomalous traffic for 90% of Netacea's clients and has a patent pending.

The project has also been a great success from an academic perspective, with a full academic publication underway. InnovateUK, which funds Knowledge Transfer Partnerships, has also certified the Intent Clustering project as "outstanding", the highest rating on its scale.

GET STARTED WITH NETACEA BOT MANAGEMENT

Intent Clustering is just one layer of defense against business logic attacks that forms the wider Netacea suite of bot management.

Powered by the Intent Analytics™ ecosystem and backed by bot experts, Netacea Bot Management is the most advanced and accurate defense against malicious bot traffic available.

Try Netacea Bot Management for your business

Do you want to find out how much more bot traffic you could be blocking? Get in touch to arrange your free trial of Netacea Bot Management.