




Understanding machine learning for fraud prevention



Contents

Introduction: Zero day and the evolution of ad fraud	3
Why machine learning should be part of your fraud defence	5
The time is now for machine learning	10
What does machine learning for fraud prevention look like	12
In Summary	15



Zero day and the evolution of ad fraud

If you think 3ve was a sophisticated operation, just imagine what all the other operations look like that continue to evade the grasp of law enforcement.

3ve reportedly made US\$29 million over the course of its operations – that is a substantial amount of money but in the context of ad fraud as an industry, it is barely a drop in the ocean. For some perspective, more money is lost to ad fraud in one day than 3ve was reported to make in 3 years.

3ve is an example of what some of the less sophisticated operations look like.

The 3ve investigation relied on an unprecedented level of cooperation, spanned 3 years and no doubt, cost a substantial amount of money to pull off. The reality is that dealing with fraud by bringing every fraudster to justice isn't feasible. But while we may not be able to catch them all, that doesn't mean we should give them our ad spend on a silver platter.

Fraud prevention is the best way to beat these organisations. Stop the fraud, make sure the fraudsters don't get paid and make the business of fraud less profitable.

Ad fraud and the sophistication of perpetrators have evolved at a rapid pace in the last 20+ years. Before 2000, it was mainly individuals exploiting loopholes in a still relatively juvenile internet advertising ecosystem. As budgets for online advertising grew, so too did the lucrativeness of ad fraud as a business opportunity, attracting more sophisticated players to the arena. Now ad fraud is an industry in its own right, worth approximately \$19 billion in 2018.

The evolution of ad fraud

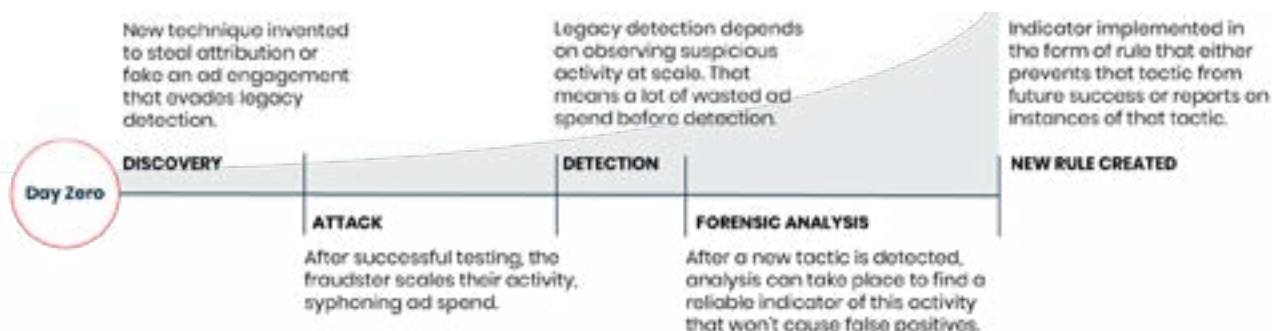


Ignorance on a grand scale allowed fraudsters to make a lot of money for a very long time and gain a substantial head start on anti-fraud providers. By the time the industry woke up and smelled the bull\$hit in their ad spend, the pockets of fraudsters were deep enough for them to adapt and conduct their own R&D to find new ways to exploit digital advertising.

The fraud prevention space has taken a similar trajectory as the cyber security space – albeit a few years behind. Initially we blacklisted (like anti-virus software) and then we wrote rules (like firewalls). But fraud continued to evolve and adapt making it harder and harder to add to blacklists and manage increasingly complex rule sets, without catching legitimate traffic in the crossfire.

What is zero day ad fraud?

In cyber security, a zero day threat is a vulnerability that there is no patch for. Borrowed from cyber security, we define zero day ad fraud as a new tactic or variation of a tactic that rules are not equipped to detect or block.



As the tricks and schemes of the last 20 years become less successful for fraudsters, it is likely we will be seeing more new types of fraud as perpetrators adapt and evolve. Using rules to address these new tactics would be a fool’s errand that perpetuates the cat and mouse game that has characterised fraud prevention until now. It is time to turn the table on the fraudster by introducing a new approach to fraud prevention that is proactive instead of reactive, removing fraud as it evolves.

Why machine learning should be part of your ad fraud defence

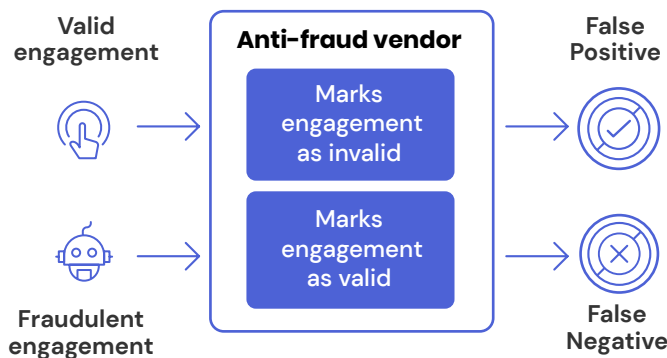
Before we discuss why machine learning (ML) is an essential component of your arsenal against zero-day fraud, we will first discount non-ML fraud mitigation in the context of zero day.



Blacklists

Blacklists are an important part of any ad fraud defence because they quickly identify sources that categorically don't have human traffic such as servers. Blacklists are a very basic first step that remove the lowest-hanging fruit but fraudsters can easily circumvent them by changing the IP addresses of their traffic. When blacklisting IPs that can have human traffic rather than isolating the fraud itself, blacklists can actually result in high volumes of false-positives.

Flow of false identification



False positives occur when a valid install is marked as fraud.

False negatives are when solutions fail to catch fraud, marking invalid traffic as valid.

Rule-based detection and mitigation

Rule-based mitigation involves identifying characteristics and thresholds that, when exceeded, block traffic or a traffic source. Rule-based mitigation is good when you know the characteristics that define a particular fraud tactic, such as an impossibly short click to install time. But, it is near impossible to formulate rules for fraud tactics that you have never encountered before.

For a rule to be created, a new fraud type must be observed at scale which means it is impacting your budget for some time before it can even be recognised. Analysts then need to find the characteristics to confidently filter it from valid traffic and distil them into rules. This process takes time – if you rely solely on rules, you are exposed and fraud is taking your ad spend until that rule exists to stop it. Rules are reactive – a new fraud type exists, then a rule is created. It is always fraud first, then rule.


To stop new fraud tactics, or zero day fraud, and stop the flow of money to fraudsters, a more proactive approach is required. That is where machine learning comes in.

What is machine learning?

Machine learning is a subset of artificial intelligence that extracts patterns and relationships from data and expresses them as a formula that can be applied to new data sets. Over time, as the data changes, new patterns are learned by the model without the need to explicitly program them.

What makes machine learning suited to combating new fraud types?

Because of the scale of data processed, insights can be more valuable and derived much faster using machine learning than by using human analysis alone.



More thorough analysis

Multi-dimensionality allows for a much deeper understanding of data, enabling machine learning to detect fraud in traffic that might initially appear valid to an analyst.

Humans are normally limited to analysing data on between 2 and 4 dimensions. Within that plane, data can be visualised and patterns easily recognised. Selecting which dimensions to explore and then running manual analysis is a time consuming process. Machine learning explores all the potential dimensions of a data set to decipher relationships not evident in 2 dimensional analysis.



Speed of analysis

Machine learning (supported by the right data architecture) can make predictions quickly, enabling fraud prevention to run in near real-time and at scale.

Human analysis is just too slow. In order to stop fraud before fraudsters get paid, fraud needs to be mitigated in real time. Analysis needs to scale as traffic fluctuates too – you can't just keep employing an endless string of analysts.



Contextual

Machine learning is contextual to make every validation based on the circumstances and characteristics of that exact transaction.

Where rules are static and based on a generalisation of normal behaviour, machine learning is much more sympathetic to the conditions of each traffic transaction.



Evolving

Continuous training, verification and self-learning ensure that machine learning models evolve in-line with changes in norms associated with valid traffic.

Human behaviour changes all the time and can vary greatly in different demographic or geographic audiences. For fraud prevention to be effective, it needs to adapt to these changes and audience nuances ensuring legitimate traffic isn't removed in the fraud mitigation process.




Proactive

Proactive by nature, machine learning effectively trained to understand valid traffic will identify invalid traffic regardless of the tactic employed (known or unknown).

With an accurate picture of what legitimate traffic looks like, it is much easier to identify invalid traffic. In order to future proof fraud mitigation, we need to stop reacting to fraud tactics and start only permitting valid traffic.

Juniper Research forecasts that by 2022, machine learning could save advertisers over \$10 billion a year in ad spend that would have been wasted on fraud.

Ad fraud savings derived from machine learning (USD)



Source: Juniper Research

“ML tools, such as those used by TrafficGuard, will enable the fight against ad fraud to move from detection to proactive mitigation in real-time.”



Fraud is a constantly moving target. Sophisticated ad fraud doesn't look like bots. On the surface, it looks like human engagement. Fraudsters adapt their processes to circumvent rules-based fraud detection and their profitability hangs on their ability to evolve.

Instead of reacting to fraud as it evolves with new rules, machine learning can be part of a proactive defence that is tactic-agnostic, more accurate and able to stop fraud before the fraudster gets paid.

The time is now for Machine Learning

Some commentators think machine learning is too new a field to be deployed for ad fraud mitigation. In the last few years, expertise and technological developments have come a long way in the field of machine learning. Widely used in the cyber security space already, ad tech seems to only just be catching on to the value of machine learning in ad fraud prevention.

Did you know the term machine learning dates back to the 1950s? So how is it that a topic older than the compact cassette tape, is one of the hottest topics in technology today?

The key drivers behind machine learning's prominence today include:

- **Scalable infrastructure** – Machine learning needs high volumes of data to reliably identify fraud. Data volumes can fluctuate depending on a variety of factors including season, so being able to scale infrastructure has driven investment in machine learning.
- **High compute power** – Ingesting high volumes of data, and in TrafficGuard's case, streaming data, requires a lot of compute power to invalidate, report and mitigate traffic in real time.
- **Affordability** – Previously, in order to access the processing power required, you would need to own your own infrastructure, making machine learning prohibitively expensive. Today, businesses can easily access powerful infrastructure on a scalable subscription model, making it more affordable.
- **Access to expertise** – As machine learning has become more feasible from a technological standpoint, data science, infrastructure and operations expertise have quickly evolved to make use of the technology.

The combination of accessible tech, affordability and skills development means that machine learning is finally at a point that it can be used to solve critical business challenges.

Naysayers that say machine learning is too new are likely the ones that have been slow to invest in the space, and are now trying to discredit the field.

As well as expertise and technology, there is one more key reason that the time for machine learning in fraud prevention is now. That is the adversaries. We are in an arms race with fraudsters. Every day new tactics are discovered and new fraud operations make the media.

Two dimensional analysis, static rules and blacklists can only take us so far. Fraudsters today are well funded operations that funnel billions out of the digital advertising industry.

At TrafficGuard, we are less concerned about the fraud-du-jour and more focused on what characterises genuine advertising engagement. While our competitors are busy talking about click injection, SDK spoofing, app install farms, our stance might seem novel. But in our view, chasing tactics is an unsustainable approach to fraud prevention. Why? Because fraud adapts – if one tactic is blocked based on a certain signature, fraudsters will innovate (like any business) in order to do their job better.

Using machine learning we bring a sustainable approach to fraud prevention to protect ad spend from tomorrow's fraud, as well as today's.



What does machine learning in fraud prevention look like?

Complex rules-based fraud prevention is no match for the sophisticated fraud stealing ad spend today. Well funded fraudsters adapt and innovate, like any successful business, when faced with new threats to their profitability. So in this arms race against fraud, what does machine learning look like?



There are 4 essential elements driving machine learning success

1 Infrastructure

From a technological standpoint, ML needs a sophisticated engine that can scale with fluctuations in volume, that can efficiently process high volumes of data, both batch and streaming and deliver actionable insights in near-real time.

2. The human element

Contrary to popular belief, machine learning isn't self sufficient. Skilled analysts need to define problems, identify appropriate technology, prepare the data, store it in task-specific locations, train machine learning models, manage and continuously verify models over time. The role of the human is integral to successful machine learning.

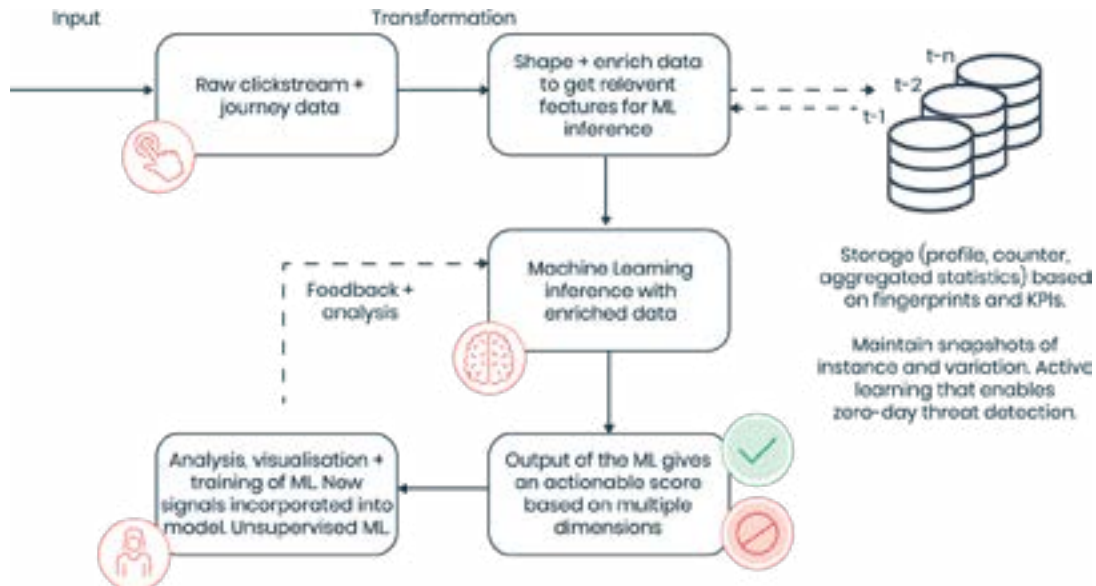
3. Data

As our head of Data Science Raigon Jolly says, "rubbish in, rubbish out". Data management involves rigorous preparation, labeling, and task-specific storage to ensure that the right data can be retrieved for specific functions. Types of data include behavioural patterns, location data, transactional data, device and network data.

4. Algorithms

It goes without saying that algorithms are an essential ingredient to machine learning but you would be surprised by the number of folks that say they have machine learning capabilities without them. In fraud prevention, machine learning needs to use computationally efficient models - deploying machine learning for tasks that can be reliably resolved with rules, only serves to slow everything down. Machine learning should be additive, not the exclusive means of fraud prevention.

TrafficGuard combines models and techniques to build unique, contextual knowledge about digital activity and user behaviour that leads to the most advanced fraud protection.



Every click, conversion and event is received by TrafficGuard along with hundreds of data points that characterise that transaction, like source IP, device, operating system, time of day etc. The transaction’s record is saved in the appropriate location and enriched by all of the other data in TrafficGuard – all the other times a device has been seen, other transactions on the same campaign, across campaigns by the same supply source etc.

With the context of that specific transaction and the trillions of data points TrafficGuard has been built on, it can confidently say whether a transaction is valid or invalid.

Machine learning models are used to validate transactions based on enriched data. In the case of zero-day threats, deep learning is particularly useful. Deep learning is the function of layers of neural networks capable of processing very large and highly dimensional data sets to uncover latent relationships in data. TrafficGuard’s neural networks ingest raw, unlabeled data to recognise patterns, cluster transactions together and assist in classification. Our neural networks are an unsupervised machine learning technique, making them critical in the fight against zero day ad fraud because they don’t rely on prior classification of fraud types. Neural networks also support other machine learning algorithms for reinforcement and regression.

Deep learning models validate transactions and the valid/invalid classification then gets fed back into the data to help future determinations of validity.

In Summary

There is no one correct way to apply machine learning and many businesses with varying levels of expertise are jumping in, in order to not be left behind. Getting all of the elements of machine learning performing reliably takes time and dedicated, experienced data science and dev ops teams. A company at the start of its machine learning journey should be treated cautiously.

At TrafficGuard, machine learning enables our award winning fraud prevention to to:

- Reduce false positives with precision fraud mitigation
- Reduce false negatives to catch fraud that other vendors or measurement platforms miss
- Mitigate fraud from known and unknown tactics
- Drive our relentless pursuit to prevent fraud at the earliest possible opportunity, supporting our efforts to stop fraudsters from getting paid



About TrafficGuard

TrafficGuard provides comprehensive ad verification, measurement and fraud protection for brands, apps, agencies and ad networks. Purpose-built specifically to fight ad fraud, TrafficGuard analyses multiple stages in the advertising journey – impressions, clicks and events such as sales or app installs – to remove invalid traffic at the earliest reliable diagnosis. By doing this, TrafficGuard improves campaign optimisation, drives ROAS and saves time usually wasted on media volume reconciliations.

Crucially, with visibility of impressions and clicks before the install, TrafficGuard is able to remove fraud earlier than measurement platforms. And with that visibility across multiple stages, attribution made by measurement platforms can be verified to prevent misattribution that skews performance data and wastes time.



Knowledge is power! Get the full picture of your performance advertising traffic quality with a free, no-obligation Traffic Quality Audit.

[CONTACT US](#)