The evolution of fraud detection models

From velocity checks to sequential transformer models

Accertify

Executive Summary

Fraud detection has evolved in response to increasingly sophisticated attack methods. What began as simple velocity checks has grown into a layered defense strategy incorporating rules engines, machine learning, graph networks, and sequential transformer models. Each stage emerged because the previous approach could no longer keep pace with fraudsters' tactics. This paper explores how each fraud detection model arose, what types of fraud it was best at catching, and why merchants need an integrated approach to stay ahead.

Introduction

Fraud Detection Through Time: From Rules to AI

The history of payment fraud is one of constant evolution. As fraudsters refine and diversify their tactics, merchants have had to adapt their defenses.

Fraud detection has progressed from simple velocity checks to rules engines, then to machine learning models (such as logistic regression or decision trees), graph networks, and now sequential transformer models. These stages arose because the previous approach could no longer keep pace with emerging threats. Each mechanism throughout the history of fraud detection has brought unique strengths to the table.

Fraud prevention is cumulative. New methods do not replace older ones entirely; they supplement them. This layered approach creates resilience against diverse attack vectors. In today's environment—where commoditized fraud tools and AI-driven schemes are widely available—merchants need defenses that combine human expertise with advanced machine intelligence.

Evolutionary Timeline

Evolutionary Timeline

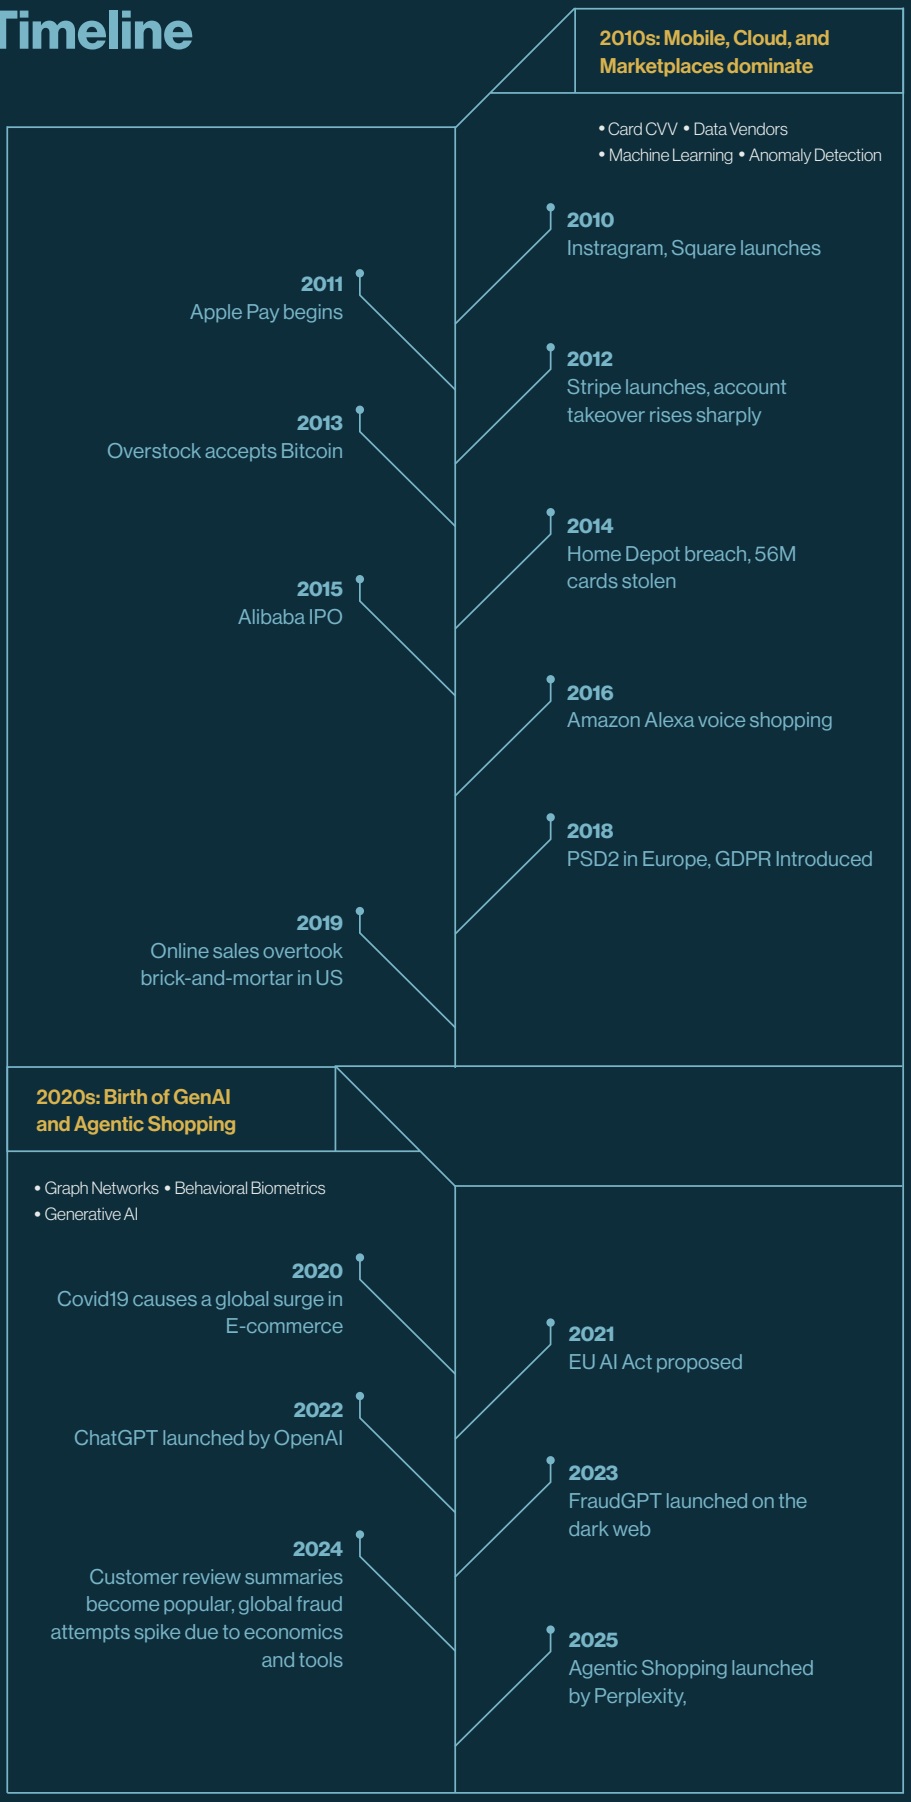

I. Velocity checks: Fraud Detection "1.0"

Why They Emerged

Early fraud detection relied on velocity checks because fraudsters using stolen cards acted fast. Merchants needed a way to spot rapid-fire transactions before losses mounted.

What are velocity checks, and how did they work?

Velocity checks are simple and fast checks that verify whether the same set of payment information is being used in rapid succession. While not highly advanced or deeply analytical, they serve an important purpose: fraudsters often attempt to exploit stolen card information by making multiple rapid transactions. Detecting this pattern early helps stop fraudulent activity before it escalates.

A velocity check, as the name implies, is concerned with speed: What's the rate that a given event occurs? These actions can include:

- Transactions within a limited span of time.
- Login attempts (especially failed login attempts).
- Password reset requests.
- Changes between different accounts by one user.
- Repeated use of the same payment.

Velocity checks: Why they worked

Once fraudsters acquire stolen credentials, they rarely hold onto them for long. Payment cards have a limited shelf life — once reported missing or flagged by another merchant, their value to criminals diminishes rapidly. As a result, attackers attempt to exploit stolen cards quickly and repeatedly, maximizing damage before detection closes the window of opportunity.

Velocity checks help identify stolen card fraud by spotting one of the most common behaviors among fraudsters: after acquiring stolen payment details, they often run quick "test" transactions—typically small purchases of digital goods like eBooks or mobile minutes—to confirm the card's validity. Once successful, they escalate rapidly to high value items like laptops or smartphones, maximizing their gains before the credentials are flagged.

Velocity checks monitor transaction frequency in real time, identifying rapid bursts of activity that signal fraud. By shutting down suspicious behavior early, merchants can prevent criminals from moving beyond small test purchases to high value items.

The Limitations of Velocity checks

Velocity checks could not serve as a merchant's sole fraud prevention method, because they only applied to a specific kind of activity. Their general weaknesses:

- They only functioned against simple, high-speed attacks on merchants.
- They didn't adapt intelligently if fraudsters modify their tactics.
- Organized fraud rings know how to circumvent velocity checks.

II. Rules Engines: Human Logic for Nuanced Patterns

Why They Emerged

Velocity checks were effective for rapid-fire attacks but too narrow to address evolving fraud tactics. As fraudsters diversified their methods—using tactics like mismatched addresses, synthetic identities, and high-value carts—merchants needed a more flexible approach. Rules engines emerged to capture these nuanced patterns through human-driven logic.

What are Rules Engines, and How Did They Work?

Rules engines rely on analysts who reviewed historical fraud data, identified recurring patterns, and codified those patterns into actionable rules. These rules were explicit conditions that triggered alerts or declines when met. Creating and implementing a new rule was relatively fast, allowing merchants to respond quickly to emerging fraud trends. Because rules were human-defined, they were inherently interpretable—merchants could easily understand why a transaction was flagged, which is critical for compliance and customer service.

Examples included:

- "Large distance between billing and shipping addresses increases fraud risk."
- "High cart value combined with a newly created email address signals potential fraud."
- "Multiple failed login attempts followed by a high-value purchase may indicate account takeover."

Rules Engines: Why They Worked

Rules engines excel at addressing point-in-time fraud trends. When fraud patterns are well understood, rules provide immediate protection without requiring complex algorithms or large datasets. They also offer transparency, which is valuable in regulated industries.

The Limitations of Rules Engines

Despite their advantages, rules engines had significant drawbacks:

- **Static nature:** Rules did not adapt automatically as fraud tactics evolved.
- **Rule accumulation:** Over time, merchants accumulated hundreds of rules, creating "rule soup" —a tangled web that was difficult to manage.
- **False positives:** Outdated or overly broad rules can block legitimate transactions, harming customer experience.
- **Reactive approach:** Rules could respond to known patterns but could not anticipate new ones.

III. Machine Learning (ML) models: Automating Rule Creation

Why ML Emerged

Manually creating and maintaining rules is laborious and error prone. As fraud tactics evolved and scaled, rules became insufficient which led to the next step in evolution: applying machine learning (ML) models to automate fraud detection. This came at a time when data volume and fraud patterns grew beyond what humans could interpret.

From Logistic Regression to Decision Trees

Early applications of ML to fraud prevention were simple. These models used logistic regression to predict whether transactions were fraudulent based on historical data, producing a binary yes/no decision to approve or reject activity.

The next advancement was the introduction of decision trees, implemented through ML libraries such as XGBoost. Decision trees operate on "90° boundaries"—clear yes/no splits based on whether a specific condition is met. By following these branching paths through multiple layers of decisions, informed by categorical fraud data, the system can identify fraud patterns and generate rules automatically.

Decision Trees: Why They Work

Decision trees are particularly effective when fraud patterns are clear. If fraudulent activity consistently exhibits specific characteristics, decision trees can model these patterns accurately.

They excel at mapping non-linear and complex patterns, handling categorical values naturally, and functioning even when datasets are incomplete. Because splits are triggered by threshold values, decision trees can manage outliers without compromising model integrity.

Another advantage is interpretability. Like human-created rules engines, decision trees produce outputs that are easy to visualize and explain. This transparency is critical for compliance in regulated industries such as finance and healthcare. Additionally, decision trees do not require feature scaling or normalization, simplifying data preparation.

The Limitations of Decision Trees

While ML-based decision trees represent a powerful alternative to manual rule creation, they have limitations:

- **Data requirements:** Decision trees require extensive training data to achieve high accuracy. This reliance on large historical datasets means they cannot catch emerging fraud patterns in real time, leaving merchants exposed until enough new data is collected and modeled.
- **Isolation bias:** These models analyze transactions individually and are not designed to track relationships between entities. This leaves merchants vulnerable to sophisticated fraud rings

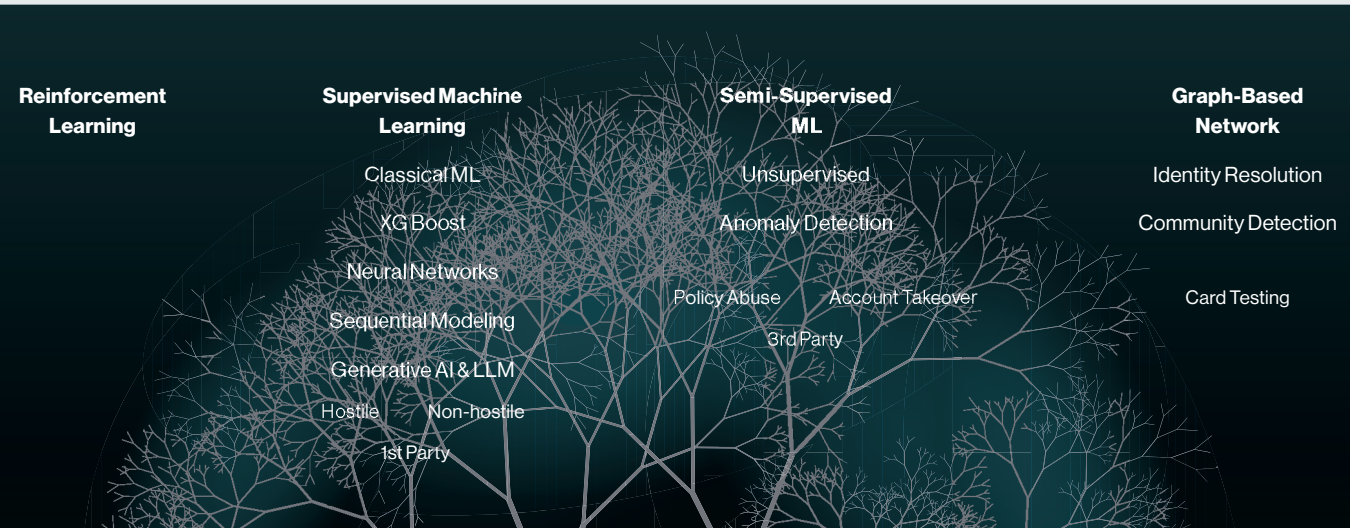

IV. Graph Networks: Uncovering Hidden Connections

Why Graph Networks Emerged

Traditional fraud detection models evaluated transactions in isolation. This left merchants vulnerable to organized fraud rings that operate across multiple identities and accounts. Graph networks address this gap by mapping relationships between transactions, enabling the detection of patterns that would otherwise remain hidden.

How Graph Networks Work

Graph networks connect attributes across transactions—such as email addresses, devices, IPs, and shipping details—within a moving time window. Think of it as building a "social network" of transactions. By analyzing these connections in real time, graph networks identify clusters of suspicious activity, even when individual transactions appear legitimate.

Modern fraudsters frequently rotate usernames, emails, and devices to evade detection. To a rules engine or decision tree, these transactions would look unrelated. Graph networks expose the underlying web of connections, making it possible to detect coordinated attacks and fraud rings before they scale.

Graph Networks and Decision Trees Are Complementary, Not Competitive

Graph networks did not replace decision tree models—they evolved alongside them. Decision trees excel at structured classification of individual transactions, while graph networks uncover relational patterns across multiple entities. Together, they form a critical layer in today's fraud prevention strategy.

Limitations of Graph Networks

- **Data complexity:** Building and maintaining graph structures requires significant computational resources and well-structured data pipelines.
- **Scalability challenges:** Real-time graph analysis can be expensive at scale, especially for merchants processing millions of transactions daily.
- **Cold-start problem:** Graph networks rely on relational data; they are less effective when transaction history or entity linkage is sparse.

V. Sequential (transformer) models: Fighting AI-generated fraud

In a foreseeable development, fraudsters have begun using generative AI (GenAI) to create synthetic identity patterns at a pace that could defeat conventional fraud detection methods. This complexity, enabled by commoditized and widely accessible fraud tools, introduces a new level of risk for merchants. Unlike earlier schemes, these AI generated identities can mimic legitimate customer behavior, adapt in real time, and scale across thousands of accounts simultaneously. This escalation raises the stakes for merchants, requiring defenses that are not only reactive, but predictive.

Sequential and transformer models can help meet this challenge by analyzing the order and timing of events, not just isolated transactions. While AI-driven fraud can imitate individual behaviors, large-scale automation often introduces subtle anomalies in sequence and timing—such as compressed activity windows or repetitive patterns across accounts. By learning what "normal" looks like and flagging deviations, these models uncover fraud that might otherwise go undetected.

What Sequential Models Add to Fraud Prevention

Sequential models go beyond temporal snapshots. While decision trees assess individual transactions in a specific moment and graph networks connect multiple transactions in the present slice of time, sequential models understand how behaviors unfold over time. This enables detection of patterns that static models miss.

Unusual activity detected by sequential models could include:

- Customer purchase habits that don't fit normal or natural patterns.
- Deviations from a user account's normal weekly or daily behavior.
- Abnormal ordering of transactions compared to expectations.

The Potential Role of Sequential Models

The ability to find new relationships and patterns within temporal sequences is a powerful one. With this capability, merchants have a potential weapon against fraud created by automated models like the commoditized and widely available "FraudGPT" toolkit used by criminals. By spotting timing irregularities and unnatural sequences, sequential models complement other detection layers and strengthen defenses against AI-driven attacks.

One of the most fascinating traits of sequential models is that their true limits and usefulness are still being researched. As both fraudsters and security professionals evolve their methodologies, there is a need to keep advancing to avoid falling behind.

While this evolution is promising, it doesn't mean that merchants should abandon other methods. A layered strategy provides the resilience needed to counter diverse and fast changing threats.

Why Merchants Need All Modeling Approaches Together

Modern fraud prevention depends on the integration of multiple models. Each method adds unique strengths: velocity checks, rules engines, decision trees, graph networks, and sequential models all detect different signals and stop different types of fraud. When combined, they create a layered defense that is far stronger than any single approach.

A consolidated fraud platform that unifies these models delivers two critical benefits: comprehensive protection and operational simplicity.

Why Eliminating Older Models Is A Mistake

Each model addresses a distinct part of the fraud spectrum:

- **Velocity checks:** Effective against rapid card testing and repeat purchases.
- **Human-defined rules:** Respond quickly to point-in-time fraud trends with clear logic.
- **ML-powered decision trees:** Stop stable, historic patterns by evaluating individual transactions.
- **Graph networks:** Detect complex fraud rings through real-time relational analysis.
- **Sequential models:** Identify AI-driven and evolving fraud by analyzing timing and sequence anomalies.

Removing older models eliminates proven defenses and creates blind spots. The future is not about replacement—it's about integration.

Unified Scoring: Turning Model Disagreement into Accuracy

When multiple models work together, each act like a different "net" cast into the water—catching fraud signals that others might miss. Velocity checks, rules engines, decision trees, graph networks, and sequential models all cover different parts of the fraud landscape. Layering them creates broader coverage and greater certainty that fraud is caught while good customers pass through.

Showing separate risk scores for each model would be confusing and counterproductive and disagreement between models isn't a flaw—it's a strength. Each model has unique blind spots, so when one flags risk and another doesn't, that tension improves overall accuracy. A unified scoring layer combines these signals using ensemble logic, weighting each model appropriately and applying overrides for critical indicators, such as high-risk graph patterns.

The frontier of fraud prevention isn't about choosing one model over another—it's about integration. By weaving multiple methodologies into a cohesive system, merchants gain the ability to respond in real time, interpret decisions confidently, and stay ahead of evolving threats.

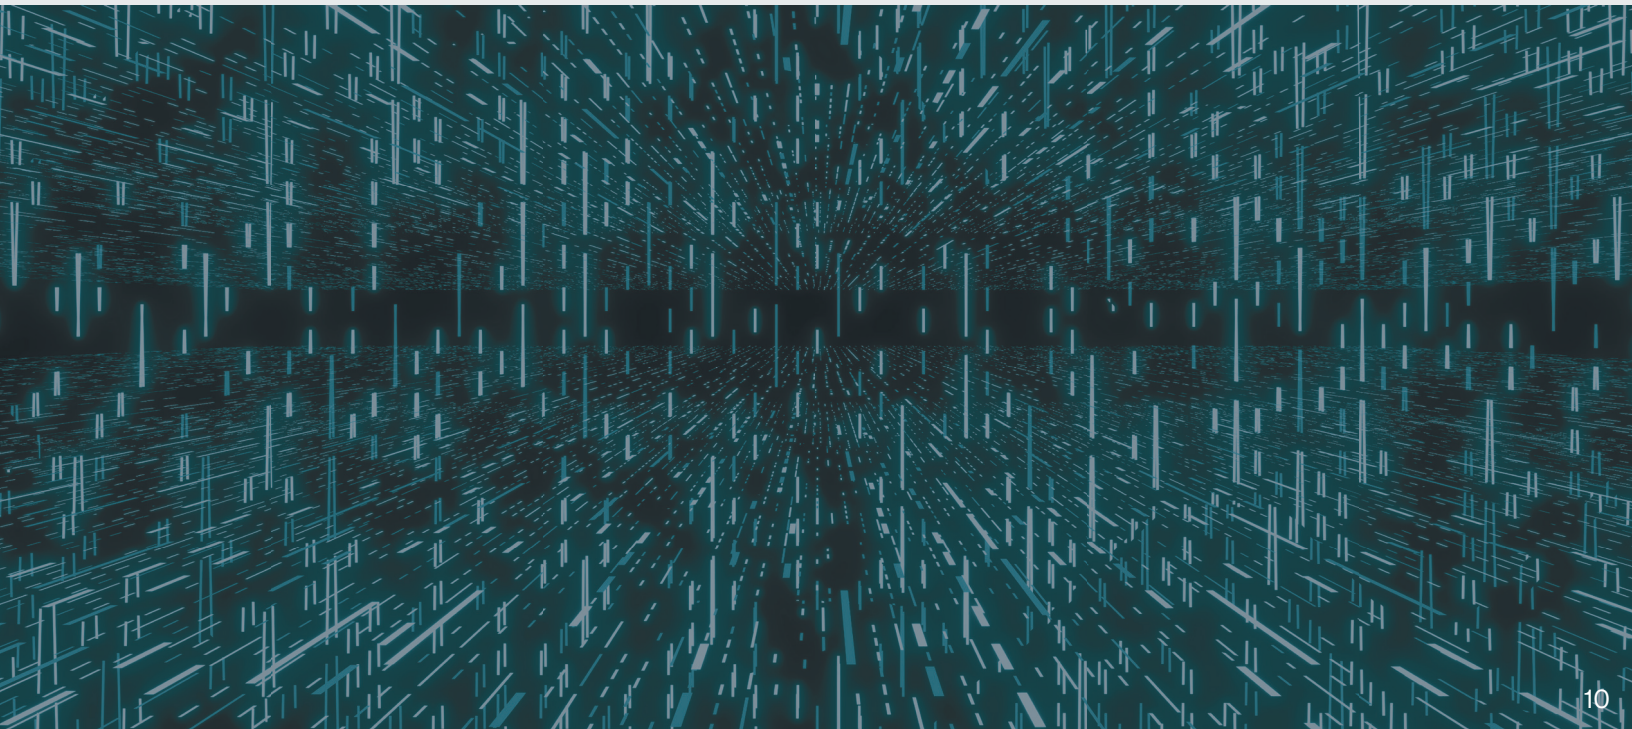

Conclusion

Fraudsters Keep Moving: Merchants Must Move Faster

Fraud has never stood still—and neither can fraud prevention. What began with simple velocity checks has evolved into rules engines, machine learning, graph networks, and now sequential models designed to counter AI-driven attacks. Each step in this evolution arose because fraud tactics changed, and each new model added a critical layer of defense.

The impact of fraud on merchants isn't just financial—it impacts reputation, customer trust, and operational stability. To stay ahead, merchants need a strategy that mirrors this history: a layered approach that combines proven methods with advanced models in a unified, adaptive system—so every layer works together to catch what others miss.

Fraudsters will keep innovating. The question is whether your defenses will evolve faster. The future belongs to merchants who embrace this layered approach, anticipate emerging threats, and turn complexity into confidence.