Fighting Financial Crimes with Artificial Intelligence

04.18 EB9994 ARTIFICIAL INTELLIGENCE

TERADATA

Introduction

Global credit fraud reached $23 billion in 2016—which is only a small fraction of the upwards of $2 trillion in risk exposure to global financial systems due to a broad spectrum of financial crimes. Beyond credit card fraud, there is check fraud, mortgage fraud, point of sale fraud, insider corporate fraud, and securities fraud. Add to these a dramatic recent growth in account and identity fraud, as well as money laundering. The diversity in types of financial crime is driven by a diverse range of perpetrators, including individuals, corporations, and organized crime groups.

Financial Crime: Cross-Channel Attacks on the Rise

One of the biggest trends in financial crime is the rise in cross-channel attack tactics. What began as a bank-led innovation to serve customers through new and multiple channels, including various mobile devices, opened a door to criminal innovation exploiting the diversity of channels and lack of data and interdiction integration across channels. Because banking institutions focus on serving each channel distinctively, they miss cross-channel activity patterns that indicate crime. Criminals take advantage of channel fragmentation by combining actions spread across multiple channels into novel attacks that banks are ill-prepared to detect.

Beyond finding success with fraud schemes that take advantage of data silos and institutional myopia, financial crime is marked by increasing collaboration between criminals. Fraudsters are working together. On black markets worldwide, they share tools and code to quickly spread attacks that work against one bank, while creating new attacks banks have not yet seen or built defenses against.

As more transactions take place across different channels—from phones, tablets, ATMs, and beyond—the diversity of systems involved introduces a challenge of retaining the whole context of a transaction. Different systems are owned and monitored by different teams and systems, and the sharing of data and predictive models between data scientists across teams can be essential to combatting financial crime—and shoring up weaknesses introduced by the diversity of channels. As the normal flow of transactions become increasingly disjointed and multi-system, anomalies and fraudulent activity are that much harder to detect. Fragmentation doesn't just slow enterprises down, it opens them up to attacks where their defenses and capacities to respond are weakest.

Fraudsters' success with collaborative, cross-channel attacks shows the weakness of silos: siloed applications, silos of data, and silos of decision. Integration of data and decisions across channels is essential to effectively defend against innovative attacks that span channels.

Traditional Crime Detection Methods No Match For Automated Attacks

Financial crime increasingly depends on processes that are largely automated. For example, when an industrial scale data breach occurs, synthetic identity creation that exploits the breach is done via automation rather than by hand. Likewise, an effective response to the scale and speed of automated attacks cannot rely on manual efforts. What is needed is automation in crime pattern detection and data preparation, as well as automation in monitoring, managing, and continuous cultivation of the ensemble of rules and models that responds to the evolving financial crimes.

The agile community of financial criminals rapidly develops new attack modes, and shares their tools and techniques with a globally dispersed, skilled, and insatiable black market. The process of manually identifying new exploit patterns is slow and error-prone, especially when it comes to the complicated patterns present in cross-channel attacks. It can take months for the existence of a new fraud scheme to be detected—and can be difficult even for trained, expert investigators to manually discern the intricate patterns that distinguish criminal from normal behavior; then to express those patterns in rules that capture only criminals.

Automated Financial Crime Detection

The automation of financial crime detection through modern algorithms, including deep learning, can dramatically accelerate the accurate detection of crime patterns. Modern, AI-enhanced detection models, appropriately constructed and tuned, will catch many more instances of crime than legacy rules-based approaches. At the same time, such models can significantly reduce the number of false positives; that is, transactions that are flagged as suspicious but are, in fact, legitimate. For some banks, false positives account for more than 99 percent of the transactions they stop. The customers and merchants involved in interrupted/blocked legitimate transactions are understandably unhappy—and unhappy customers may not remain customers for long.

Beyond the detection of any particular pattern, however, is the importance of automating the deployment and management of all crime pattern classifiers with a bank's authorization system. The right solution will support the

Financial Crime: Outpacing Banks' Efforts to Fight it

Countering the many types of financial crime across diverse transactional platforms and products has been a constant struggle for banks and other financial institutions. Financial criminals are relentless in devising new ways to attack—battering against old defenses and, all too often, slipping by undetected. Despite the best efforts of skilled teams focused on varied aspects of the problem, financial crime remains a major cost to business. It doesn't have to be this way.

Banks have slowly institutionalized past learning in legacy rules-based engines, but have largely failed to develop the analytic agility to keep pace with fast-evolving criminal schemes. For example, consider fraud: the old way of fraud detection is based on long lists of rules that hard-code logic that transactions must follow to be authorized as valid. Methods for maintaining the ensemble of rules can be haphazard and ad hoc.

Not only do new and better rules not make their way into the ensemble fluidly and systematically, but the effectiveness of existing rules can decay without detection. Using rules-based engines—informed by light analytics as the only mechanism for detecting and thwarting financial crimes—is an antiquated and inadequate response. As a result, banks can't keep up. They suffer losses that could be avoided if only they invested in the right enhancements to their existing systems.

automation of model operationalization, reducing the time it takes for banks to go from systematic crime pattern identification to effective response deployment.

Sophisticated modern algorithms can automatically process massive amounts of complicated, interconnected data from many applications and source systems—quickly pinpointing anomalies easily missed by the human investigator.

Accelerating the Speed of Analytics Innovation

It's difficult to encourage teams to come up with new crime mitigation strategies when it's hard to test new strategies against the old—and difficult to integrate them with existing authorization and crime interdiction infrastructure. The last thing enterprises want is another point solution for one type of fraud, or one type of data, that doesn't integrate with their battle-hardened legacy systems for authorization.

On the other hand, with a proven, tailored system that is built to incorporate ongoing changes in a fast, reliable, auditable fashion, enterprises can accelerate the speed of analytics innovation. Analysts and data scientists know their efforts can be deployed to stop crime and make a positive difference today, rather than weeks or months from now. Designed for rapid iterative learning, an effective solution accelerates the integration into operational systems of improved models built by data scientists, together with models that have been improved algorithmically without human intervention.

Sophisticated deep learning models can isolate patterns to complement human identification of fraud patterns. Those models—whether human or machine-generated—can be tested automatically to see whether adding them will catch more fraud than the set of rules and models already in place, and what the expected cost of false positives might be. A modern AI-enhanced crime detection system thus protects banks from criminal losses with a system of rules and statistical/deep learning models. The system strengthens itself with time and experience as an evolving pool of candidates that could replace or supplement the production ensemble and tested for inclusion in the operational defenses.

A modern solution can continuously integrate new data to refresh models, and the rules set so that new criminal patterns are recognized early—and stopped cold. Rather than a system that grows stale and ever less effective between manual model refreshes, it's better to have one that feeds on all data to constantly increase the strength of its defenses with each transaction processed, and each customer served.

Modern Techniques

Though rules will always have an important role to play throughout financial services and transaction processing, successful crime prevention today goes beyond rules alone to incorporate artificial intelligence (AI) and other advanced machine learning algorithms. Existing systems can be supplemented with AI, rapidly identifying fraud patterns of limitless complexity—beyond what manual analysis can identify and articulate as a logical rule. Modern algorithmic techniques have been proven to stop financial crimes at top banks, and have been tailored for distinct use cases in mobile banking, online payments, anti-money laundering (AML), and beyond.

A human-centered artificial intelligence accelerator to fight financial crime

As criminals create new schemes that rely on multiple channels and applications and accounts—for example, in money laundering schemes—there are challenges in detecting patterns across the variety of touchpoints and data sources. A graph analytic approach, together with deep learning, can be especially powerful for AML. For identity account fraud and account takeover, patterns in account signup and account information change can indicate that a fraudster is at work. In these cases, important clues to criminal activity are present in the unstructured textual data and the event series data. These and other methods squeeze as much information as possible out of all the data—whatever the type or size—using diverse, use case-aligned techniques including deep learning algorithms, that have been tested and deployed in major banks worldwide.

To the limited extent that other companies and consultancies provide AI solutions at all, the solutions they offer tend to be black boxes. Black box solutions emit predictions—predictions that may be more accurate than ones based in rules—but they come with no regulations-compliant explanations. In a tightly regulated environment, such an approach would only introduce another set of headaches: Black box solutions are non-starters. What is needed instead are the reasons why a transaction was rejected, reasons a human can understand, not only for regulatory compliance, but also because such information makes it possible to improve operational and management decisions going forward.

Enabling Collaboration

By supporting data and model sharing, subject to enterprise-grade user access controls—and including clear, easy-to-understand explanations of model-based predictions—an enterprise-wide financial crime solution enhances collaboration between data scientists on different teams, and between data scientists and the IT professionals involved in maintaining operational systems. The solution can integrate fraud detection events and processes across a variety of channels, departments, and user types. By supporting a unified forum for data science model comparison across teams, an integrated financial crime analytics solution can help the AML group with its traditional focus on compliance and law, complement the operational focus of the fraud group, and provide a common clearinghouse of data, data engineering pipelines, and models to support better decisions at every encounter with criminals. Such a solution facilitates swift, coordinated reactions to the incidence of new financial crime.

The architecture of an effective solution maximizes the value of existing investments in legacy authorization infrastructure, as well as facilitates the development of better detection models that can use any and all crime-relevant data. Additionally, it supports coordination of detection and interdiction efforts across diverse teams, ending the infrastructural and operational silos that leave the bank vulnerable to sophisticated cross-channel attacks. It is essential to create a unified, cohesive front against the relentless assault of fraudsters. Though it may make sense to have separate business and application teams run different channels and related applications, the authorization decisions made on those channels need to reflect an overarching view of the bank.

The Teradata Financial Crimes Accelerator enables the rapid development of tailored, best-in-class solutions for financial crime detection, with the following characteristics:

- **Cross-channel AI-enhanced detection models:** Enables the development and deployment of deep learning and AI models utilizing data from any source system, across all channels
- **Automation:** Uses judicious automation to accelerate the development, deployment, monitoring, and management of detection models
- **Robust operations:** Provides a robust analytic ops environment; easy for data scientists to use, and familiar/trustworthy to IT and data engineers
- **Integration with existing systems:** Designed for seamless integration with rules engines and operational transaction processing systems
- **Collaboration:** Supports effective collaboration between data scientists across teams, and between data scientists and data engineers

The Teradata Financial Crimes Accelerator

The Teradata Financial Crimes Accelerator puts continuous analytic innovation first, enabling our customers to dramatically improve legacy systems and processes—and keep pace in the fight against relentlessly evolving attacks. It's composed of best practices, code, IP, and proven design patterns to help accelerate deployment and ensure quick ROI.

Our combination of technical IP and services enable customers to quickly create a custom solution that enhances existing systems, and detects fraud in fundamentally different and better ways to support real-time approve/decline decisioning, and bring transactions within risk tolerance.

In large US and global banks, we have integrated our solution with pre-existing rules-based authorization systems to achieve impressive results. For example,

with help from Teradata IP, services, and technologies, one large bank realized a 60 percent reduction in fraud detection false positives, while increasing true positive detection rates by 50 percent. With our services and extensive IP, our customers achieve substantial bottom-line improvements, including reduced fraud loss, reduced false positive loss, improved customer experience and greater trust, and higher net promoter score (NPS).

The Financial Crimes Accelerator can take full advantage of diverse data across diverse products. Cross-product and multichannel, our solution supports detection methods that incorporate endpoint and navigational data, as well as data from users and accounts, among others. Credit card, debit card, online, branch banking, ATM, wire transfer, call center—deep views into all banking products can be brought together so nothing goes missing. By supporting the development of models for crime detection that incorporate a cohesive view of a bank's universe of data, the Accelerator puts an end to the myopic silos of data and decisions that threat actors have proved all too able to exploit. By accelerating the deployment of human-centered AI solutions, the Financial Crimes Accelerator gives financial institutions the robust defenses they need to protect themselves and their customers against modern and evolving attacks.

The Accelerator integrates new AI-based analytic approaches with existing systems. Starting from our IP, including pre-configured detection models and solutions for managing deployed models, we can rapidly tailor the system to the customer's particular infrastructure, data, and use cases. The Teradata approach is not to "rip and replace" legacy enterprise systems; instead, it is to augment them with new capabilities.

The Accelerator can comply with aggressive corporate cloud mandates, or be deployed on site. Our customers experience fast time-to-value, even in large, complex deployments. We offer cloud-based deployments that facilitate fast procurement; subscription-based rather than CAPEX pricing; and provide the convenience of managed infrastructure that frees customers' in-house talent to focus on unique value-add.

AI Services Supporting the Teradata Financial Crimes Accelerator

AI Strategy Service

Teradata helps identify and recommend a series of practical AI use cases that are aligned to the strategic business goals of the enterprise, as it relates to financial crime detection and mitigation. We review key enterprise AI capabilities, and recommend next steps to help our customers successfully get value from AI.

AI Rapid Analytic Consulting Engagement (RACE)

With our RACE methodology, we can quickly demonstrate the value of AI-enhanced approaches to test use cases, enabling customers to gain buy-in from relevant stakeholders across the organization.

AI Foundation Service/Analytic Ops Accelerator

We help customers build and deploy a machine and deep learning solution, draw in the enterprise data sources that fuel model training, build AI models that deliver tangible business outcomes such as reduced charge-offs, and integrate the new AI techniques with existing operational systems. Our experience informs an array of proven code, design patterns, and best practices we use to accelerate value and reduce implementation risk.

The Analytic Ops Accelerator adapts software engineering best practices to advanced analytics and data science. It enables fast, repeatable, reliable deployment of new insights; model management; leading tools and interface for advanced data science; and curated pipelines for managing cross-channel data flows across multiple systems.

AI Analytics As-a-Service

Teradata helps customers design and oversee mechanisms to optimize and improve existing financial crime detection, interdiction, and investigation efforts using AI. Our team of world-class data scientists and engineers will manage an iterative, stage-gate process for analytic models from development to operational handover.

About Teradata

Teradata empowers companies to achieve high-impact business outcomes. With a portfolio of business analytics solutions, architecture consulting, and industry-leading big data and analytics technology, Teradata unleashes the potential of great companies.

To learn more about Teradata and Artificial Intelligence, visit Teradata.com/AI.

10000 Innovation Drive, Dayton, OH 45342 **Teradata.com**

Teradata and the Teradata logo are registered trademarks of Teradata Corporation and/or its affiliates in the U.S. and worldwide. Teradata continually improves products as new technologies and components become available. Teradata, therefore, reserves the right to change specifications without prior notice. All features, functions, and operations described herein may not be marketed in all parts of the world. Consult your Teradata representative or Teradata.com for more information.

Copyright © 2018 by Teradata Corporation All Rights Reserved. Produced in U.S.A.

04.18 EB9994