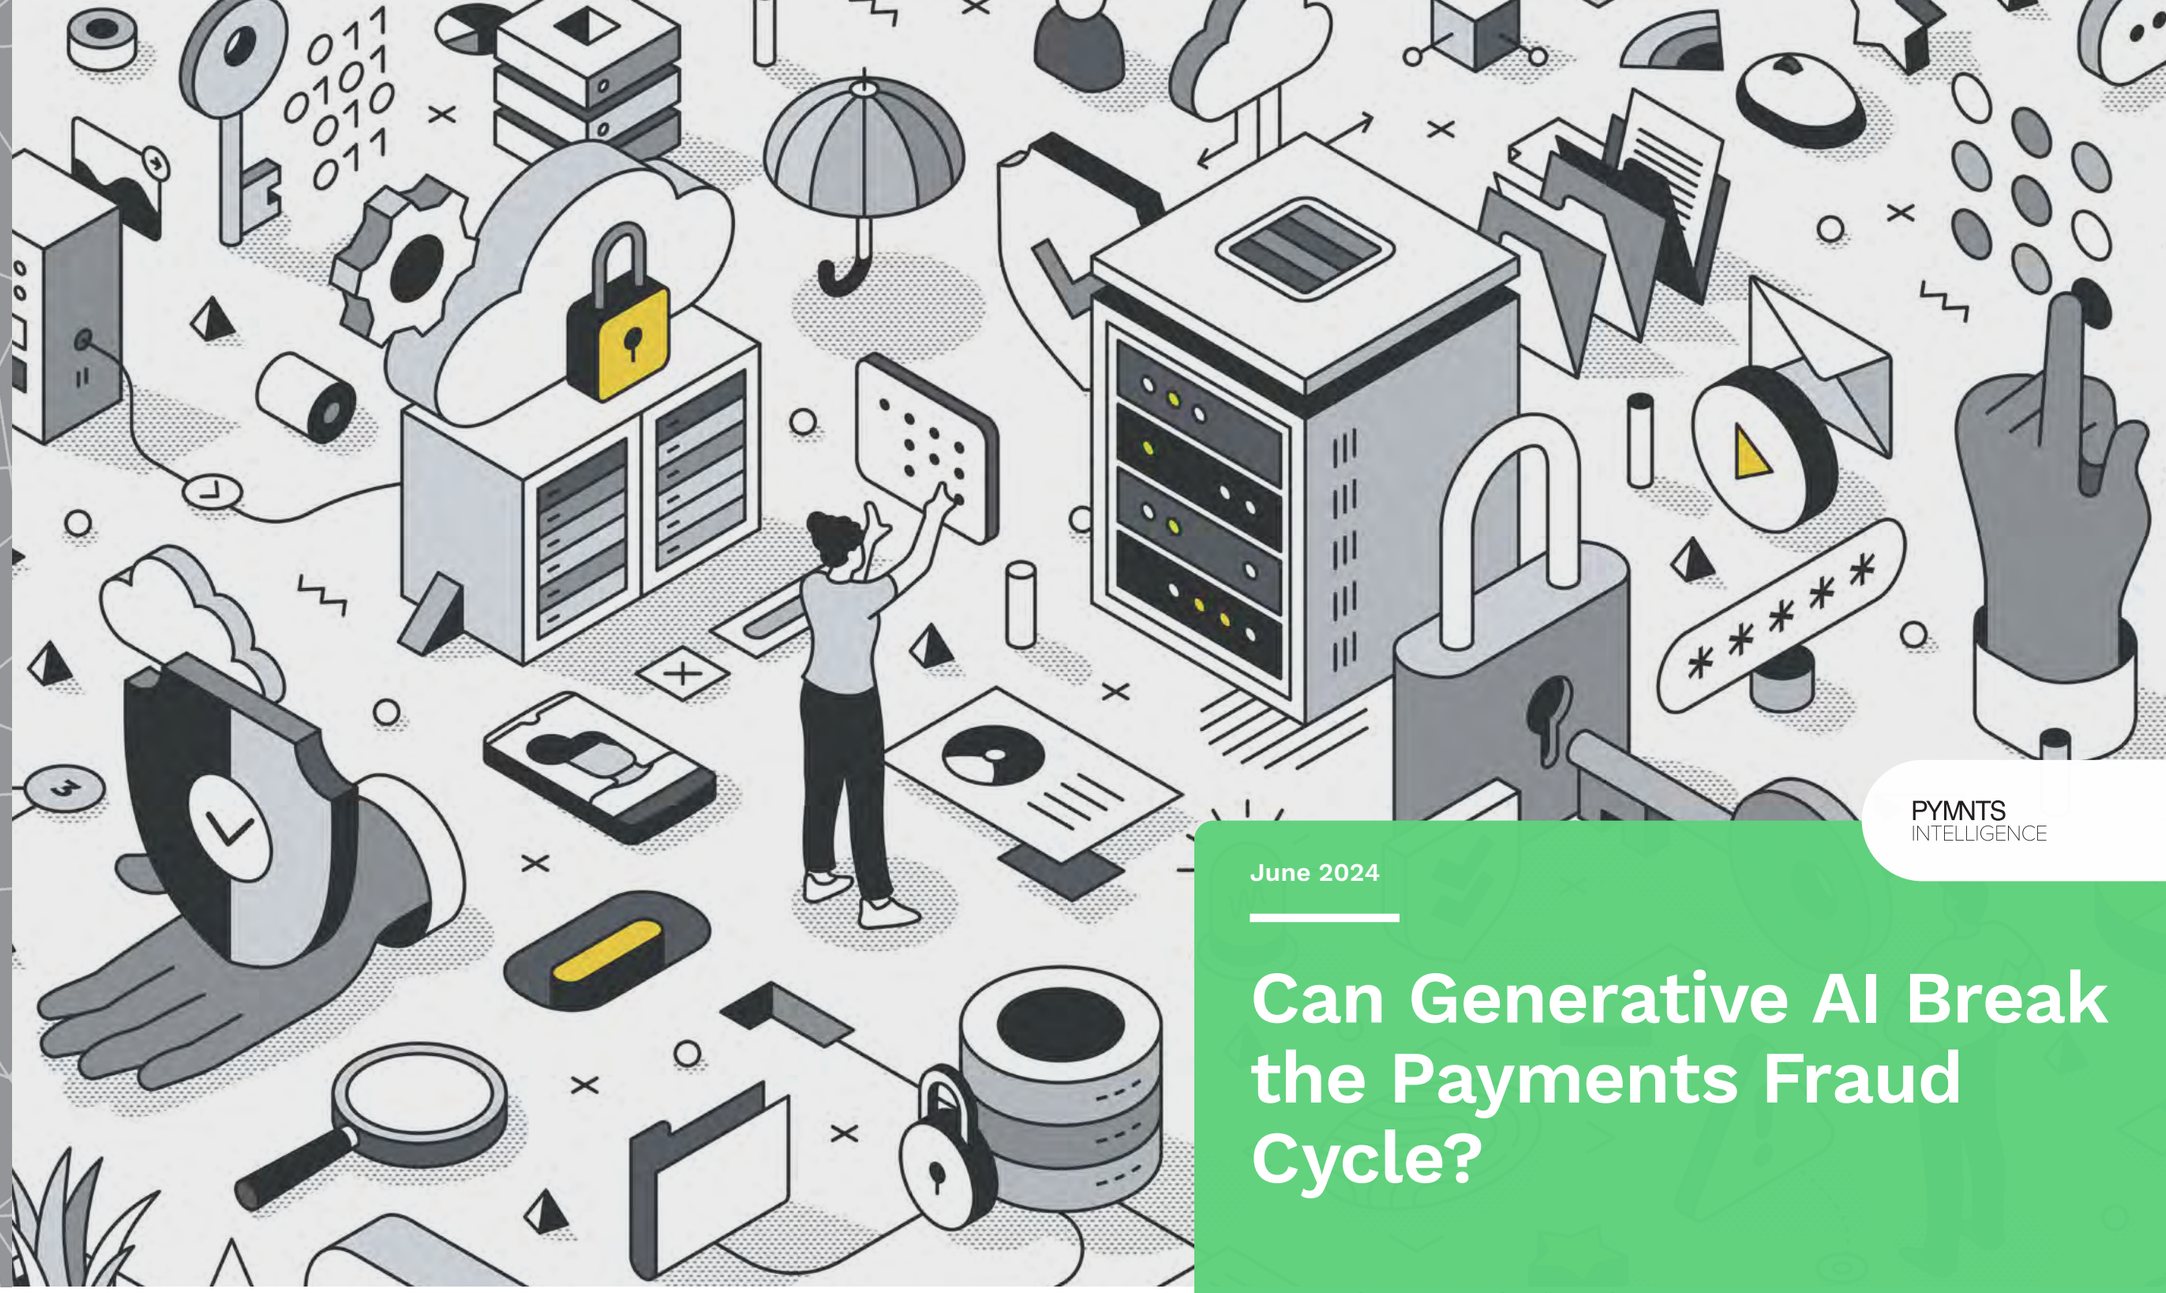

PYMNTS
INTELLIGENCE

June 2024

Can Generative AI Break the Payments Fraud Cycle?

Generative AI Tracker® Series

What's Inside

PYMNTS
INTELLIGENCE

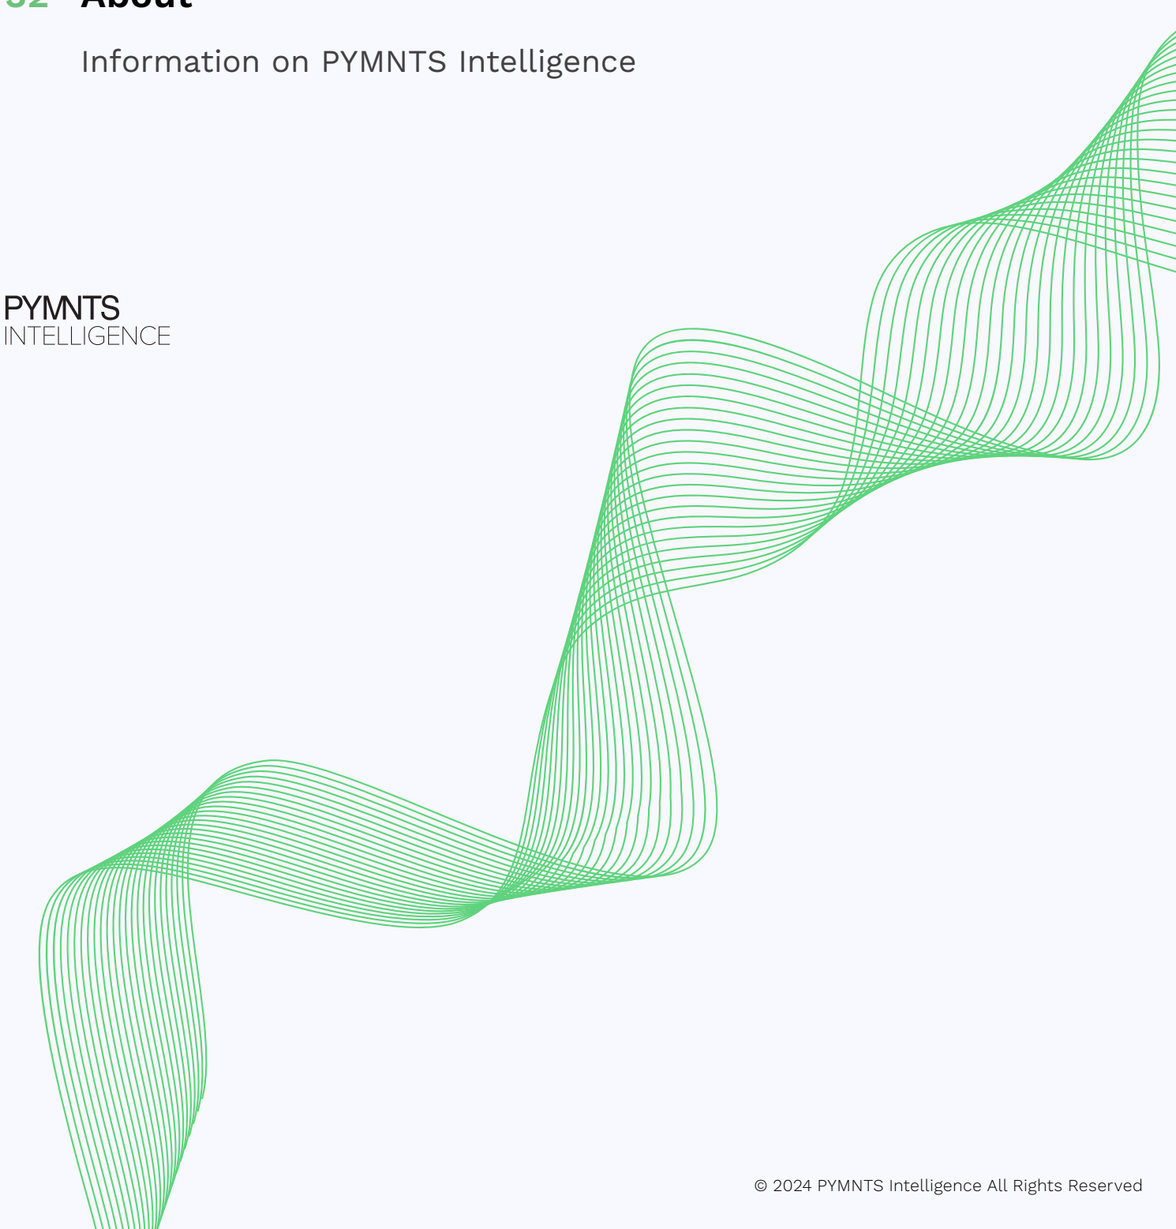

© 2024 PYMNTS Intelligence All Rights Reserved

Background

Payments Fraud Detection Is a Key Use Case for Generative AI

Limitations of conventional methods

Traditional rules-based systems for detecting payments fraud are increasingly inadequate against fraudsters devising complex schemes that rely on advanced technologies. The traditional rules-based approach uses predefined, static rules to identify suspicious activities based on known fraud patterns and expert knowledge, yet it requires frequent manual updates to rules, is prone to high false positive rates and has a limited ability to adapt to new fraud schemes.

Predictive artificial intelligence (AI), or machine learning (ML), is an advanced technique that uses supervised learning models trained on labeled historical data. It is more effective at reducing false positives and can adapt to new fraud schemes through retraining on new data. This technology has achieved nearly prerequisite status in recent years for modern fraud detection.

However, the relative newcomer of generative AI offers the vast potential to supplement existing methods through its ability to detect subtle patterns in payments data that elude both rules-based systems and predictive AI alternatives. Using unsupervised or semi-supervised learning techniques, generative AI excels at detecting novel and evolving fraud patterns by analyzing unstructured data, thereby augmenting traditional ML models for cutting-edge fraud detection.

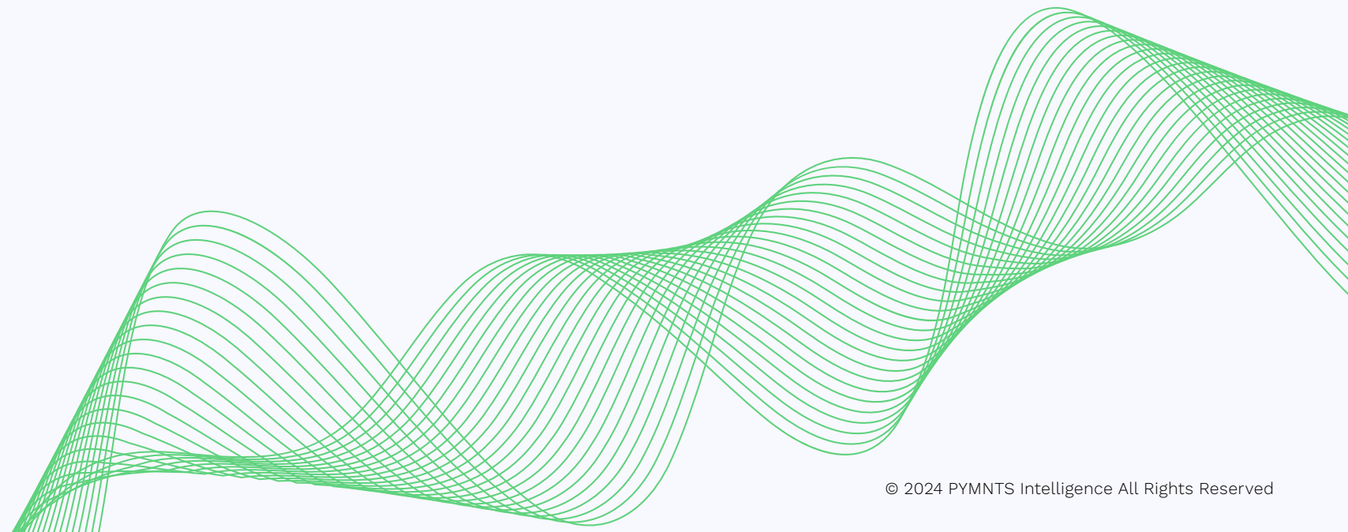

Background

Need for real-time learning

Modern payments fraud demands a solution capable of adapting in real time and at scale. Generative AI, with its capacity for continuous learning, offers the unprecedented advantages of rapidly refining and adapting its understanding of patterns to distinguish more accurately between legitimate and fraudulent payments.

Background

Promise of synthetic data

Utilizing real-world financial data for training fraud detection models often raises privacy concerns. Generative AI offers an alternative through the production of synthetic datasets that mimic the characteristics of genuine data, an approach that allows for training robust fraud detection models without jeopardizing privacy rights or compliance.

Better consumer experiences

False positives in fraud detection often frustrate customers when legitimate transactions are incorrectly flagged. Generative AI reduces these false positives by more accurately distinguishing between genuine and fraudulent purchasing behaviors. This precision helps to ensure smoother transaction experiences and fewer customer frustrations.

Companies of Note

Leading the Way in Generative AI for Payments Fraud Detection

Generative AI's rise has not gone unnoticed by the financial industry. Indeed, two industry stalwarts — Visa and Mastercard — have already built and deployed their own in-house generative AI payments fraud detection tools. While the technology's use by the industry remains very much in its infancy, these companies provide a glimpse at the various ways generative AI can be used to combat payments-related fraud — and early adopters are already seeing tangible benefits.

Early types of generative AI fraud detection solutions:

Real-time transaction scoring: Tools that assign scores to transaction risks and reduce the number of false positives

Proactive fraud prevention: Tools that identify at-risk merchants and secure payments networks

Scaling advanced fraud detection: Tools that speed up fraud detection training processes and drive system scalability

Companies of Note

PYMNTS Intelligence spotlights three financial services providers leading the vanguard:

Real-time transaction scoring

Visa recently introduced its Visa Account Attack Intelligence (VAAI) Score, a solution leveraging generative AI to identify and score enumeration attacks in real time. Enumeration attacks are a type of cyberattack in which a malicious actor systematically gathers information such as usernames, email addresses and account details from a target system or network by exploiting its responses, potentially leading to unauthorized access or further attacks. The VAAI Score assigns risk scores to card-not-present transactions, allowing card issuers to make more informed decisions about blocking suspicious activity. The ability to learn and discern typical from atypical transaction patterns is driving down the rates of false positives, which are now 85% lower compared to other models. This type of solution not only enhances consumer satisfaction but also directly contributes to reducing financial losses for issuers.

Companies of Note

Proactive fraud prevention

Mastercard is deploying generative AI to double the speed of detecting compromised cards and better predict their full details on its network. This approach enables faster blocking and replacement of compromised cards, thereby reducing the window for fraudulent activities. Additionally, the model accelerates the identification of at-risk or compromised merchants by 300%, thus enhancing the security of the digital payments ecosystem.

Scaling advanced fraud detection

European FinTech Bunq is utilizing generative AI to streamline and automate its transaction-monitoring system. Employing both supervised and unsupervised learning, the digital bank has developed a generative AI-powered system that is fully automated and easily scalable. This approach has notably accelerated its data processing pipeline by more than five times and increased fraud detection model training speed by nearly 100 times compared to previous methods. This improvement allows Bunq to refine its detection algorithms continuously, resulting in more effective fraud prevention and increased operational efficiency.

Innovation and Use Cases

Generative AI Could Cut Through the Fraud Detection Noise

The potential of generative AI to alleviate the costly headache of payments fraud has garnered considerable attention in the financial industry. As this technology continues to mature and its adoption gains traction, generative AI could become a cornerstone of modern payments fraud prevention strategies, promising significant improvements in accuracy, efficiency and cost savings.

Excitement stems largely from generative AI's potential to overcome inherent limitations of traditional fraud detection systems. Its capabilities hold the potential to supplement current methods with real-time identification and neutralization of payments fraud, a prospect with significant implications for further safeguarding the purchasing experience and improving the bottom line of financial institutions (FIs) and businesses.

Innovation and Use Cases

The foundations of predictive AI in fraud detection

To better understand the ways in which generative AI could reshape payments fraud detection, aoverview of predictive AI — the current go-to method — is needed.

Predictive AI works by leveraging a combination of historical data, statistical modeling, data mining techniques and machine learning (ML) to predict future outcomes. Applied to detecting payments frauspecifically, predictive AI analyzes historical transaction data to distinguish between fraudulent and legitimate payment patterns. Having established these baseline patterns, the model can flag suspicious transactions thadeviate from established trends, thereby enabling banks and FIs to identify potentially fraudulent activity in real time.

© 2024 PYMNTS Intelligence All Rights Reserved

Innovation and Use Cases

Potential applications of generative AI in payments fraud detection

Building on this, PYMNTS Intelligence offers a comparative breakdown of predictive AI and generative AI for several key applications in payments fraud detection and prevention, highlighng the distinct strengths of generative AI.

DETECTING FRAUD

NUANCED PATTERNS

Predictive AI: Utilizes historical data and predefined features to predict and flag transactions based on established patterns. Effective in recognizing known fraud scenarios but potentially misses novel threats.

Generative AI: Excels at identifying complex and subtle patterns missed by traditional rules-based systems, such as unusual transaction amounts and atypical spending behaviors.

BASELINES PAYMENTS ACTIVITY

Predictive AI: Relies on historical data to define baselines, which require periodic updates to remain effective.

Generative AI: Continuously adapts to evolving payments behavior, establishing dynamic baselines and instantly identifying anomalies like unusual login times and transaction patterns.

SOPHISTICATED FRAUD SCHEMES

Predictive AI: Efficiently processes large volumes of data to identify patterns indicative of fraud but may require manual intervention to adapt to new types of fraud.

Generative AI: Analyzes vast datasets of both legitimate and fraudulent transactions to detect complex fraud schemes involving multiple accounts or channels.

Innovation and Use Cases

PREVENTING FRAUD

REALISTIC SYNTHETIC DATA

Predictive AI: Relies on genuine historical data for training, the use of which can be restricted by privacy regulations and limited by diversity and variability of actual fraud data within these datasets.

Generative AI: Generates synthetic payments data for model training, reducing or eliminating reliance on real-world data and addressing privacy concerns while boosting fraud detection accuracy.

REAL-TIME ANALYSIS

Predictive AI: Capable of effective real-time monitoring but, again, not designed to learn and adapt in real time to novel fraud tactics, limiting analysis to known threats.

Generative AI: Facilitates real-time analysis of transactions, network activities and interactions among involved parties, enabling near-instant responses to potential fraud and minimizing financial losses.

SYNTHETIC PAYMENTS FRAUD

Predictive AI: Effective at detecting fraud based on known indicators but less capable at detecting novel synthetic fraud without retraining.

Generative AI: Identifies patterns associated with fictitious identities and flags suspicious transactions, leveraging its ability to generate and compare synthetic identities.

FALSE POSITIVES

Predictive AI: Can achieve low false positive rates with well-tuned models but requires more frequent updates to maintain accuracy and efficacy.

Generative AI: Continuously learns and adapts, exhibiting a nuanced understanding of legitimate payments behavior that reduces false positives, thus minimizing purchase disruptions.

Potential applications of generative AI in the payments industry: A quick guide

1. Advanced pattern detection
2. Dynamic behavior modeling
3. Multi-account fraud analysis
4. Synthetic data generation
5. Real-time transaction analysis
6. Synthetic identity fraud detection
7. False-decline reduction

Predictive AI has proven to be a powerful tool, and current iterations of generative AI would complement rather than replace it. However, as these technologies converge and model efficacy improves, the payments landscape will move closer to a watershed moment when the opportunity cost of even attempting payments fraud will likely far outweigh the potential gains. The more immediate questions are how quickly the financial industry can navigate the technical hurdles to make this a reality and whether the regulatory landscape will evolve quickly enough to support this transformation.

© 2024 PYMNTS Intelligence All Rights Reserved

Issues and Challenges

The Tough Road Ahead for Generative AI to Overtake Fraud

The breadth and impact of these use cases help explain the financial industry's excitement about generative AI's potential to fight fraud, with 83% of FIs already eyeing its use for these purposes. However, this widespread enthusiasm is tempered by a sobering reality: The same sophistication that makes generative AI such a powerful sentinel against fraud also poses significant obstacles to its industrywide adoption.

Generative AI's need for vast amounts of financial data puts it at odds with privacy laws.

The privacy-utility paradox

A central concern regarding the adoption of generative AI for fraud detection lies at the intersection of data utility and privacy protection. Generative AI models require vast datasets for training, but using real-world financial data poses critical ethical and privacy concerns. Sensitive financial information is governed by strict laws in the United States, such as the Gramm-Leach-Bliley Act and the Bank Secrecy Act, as well as by industry best practices, such as the Payment Card Industry Data Security Standard (PCI DSS). The current oversight landscape, however, is struggling to keep pace with rapid advancements in generative AI, creating regulatory gaps and leaving the financial industry without a clear roadmap for navigating this transition.

Fairness of generative AI outputs is under close watch.

This presents FIs with a difficult paradox: While generative AI is a novel technology with enormous potential to strengthen fraud prevention and detection methods, it requires sensitive data to learn, raising concerns about potential exposure or misuse of this information. Data anonymization or creating synthetic alternatives that retain training utility seem like obvious solutions. However, the limited number of financial services providers launching generative AI fraud detection tools suggests that significant challenges remain in this regard. Even minor missteps in data handling could jeopardize the role of generative AI, particularly if training data is found to subtly channel bias.

How, then, can FIs confidently deploy this technology without risk of compromising privacy and ethical norms? As generative AI continues to evolve and improve, will we witness a version capable of protecting privacy? Or, will these challenges prove insurmountable, potentially stalling the technology's broader adoption for this use case? Regulation will likely be key for signposting a path forward.

© 2024 PYMNTS Intelligence All Rights Reserved

Issues and Challenges

The specter of bias

Bias remains a persistent and pronounced drawback of generative AI. When trained on biased data, generative AI can yield results that reproduce and amplify existing inequities, potentially leading to discriminatory outcomes that disproportionately impact certain groups. In the context of payments fraud detection, this bias could manifest in several ways, such as higher false positive rates for transactions involving individuals from specific demographic groups. This could lead to unnecessary friction in consumer purchase experiences or, worse, unwarranted scrutiny and potential financial harm.

The consequences of biased outcomes extend beyond the ethical implications of discrimination to include considerable business risks. FIs could face increased operational costs investigating these arbitrary false positives and potential legal liabilities arising from discriminatory practices. Given that trust is the cornerstone of the financial industry, deploying untrustworthy fraud detection models could undermine this trust and threaten an FI's market share, particularly with Big Tech making inroads into the industry's territory.

As fraudsters increasingly lean on generative AI to create elaborate payments fraud, the promise of this technology to detect and prevent it largely remains unfulfilled. Will the financial industry, technology providers, regulators and society discover an acceptable balance between leveraging generative AI capabilities and mitigating the risks, or will the specter of bias cast a long shadow over the future of generative AI in fraud detection and prevention?

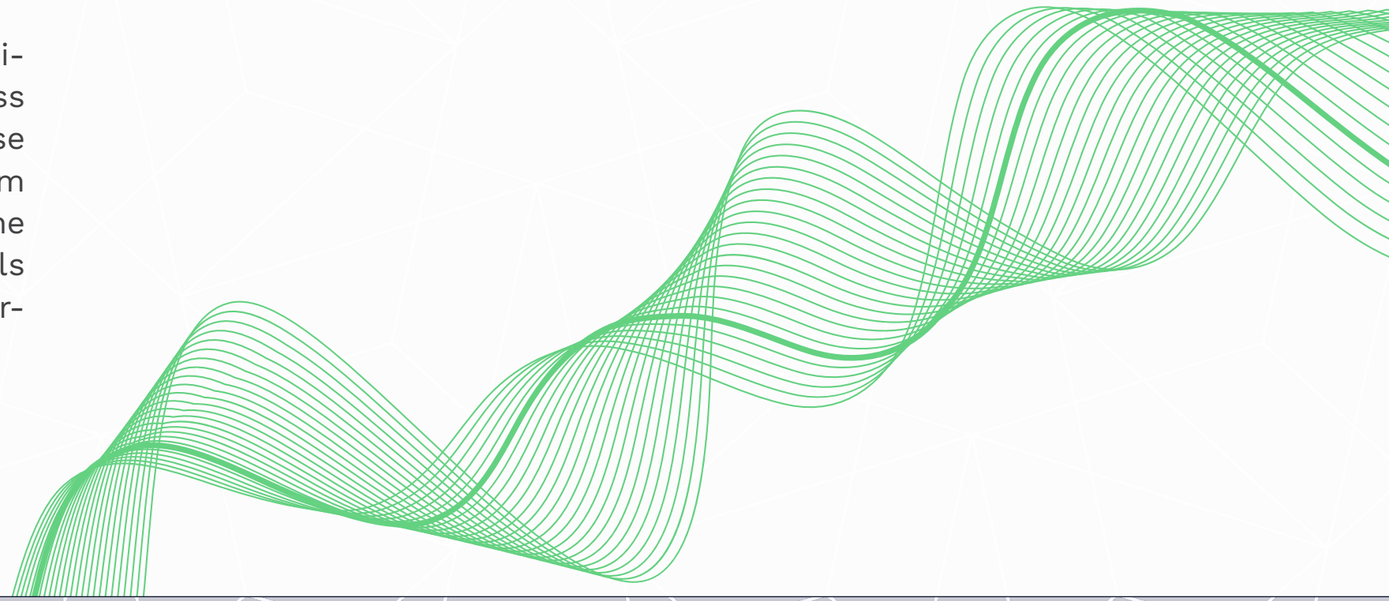

© 2024 PYMNTS Intelligence All Rights Reserved

About

PYMNTS INTELLIGENCE

PYMNTS Intelligence is a leading global data and analytics platform that uses proprietary data and methods to provide actionable insights on what's now and what's next in payments, commerce and the digital economy. Its team of data scientists include leading economists, econometricians, survey experts, financial analysts and marketing scientists with deep experience in the application of data to the issues that define the future of the digital transformation of the global economy. This multilingual team has conducted original data collection and analysis in more than three dozen global markets for some of the world's leading publicly traded and privately held firms.

The PYMNTS Intelligence team that produced this Tracker:

Aitor Ortiz
Managing Director

Randall Brown
Senior Writer

Disclaimer

The Generative AI Tracker® Series may be updated periodically. While reasonable efforts are made to keep the content accurate and up to date, PYMNTS MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND, EXPRESS OR IMPLIED, REGARDING THE CORRECTNESS, ACCURACY, COMPLETENESS, ADEQUACY, OR RELIABILITY OF OR THE USE OF OR RESULTS THAT MAY BE GENERATED FROM THE USE OF THE INFORMATION OR THAT THE CONTENT WILL SATISFY YOUR REQUIREMENTS OR EXPECTATIONS. THE CONTENT IS PROVIDED "AS IS" AND ON AN "AS AVAILABLE" BASIS. YOU EXPRESSLY AGREE THAT YOUR USE OF THE CONTENT IS AT YOUR SOLE RISK. PYMNTS SHALL HAVE NO LIABILITY FOR ANY INTERRUPTIONS IN THE CONTENT THAT IS PROVIDED AND DISCLAIMS ALL WARRANTIES WITH REGARD TO THE CONTENT, INCLUDING THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, AND NON-INFRINGEMENT AND TITLE. SOME JURISDICTIONS DO NOT ALLOW THE EXCLUSION OF CERTAIN WARRANTIES, AND, IN SUCH CASES, THE STATED EXCLUSIONS DO NOT APPLY. PYMNTS RESERVES THE RIGHT AND SHOULD NOT BE LIABLE SHOULD IT EXERCISE ITS RIGHT TO MODIFY, INTERRUPT, OR DISCONTINUE THE AVAILABILITY OF THE CONTENT OR ANY COMPONENT OF IT WITH OR WITHOUT NOTICE.

PYMNTS SHALL NOT BE LIABLE FOR ANY DAMAGES WHATSOEVER, AND, IN PARTICULAR, SHALL NOT BE LIABLE FOR ANY SPECIAL, INDIRECT, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, OR DAMAGES FOR LOST PROFITS, LOSS OF REVENUE, OR LOSS OF USE, ARISING OUT OF OR RELATED TO THE CONTENT, WHETHER SUCH DAMAGES ARISE IN CONTRACT, NEGLIGENCE, TORT, UNDER STATUTE, IN EQUITY, AT LAW, OR OTHERWISE, EVEN IF PYMNTS HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

SOME JURISDICTIONS DO NOT ALLOW FOR THE LIMITATION OR EXCLUSION OF LIABILITY FOR INCIDENTAL OR CONSEQUENTIAL DAMAGES, AND IN SUCH CASES SOME OF THE ABOVE LIMITATIONS DO NOT APPLY. THE ABOVE DISCLAIMERS AND LIMITATIONS ARE PROVIDED BY PYMNTS AND ITS PARENTS, AFFILIATED AND RELATED COMPANIES, CONTRACTORS, AND SPONSORS, AND EACH OF ITS RESPECTIVE DIRECTORS, OFFICERS, MEMBERS, EMPLOYEES, AGENTS, CONTENT COMPONENT PROVIDERS, LICENSORS, AND ADVISERS.

Components of the content original to and the compilation produced by PYMNTS is the property of PYMNTS and cannot be reproduced without its prior written permission.

The Generative AI Tracker® Series is a registered trademark of What's Next Media & Analytics, LLC ("PYMNTS").