International Journal of Advance and Applied Research
www.ijaar.co.in
ISSN – 2347-7075 Impact Factor – 8.141
Peer Reviewed Bi-Monthly
Vol.7 No.22 March –April- 2026

# Fraud Detection in Banking Transactions

**Patiksha Somnath Udawant**
*Department of Computer Science,*
*Dr. D. Y. Patil Arts, Commerce and Science College, Akurdi, Pune – 44*
***Corresponding Author – Patiksha Somnath Udawant***
**DOI - 10.5281/zenodo.19345462**

***Abstract:***

*Machine Learning (ML) is fundamentally transforming banking fraud detection, offering unprecedented opportunities for high-volume scalability and the identification of complex, multi channel criminal networks, such as Synthetic Identity Fraud. However, this shift introduces significant operational vulnerabilities, notably the economic burden of high false positive rates, which can account for 19% of the total cost of fraud , and the risk of targeted adversarial evasion attacks against predictive models. This paper synthesizes recent findings to explore the dual-edged nature of ML in financial security, examining how deep learning architectures (e.g., CNN-LSTM) demonstrate superior performance while highlighting systemic challenges like data silos and regulatory conflicts (GDPR vs. AML). Drawing on strategic insights, it discusses the imperative for risk-aligned metric optimization, specifically recommending the adoption of the weighted F-beta score to prioritize recall , alongside technological safeguards such as adversarial training to enhance model robustness. Findings underscore the critical need for unified data architectures and collaborative industry benchmarking to build a resilient defense against evolving financial crime.*

## Introduction:

The modern financial landscape is characterized by accelerating digital transformation, which, while beneficial for customers, provides a sophisticated environment for organized financial crime. The nature of financial crime has fundamentally shifted; the core challenge is no longer about stopping individual bad actors but rather dismantling complex, multi-channel criminal rings that exploit systemic gaps across different banking products and platforms—ranging from online portals to payment processing and loan origination. Criminals fabricate entirely new identities by combining real and stolen information (Synthetic Identity Fraud), orchestrating complex webs of activity that are nearly impossible to trace using conventional methods.

Legacy fraud detection systems, which are typically static, rule-based, and often operate within isolated data silos, are fundamentally ill-equipped for this new reality. These older systems are known to generate overwhelming false positives and lack the contextual intelligence needed for comprehensive threat analysis. Machine learning and Artificial Intelligence (AI) are emerging as the necessary transformative force, offering adaptive, proactive defense capabilities that can manage high transaction volumes and dynamically respond to evolving tactics. This paper synthesizes recent technical and operational literature to explore the efficacy and feasibility of ML in securing banking transactions, addressing its capacity for scalable threat mitigation, its intrinsic vulnerabilities, and the required strategic frameworks for achieving

sustainable operational resilience in an increasingly digital and adversarial environment.

**Operational Opportunities of ML in Banking Fraud Detection:**

ML systems present compelling advantages over conventional methods, largely centered on dynamic adaptability, scalability, and enhanced precision in highly complex and high-volume environments.

**Fraud Detection Process Flow:**

Modern fraud detection in banking typically follows this sequence:

1. Transaction initiated by the customer.
2. Data collected from transaction and account behavior.
3. Preliminary rule-based checks for obvious anomalies.
4. Machine learning models analyze transactions for deeper patterns.
5. Suspicious activity flagged automatically.
6. Human review for flagged transactions, if required.
7. Final action: approve, block, or monitor transaction further.

**Process Flow Diagram:**

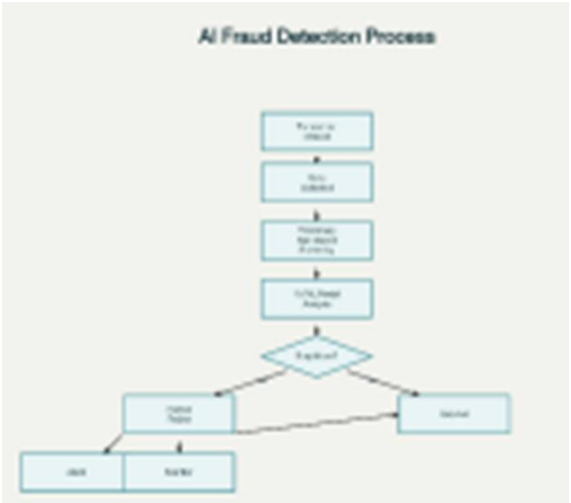


**AI-Based Fraud Detection Process in Banking:**

**Types of Banking Fraud:**

AI tools help detect a variety of frauds in the banking sector. The most frequent types and their prevalence (approximate) are:

- Phishing (25%)
- Identity Theft (20%)
- Card Skimming (15%)
- Account Takeover (20%)
- Internal Fraud (10%)
- Money Laundering (10%)

**Fraud Types Diagram:**

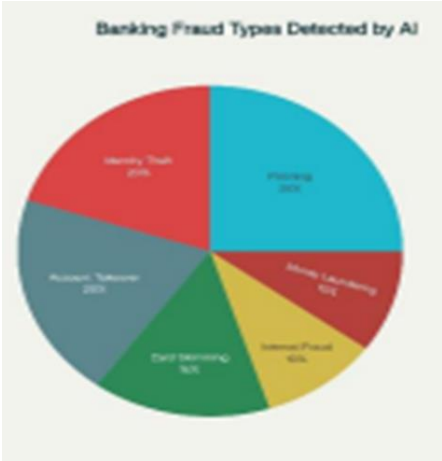


**Real-Time Scalability and High-Volume Processing Efficiency:**

ML-powered fraud protection offers inherent scalability, capably handling massive transaction volumes during peak periods without requiring the constant addition of new, complex rules characteristic of legacy systems. This operational advantage is critical in digital banking where users expect immediate feedback and decisions must often be made in milliseconds. The ability of ML models to manage peak transaction volumes, such as during high-traffic sales events or product launches, provides a stability that rule-based systems cannot match.

Furthermore, ML increases efficiency by providing automated 'approve' or 'decline' decisions, accelerating the order flow and fulfillment process in payment systems. Benchmarking key operational metrics, such as average detection time and case closure efficiency, allows financial institutions (FIs) to evaluate how effectively their AI tools perform compared to industry standards. This continuous

evaluation is necessary to fine-tune algorithms for optimal processing speed, ensuring that the enhanced accuracy of complex models does not inadvertently translate into unacceptable user friction due to increased latency.

**Superior Anomaly Detection and Deep Learning Architectures:**

ML models leverage historical, behavioral, and transaction data to identify subtle, unusual spending patterns and high-risk geographies that defy static rules. This capability enables the identification of sophisticated fraud schemes where criminals intentionally modify standard transaction features to evade detection.

Algorithmic Superiority of Deep Learning: Advanced deep learning (DL) architectures demonstrate superior effectiveness in identifying fraudulent transactions compared to traditional ML techniques. Empirical studies have shown that sequential and contextual models, such as Convolutional Neural Networks combined with Long Short-Term Memory networks (CNN LSTM), achieved higher overall performance across standard metrics, including accuracy, precision, and recall, compared to models like Random Forest. This performance enhancement is attributed to the deep learning models' capacity to capture complex temporal and spatial dependencies inherent in organized fraud. Additionally, Graph Neural Networks (GNNs) are recognized for robust generalization when analyzing transaction data, particularly when modeling relationships between entities like users, merchants, and devices. The quantitative superiority of these deep learning models is a direct, algorithmic necessity for addressing the qualitative shift in financial crime toward organized, multi-channel exploitation. These architectures are uniquely positioned to capture the interconnectedness of complex, sequential activities—linking separate fraudulent actions across platforms to dismantle the criminal web that traditional methods cannot perceive.

**Handling Extreme Class Imbalance:**

The successful application of ML in banking is contingent upon effective handling of the severe class imbalance inherent in transaction data, where fraudulent transactions typically comprise less than 0.17% of the total dataset. This imbalance biases models toward the majority (non-fraudulent) class, leading to high precision but catastrophically low recall (sensitivity to actual fraud).

To counteract this, sophisticated pre processing is mandatory. The Synthetic Minority Over-sampling Technique (SMOTE) is highly effective, generating synthetic minority class samples by interpolating between existing fraudulent instances. This approach creates a balanced training dataset, often targeting a 1:1 ratio, which is critical for achieving enhanced model sensitivity (Recall) to the fraudulent class. The cost of a False Negative (failing to detect a fraudulent transaction) is exceptionally high, resulting in "significant financial loss". Therefore, any methodology, such as SMOTE, that enhances the model's ability to identify the minority class is an indispensable operational requirement, justifying the added computational complexity during the training phase.

**Critical Challenges and Vulnerabilities in ML-Driven Fraud Detection:**

Despite the algorithmic opportunities presented by machine learning, the adoption of these systems in banking introduces severe operational friction points, systemic vulnerabilities, and entirely new adversarial attack surfaces.

**The Financial and Operational Burden of False Positives (FPs):**

One of the most immediate and costly

non algorithmic challenges is the management of False Positives (FPs), which are legitimate transactions erroneously declined due to overly sensitive fraud detection parameters. This is not merely an inconvenience; it represents a multi-billion dollar problem. False positive losses are estimated to account for 19% of the total cost of fraud, a figure substantially higher than the 7% loss typically attributed to actual fraud.

This staggering ratio reveals a critical contradiction: systems designed to prevent financial loss are simultaneously causing a larger indirect financial loss due to a metric imbalance. High FP rates, which can exceed 95% in poorly managed or legacy rule-based systems , result in rising operational costs due to necessary manual reviews, wasted analyst time, and critically, "broken customer journeys". Such issues erode customer trust and cause significant frustration. Therefore, the modern imperative for fraud detection is shifting from simple fraud avoidance to risk-managed customer experience, forcing institutions to adopt cost-sensitive metrics that explicitly weigh the expense of false declines against the risk of undetected fraud.

**Latency-Accuracy Trade-offs in Real-Time Systems:**

The necessity of real-time detection creates an inherent conflict between model complexity (accuracy) and processing time (latency). In financial applications, decisions must often be instantaneous to ensure efficiency and a seamless customer experience. Low latency is therefore a critical operational requirement.

However, a direct trade-off exists: the complex deep learning models necessary to achieve high accuracy in detecting sophisticated fraud require extensive computational resources and time to process inputs, inevitably increasing latency. Conversely, simplifying models to reduce response time inherently lowers the prediction accuracy and may increase the risk of missing sophisticated fraud patterns. Managing this constraint is vital, as the acceptable balance between these two factors varies based on the specific application, such as high-frequency trading versus loan origination.

**Data Fragmentation, Silos, and Contextual Blindness:**

The financial services industry frequently operates with isolated data architectures, which acts as a profound organizational barrier to effective fraud detection. Typically, specialized fraud teams work with transaction events, account profiles, and histories, while cybersecurity teams manage authentication logs, firewall data, and Security Information and Event Management (SIEM) telemetry.

This systemic data fragmentation prevents the necessary data fusion required for holistic behavioral pattern analysis. Organized financial crime explicitly exploits gaps between systems and product lines. If the data streams are segregated, the model lacks the contextual intelligence required to link cyber-level events with transactional events, making effective feature engineering impossible. The best ML models are thus rendered ineffective because they lack the necessary cross-channel feature inputs needed to uncover the "hidden connections" and trace the complex, multi channel exploitation paths utilized by sophisticated criminal networks. This structural challenge transforms fraud detection from a purely machine learning optimization problem into a fundamental enterprise architecture problem.

**Model Robustness and Adversarial Evasion Attacks:**

The reliance on predictive ML models introduces a new attack surface: model vulnerability to targeted manipulation. ML systems are vulnerable to evasion attacks, a type of cyber threat where malicious actors subtly

manipulate the input data (the transaction parameters) to cause the model to make incorrect predictions and bypass detection.

In the financial context, this involves attackers altering features such as transaction timing, amount, or frequency in calculated ways to ensure the fraud detection model produces an output favorable to them, such as approving an unauthorized transaction. The model itself continues to operate normally, but its inherent limitations are exploited. These targeted manipulations—which essentially involve fooling the AI—directly impact the reliability of the system and, without mitigation, could paralyze financial security measures.

**Regulatory Compliance and Data Governance Constraints:**

Financial institutions operating globally face the immense challenge of balancing the necessity of robust transaction monitoring for Anti-Money Laundering (AML) and Know Your Customer (KYC) compliance with the strict mandates of consumer data privacy, particularly the General Data Protection Regulation (GDPR) in the European Union (EU).

AML regulations require FIs to actively monitor suspicious transactions and maintain records. Conversely, GDPR emphasizes personal data protection, requiring institutions to establish legal grounds for processing personal data (such as transaction logs and customer identification data) and demanding transparency. Furthermore, GDPR places the burden of proof on the data controller, requiring banks to demonstrate "full control" over all personal data flows. This introduces significant friction when deploying complex ML models that require vast, centralized pools of personal data for training and inference, potentially conflicting with data retention policies and the principle of data minimization.

**Strategic Mitigation and Achieving Operational Resilience:**

Achieving resilience against sophisticated financial crime requires integrated strategies that optimize performance metrics for actual risk, unify disparate data architectures, and proactively fortify models against adversarial threats.

Optimizing Performance Metrics with Weighted F-Beta Scoring Traditional ML evaluation, often relying on the standard F1 score, applies equal weight ($\beta = 1$) to Precision and Recall.7 In fraud detection, however, the financial cost of a False Negative (missed fraud) is catastrophic, resulting in "significant financial loss" 8, while the cost of a False Positive is high but manageable (19% of total cost of fraud).2 Therefore, a cost-sensitive and risk-aligned metric is required. The weighted F-beta score ($\text{F}_{\beta}$) is the definitive solution, allowing institutions to strategically align the model's performance with organizational risk tolerance.7 By implementing a factor where $\beta > 1$ (e.g., $\beta=2$ if recall is considered twice as important as precision), the metric heavily weights Recall, thereby minimizing the occurrence of the most catastrophic error: undetected fraud.7 This metric optimization is an essential strategic defense layer, ensuring the system's primary objective function is rooted in minimizing financial risk rather than simply maximizing statistical accuracy.7.2. Architecting for Low-Latency and High-Accuracy To resolve the inherent conflict between processing speed and prediction quality 11, institutions must deploy hybrid detection architectures.22 A purely monolithic, high-complexity model is impractical for real-time applications. Instead, a tiered or hybrid approach allows a fast, low-accuracy algorithm (such as a simple Logistic Regression or rule check) to handle initial screening and low risk transactions rapidly.22 Concurrently, a slower, more

sophisticated, high-accuracy deep learning model (e.g., CNN-LSTM) is reserved for refining decisions concerning high-risk or ambiguous flags in the background.22 This framework allows for dynamic thresholding, adjusting accuracy demands based on current system load or latency constraints, thus ensuring real-time speed for the majority of transactions while dedicating necessary computational resources to high-risk events.The effectiveness of these strategies is underpinned by the implementation of an integrated, real-time data and model pipeline, as conceptualized in Figure 3. This architecture merges temporal and relational analysis capabilities to provide a holistic view of the transaction environment.

**Proactive Model Defense via Adversarial Training:**

To secure ML models against deliberate manipulation through evasion attacks , passive monitoring is insufficient; hybrid defense mechanisms that integrate security into the model lifecycle are mandatory. This involves recognizing that AI security is an extension of traditional software development.

Studies have confirmed the superiority of adversarial training over standalone detection algorithms. Adversarial training involves intentionally augmenting the training dataset with malicious, perturbed examples that mimic the input manipulations used by fraudsters. This process fortifies the model's decision boundary, enhancing its inherent robustness against unforeseen malicious inputs and providing a substantial robustness gain (e.g., 23.29% increase reported in countermeasure evaluations). Resilience against targeted manipulation is best achieved by integrating this adversarial training regime with continuous real-time anomaly detection.

**Data Unification and Enhanced Feature Engineering:**

Overcoming the data silo problem requires an organizational and technical commitment to data fusion. Institutions must implement investigative platforms capable of unifying data from traditionally separate domains, specifically integrating cybersecurity and network telemetry data (e.g., SIEM logs, EDR data) with financial transaction systems and customer profiles.

This systemic integration enables sophisticated feature engineering, which is the process of creating complex contextual features necessary for effective supervised learning models. Without fusing these disparate streams, it is impossible to generate the comprehensive behavioral features required to identify the specific characteristics of organized, cross-channel fraud. Furthermore, the legal necessity under GDPR to audit and demonstrate "full control" over personal data flows unintentionally acts as a strategic driver for this architectural cleanup, forcing institutions to map and manage data streams which is a necessary precondition for effective unification and advanced feature enrichment.

**Future Trends and Industry Resilience (Outlook 2026):**

The future of banking fraud detection will be defined by the institutional response to AI powered threats, characterized by enhanced collaboration and continuous regulatory adaptation.

**Collaborative Fraud Benchmarking and Shared Intelligence:**

As criminal tactics become more generalized and global, the internal effectiveness of proprietary AI systems must be validated against externally evolving threats through mandatory fraud benchmarking. The future

mandates industry collaboration, similar to frameworks established by organizations like the Financial Crimes Enforcement Network (FinCEN), where anonymized benchmarks related to fraud detection rates, average fraud loss rates, and operational metrics are shared. This collective intelligence provides firms with generalized threat insights, allowing them to measure their operational resilience against common campaigns and adapt their models more quickly than isolated internal analyses permit.

**Continuous Regulatory Evolution:**

Regulatory bodies, including the Financial Action Task Force (FATF), will continue to evolve international guidelines, specifically urging risk-based assessments and active monitoring of suspicious transactions. This compliance pressure ensures that fraud detection solutions remain aligned with critical KYC and AML mandates. Crucially, as ML models automate high-stakes decisions, future systems must incorporate explainability (XAI) features to ensure that automated decisions can be audited, justified, and demonstrated to comply with personal data processing rules under regulations like GDPR, particularly the requirement for demonstrated control over data processing.

**Conclusion:**

The integration of Machine Learning into banking security is an essential evolutionary step, providing the necessary scalability and complexity to counter organized financial crime and sophisticated synthetic identity schemes. By utilizing advanced deep learning architectures such as CNN-LSTM and GNNs, institutions can achieve superior anomaly detection capabilities compared to legacy systems.

However, this paper concludes that systemic vulnerabilities pose substantial threats to sustained operational success. Specifically, the exorbitant and often overlooked economic cost of false positives (19% of total fraud cost) , the inhibiting effect of organizational data silos , and the emerging operational risk of targeted adversarial evasion attacks must be proactively addressed.

Resilience is attainable only through the disciplined adoption of strategic mitigation efforts. These include utilizing risk-aligned metrics like the F-beta score ($\beta>1$) to strategically minimize catastrophic False Negatives , implementing hybrid detection architectures (as illustrated in Figure 3) to effectively manage the latency-accuracy conflict , and employing proactive AI security measures, such as adversarial training, to fortify models against manipulation. Ultimately, collaborative efforts involving industry benchmarking and rigorous adherence to unified data governance standards (driven partially by mandates like GDPR/AML) are essential to ensure ML empowers, rather than endangers, the security and integrity of the global financial ecosystem.

**References:**

1. Financial Crime Academy. (n.d.). GDPR and AML: Transaction Monitoring.
2. Sift. (n.d.). FIBR: Fraud Industry Benchmarking Resource.
3. JPMorgan. (n.d.). CNP fraud prevention: combat chargebacks.
4. DataWalk. (n.d.). Fraud Detection in Banking 2026: Future Trends Predictions.
5. Alkami. (n.d.). Cybersecurity in Digital Banking: Breaking Silos, Building Fusion.
6. Fintech Series. (n.d.). Fraud Benchmarking with FinTech..
7. Wikipedia. (n.d.). Precision and recall.
8. Towards Data Science. (n.d.). Performance Metrics: Confusion Matrix, Precision, Recall, and F1 Score.
9. Sunandoroy. (2025). Fraud Detection in Financial Transactions: Challenges and

Innovations.
10. Bugfree.ai. (n.d.). Latency vs Accuracy Trade-offs in Real-Time Systems.
11. Milvus. (n.d.). What are the tradeoffs between latency and accuracy.
12. BuiltIn. (n.d.). Feature Engineering.
13. ResearchGate. (n.d.). Fraud Detection
14. Algorithms: Supervised vs Unsupervised Learning.
15. ArXiv. (n.d.). Performance comparison of LSTM, GNN, and Random Forest.
16. ResearchGate. (n.d.). Comparison of Classification Metrics for CNN, LSTM, and Random Forest Models.
17. PwC. (2017). GDPR Banking Industry Report.
18. Signifyd. (n.d.). Rules-Based vs Machine Learning Fraud Protection.
19. PayPal. (n.d.). Fraud Prevention with Rules vs Machine Learning.
20. Sopra Steria. (n.d.). False Positives Are a Cost Centre.
21. MDPI. (n.d.). Illustrative data for table comparing fraud detection strategies.
22. Frontiers in Artificial Intelligence. (2025). Its effectiveness was evaluated by comparing model performance before and after resampling, with particular attention to recall and F1 score.
23. Neptune.ai. (n.d.). Adversarial Machine Learning Defense Strategies.
24. Journal JERR. (2025). Adversarial Threats to AI Driven Systems: Exploring the Attack Surface of Machine Learning Models and Countermeasures.
25. Nightfall.ai. (n.d.). Evasion Attacks.
26. Startup Defense. (n.d.). Evasion Attacks ML.
27. International Journal of Advance and Applied Research. (2025). Balancing AI's Power in Cybersecurity: Opportunities and Vulnerabilities