ACTA SCIENTIAE

ISSN: 2178-7727

AI-Based Transaction And Fraud Detection Observation Systems

Furqaan Mujtahid

1051 Amador St Claremont, CA 91711

ABSTRACT

The exponential growth of digital transactions has created unprecedented opportunities for fraudulent activities, necessitating advanced detection mechanisms beyond traditional rule-based systems. Artificial intelligence and machine learning technologies have emerged as powerful tools for identifying fraudulent patterns in real-time transaction environments. This paper examines the current state of AI-based fraud detection systems, exploring various machine learning algorithms, deep learning architectures, and hybrid approaches employed in financial transaction monitoring. We analyze the challenges associated with imbalanced datasets, evolving fraud patterns, and the need for explainable AI in regulatory environments. The study discusses implementation frameworks, performance metrics, and practical considerations for deploying AI-driven fraud detection systems across different financial sectors. Additionally, we address emerging trends including federated learning for privacy-preserving fraud detection, graph neural networks for relationship analysis, and real-time adaptive systems that continuously learn from new fraud patterns. Our findings suggest that while AI-based systems significantly outperform traditional methods in detection accuracy and speed, successful implementation requires careful consideration of data quality, model interpretability, and continuous system adaptation to combat increasingly sophisticated fraud schemes.

KEYWORDS: Artificial Intelligence, Fraud Detection, Machine Learning, Transaction Monitoring, Anomaly Detection, Financial Security

INTRODUCTION

Financial fraud has become one of the most pressing challenges facing the global economy, with losses reaching billions of dollars annually across various sectors including banking, insurance, e-commerce, and digital payments. The traditional approach to fraud detection relied heavily on predefined rules and manual review processes, which proved inadequate against the sophistication and volume of modern fraudulent schemes (Abdallah et al., 2016). As digital transformation accelerates across financial services, the sheer volume of transactions has grown exponentially, making manual oversight practically impossible and creating opportunities for fraudsters to exploit system vulnerabilities.

The integration of artificial intelligence into fraud detection represents a paradigm shift from reactive to proactive security measures. Machine learning algorithms can analyze vast amounts of transactional data, identify subtle patterns invisible to human analysts, and adapt to evolving fraud tactics in ways that static rule-based systems cannot (Hilal et al., 2022). These capabilities have made AI-based systems essential components of modern financial infrastructure, protecting both institutions and consumers from increasingly sophisticated attacks.

However, implementing effective AI-based fraud detection systems involves numerous technical and operational challenges. The highly imbalanced nature of fraud data, where legitimate transactions vastly outnumber fraudulent ones, creates significant difficulties for model training and evaluation. Additionally, fraudsters continuously evolve their tactics to evade detection, requiring systems that can adapt quickly without extensive retraining (Randhawa et al., 2018). The need for real-time processing in transaction environments adds another layer of complexity, as detection systems must make accurate decisions within milliseconds to prevent fraud without disrupting legitimate customer activities.

This paper provides a comprehensive examination of AI-based transaction and fraud detection systems,

DOI: 10.22178/acta.27.1.4

ISSN: 2178-7727

focusing on practical implementations and emerging technologies. We explore various algorithmic approaches, discuss performance considerations, and analyze the challenges and opportunities in this rapidly evolving field. Our goal is to provide insights that can guide both researchers and practitioners in developing more effective fraud detection solutions.

EVOLUTION OF FRAUD DETECTION METHODOLOGIES

The history of fraud detection reflects the ongoing arms race between financial institutions and fraudsters. Early systems relied entirely on manual review processes where human analysts examined suspicious transactions based on simple threshold rules. If a transaction exceeded a certain amount or occurred in an unusual location, it would be flagged for manual investigation. While straightforward, this approach was labor-intensive, slow, and easily circumvented by fraudsters who learned to keep individual transactions below detection thresholds.

The introduction of rule-based expert systems in the 1990s marked the first significant technological advancement. These systems encoded domain expertise into if-then rules that could automatically flag suspicious activities. For example, rules might trigger alerts for multiple transactions from different geographic locations within short timeframes or unusual purchasing patterns for specific account types. While more efficient than purely manual processes, rule-based systems suffered from significant limitations including high false positive rates, inability to detect novel fraud patterns, and constant need for manual rule updates as fraud tactics evolved (Makki et al., 2019).

Statistical methods represented the next evolutionary step, introducing probabilistic modeling to fraud detection. Techniques like logistic regression and decision trees allowed systems to learn from historical data and identify statistical anomalies that might indicate fraud. These approaches improved detection rates and reduced false positives compared to simple rule-based systems, but they still struggled with the dynamic nature of fraud and required substantial feature engineering effort from domain experts.

The emergence of machine learning marked a fundamental shift toward data-driven fraud detection. Instead of relying on explicitly programmed rules, ML algorithms could automatically discover patterns in transaction data that distinguished fraudulent from legitimate behavior. Early ML applications used algorithms like support vector machines and random forests, achieving significantly better performance than previous approaches. However, these traditional ML methods still required careful feature selection and struggled with the complexity and high dimensionality of modern transaction data (Taha and Malebary, 2020).

Deep learning and neural networks have pushed fraud detection capabilities to new levels in recent years. These sophisticated architectures can automatically extract relevant features from raw transaction data, detect complex non-linear patterns, and handle massive datasets with millions of transactions. Recurrent neural networks excel at capturing temporal patterns in transaction sequences, while autoencoders prove effective for anomaly detection in high-dimensional spaces. The combination of increased computational power, availability of large datasets, and algorithmic innovations has made deep learning the current frontier in fraud detection technology (Prusti et al., 2020).

MACHINE LEARNING ALGORITHMS FOR FRAUD DETECTION

Various machine learning algorithms have been successfully applied to fraud detection, each with distinct strengths and limitations that make them suitable for different scenarios and requirements.

Supervised learning algorithms form the backbone of most fraud detection systems. These algorithms learn from labeled historical data where transactions are marked as either fraudulent or legitimate. Logistic regression, despite its simplicity, remains widely used due to its interpretability and computational efficiency, making it suitable for real-time applications where decision explanations are required for regulatory compliance. Random forests and gradient boosting machines have become extremely popular due to their ability to handle complex non-linear relationships, built-in feature importance metrics, and robustness to

DOI: 10.22178/acta.27.1.4

ISSN: 2178-7727

overfitting (Rtayli and Enneya, 2020).

Support vector machines work by finding optimal boundaries that separate fraudulent from legitimate transactions in high-dimensional feature spaces. While computationally intensive for large datasets, SVMs excel in scenarios with clear separation between classes and can capture complex decision boundaries through kernel functions. Neural networks, particularly deep learning architectures, have shown remarkable success in fraud detection by automatically learning hierarchical feature representations from raw data without extensive feature engineering.

Unsupervised learning approaches address the challenge of limited labeled fraud data by detecting anomalies without requiring prior fraud examples. Clustering algorithms like K-means and DBSCAN group similar transactions together, identifying outliers that deviate significantly from normal behavior patterns. Isolation forests specifically designed for anomaly detection work by isolating observations that are few and different, making them particularly effective for fraud detection where fraudulent transactions represent rare events. One-class SVM and autoencoders learn the characteristics of normal transactions and flag anything that deviates significantly from this learned normal behavior (Saia and Carta, 2017).

Table 1: Comparison of Machine Learning Algorithms for Fraud Detection

Algorithm	Strengths	Limitations	Best Use Case
Logistic Regression	Fast, interpretable, low computational cost	Limited to linear patterns	Real-time systems requiring explainability
Random Forest	Handles non-linear patterns, robust to overfitting	Black box nature, slower prediction	Batch processing with complex patterns
Gradient Boosting	High accuracy, feature importance	Prone to overfitting, computationally intensive	Scenarios prioritizing accuracy over speed
Neural Networks	Learns complex patterns automatically	Requires large data, black box	Large-scale systems with abundant data
Isolation Forest	No labeled data needed, scales well	May miss subtle fraud patterns	Situations with limited labeled fraud cases
Autoencoders	Detects novel fraud types, unsupervised	Requires careful threshold tuning	Detecting previously unseen fraud patterns

Ensemble methods combine multiple algorithms to achieve better performance than individual models. By aggregating predictions from diverse models, ensemble approaches reduce both bias and variance while improving robustness. Stacking techniques that use one model's outputs as inputs to another have shown particular promise in fraud detection by leveraging complementary strengths of different algorithms (Zareapoor and Shamsolmoali, 2015).

DEEP LEARNING ARCHITECTURES FOR TRANSACTION MONITORING

Deep learning has revolutionized fraud detection by enabling systems to automatically extract meaningful features from raw transaction data and detect increasingly subtle fraud patterns that traditional methods miss. Convolutional neural networks, traditionally used for image processing, have found applications in fraud detection by treating transaction sequences as temporal patterns similar to images. CNNs can identify local patterns in transaction data and combine them hierarchically to recognize complex fraud signatures. Their ability to share weights across different positions makes them efficient for processing large transaction datasets.

Recurrent neural networks, particularly Long Short-Term Memory (LSTM) networks, excel at analyzing sequential transaction data where temporal relationships matter. Unlike traditional methods that treat each transaction independently, LSTMs maintain memory of previous transactions, allowing them to detect fraud patterns that emerge over time. For example, account takeover fraud often involves a gradual escalation of suspicious activities that LSTM networks can identify by analyzing the entire transaction sequence (Prusti et al., 2020).

DOI: 10.22178/acta.27.1.4

ISSN: 2178-7727

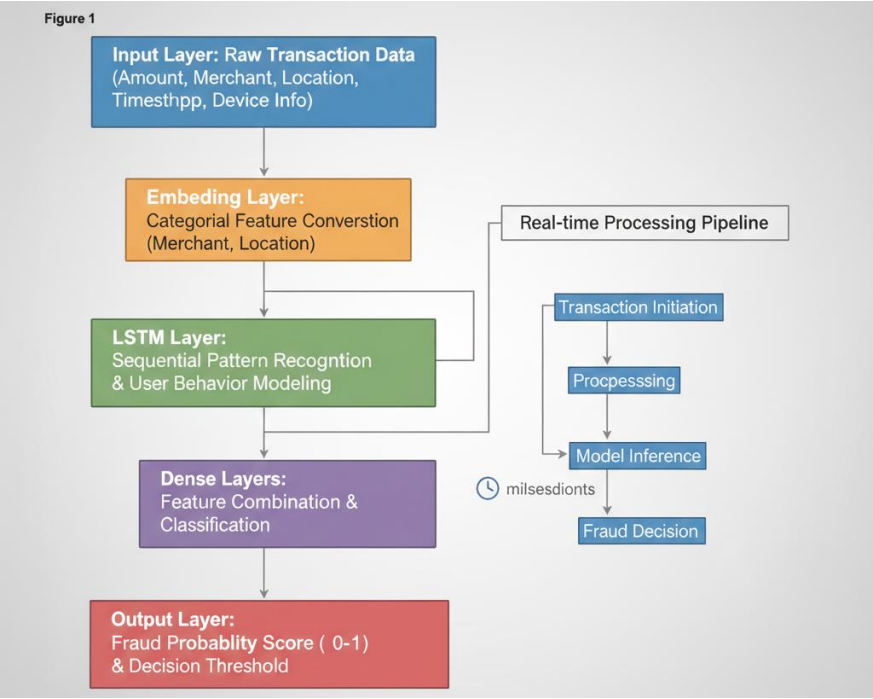

Figure 1: Deep Learning Architecture for Real-Time Fraud Detection

This figure illustrates a multi-layer deep learning system for transaction fraud detection:

- Input Layer: Raw transaction data including amount, merchant category, location, timestamp, device information
- Embedding Layer: Converts categorical features (merchant, location) into dense vector representations
- LSTM Layer: Processes transaction sequences to capture temporal patterns and user behavior over time
- Attention Layer: Focuses on most relevant transactions in the sequence for fraud prediction
- Dense Layers: Combines learned features for final classification
- Output Layer: Fraud probability score (0-1) with decision threshold
- Real-time Processing Pipeline: Shows data flow from transaction initiation through preprocessing, model inference, and fraud decision within milliseconds

Autoencoders function by learning compressed representations of normal transaction patterns and reconstructing input data from these representations. During fraud detection, transactions that cannot be accurately reconstructed from the learned normal patterns are flagged as potentially fraudulent. Variational autoencoders extend this approach by learning probabilistic distributions of normal behavior, providing more nuanced anomaly detection capabilities.

Graph neural networks represent one of the most exciting recent developments in fraud detection. Financial transactions naturally form network structures where accounts, merchants, and devices are nodes connected by transaction edges. GNNs can analyze these network structures to identify suspicious patterns like fraud rings where multiple accounts coordinate fraudulent activities. By propagating information through the transaction graph, GNNs detect fraud patterns that manifest in relationships rather than individual transactions alone (Hilal et al., 2022).

Attention mechanisms have enhanced deep learning models by allowing them to focus on the most relevant parts of transaction data when making predictions. In fraud detection, attention helps models identify which specific transactions or features most strongly indicate fraud, improving both accuracy and interpretability. Transformer architectures combining attention mechanisms with parallel processing have shown promising results in handling long transaction sequences efficiently.

DOI: 10.22178/acta.27.1.4

ISSN: 2178-7727

HANDLING IMBALANCED DATASETS

The class imbalance problem represents one of the most significant challenges in fraud detection, where fraudulent transactions typically account for less than 1% of all transactions. Standard machine learning algorithms trained on such imbalanced data tend to develop a bias toward the majority class, achieving high overall accuracy while failing to detect the minority fraud cases that matter most.

Several approaches address this imbalance challenge at the data level. Oversampling techniques increase the representation of fraudulent transactions in training data through duplication or synthetic generation. The Synthetic Minority Over-sampling Technique (SMOTE) creates synthetic fraud examples by interpolating between existing fraudulent transactions, providing more diverse training examples. However, oversampling risks overfitting if synthetic examples don't represent realistic fraud patterns (Makki et al., 2019).

Undersampling reduces the number of legitimate transactions to balance class distribution. While computationally efficient, undersampling discards potentially valuable information about normal behavior patterns. Hybrid approaches combining both oversampling and undersampling attempt to achieve optimal balance while preserving important information from both classes.

Table 2: Techniques for Addressing Class Imbalance in Fraud Detection

Technique	Approach	Advantages	Disadvantages
Random Oversampling	Duplicate fraud cases	Simple, preserves all fraud examples	Risk of overfitting
SMOTE	Generate synthetic fraud cases	Creates diverse examples	May introduce noise
Random Undersampling	Remove legitimate transactions	Fast training	Loses information
Ensemble Sampling	Combine multiple balanced subsets	Preserves information	Computationally intensive
Cost-sensitive Learning	Assign higher cost to fraud misclassification	No data modification needed	Requires careful cost tuning
Anomaly Detection	Treat fraud as anomaly	Works with extreme imbalance	May miss organized fraud

Algorithm-level approaches modify learning algorithms to better handle imbalance. Cost-sensitive learning assigns higher misclassification costs to fraudulent transactions, forcing models to pay more attention to correctly identifying fraud even at the expense of some false positives. Ensemble methods train multiple models on different balanced subsets of data and combine their predictions, leveraging the entire dataset while maintaining balance during training (Rtayli and Enneya, 2020).

Evaluation metrics must be carefully selected for imbalanced fraud detection scenarios. Accuracy, which measures overall correctness, becomes misleading when 99% accuracy can be achieved by simply classifying everything as legitimate. Instead, practitioners focus on precision (accuracy of fraud predictions), recall (percentage of frauds detected), and F1-score (harmonic mean of precision and recall). The area under the Receiver Operating Characteristic curve (AUC-ROC) provides a threshold-independent measure of model performance across different operating points.

REAL-TIME FRAUD DETECTION SYSTEMS

The effectiveness of fraud detection systems depends not just on accuracy but also on the speed of detection and response. Real-time systems must analyze transactions and make fraud decisions within milliseconds to prevent fraudulent transactions while avoiding delays that frustrate legitimate customers.

System architecture for real-time fraud detection typically involves multiple layers working in parallel. A fast preliminary screening layer using simple rules or lightweight models quickly filters obvious legitimate

DOI: 10.22178/acta.27.1.4

ISSN: 2178-7727

transactions, allowing only suspicious cases to proceed to more sophisticated ML analysis. This tiered approach balances thoroughness with speed, ensuring that complex models focus computational resources on transactions that actually need deeper analysis (Abdallah et al., 2016).

Feature engineering for real-time systems requires careful consideration of latency constraints. While aggregated features based on long-term transaction history provide valuable fraud signals, computing these features in real-time can introduce unacceptable delays. Pre-computed feature stores that maintain up-to-date aggregate statistics allow real-time systems to access complex features instantly without on-the-fly computation. Stream processing frameworks enable continuous updating of these features as new transactions occur.

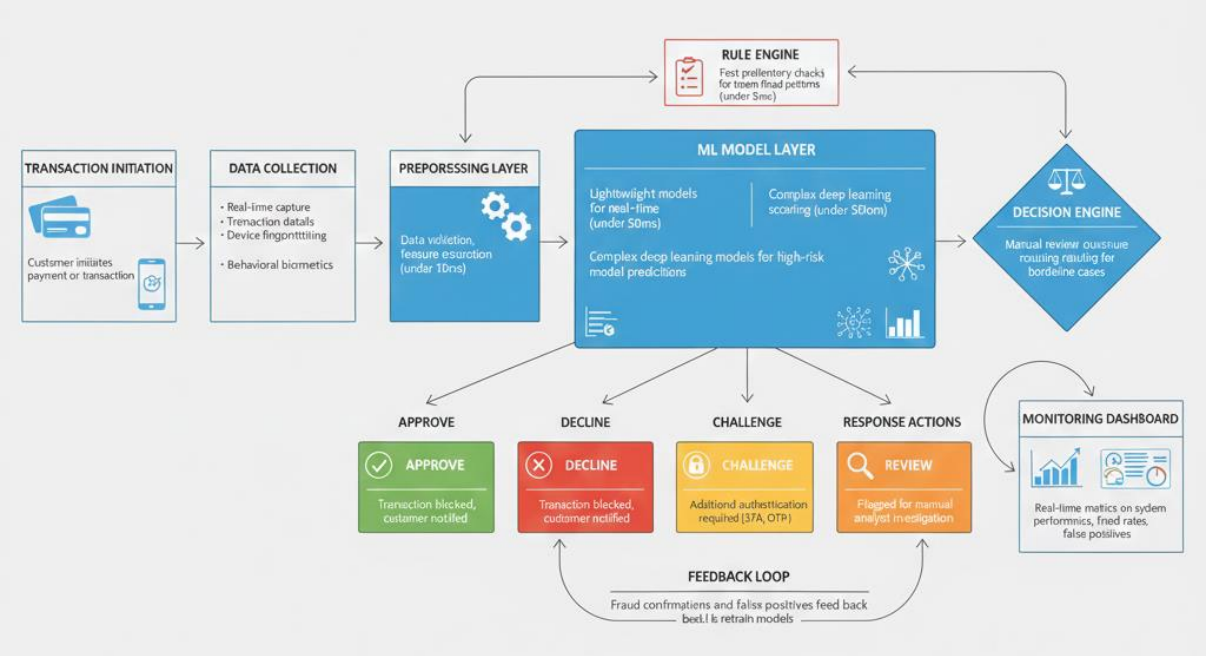

Figure 2: Real-Time Fraud Detection System Architecture

This figure depicts the end-to-end architecture of a production fraud detection system:

- Transaction Initiation: Customer initiates payment or transaction
- Data Collection: Real-time capture of transaction details, device fingerprinting, behavioral biometrics
- Preprocessing Layer: Data validation, normalization, feature extraction (under 10ms)
- Rule Engine: Fast preliminary checks for known fraud patterns (under 5ms)
- ML Model Layer:
    - Lightweight models for real-time scoring (under 50ms)
    - Complex deep learning models for high-risk transactions
    - Ensemble aggregation of multiple model predictions
- Decision Engine:
    - Risk scoring and threshold comparison
    - Business rules application
    - Manual review queue routing for borderline cases
- Response Actions:
    - Approve: Transaction proceeds normally
    - Decline: Transaction blocked, customer notified
    - Challenge: Additional authentication required (2FA, OTP)
    - Review: Flagged for manual analyst investigation
- Feedback Loop: Fraud confirmations and false positives feed back to retrain models
- Monitoring Dashboard: Real-time metrics on system performance, fraud rates, false positives

Model deployment strategies impact real-time performance significantly. Some systems deploy models as microservices that can be independently scaled based on transaction volume. Others use model compilation

DOI: 10.22178/acta.27.1.4

ISSN: 2178-7727

techniques to optimize inference speed, converting trained models into highly efficient execution graphs. Edge deployment, where models run closer to the transaction source, reduces network latency but complicates model updates and monitoring.

Handling model updates in production systems requires careful coordination to avoid service disruptions. A/B testing frameworks allow gradual rollout of new models, comparing their performance against existing systems on live traffic before full deployment. Shadow mode deployment runs new models in parallel with production systems, logging predictions without affecting actual decisions, enabling thorough validation before going live (Hilal et al., 2022).

EXPLAINABLE AI AND REGULATORY COMPLIANCE

The black-box nature of many advanced AI models creates challenges in regulated financial environments where institutions must explain why specific transactions were flagged or declined. Regulatory requirements in many jurisdictions mandate that customers receive explanations for adverse actions, making model interpretability crucial for fraud detection systems.

SHAP (SHapley Additive exPlanations) values have emerged as a popular technique for explaining individual fraud predictions. By calculating each feature's contribution to a specific prediction, SHAP provides human-understandable explanations of why particular transactions were flagged as suspicious. For example, a SHAP explanation might indicate that a transaction was flagged due to unusual transaction amount (40% contribution), foreign location (30%), and time of day (20%), with other factors contributing the remaining 10% (Taha and Malebary, 2020).

LIME (Local Interpretable Model-agnostic Explanations) generates explanations by training simple, interpretable models locally around specific predictions. These local models approximate the complex model's behavior for particular instances, providing understandable explanations without sacrificing the global model's sophisticated pattern recognition capabilities.

Attention mechanisms in deep learning models provide inherent interpretability by highlighting which parts of input data most influenced predictions. In transaction fraud detection, attention weights can show which specific transactions in a sequence or which features most strongly indicated fraud, offering transparency into model decision-making processes.

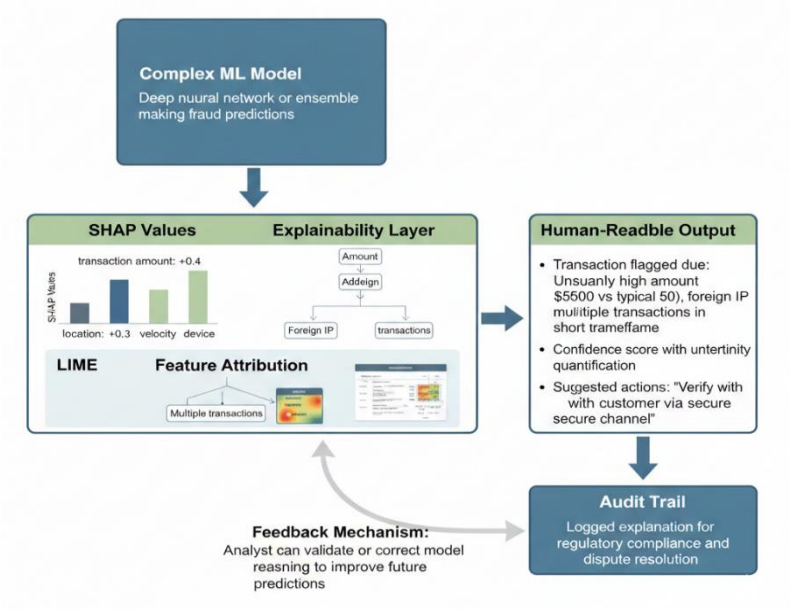

Figure 3: Model Explainability Framework for Fraud Detection

DOI: 10.22178/acta.27.1.4

ISSN: 2178-7727

This figure illustrates how explainable AI techniques provide transparency in fraud detection decisions:

- Complex ML Model: Deep neural network or ensemble making fraud predictions
- Explainability Layer:
  - SHAP Values: Bar chart showing feature importance for specific transaction (transaction amount: +0.4, location: +0.3, velocity: +0.2, device: +0.1)
  - LIME Explanation: Simplified decision tree showing key factors
  - Feature Attribution: Heatmap highlighting most influential transaction attributes
- Human-Readable Output:
  - "Transaction flagged due to: Unusually high amount ($5000 vs typical $50), foreign IP address, multiple transactions in short timeframe"
  - Confidence score with uncertainty quantification
  - Suggested actions: "Verify with customer via secure channel"
- Audit Trail: Logged explanation for regulatory compliance and dispute resolution
- Feedback Mechanism: Analyst can validate or correct model reasoning to improve future predictions

Rule extraction techniques attempt to distill complex models into interpretable rule sets that approximate their behavior. While the extracted rules may not capture every nuance of the original model, they provide actionable insights that human analysts and regulators can understand and verify. This approach allows organizations to benefit from sophisticated AI while maintaining transparency.

Model documentation and governance frameworks ensure that fraud detection systems meet regulatory requirements. Comprehensive documentation of model development, validation, and monitoring processes demonstrates due diligence. Regular model audits verify that systems operate as intended and don't introduce unfair biases. Version control and change management procedures maintain accountability for model updates and their impacts on fraud detection decisions (Prusti et al., 2020).

EMERGING TRENDS AND FUTURE DIRECTIONS

The landscape of AI-based fraud detection continues evolving rapidly with several promising technologies and approaches on the horizon.

Federated learning enables multiple financial institutions to collaboratively train fraud detection models without sharing sensitive customer data. Each institution trains models locally on their data, sharing only model updates rather than raw transactions. This approach allows the industry to benefit from collective fraud intelligence while preserving privacy and competitive confidentiality. Cross-institutional patterns that would be invisible to individual institutions become detectable through federated learning (Randhawa et al., 2018).
Adversarial machine learning addresses the arms race between fraud detection systems and increasingly sophisticated fraudsters. By anticipating how fraudsters might adapt to detection systems, adversarial training creates more robust models that resist evasion attempts. Generative adversarial networks can synthesize realistic fraud scenarios for testing detection systems against attacks that haven't yet occurred in the real world.
Behavioral biometrics analyzes how users interact with devices and applications, creating unique behavioral fingerprints that supplement traditional authentication. Machine learning models can detect account takeover by recognizing changes in typing patterns, mouse movements, or touchscreen interactions that indicate someone other than the legitimate user is conducting transactions. These behavioral signals are difficult for fraudsters to replicate even when they possess stolen credentials.

Graph analytics and network analysis techniques are becoming increasingly sophisticated at identifying organized fraud rings and money laundering schemes. By analyzing the complete network of transactions, accounts, and entities, these systems detect suspicious patterns like circular money flows, layering transactions, or coordinated account activities that individual transaction analysis would miss (Saia and Carta, 2017).

Quantum computing, while still largely theoretical for fraud detection applications, promises to revolutionize pattern recognition and optimization problems. Quantum algorithms could potentially analyze vastly larger

DOI: 10.22178/acta.27.1.4

ISSN: 2178-7727

datasets and more complex patterns than classical computing allows, though practical quantum fraud detection systems likely remain years in the future.

CHALLENGES AND LIMITATIONS

Despite impressive advances, AI-based fraud detection systems face ongoing challenges that require continued research and innovation.

Concept drift, where the statistical properties of fraud change over time, remains a fundamental challenge. Fraudsters continuously adapt their tactics, meaning that models trained on historical data gradually become less effective. Continuous learning systems that automatically update models as new fraud patterns emerge help address this issue, but they introduce risks of inadvertently learning from fraudulent data or reducing detection of established fraud types (Zareapoor and Shamsolmoali, 2015).

Adversarial attacks specifically designed to fool ML models represent an escalating threat. Fraudsters with knowledge of detection systems can craft transactions that evade detection while still achieving fraudulent objectives. Robust ML techniques that maintain performance under adversarial conditions are essential but add complexity and computational overhead to fraud detection systems.

Privacy concerns arise from the extensive data collection and analysis required for effective fraud detection. Balancing fraud prevention with customer privacy rights requires careful consideration, especially under regulations like GDPR that restrict how personal data can be collected and used. Privacy-preserving machine learning techniques that minimize data exposure while maintaining detection effectiveness represent an important research direction.

False positives remain a persistent problem even in advanced AI systems. Each legitimate transaction incorrectly flagged as fraud creates customer friction, potentially damaging relationships and driving customers to competitors. Optimizing the trade-off between catching fraud and minimizing false alarms requires careful threshold tuning and ongoing monitoring (Rtayli and Enneya, 2020).

CONCLUSION

AI-based transaction and fraud detection systems have fundamentally transformed how financial institutions protect themselves and their customers from fraudulent activities. Machine learning and deep learning approaches offer substantial improvements over traditional rule-based systems in detection accuracy, adaptability, and scalability. The ability to automatically learn from data, detect complex patterns, and adapt to evolving fraud tactics makes AI essential for modern fraud prevention.

However, successful implementation requires careful attention to numerous technical and operational challenges. The imbalanced nature of fraud data, need for real-time processing, requirement for explainable decisions, and constantly evolving fraud landscape all demand sophisticated solutions that go beyond simply applying off-the-shelf algorithms. Organizations must invest in data infrastructure, modeling expertise, and continuous monitoring to maintain effective fraud detection capabilities.

Looking forward, emerging technologies like federated learning, behavioral biometrics, and advanced graph analytics promise to further enhance fraud detection capabilities. The ongoing competition between fraud detection systems and increasingly sophisticated fraudsters will continue driving innovation in AI techniques and their applications. Success will require not just technological advancement but also thoughtful consideration of privacy, fairness, and the human elements of fraud prevention. As AI technologies mature and best practices emerge, we can expect fraud detection systems to become even more effective at protecting the integrity of financial systems while minimizing disruption to legitimate customers.

DOI: 10.22178/acta.27.1.4

ISSN: 2178-7727

REFERENCES

1. Abdallah, A., Maarof, M.A., & Zainal, A. (2016). Fraud detection system: A survey. Journal of Network and Computer Applications, 68, 90-113.
2. Hilal, W., Gadsden, S.A., & Yawney, J. (2022). Financial fraud: A review of anomaly detection techniques and recent advances. Expert Systems with Applications, 193, 116429.
3. Makki, S., Assaghir, Z., Taher, Y., Haque, R., Hacid, M.S., & Zeineddine, H. (2019). An experimental study with imbalanced classification approaches for credit card fraud detection. IEEE Access, 7, 93010-93022.
4. Sylvester, Prof. C. J., & Abhilashi, Dr. R. K. (2025). 6th International Conference on "Applications of Artificial Intelligence" (J. Maheshkar, Ed.; 1st ed.). Geh Press. https://www.researchgate.net/publication/398044792_Applications_of_Artificial_Intelligence_Opportunity_of_Start-up_Skill_India
5. Maheshkar, J. A. (2025). Bridging the Gap: A Systematic framework for Agentic AI root cause analysis in hybrid distributed systems. Acta Scientiae, 26(1), 228–245. 10.13140/RG.2.2.32368.52485
6. Rtayli, N., & Enneya, N. (2020). Enhanced credit card fraud detection based on SVM-recursive feature elimination and hyper-parameters optimization. Journal of Information Security and Applications, 55, 102596.
7. Saia, R., & Carta, S. (2017). A frequency-based approach for detecting anomalies in credit card transactions. In Proceedings of the International Conference on Security and Cryptography (pp. 386-391). SciTePress.
8. Taha, A.A., & Malebary, S.J. (2020). An intelligent approach to credit card fraud detection using an optimized light gradient boosting machine. IEEE Access, 8, 25579-25587.
9. Zareapoor, M., & Shamsolmoali, P. (2015). Application of credit card fraud detection: Based on bagging ensemble classifier. Procedia Computer Science, 48, 679-685.
10. Dornadula, V.N., & Geetha, S. (2019). Credit card fraud detection using machine learning algorithms. Procedia Computer Science, 165, 631-641.
11. Varmedja, D., Karanovic, M., Sladojevic, S., Arsenovic, M., & Anderla, A. (2019). Credit card fraud detection - machine learning methods. In 18th International Symposium INFOTEH-JAHORINA (pp. 1-5). IEEE.
12. Lucas, Y., Portier, P.E., Laporte, L., He-Guelton, L., Caelen, O., Granitzer, M., & Calabretto, S. (2020). Towards automated feature engineering for credit card fraud detection using multi-perspective HMMs. Future Generation Computer Systems, 102, 393-402.
13. Bahnsen, A.C., Aouada, D., Stojanovic, A., & Ottersten, B. (2016). Feature engineering strategies for credit card fraud detection. Expert Systems with Applications, 51, 134-142.
14. Kho, J.M., & Rogers, M.K. (2021). A review on machine learning algorithms for fraud detection in financial transactions. Journal of Financial Crime, 28(4), 1277-1293.
15. Ghobadi, F., & Rohani, M. (2021). Cost sensitive learning for improving fraud detection systems. Expert Systems with Applications, 169, 114472.
16. Jaykumar Ambadas Maheshkar. (2025). Bridging the Gap: A Systematic Framework for Agentic AI Root Cause Analysis in Hybrid Distributed Systems. Acta Scientiae, 26(1), 228–245. Retrieved from https://www.periodicos.ulbra.org/index.php/acta/article/view/502
17. Jaykumar Ambadas Maheshkar. (2024). Intelligent CI/CD Pipelines Using AI-Based Risk Scoring for FinTech Application Releases. Acta Scientiae, 25(1), 90–108. Retrieved from https://www.periodicos.ulbra.org/index.php/acta/article/view/532
18. Maheshkar, J. A. (2024c). AI-POWERED PAYMENT FRAUD SIGNATURE GENERATION AND CONTINUOUS RETRAINING METHODS. Power System Protection and Control, 52(4), 75–93. https://doi.org/10.46121/pspc.52.4.7
19. Maheshkar, J. A. (2025b). AUTONOMOUS CLOUD RESOURCE OPTIMIZATION USING REINFORCEMENT LEARNING FOR FINTECH MICROSERVICES. Power System Protection and Control, 53(3), 231–246. https://doi.org/10.46121/pspc.53.3.15
20. Maheshkar, J. A. (2024b, September 20). AI-Driven FinOps: Intelligent Budgeting and Forecasting in Cloud Ecosystems. https://eudoxuspress.com/index.php/pub/article/view/4128

DOI: 10.22178/acta.27.1.4

ISSN: 2178-7727

21. Maheshkar, J. A. (2023). AI-Assisted Infrastructure as Code (IAC) validation and policy enforcement for FinTech systems. Academic Social Research, 9(4), 20–44. https://doi.org/10.13140/rg.2.2.26249.92002
22. Maheshkar, J. A. (2023). Automated code vulnerability detection in FinTech applications using AI-Based static analysis. Academic Social Research, 9(3), 1–24. https://doi.org/10.13140/RG.2.2.32960.80648
23. Sylvester, Prof. C. J., & Abhilashi, Dr. R. K. (2025). 6th International Conference on "Applications of Artificial Intelligence" (J. Maheshkar, Ed.; 1st ed.). Geh Press. https://www.researchgate.net/publication/398044792_Applications_of_Artificial_Intelligence_Opportunity_of_Start-up_Skill_India
24. Maheshkar, J. A. (n.d.). System and Method for Secure AI-Based Financial Technology Governance and Risk Management (US Patent No. 19,391,736) U.S. Patent and Trademark Office.
25. Maheshkar, J. A. (2025). Software Testing Device. UK Intellectual Property Office Patent no. GB6488596. Available at: https://www.search-for-intellectual-property.service.gov.uk/
26. Maheshkar, J. A. (2026). AI-driven cloud engineering migrating and modernizing legacy applications with security, observability, and SRE. Pearson Education. ISBN: 978-1970596311. ASIN: B0GF1NLZX4 https://a.co/d/8DjLAEX
27. Anisha Gupta. (2025-05-24). Business Intelligence Ecosystems: Data Analysts, Visualization Platforms, and Data-Driven Decision Systems (ACTA SCIENTIAE), Vol. 26 No. 2 (2025): ACTA SCIENTIAE, 690-700. Retrieved from
28. https://www.periodicos.ulbra.org/index.php/acta/article/view/562
29. Anisha Gupta. (2024-05-31) . EVALUATING THE IMPACT OF TECHNICAL RECRUITING, TECHNICAL DOCUMENTATION, AND HR ANALYTICS SYSTEMS ON WORKFORCE OPTIMIZATION(Journal of Computational Analysis and Applications (JoCAAA)),3474-3491, Vol. 33 No. 05 (2024): JOCAAA, Retrieved from https://eudoxuspress.com/index.php/pub/article/view/4612
30. Anisha Gupta. (2023-05-31), DATA ANALYTICS LIFECYCLE: OBSERVATION METHODS & PROCESSING FRAMEWORK(Journal of Computational Analysis and Applications (JoCAAA)), Vol. 31 No. 3 (2023): JOCAAA, (851-866), Retrieved from https://eudoxuspress.com/index.php/pub/article/view/4613

DOI: 10.22178/acta.27.1.4