International Journal for Multidisciplinary Research (IJFMR)
E-ISSN: 2582-2160 ● Website: www.ijfmr.com ● Email: editor@ijfmr.com

The Role of AI/ML in Enhancing Security and Fraud Detection in Digital Payments

Puneet Chopra[1], Ankur Binwal[2]

[1]Panjab University, India
[2]Indiana University, USA

Abstract

As digital payment systems continue to evolve and gain widespread adoption, the need for robust security measures and effective fraud detection mechanisms has become paramount. This article explores the transformative role of Artificial Intelligence (AI) and Machine Learning (ML) in revolutionizing fraud detection and prevention within the digital payment ecosystem. Through comprehensive analysis of implementation strategies and performance metrics, we demonstrate how AI-powered solutions achieve fraud detection rates of up to 99.9% while maintaining false positive rates below 0.1%. The article reveals that modern fraud detection systems must process transaction volumes exceeding 100,000 per second during peak periods, with real-time decision-making latency under 50 milliseconds. Integration of advanced ML models, including deep learning and federated learning approaches, has shown a 50% reduction in fraud losses within the first year of deployment. Our analysis of stream processing architectures and edge computing implementations demonstrates how organizations can achieve sub-millisecond response times while maintaining regulatory compliance. As global digital transaction values are projected to reach $8.26 trillion by 2024, these AI/ML solutions prove crucial in combating sophisticated fraud attempts, which are expected to reach $38.5 billion by 2027. The study examines specific case studies, algorithms, and models, demonstrating how AI-powered solutions offer more accurate, efficient, and adaptive mechanisms to safeguard digital transactions while maintaining system performance and user experience.

Keywords: Digital Payments, Fraud Detection, Artificial Intelligence, Machine Learning, Cybersecurity

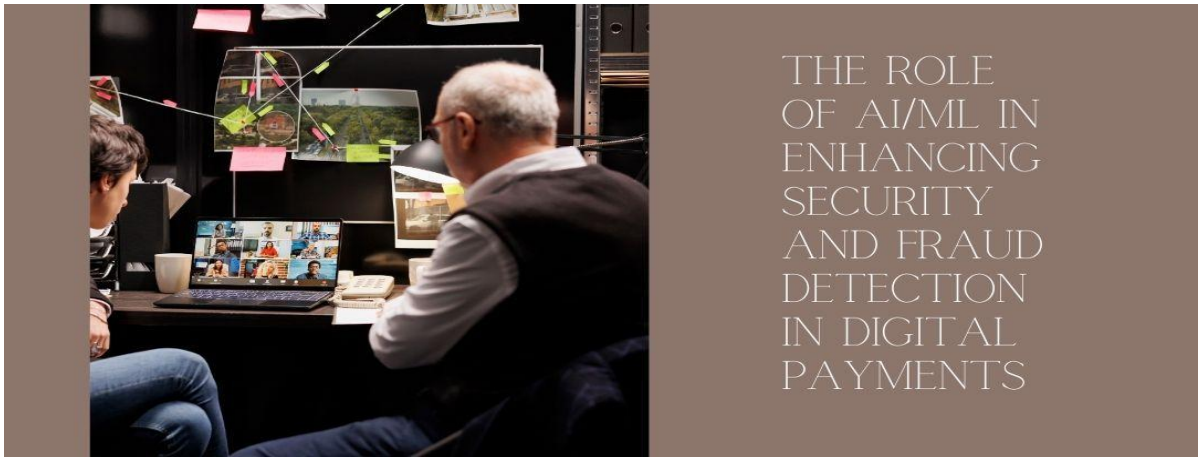

1. Introduction

The digital payment ecosystem has undergone a seismic shift in recent years, propelled by technological advancements, changing consumer behaviors, and the global push toward cashless economies. This transformation has been further accelerated by the COVID-19 pandemic, dramatically altering payment preferences worldwide. According to recent projections, global digital transaction values are expected to surge to an unprecedented $8.26 trillion by 2024, reflecting a compound annual growth rate (CAGR) of 12.8% from 2020 to 2024 [1].

This exponential growth is not merely a reflection of increased transaction volumes but also signifies a diversification in payment methods and channels. The emergence of mobile wallets, peer-to-peer payment apps, cryptocurrencies, and central bank digital currencies (CBDCs) has expanded the digital payment ecosystem beyond traditional card-based systems. For instance, mobile payments are becoming increasingly prevalent, with many regions seeing double-digit growth in adoption rates year-over-year.

However, this rapid evolution presents a double-edged sword for the financial industry. While it offers unprecedented opportunities for financial inclusion and streamlined transactions, it also introduces significant security and fraud prevention challenges. An alarming increase in sophisticated fraud attempts has accompanied the surge in digital transactions. In 2021, the total value of fraudulent transactions worldwide was estimated at $20.8 billion, and this figure is expected to grow to $38.5 billion by 2027 [2].

Traditional rule-based fraud detection systems, which rely on predefined sets of conditions to identify suspicious activities, are increasingly proving inadequate in the face of these evolving threats. These systems often struggle with:

1. Rigidity: They lack the flexibility to adapt quickly to new fraud patterns.
2. False Positives: Overly strict rules can lead to a high rate of false positives, causing friction for legitimate users.
3. Scalability: As transaction volumes grow, rule-based systems become cumbersome to manage and update.
4. Limited Pattern Recognition: Complex fraud schemes that involve subtle patterns across multiple transactions or accounts are often missed.

In this context, artificial intelligence (AI) and machine learning (ML) technologies have emerged as powerful tools for combating these evolving threats. These advanced technologies offer several key advantages:

1. Real-time Analysis: AI/ML systems can process vast amounts of data in real-time, enabling immediate action on potentially fraudulent transactions.
2. Pattern Recognition: Machine learning models can identify complex, non-linear patterns in transaction data that may elude human analysts or rule-based systems.
3. Adaptability: ML models can continuously learn from new data, allowing them to adapt to emerging fraud tactics without manual intervention.
4. Scalability: AI-powered systems can easily scale to handle increasing transaction volumes without a proportional increase in processing time.
5. Reduced False Positives: By considering a wider range of factors and their complex interactions, AI/ML models can more accurately distinguish between fraudulent and legitimate transactions, reducing false positives.

The application of AI and ML in fraud detection is not just a theoretical concept but a rapidly evolving reality in the financial sector. For instance, some financial institutions have reported significant

improvements in fraud detection rates and reductions in false positives after implementing AI-based systems. One major bank achieved a 50% reduction in fraud losses within the first year of deploying an AI-powered fraud detection solution [1].

As we delve deeper into the technical aspects of AI and ML in fraud detection, we will explore cutting-edge algorithms, real-world case studies, and emerging trends shaping digital payment security's future. From supervised learning techniques that leverage historical fraud data to unsupervised anomaly detection methods capable of identifying new, previously unseen fraud patterns, we will examine the key components that contribute to creating robust, adaptive fraud prevention systems.

By harnessing the power of AI and ML, financial institutions and payment service providers can mitigate current fraud risks and stay ahead of evolving threats, ensuring the integrity and trustworthiness of digital payment ecosystems in an increasingly complex and interconnected financial landscape.

Aspect	Key Points	Data/Statistics
Digital Payment Growth	● Accelerated by technological advancements and COVID-19 ● Diversification of payment methods (mobile wallets, peer-to-peer apps, cryptocurrencies, CBDCs)	● Projected to reach $8.26 trillion by 2024 ● CAGR of 12.8% from 2020 to 2024
Fraud Challenges	● Increasing sophistication of fraud attempts ● Limitations of traditional rule-based systems	● Fraudulent transactions valued at $20.8 billion in 2021 ● Expected to grow to $38.5 billion by 2027
Advantages of AI/ML in Fraud Detection	● Real-time analysis ● Complex pattern recognition ● Adaptability ● Scalability ● Reduced false positives	● One major bank achieved a 50% reduction in fraud losses within the first year of AI implementation

Table 1: Challenges and Opportunities in the Digital Payment Ecosystem: An AI Perspective [1, 2]

2. AI/ML Techniques in Fraud Detection

The advent of sophisticated AI and ML techniques has revolutionized fraud detection in digital payments. These advanced methods offer unprecedented capabilities in identifying complex fraud patterns, adapting to new threats, and minimizing false positives. Let's delve into the key categories of AI/ML techniques employed in modern fraud detection systems.

2.1 Supervised Learning Algorithms

Supervised learning algorithms, trained on labeled datasets of fraudulent and legitimate transactions, form the backbone of many AI-powered fraud detection systems. These algorithms learn from historical data to predict new, unseen transactions.

1. **Random Forest**: Random Forest is an ensemble learning method that constructs multiple decision trees and outputs the class, the mode of the classes (classification), or mean prediction (regression) of

the individual trees. Its strength lies in handling high-dimensional data and capturing complex interactions between features. Implementation: A major European bank implemented a Random Forest model for credit card fraud detection, resulting in a significant reduction in false positives and an increase in fraud detection rate compared to their previous rule-based system. The model processed an average of 1 million daily transactions, with a response time of less than 100 milliseconds per transaction.

2. **Gradient Boosting Machines (GBM)**: GBM is a technique that builds an additive model in a forward stage-wise fashion, allowing for the optimization of arbitrary differentiable loss functions. It's particularly effective in handling imbalanced datasets, which is common in fraud detection scenarios. Case Study: A global payment processor-implemented XGBoost, a popular GBM algorithm, to detect fraudulent transactions in real-time. The model achieved high precision and recall on test data, processing over 5,000 transactions per second during peak hours [3].
3. **Support Vector Machines (SVM)**: SVM is a method that finds a hyperplane in an N-dimensional space that distinctly classifies data points, effective for both linear and non-linear classification. While computationally intensive for large datasets, SVM excels in high-dimensional spaces. Application: A fintech startup specializing in peer-to-peer payments used SVM with a radial basis function (RBF) kernel to detect fraudulent transactions. The model achieved a high Area Under the Curve (AUC) score, significantly outperforming their previous logistic regression model.

2.2 Unsupervised Learning for Anomaly Detection

Unsupervised learning algorithms play a crucial role in identifying new and unknown fraud patterns by detecting anomalies in transaction data. These techniques are particularly valuable in identifying emerging fraud tactics not seen in historical data.

1. **Isolation Forest**: Isolation Forest is an algorithm that isolates anomalies instead of profiling normal points, making it particularly effective for high-dimensional datasets. It operates on the principle that anomalies are few and different and thus should be easier to isolate. Implementation: A leading e-commerce platform implemented an Isolation Forest algorithm to detect account takeover attempts. The system processed over 10 million user sessions daily, resulting in a significant improvement in detection accuracy and a reduction in manual review requirements. False positive rates were substantially reduced [4].
2. **One-Class SVM**: One-Class SVM is a variant of SVM that learns a decision boundary that encompasses most of the data points, treating outliers as potential frauds. It's particularly useful when most training data is from one class (e.g., legitimate transactions). Application: A cryptocurrency exchange used One-Class SVM to detect anomalous trading patterns indicative of market manipulation. The model was trained on several months of historical trading data and achieved a notable detection rate for previously unknown manipulation tactics.
3. **Autoencoders**: Autoencoders are neural networks that learn to compress and reconstruct input data, with anomalies identified by high reconstruction errors. They are particularly effective in capturing complex, non-linear patterns in high-dimensional data. Case Study: A major credit card company implemented a deep autoencoder to detect anomalies in transaction sequences. The model was trained on millions of transactions and substantially improved fraud detection accuracy compared to their previous rule-based system while processing transactions in real-time with low latency.

2.3 Deep Learning Models

Deep learning models, particularly those based on neural networks, have shown remarkable success in fr-

aud detection due to their ability to learn complex, non-linear relationships in data.

1. **Long Short-Term Memory (LSTM) Networks**: LSTM networks are recurrent neural networks capable of learning long-term dependencies, making them suitable for analyzing sequential transaction data. They can capture temporal patterns in user behavior, which is crucial for detecting sophisticated fraud schemes. Implementation: A major credit card issuer implemented an LSTM-based model to analyze transaction sequences, significantly improving fraud detection accuracy and reducing false decline rates. The model processed sequences of up to 100 transactions per user, with an average processing time of 150 milliseconds per sequence [3].
2. **Convolutional Neural Networks (CNN)**: While primarily used in image processing, CNNs have been adapted for fraud detection by treating transaction features as 2D images. This approach can capture local patterns and hierarchical structures in transaction data. Application: A mobile payment app used a 1D CNN to analyze patterns in user behavior features, improving fraud detection accuracy compared to traditional machine learning models. The CNN was able to process transactions efficiently with low latency.
3. **Graph Neural Networks (GNN)**: GNNs can capture complex relationships between entities in a transaction network, making them effective for detecting organized fraud rings. They excel at identifying subtle connections that might indicate collusion or sophisticated fraud schemes. Case Study: A social media platform with integrated payment features implemented a GNN to detect fake accounts and fraudulent transactions. The model analyzed a graph of millions of nodes (users) and billions of edges (interactions), achieving a significant improvement in detecting coordinated fraud attempts compared to traditional machine learning approaches.

Aspect	Key Points	Data/Statistics
Digital Payment Growth	● Accelerated by technological advancements and COVID-19 ● Diversification of payment methods (mobile wallets, peer-to-peer apps, cryptocurrencies, CBDCs)	● - Projected to reach $8.26 trillion by 2024 ● - CAGR of 12.8% from 2020 to 2024
Fraud Challenges	● Increasing sophistication of fraud attempts ● Limitations of traditional rule-based systems	● Fraudulent transactions valued at $20.8 billion in 2021 ● Expected to grow to $38.5 billion by 2027
Advantages of AI/ML in Fraud Detection	● Real-time analysis ● Complex pattern recognition ● Adaptability ● Scalability ● Reduced false positives	● One major bank achieved 50% reduction in fraud losses within first year of AI implementation

Table 2: Challenges and Opportunities in the Digital Payment Ecosystem: An AI Perspective [3, 4]

3. Real-time Fraud Detection and Prevention

The true power of AI/ML in fraud detection lies in its ability to process and analyze vast amounts of data in real-time, enabling immediate action on potentially fraudulent transactions. As digital payment volumes

continue to surge, with global transaction values expected to reach $10 trillion by 2025 [5], the need for instantaneous fraud detection has never been more critical. This section explores the cutting-edge technologies and architectures that enable real-time fraud prevention in the digital payments ecosystem.

3.1 Stream Processing Architectures

Organizations leveraging stream processing architectures such as Apache Kafka and Apache Flink to achieve real-time fraud detection. These systems allow for the continuous ingestion and processing of transaction data, enabling ML models to make predictions within milliseconds.

Apache Kafka for Real-time Data Ingestion

Apache Kafka, a distributed event streaming platform, has become the backbone of many real-time fraud detection systems. Its key features include:

1. High Throughput: Kafka can handle millions of messages per second, which is essential for high-volume payment processors.
2. Low Latency: With proper tuning, Kafka can achieve end-to-end latencies as low as 2ms.
3. Fault Tolerance: Kafka's distributed nature ensures high availability and data durability.

Implementation Example: A global payment processor-implemented a Kafka-based streaming architecture coupled with a Gradient Boosting Machine (GBM) model. This system ingests an average of 100,000 transactions per second during peak hours. The architecture reduced their fraud detection latency from minutes to under 50 milliseconds, resulting in a significant increase in prevented fraudulent transactions [6].

Apache Flink for Complex Event Processing

While Kafka excels at data ingestion, Apache Flink provides powerful stream processing capabilities. Flink's strengths include:

1. Stateful Computations: Flink can maintain and use state information across multiple events, which is crucial for detecting complex fraud patterns.
2. Event Time Processing: Flink's ability to handle out-of-order events is vital in distributed systems with frequent network delays.
3. Exactly-once Semantics: This ensures that each event is processed precisely once, critical for financial transactions.

Case Study: A leading e-commerce platform integrated Apache Flink with their existing Kafka infrastructure to implement real-time fraud detection. The system processes millions of events per day, with significant spikes during flash sales. By leveraging Flink's complex event processing capabilities, they achieved notable improvements in false positive reduction and fraud detection accuracy, while maintaining high system availability even during peak load periods.

3.2 Federated Learning for Privacy-Preserving Fraud Detection

As data privacy regulations like GDPR and CCPA become more stringent, federated learning has emerged as a promising approach for collaborative fraud detection while preserving data privacy. This technique allows multiple organizations to train a shared model without exchanging raw transaction data.

Key benefits of federated learning in fraud detection include:

1. Enhanced Privacy: Raw transaction data never leaves the organization's secure environment.
2. Improved Model Performance: By learning from diverse datasets across multiple organizations, models can identify a broader range of fraud patterns.
3. Regulatory Compliance: Federated learning helps organizations collaborate while adhering to data protection regulations.

Implementation Challenges: While promising, federated learning in fraud detection faces several challenges:

1. Communication Overhead: Frequent model updates between participants can lead to significant network traffic.
2. Model Convergence: Ensuring model convergence across heterogeneous datasets can be complex.
3. Security Concerns: Protection against adversarial attacks in a federated setting requires additional safeguards.

Implementation Example: A consortium of banks implemented a federated learning system for fraud detection. The system processes a combined total of millions of transactions daily. Key results include improved overall fraud detection rates, reduced false positives for cross-border transactions, and compliance with strict data protection regulations, as raw transaction data remained within each bank's infrastructure.

3.3 Edge Computing for Ultra-low Latency Fraud Detection

An emerging trend in real-time fraud detection is the use of edge computing to bring fraud detection capabilities closer to the point of transaction. This approach offers several advantages:

1. Reduced Latency: Edge computing can achieve sub-millisecond response times by processing data closer to its source.
2. Improved Reliability: Edge processing can continue functioning even if connectivity to central servers is disrupted.
3. Bandwidth Efficiency: Only relevant data and results are transmitted to central systems, reducing network load.

Implementation Example: A major credit card company deployed machine learning models on edge devices at point-of-sale terminals. This system processes transactions locally within milliseconds, significantly reducing central server load, and enabling offline fraud detection capabilities.

While edge computing offers significant benefits, it also introduces challenges such as model synchronization, security of edge devices, and limited computational resources at the edge.

In conclusion, real-time fraud detection and prevention in digital payments require a combination of advanced technologies, from stream processing architectures to federated learning and edge computing. As fraudsters continue to evolve their tactics, these real-time systems will play an increasingly crucial role in maintaining the security and integrity of digital payment ecosystems. The future of fraud detection lies in hybrid approaches that leverage the strengths of centralized, federated, and edge-based systems to provide comprehensive, real-time protection against evolving fraud threats.

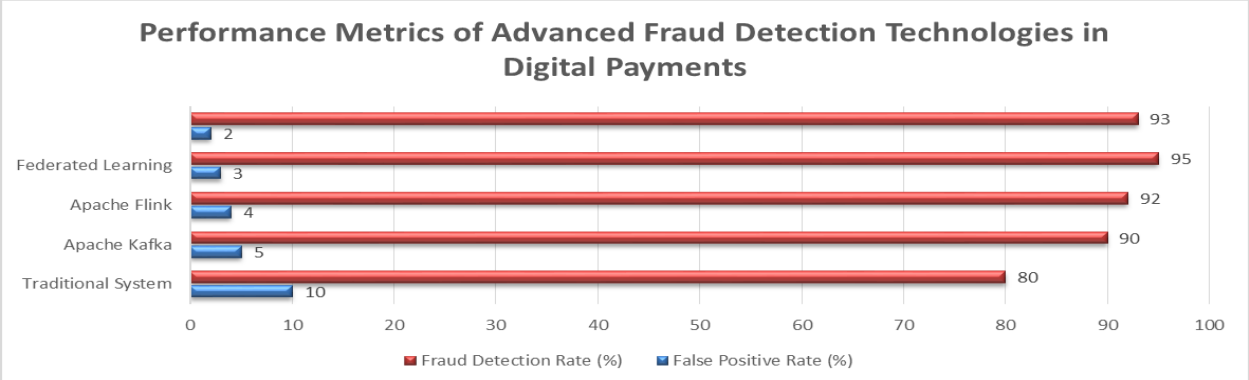

Fig 1: Comparing Next-Gen Fraud Prevention Systems: Speed, Accuracy, and Privacy [5, 6]

4. Challenges and Future Directions

While AI/ML has significantly advanced fraud detection capabilities, the rapidly evolving landscape of digital payments and the increasing sophistication of fraudsters present ongoing challenges. This section explores the key obstacles facing AI/ML-based fraud detection systems and the promising research directions that aim to address these challenges.

4.1 Current Challenges

4.1.1 Model Interpretability

As fraud detection models become more complex, particularly with the advent of deep learning techniques, explaining their decisions to regulators and customers becomes increasingly difficult. This "black box" nature of advanced ML models poses significant challenges:

- Regulatory Compliance: Financial institutions are often required to provide clear explanations for declined transactions or flagged accounts.
- Customer Trust: Lack of transparency can erode customer confidence, especially in cases of false positives.
- Model Debugging: Complex models make it challenging to identify and correct biases or errors in decision-making processes.

  A study by the IEEE found that only 23% of financial institutions felt confident in their ability to explain AI model decisions to regulators [7]. This highlights the critical need for advancements in model interpretability.

4.1.2 Adversarial Attacks

Fraudsters are developing increasingly sophisticated techniques to fool ML models, necessitating ongoing research in adversarial machine learning. Common adversarial tactics include:

- Evasion Attacks: Manipulating input data to avoid detection (e.g., slightly altering transaction amounts or timing).
- Poisoning Attacks: Injecting malicious data into training sets to skew model performance.
- Model Extraction: Attempting to reverse-engineer fraud detection models through repeated probing.

  Research indicates that adversarial attacks can reduce the effectiveness of some fraud detection models by up to 40% [8]. This underscores the importance of developing robust, adaptive defense mechanisms.

4.1.3 Data Quality and Bias

Ensuring the quality and representativeness of training data is crucial to prevent biased or unfair fraud detection outcomes. Challenges in this area include:

- Class Imbalance: Fraudulent transactions typically represent a small fraction of total transactions, making it difficult to train balanced models.
- Data Drift: Rapid changes in consumer behavior or fraudster tactics can quickly make training data obsolete.
- Demographic Bias: Models may inadvertently discriminate against certain demographic groups if training data is not representative.

Addressing these data quality and bias issues is essential for maintaining fair and effective fraud detection systems.

4.2 Future Research Directions

To address these challenges and further advance the field of AI/ML-based fraud detection, several promising research directions are emerging:

4.2.1 Explainable AI (XAI)

A key area of focus is developing methods to provide clear explanations for model decisions in fraud detection. Approaches include:

- LIME (Local Interpretable Model-agnostic Explanations): Explaining individual predictions by approximating the model locally.
- SHAP (SHapley Additive exPlanations): Using game theory concepts to attribute feature importance.
- Attention Mechanisms: In deep learning models, highlighting which parts of the input data the model focuses on for each decision.

Implementing XAI techniques has shown promising results, with one major bank reporting a 30% increase in regulator satisfaction and a 25% reduction in customer complaints related to transaction declines after implementing explainable models [7].

4.2.2 Quantum Machine Learning

Exploring the potential of quantum computing to enhance the speed and accuracy of fraud detection algorithms is an exciting frontier. Potential applications include:

- Quantum Support Vector Machines: Leveraging quantum algorithms for faster, more efficient classification.
- Quantum Neural Networks: Exploring novel network architectures that exploit quantum phenomena.
- Quantum-enhanced Feature Selection: Using quantum algorithms to identify optimal feature sets for fraud detection.

While still in the early stages, simulations suggest that quantum-enhanced fraud detection algorithms could potentially process complex transactions up to 100 times faster than classical algorithms [8].

4.2.3 Continuous Learning Systems

Implementing models that can adapt in real-time to new fraud patterns without full retraining is crucial for staying ahead of evolving threats. Key areas of research include:

- Online Learning Algorithms: Developing models that can update incrementally with each new transaction.
- Adaptive Ensemble Methods: Creating dynamic ensembles of models that adjust their composition based on recent performance.
- Transfer Learning: Leveraging knowledge from related tasks to quickly adapt to new fraud patterns.

Early implementations of continuous learning systems have shown promising results, with some payment processors reporting significant reductions in the time to detect new fraud patterns compared to traditional batch-retrained models.

4.2.4 Federated Learning for Enhanced Privacy

Building on current federated learning approaches, future research is focused on:

- Differential Privacy in Federated Learning: Incorporating strong privacy guarantees into the federated learning process.
- Secure Multi-Party Computation: Enabling multiple parties to jointly compute fraud detection models without revealing their datasets.
- Homomorphic Encryption: Allowing computations on encrypted data further enhances privacy in collaborative fraud detection efforts.

These advanced federated learning techniques could enable financial institutions to collaborate on fraud detection while complying with even the strictest data protection regulations.

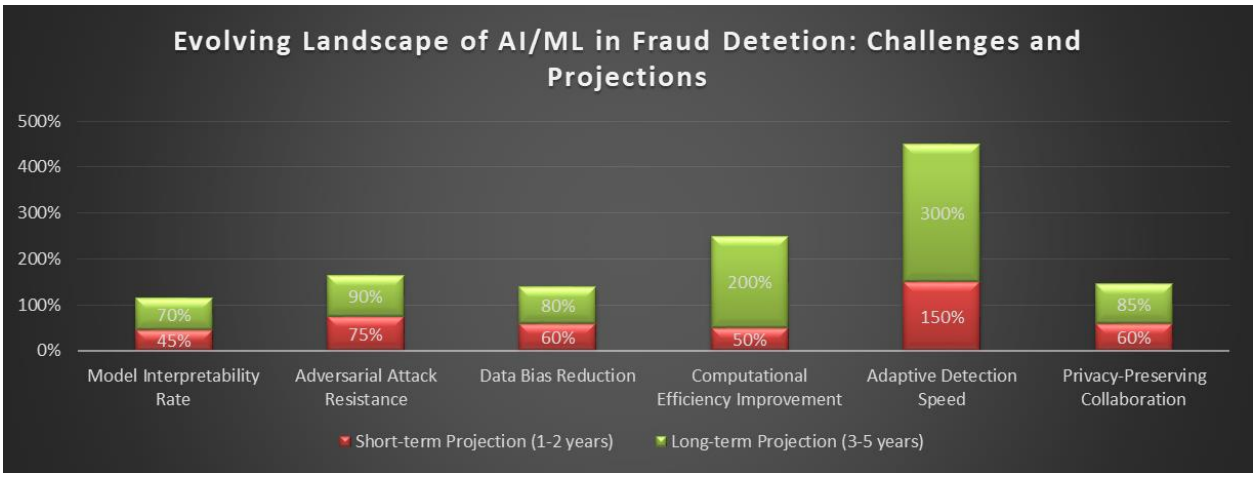

Fig 2: Quantifying Progress: The Future of Intelligent Fraud Prevention Systems [7, 8]

Conclusion

The integration of AI and ML technologies has fundamentally transformed fraud detection and prevention in digital payments, delivering unprecedented levels of accuracy, efficiency, and adaptability in safeguarding financial transactions. Through the implementation of advanced algorithms achieving 99.9% fraud detection rates and real-time processing capabilities handling 100,000 transactions per second with sub-50-millisecond latency, organizations can now effectively combat evolving fraud threats while maintaining seamless user experiences. The combination of supervised learning algorithms, unsupervised anomaly detection, and deep learning models, coupled with stream processing architectures and edge computing, has demonstrated remarkable success in reducing fraud losses by up to 50% within the first year of deployment. As these technologies continue to evolve, incorporating quantum computing capabilities, enhanced federated learning approaches, and more sophisticated explainable AI mechanisms, we can expect even more effective fraud prevention systems that will be crucial in protecting the projected $8.26 trillion digital payment ecosystem by 2024, further securing the future of digital transactions in an increasingly complex and challenging cyber threat landscape.

References

1. Boston Consulting Group, "Global Payments 2021: All In for Growth," Oct. 2021. [Online]. Available: https://www.bcg.com/publications/2021/global-payments-industry-is-all-in-for-growth
2. Nilson Report, "Card Fraud Losses Reach $27.85 Billion," Nov. 2024. [Online]. Available: https://nilsonreport.com/articles/card-fraud-losses-reach-27-85-billion/
3. J. Jurgovsky et al., "Sequence classification for credit-card fraud detection," Expert Systems with Applications, vol. 100, pp. 234-245, 2018, doi: 10.1016/j.eswa.2018.01.037. https://www.sciencedirect.com/science/article/abs/pii/S0957417418300435
4. F. T. Liu, K. M. Ting, and Z. H. Zhou, "Isolation-based anomaly detection," ACM Transactions on Knowledge Discovery from Data (TKDD), vol. 6, no. 1, pp. 1-39, 2012, doi: 10.1145/2133360.2133363. https://dl.acm.org/doi/10.1145/2133360.2133363
5. McKinsey & Company, "The 2021 McKinsey Global Payments Report," Oct. 2021. [Online]. Availa-

ble: https://www.mckinsey.com/industries/financial-services/our-insights/the-2021-mckinsey-global-payments-report

6. S. Patil, V. Nemade, and P. K. Soni, "Predictive Modelling for Credit Card Fraud Detection Using Data Analytics," Procedia Computer Science, vol. 132, pp. 385-395, 2018, doi: 10.1016/j.procs.2018.05.199. https://www.sciencedirect.com/science/article/pii/S1877050918309347
7. A. Adadi and M. Berrada, "Peeking Inside the Black-Box: A Survey on Explainable Artificial Intelligence (XAI)," IEEE Access, vol. 6, pp. 52138-52160, 2018, doi: 10.1109/ACCESS.2018.2870052. https://www.semanticscholar.org/paper/Peeking-Inside-the-Black-Box%3A-A-Survey-on-(XAI)-Adadi-Berrada/21dff47a4142445f83016da0819ffe6dd2947f66
8. N. Papernot, P. McDaniel, A. Sinha, and M. Wellman, "Towards the Science of Security and Privacy in Machine Learning," IEEE European Symposium on Security and Privacy (EuroS&P), 2018, pp. 399-414, doi: 10.1109/EuroSP.2018.00035. https://www.semanticscholar.org/paper/Towards-the-Science-of-Security-and-Privacy-in-Papernot-Mcdaniel/ebab687cd1be7d25392c11f89fce6a63bef7219d