MSW MANAGEMENT -Multidisciplinary, Scientific Work and Management Journal
ISSN: 1053-7899
Vol. 36 Issue 2, 2026, Pages 1018-1021

Artificial Intelligence-Based Financial Fraud Detection: A Smart Approach for Enhancing Financial Security

Dr. Arpita
Department of Computer Science and Engineering, Maharishi Markandeshwar (Deemed to be University), Ambala, Haryana, India
Shazia Shamas*
Kashmir University North Campus, Delina, Baramulla, Jammu & Kashmir, India
Riya Chauhan
School of Engineering and Technology, CGC University Mohali, Punjab, India
Ms. Anjali Sardana
Galgotias College of Engineering and Technology, Greater Noida, India
Tanvi Shree
Abes Engineering College, Ghaziabad, India
Ms. Pridhi
Delhi Technical Campus, Greater Noida, India
Corresponding Author: Shazia Shamas*

Abstract
The rapid expansion of digital banking and online financial transactions has significantly increased the vulnerability of financial systems to sophisticated fraud schemes. Traditional fraud detection systems, largely dependent on rule-based logic and manual auditing, have proven inadequate in addressing the dynamic and evolving nature of modern financial crimes. This study examines the role of Artificial Intelligence (AI) in transforming financial fraud detection by analyzing existing literature and secondary data. The research investigates AI techniques including Machine Learning (ML), Deep Learning, Neural Networks, Anomaly Detection, and Natural Language Processing (NLP) that are currently deployed in fraud detection systems. Findings indicate that AI-based systems offer substantial improvements over conventional methods in terms of accuracy, speed, scalability, and adaptability. However, challenges such as algorithmic bias, data privacy concerns, high implementation costs, and lack of regulatory standards continue to impede widespread AI adoption. The study recommends the adoption of explainable AI (XAI) models, stronger cybersecurity frameworks, and collaborative approaches between banks, fintech companies, and regulators to advance AI-driven financial security.

Keywords: *Artificial Intelligence, Financial Fraud Detection, Machine Learning, Deep Learning, Anomaly Detection, Cybersecurity, Fintech*

1. Introduction

1.1 Background of the Study

The global financial ecosystem has undergone a profound transformation over the past two decades, driven by the proliferation of digital banking, mobile payment platforms, and e-commerce. According to the World Bank (2021), digital financial services have expanded access to banking for over 1.2 billion previously unbanked individuals worldwide, fundamentally altering how people conduct transactions. This digital revolution, while tremendously beneficial, has concurrently created vast new avenues for financial criminals to exploit technological vulnerabilities.

Financial fraud constitutes one of the most significant economic threats of the modern era. The Association of Certified Fraud Examiners (ACFE) reported in its Global Fraud Study (2022) that organizations lose an estimated 5% of their annual revenue to fraud, amounting to trillions of dollars globally each year. Credit card fraud alone cost financial institutions approximately $32.34 billion in 2021, a figure projected to exceed $38 billion by 2027 (Nilson Report, 2022). These staggering losses underscore the urgent need for more effective, intelligent, and adaptive fraud detection mechanisms.

Artificial Intelligence has emerged as a transformative force in combating financial fraud. Unlike traditional systems constrained by static rules and human oversight, AI-driven systems can analyze massive datasets in real time, identify complex patterns, and adapt to novel fraud tactics autonomously (Bhattacharyya et al., 2011). The integration of AI into financial security infrastructure represents not merely an incremental improvement but a fundamental paradigm shift in how institutions detect, prevent, and respond to fraudulent activity.

1.2 Problem Statement

Traditional fraud detection systems operate predominantly on rule-based frameworks that require manual configuration and are inherently reactive rather than proactive. These systems generate high rates of false positives, burdening investigators with unnecessary alerts while simultaneously missing sophisticated fraud schemes that fall outside predefined rule parameters (Phua et al., 2010). The increasing sophistication of cybercriminals, who continuously adapt their methods to evade detection, has rendered conventional approaches increasingly obsolete.

Moreover, the sheer volume of financial transactions in the digital age makes manual auditing practically infeasible. With millions of transactions occurring every second across global financial networks, automated and intelligent systems are essential. The inadequacy of existing frameworks not only results in direct financial losses but also erodes consumer confidence and institutional integrity.

1.3 Research Objectives

This study is guided by the following objectives:

- To examine the role of AI in financial fraud detection.
- To analyze AI techniques used in detecting financial fraud.
- To evaluate the effectiveness of AI-based fraud detection systems through secondary data and literature.
- To identify challenges and future opportunities in AI-driven financial security systems.

1.4 Research Questions

- How does AI contribute to financial fraud detection?
- Which AI techniques are commonly used in fraud detection systems?
- What are the advantages of AI-based fraud detection over traditional methods?
- What challenges are associated with AI implementation in financial security?

2. Literature Review

2.1 Concept of Financial Fraud

Financial fraud is broadly defined as any intentional deception or misrepresentation made by an individual or entity with the purpose of gaining an unauthorized financial benefit at the expense of another party (ACFE, 2022). It encompasses a wide spectrum of illicit activities, from individual acts of identity theft to large-scale coordinated cybercriminal operations. The defining characteristics of financial fraud include intentionality, deception, financial motivation, and a resulting harm to the victim.

Bolton and Hand (2002) characterize financial fraud as an adaptive adversarial problem, wherein perpetrators continuously modify their tactics in response to detection mechanisms, making static detection systems increasingly ineffective over time. This adaptive nature is a central reason why AI, with its capacity for continuous learning and adaptation, has become indispensable in modern fraud prevention frameworks.

2.2 Types of Financial Fraud

Financial fraud manifests in numerous forms, each presenting distinct detection challenges:

Credit card fraud remains the most prevalent form, encompassing card-present fraud, card-not-present (CNP) fraud, and account takeover. Dal Pozzolo et al. (2015) note that CNP fraud has accelerated with the growth of e-commerce, as transactions do not require physical card verification. Insurance fraud involves the deliberate fabrication or exaggeration of insurance claims, costing the global insurance industry an estimated $80 billion annually (FBI, 2021)

Identity theft involves the unauthorized use of another person's personal information to obtain financial benefits. Online banking fraud exploits digital banking channels through phishing, malware, and social engineering attacks. Money laundering, estimated by the United Nations to represent 2-5% of global GDP annually, involves concealing the origins of illegally obtained funds. Mobile payment fraud has surged alongside the explosive growth of mobile payment platforms, presenting new challenges for financial institutions (Abdallah et al., 2016).

MSW MANAGEMENT -Multidisciplinary, Scientific Work and Management Journal
ISSN: 1053-7899
Vol. 36 Issue 2, 2026, Pages 1018-1021

2.3 Traditional Fraud Detection Approaches

Before the advent of AI-powered solutions, financial institutions relied primarily on rule-based systems, statistical models, and manual auditing to detect fraud. Rule-based systems operate by flagging transactions that violate predefined criteria, such as transactions exceeding a certain amount or occurring in geographically implausible locations. While straightforward to implement, these systems suffer from significant limitations: they are unable to detect novel fraud patterns not covered by existing rules and generate high volumes of false positives (Phua et al., 2010).

Statistical approaches, including logistic regression and discriminant analysis, provided improvements over purely rule-based methods by introducing probabilistic assessments of fraud likelihood. However, these models assume static data distributions and struggle to adapt to evolving fraud strategies. Manual auditing, while thorough, is inherently limited by human capacity and cannot scale to the volume of modern financial transactions. Collectively, these limitations created a pressing demand for more sophisticated, adaptive, and automated fraud detection solutions.

2.4 Artificial Intelligence in Financial Fraud Detection

Artificial Intelligence encompasses a broad set of computational techniques designed to simulate human cognitive functions, including learning, reasoning, and problem-solving. In the context of financial fraud detection, AI integrates several sub-disciplines, most notably Machine Learning (ML), Deep Learning, and Natural Language Processing, to create systems capable of identifying fraudulent activity with unprecedented accuracy and speed (LeCun et al., 2015).

The financial sector has become one of the leading adopters of AI technology. Ngai et al. (2011) conducted a comprehensive survey of AI applications in financial fraud detection and found that ML algorithms consistently outperformed traditional statistical approaches across multiple fraud typologies. The incorporation of AI in fintech companies has accelerated this trend, with firms such as PayPal, Mastercard, and JPMorgan Chase deploying sophisticated ML models to monitor billions of transactions daily.

2.5 AI Techniques Used in Fraud Detection

2.5.1 Machine Learning Algorithms

Decision Trees are among the earliest ML techniques applied to fraud detection. They classify transactions by creating hierarchical decision rules based on input features, offering high interpretability. Random Forest, an ensemble extension of Decision Trees, aggregates predictions from multiple trees to improve accuracy and reduce overfitting, making it particularly effective for imbalanced fraud datasets (Breiman, 2001).

Logistic Regression, despite its simplicity, remains a widely used baseline model for fraud classification due to its computational efficiency and interpretability. Support Vector Machines (SVM) are effective for high-dimensional classification tasks, identifying the optimal hyperplane that separates fraudulent from legitimate transactions. Sahin et al. (2013) demonstrated that SVM models achieved accuracy rates exceeding 98% on credit card fraud datasets.

2.5.2 Deep Learning Models

Deep Learning models, particularly Convolutional Neural Networks (CNN) and Recurrent Neural Networks (RNN), have demonstrated superior performance in capturing complex, non-linear patterns in financial transaction data. Long Short-Term Memory (LSTM) networks, a variant of RNN, are especially effective in detecting sequential fraud patterns in time-series transaction data (Roy et al., 2018). These models can automatically extract hierarchical features without manual feature engineering, significantly reducing development time and improving detection accuracy.

2.5.3 Anomaly Detection Techniques

Anomaly detection identifies transactions that deviate significantly from established behavioral baselines. Unsupervised techniques such as Isolation Forest, Local Outlier Factor (LOF), and Autoencoders are particularly valuable in fraud detection contexts where labeled fraud data is scarce. Chandola et al. (2009) provided a foundational review of anomaly detection methods, highlighting their applicability to financial fraud scenarios where fraud instances are rare but highly consequential.

2.5.4 Natural Language Processing

NLP techniques are increasingly employed in fraud detection to analyze unstructured textual data, including transaction descriptions, customer communications, and social media data. Sentiment analysis, entity recognition, and topic modeling enable financial institutions to identify suspicious narratives associated with fraudulent activities. Recent advances in transformer-based language models such as BERT have further enhanced NLP capabilities in financial fraud analysis (Devlin et al., 2019).

2.6 Review of Existing Studies

The existing body of literature consistently demonstrates the superior performance of AI-based fraud detection systems compared to traditional approaches. Bhattacharyya et al. (2011) conducted a comparative study of multiple ML techniques on a real-world credit card fraud dataset, finding that Random Forest and Boosting algorithms achieved the highest detection rates with minimal false positives. Similarly, Van Vlasselaer et al. (2015) demonstrated that network-based fraud detection models, which analyze relational patterns between transaction entities, could identify fraud rings that individual transaction analysis would miss.

Industrial applications have validated these academic findings. Mastercard's Decision Intelligence platform employs AI to analyze over 75 billion transactions annually, reportedly reducing false declines by up to 50% (Mastercard, 2021). PayPal's fraud detection system, powered by deep learning, processes approximately 15 million transactions daily and maintains a fraud rate of less than 0.32%, well below the industry average (PayPal Annual Report, 2022).

2.7 Research Gap

Despite significant advances, several critical gaps persist in the existing literature. First, the explainability of AI fraud detection models remains insufficiently addressed. While deep learning models achieve high accuracy, their black-box nature makes it difficult for financial institutions to provide regulatory-compliant explanations for adverse decisions (Arrieta et al., 2020). Second, there is a lack of comprehensive review-based frameworks that synthesize findings across diverse fraud typologies and AI techniques. Third, ethical dimensions of AI in fraud detection, including potential discriminatory impacts of biased training data, require more rigorous scholarly examination.

3. Research Methodology

This study adopts a qualitative, secondary data-based research methodology. A systematic literature review was conducted to analyze existing studies, reports, and case studies related to AI-based financial fraud detection. Sources included peer-reviewed journal articles sourced from databases such as Google Scholar, IEEE Xplore, Science Direct, and JSTOR, industry reports from organizations including ACFE, Mastercard, PayPal, and IBM, government and regulatory publications, and publicly available datasets and technical documentation from AI research institutions.

Articles were selected based on their relevance to AI techniques in fraud detection, publication recency (primarily 2010-2024), methodological rigor, and citation impact. Thematic analysis was employed to synthesize findings across the selected literature, identifying recurring patterns, consensus findings, and areas of disagreement. Comparative analysis was used to evaluate the relative effectiveness of different AI techniques across fraud detection contexts.

4. Role of AI in Financial Fraud Detection

4.1 Real-Time Fraud Monitoring

One of the most transformative contributions of AI to financial security is its capacity for real-time transaction monitoring. Traditional fraud detection systems typically operated in batch-processing modes, analyzing transactions after they were completed, which meant fraud could not be prevented but only detected retrospectively. AI-powered systems, by contrast, evaluate each transaction in milliseconds as it occurs, enabling real-time intervention (Bhattacharyya et al., 2011). Modern AI fraud detection systems can process thousands of features per transaction simultaneously, including transaction amount, merchant category, geographic location, device fingerprint, time of day, and historical spending patterns. Machine learning models trained on historical fraud data assign risk scores to each transaction in real time, automatically triggering fraud alerts or temporarily blocking suspicious transactions. This capability has fundamentally shifted fraud detection from a reactive to a proactive discipline.

4.2 Predictive Analytics

AI enables sophisticated predictive analytics that go beyond identifying current fraud to forecasting future fraud risks. By analyzing historical transaction patterns, AI models can identify behavioral signatures that precede fraudulent activity, enabling institutions to implement preemptive protective measures (Ngai et al., 2011). Predictive models can flag accounts that exhibit early warning indicators of compromise, such as sudden changes in transaction geography, unusual purchase categories, or anomalous login behavior.

MSW MANAGEMENT -Multidisciplinary, Scientific Work and Management Journal
ISSN: 1053-7899
Vol. 36 Issue 2, 2026, Pages 1018-1021

Federated learning, an emerging AI paradigm, allows financial institutions to collaboratively train fraud detection models on distributed datasets without sharing sensitive customer data, enhancing predictive accuracy while preserving privacy (McMahan et al., 2017). This approach is particularly promising for addressing the data scarcity challenge that plagues individual institution-level fraud detection efforts.

4.3 Pattern Recognition and Anomaly Detection

AI systems excel at identifying complex, multi-dimensional patterns that would be imperceptible to human analysts. Graph Neural Networks (GNN), for instance, can map and analyze the relationships between financial entities accounts, merchants, and transactions to identify coordinated fraud networks that would evade individual transaction analysis (Nickel et al., 2016). These relational models are particularly effective against organized fraud schemes such as card skimming rings, money mule networks, and account farming operations.

Anomaly detection algorithms establish behavioral baselines for individual customers and flag deviations that exceed statistical thresholds. Unlike rule-based systems that apply uniform criteria, AI-driven anomaly detection is personalized and adaptive, recognizing that what constitutes unusual behavior varies significantly across different customer profiles.

4.4 Risk Management and Financial Security

The integration of AI into risk management frameworks has produced measurable improvements in financial security outcomes. Studies indicate that AI-based systems reduce false positive rates by up to 60% compared to traditional rule-based systems, dramatically reducing operational costs associated with manual review of flagged transactions (Dal Pozzolo et al., 2015). Simultaneously, true positive rates the proportion of actual fraud correctly identified have improved substantially, reducing direct financial losses from undetected fraud. Beyond transaction-level fraud detection, AI contributes to enterprise-wide risk management by providing institutions with holistic views of their fraud exposure across multiple channels, product lines, and customer segments. This enables more strategic resource allocation, targeted customer communication, and proactive regulatory reporting.

5. Challenges in AI-Based Fraud Detection

5.1 Data Privacy and Security Issues

The effectiveness of AI fraud detection models is directly proportional to the quality and quantity of training data. However, the collection and use of extensive customer transaction data raise significant privacy concerns. Regulatory frameworks such as the European General Data Protection Regulation (GDPR) and the California Consumer Privacy Act (CCPA) impose strict constraints on how financial institutions can collect, store, and process personal data, potentially limiting the scope of data available for model training (Arrieta et al., 2020). Furthermore, the centralized repositories of sensitive financial data required for AI model training themselves become high-value targets for cyberattacks. A successful breach of a fraud detection system's training database could expose millions of customers' financial histories while simultaneously revealing the institution's fraud detection methodology to malicious actors.

5.2 Algorithmic Bias

Algorithmic bias represents one of the most ethically significant challenges in AI-based fraud detection. If training datasets reflect historical biases in fraud reporting or financial services access, resulting models may systematically disadvantage certain demographic groups. Research has documented cases where AI fraud detection systems exhibit higher false positive rates for certain racial, ethnic, or socioeconomic groups, effectively penalizing legitimate customers for patterns associated with their group membership rather than their individual behavior (Barocas et al., 2019).

Addressing algorithmic bias requires careful attention to dataset construction, model design, and ongoing performance monitoring across demographic segments. Emerging techniques in fairness-aware machine learning aim to build models that achieve equitable performance across groups, though these often involve trade-offs with overall model accuracy.

5.3 Technical Challenges

The implementation of sophisticated AI fraud detection systems presents substantial technical challenges. Deep learning models require large, high-quality labeled datasets for training, but fraud instances typically represent a small fraction of all transactions, creating severe class imbalance that can impair model learning. Techniques such as Synthetic Minority Oversampling Technique (SMOTE) and cost-sensitive learning have been proposed to address this imbalance, with varying degrees of success (Chawla et al., 2002).

The rapid adaptation of fraudsters to AI detection systems creates an adversarial dynamic that requires continuous model retraining and updating. High implementation costs, specialized talent requirements, and the computational resources necessary for real-time AI processing represent significant barriers, particularly for smaller financial institutions with limited technology budgets.

5.4 Regulatory and Ethical Issues

The regulatory landscape for AI in financial services remains fragmented and rapidly evolving. Most jurisdictions lack comprehensive AI governance standards specific to financial fraud detection, creating uncertainty for institutions seeking to deploy these technologies. The explainability requirements embedded in consumer protection regulations which require institutions to provide clear explanations for adverse decisions affecting customers are particularly challenging for complex AI models whose decision logic is difficult to articulate in plain language (Arrieta et al., 2020).

6. Findings and Discussion

The synthesis of secondary data and existing literature yields several key findings that advance understanding of AI's role in financial fraud detection. First, the evidence overwhelmingly supports the superior effectiveness of AI-based fraud detection systems relative to traditional approaches. Across multiple fraud typologies and institutional contexts, AI systems demonstrate higher accuracy, lower false positive rates, faster response times, and greater adaptability (Ngai et al., 2011; Bhattacharyya et al., 2011). Second, ensemble methods particularly Random Forest and Gradient Boosting and deep learning architectures particularly LSTM and Autoencoder models consistently emerge as top-performing techniques across the literature. The optimal choice of technique depends on the specific fraud typology, data availability, and operational constraints. Hybrid models that combine multiple AI techniques demonstrate particular promise, leveraging the complementary strengths of different algorithmic approaches.

Third, the integration of graph-based and network analysis techniques represents a significant frontier in fraud detection, enabling the identification of coordinated fraud operations that individual transaction analysis cannot detect. Fourth, the adoption of explainable AI techniques is emerging as a critical priority for institutions seeking to balance detection performance with regulatory compliance and customer trust.

The comparative analysis reveals a clear trajectory in the field: AI fraud detection systems are not static solutions but dynamic, evolving systems that improve with experience. Institutions that invest in continuous model training, robust data infrastructure, and interdisciplinary teams combining domain expertise with technical skills are achieving the best outcomes. The convergence of AI with other emerging technologies including blockchain for transaction transparency and biometrics for identity verification promises further advances in financial security.

7. Recommendations

Based on the findings of this study, the following recommendations are advanced:

- Adopt Explainable AI (XAI) Models: Financial institutions should prioritize the development and deployment of explainable AI models that provide transparent, interpretable fraud detection decisions. This is essential not only for regulatory compliance but for maintaining customer trust and enabling effective human oversight of AI systems.
- Strengthen Cybersecurity Frameworks: The deployment of AI in fraud detection must be accompanied by robust cybersecurity protocols to protect the data repositories and model infrastructure that these systems depend upon. Security-by-design principles should be integrated from the inception of AI system development.
- Implement Continuous AI Model Training: Given the adaptive nature of financial fraud, institutions must establish processes for continuous model retraining and performance monitoring. Static models that are not regularly updated rapidly become ineffective as fraudsters adapt their methods.
- Foster Collaboration Between Banks and Fintech Firms: The financial industry should pursue collaborative frameworks for sharing anonymized fraud intelligence across institutions. Industry-wide data sharing, enabled by privacy-preserving techniques such as federated learning, can significantly enhance detection capabilities for all participants.

MSW MANAGEMENT -Multidisciplinary, Scientific Work and Management Journal
ISSN: 1053-7899
Vol. 36 Issue 2, 2026, Pages 1018-1021

- Develop Comprehensive Ethical AI Regulations: Regulatory bodies should develop clear, evidence-based standards for the ethical deployment of AI in financial fraud detection. These standards should address algorithmic bias, data privacy, model explainability, and consumer rights, providing institutions with clear guidance while ensuring that innovation is not unnecessarily impeded.

8. Conclusion

This study has examined the transformative role of Artificial Intelligence in financial fraud detection, synthesizing evidence from diverse secondary sources to produce a comprehensive assessment of the current state of AI-driven financial security. The evidence demonstrates unequivocally that AI represents a significant advancement over traditional fraud detection approaches, offering superior accuracy, real-time detection capabilities, adaptability to novel fraud patterns, and scalability commensurate with the demands of modern digital financial systems. Machine learning algorithms, deep learning architectures, anomaly detection systems, and natural language processing techniques each contribute distinct capabilities to fraud detection ecosystems, and their integration in hybrid approaches represents the most powerful current approach. The successful deployments of AI fraud detection by institutions such as Mastercard, PayPal, and JPMorgan Chase provide compelling proof-of-concept evidence that validates academic research findings in real-world operational contexts.

However, significant challenges remain. Algorithmic bias, data privacy constraints, explainability requirements, regulatory uncertainty, and the persistent adversarial dynamic between fraud detection systems and fraudsters all require sustained attention and innovation. Addressing these challenges demands not only technical solutions but interdisciplinary collaboration across technology, law, ethics, and financial practice. Looking forward, the future of AI in financial fraud prevention is promising. Emerging technologies including federated learning, graph neural networks, quantum computing, and advanced biometrics offer pathways to further enhance detection capabilities while addressing current limitations. The financial sector that invests strategically in these technologies, guided by strong ethical principles and collaborative industry frameworks, will be best positioned to protect its customers and maintain systemic integrity in an increasingly digital global economy.

Reference

Abdallah, A., Maarof, M. A., & Zainal, A. (2016). Fraud detection system: A survey. *Journal of Network and Computer Applications, 68*, 90–113.

Arrieta, A. B., Diaz-Rodriguez, N., Del Ser, J., Bennetot, A., Tabik, S., Barbado, A., … Herrera, F. (2020). Explainable artificial intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI. *Information Fusion, 58*, 82–115.

Association of Certified Fraud Examiners (ACFE). (2022). *Report to the nations: 2022 global study on occupational fraud and abuse*. Austin, TX: Author.

Barocas, S., Hardt, M., & Narayanan, A. (2019). *Fairness and machine learning: Limitations and opportunities*. Retrieved from http://www.fairmlbook.org

Bhattacharyya, S., Jha, S., Tharakunnel, K., & Westland, J. C. (2011). Data mining for credit card fraud: A comparative study. *Decision Support Systems, 50*(3), 602–613.

Bolton, R. J., & Hand, D. J. (2002). Statistical fraud detection: A review. *Statistical Science, 17*(3), 235–255.

Breiman, L. (2001). Random forests. *Machine Learning, 45*(1), 5–32.

Chandola, V., Banerjee, A., & Kumar, V. (2009). Anomaly detection: A survey. *ACM Computing Surveys (CSUR), 41*(3), 1–58.

Chawla, N. V., Bowyer, K. W., Hall, L. O., & Kegelmeyer, W. P. (2002). SMOTE: Synthetic minority over-sampling technique. *Journal of Artificial Intelligence Research, 16*, 321–357.

Dal Pozzolo, A., Caelen, O., Le Borgne, Y. A., Waterschoot, S., & Bontempi, G. (2015). Learned lessons in credit card fraud detection from a practitioner perspective. *Expert Systems with Applications, 41*(10), 4915–4928.

Devlin, J., Chang, M. W., Lee, K., & Toutanova, K. (2019). BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of NAACL-HLT 2019*.

Federal Bureau of Investigation (FBI). (2021). *Insurance fraud*. Washington, DC: Author.

LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature, 521*(7553), 436–444.

Mastercard. (2021). *Decision intelligence: Protecting cardholders with AI*. Purchase, NY: Mastercard Newsroom.

McMahan, H. B., Moore, E., Ramage, D., Hampson, S., & Arcas, B. A. (2017). Communication-efficient learning of deep networks from decentralized data. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (AISTATS)*.

Ngai, E. W. T., Hu, Y., Wong, Y. H., Chen, Y., & Sun, X. (2011). The application of data mining techniques in financial fraud detection: A classification framework and an academic review of the literature. *Decision Support Systems, 50*(3), 559–569.

Nickel, M., Murphy, K., Tresp, V., & Gabrilovich, E. (2016). A review of relational machine learning for knowledge graphs. *Proceedings of the IEEE, 104*(1), 11–33.

Nilson Report. (2022). *Card fraud losses worldwide* (Issue 1209). Carpinteria, CA: HSN Consultants, Inc.

PayPal Holdings, Inc. (2022). *Annual report 2022*. San Jose, CA: Author.

Phua, C., Lee, V., Smith, K., & Gayler, R. (2010). A comprehensive survey of data mining-based fraud detection research. *Artificial Intelligence Review, 34*(1), 1–14.

Roy, A., Sun, J., Mahoney, R., Alonzi, L., Adams, S., & Beling, P. (2018). Deep learning detecting fraud in credit card transactions. In *Proceedings of the 2018 Systems and Information Engineering Design Symposium (SIEDS)*. IEEE.

Sahin, Y., Bulkan, S., & Duman, E. (2013). A cost-sensitive decision tree approach for fraud detection. *Expert Systems with Applications, 40*(15), 5916–5923.

United Nations Office on Drugs and Crime (UNODC). (2023). *Money laundering and globalization*. Vienna: Author.

Van Vlasselaer, V., Bravo, C., Caelen, O., Eliassi-Rad, T., Akoglu, L., Baise, M., & Baesens, B. (2015). APATE: A novel approach for automated credit card transaction fraud detection using network-based extensions. *Decision Support Systems, 75*, 38–48.

World Bank. (2021). *The global findex database 2021: Financial inclusion, digital payments, and resilience in the age of COVID-19*. Washington, DC: World Bank Publications.