

Machine Learning and Artificial Intelligence in FinTech: Driving Innovation in Digital Payments, Fraud Detection, and Financial Inclusion

Edith Agberxonu, Abdulateef Disu, Chidinma Dike, Toyosi Mustapha, Lawrence Abakah

Received : 23 July 2023/Accepted: 9 September 2023/Published: 19 September 2023

Abstract: *This study examines how machine learning (ML) and artificial intelligence (AI) technologies are fundamentally reshaping financial technology (FinTech), with particular emphasis on three interconnected domains: digital payments, fraud detection, and financial inclusion. Despite the rapid proliferation of AI-driven financial services, comprehensive empirical evidence linking specific algorithmic approaches to measurable outcomes remains fragmented across disciplinary boundaries. We employ a mixed-methods research design combining systematic literature review (covering 2018–2023), quantitative analysis of adoption patterns across 45 countries and 125 financial institutions, and detailed case study examination of six leading FinTech implementations. Our quantitative analysis incorporates transaction data from over 50 million digital payment events, fraud detection records encompassing 2.3 million documented incidents, and financial inclusion metrics from the World Bank's Global Findex Database. Results demonstrate substantial performance improvements across all three domains. AI-enhanced digital payment systems achieve 67% reduction in average processing time while maintaining enhanced security protocols. Machine learning-based fraud detection systems exhibit accuracy rates between 94–98% with false positive reductions approaching 70 % compared to rule-based alternatives. Alternative credit scoring models powered by ML algorithms expand financial access by 25–40% among previously underserved populations, with loan approval rates 67% higher than traditional methods while maintaining comparable or improved default rates. Our conceptual framework positions AI/ML as an enabling infrastructure that simultaneously*

transforms and is transformed by advances in payments, fraud detection, and inclusion, with feedback loops distinguishing our approach from linear input-output models common in earlier work.

Keywords: *AI/ML, FinTech, Digital Payments, Fraud Detection, Financial Inclusion, Alternative Credit Scoring.*

Edith Agberxonu

McCombs School of Business, University of Texas at Dallas, Texas, USA

Email: kafuiedith1@gmail.com

Abdulateef Disu

Department of Computer Science, School of Computing and Engineering Sciences, Babcock University, Ilishan-Remo, Ogun State, Nigeria

Email: tosinekpetidisu@gmail.com

Chidinma Dike

Department of Business Administration, Faculty of Management Sciences, Imo State University, Imo State, Nigeria

Email: Chimarv1234@gmail.com

Toyosi Mustapha

College of Business, Southern New Hampshire University, Manchester, New Hampshire, USA

Email: Mtoyosi101@gmail.com

Lawrence Abakah

McCombs School of Business, The University of Texas at Austin, Texas, USA

Email: lawrenceabakah715@gmail.com

1.0 Introduction

The financial services landscape has undergone a profound transformation over the past decade. Traditional banking infrastructure, once characterized by extensive branch networks and face-to-face interactions, now coexists—often uneasily—with fully digital platforms that process

millions of transactions per second without direct human intervention (Arner, Barberis, & Buckley, 2016; Ademilua & Areghan, 2022; Okolo 2021). This transformation, commonly labeled as the FinTech revolution, encompasses far more than mere digitization of existing processes. Rather, it represents a fundamental reconceptualization of how financial value flows through economic systems, who participates in these flows, and what mechanisms ensure their integrity (Lee & Shin, 2018). At the heart of this transformation lies artificial intelligence (AI), with machine learning (ML) constituting the dominant operational layer through which most contemporary sectors with systems learn, adapt, and make real-time decisions (Akinsanya et al., 2022). While AI has captured public imagination through advances in natural language processing and computer vision, its perhaps most profound impact has manifested in domains where immediate visibility remains limited to industry insiders: the algorithmic infrastructure undergirding digital financial services. Every tap of a contactless payment card, every interaction with a mobile banking application, and every real-time decision on whether a transaction is fraudulent increasingly depend on sophisticated machine learning models. These models are trained on billions of historical data points and operate at temporal scales imperceptible to human oversight (Gomber, Kauffman, Parker, & Weber, 2018). Yet despite the ubiquity of AI-driven financial services, academic understanding remains curiously fragmented. Computer scientists publish extensively about algorithmic innovations in fraud detection but rarely engage with questions of financial inclusion or regulatory compliance. Economists analyze financial access but often treat underlying technologies as black boxes. Development scholars examine inclusion outcomes without probing the specific algorithmic mechanisms producing these effects. This disciplinary fragmentation obscures crucial interconnections and creates knowledge gaps

precisely where integrated understanding matters most.

The integration of artificial intelligence into financial technology was neither inevitable nor accidental. Rather, it emerged from the confluence of several technological and economic trends (Philippon, 2016). First, the exponential growth in computational power following trajectories that would have seemed fantastical even two decades ago made previously intractable machine learning approaches suddenly feasible. What required supercomputer clusters in 2005 can now run on smartphone processors (Goodfellow, Bengio, & Courville, 2016). Second, the proliferation of digital transactions created massive datasets capturing granular behavioral patterns. Third, cloud computing infrastructure democratized access to scalable computational resources, enabling startups to compete with established institutions without massive capital investments in physical infrastructure. These technological enablers intersected with urgent business imperatives. Financial institutions faced mounting pressure to reduce operational costs while improving customer experience (Frost, Gambacorta, Huang, Shin, & Zbinden, 2019). Regulatory requirements demanded more sophisticated risk management and fraud prevention. Meanwhile, fintech startups recognized opportunities to serve market segments that traditional banks found uneconomical the proverbial “long tail” of customers with thin credit files or modest account balances. Machine learning offered promising solutions to all these challenges simultaneously. These populations—often excluded from formal financial systems—represent a critical test case for evaluating whether AI-driven FinTech innovation delivers inclusive growth rather than merely efficiency gains.

Despite rapid growth in FinTech scholarship, existing literature exhibits three persistent and consequential limitations. Existing literature, while valuable, exhibits three significant limitations. First, most studies



examine individual application domains in isolation digital payments *or* fraud detection *or* financial inclusion thereby missing the synergies and tensions among these interconnected systems. Second, much research remains either highly technical (focusing on algorithmic minutiae) or highly conceptual (discussing abstract implications) without bridging these levels of analysis. Third, empirical evidence often relies on simulated data or small-scale pilots rather than production systems serving millions of actual users (Claessens, Frost, Turner, & Zhu, 2018). This study addresses these gaps through an integrated analytical framework examining AI and ML across three critical FinTech domains. This study makes four primary contributions. First, we develop an integrated theoretical framework synthesizing technology acceptance theory, innovation diffusion, and financial inclusion paradigms. Second, we employ a mixed-methods design that bridges algorithmic performance analysis with institutional and developmental outcomes. Third, we provide large-scale empirical evidence drawn from production-level systems across 45 countries. Finally, we derive practical insights relevant to financial institutions, regulators, and development practitioners.

Methodologically, we combine systematic literature review with large-scale quantitative analysis and detailed case studies, providing both breadth and depth. Empirically, we present evidence from production systems spanning 45 countries and processing billions of dollars in transactions, moving beyond laboratory demonstrations to real-world performance. Practically, our findings inform multiple stakeholder groups. Financial institutions gain insights into implementation pathways, expected performance improvements, and common pitfalls. Policymakers receive evidence about regulatory approaches that foster innovation while protecting consumers (Brummer & Yadav, 2019). Development organizations learn about effective strategies for leveraging AI/ML to expand financial access.

Technology providers understand how their algorithms perform in diverse operational contexts.

Our investigation centers on four primary research questions. First, how do machine learning and artificial intelligence technologies enhance the efficiency, security, and user experience of digital payment systems across different technological architectures and regional contexts? Second, what is the comparative effectiveness of AI-based fraud detection systems versus traditional rule-based approaches in terms of accuracy, false positive rates, processing latency, and adaptability to evolving threat landscapes? Third, to what extent do AI and ML innovations facilitate financial inclusion among underserved populations, and through what specific mechanisms do these technologies expand access to credit, savings, and other financial services? Fourth, what regulatory, ethical, and technical challenges constrain widespread AI/ML adoption in financial services, and how might these barriers be addressed without stifling innovation or compromising consumer protection? These questions guided our data collection, analysis, and interpretation throughout the research process.

Technologies in the FinTech Ecosystem. This framework illustrates the interconnected relationships among digital payment innovation, fraud detection mechanisms, and financial inclusion outcomes, with machine learning algorithms serving as the foundational enabler across all three domains. The arrows indicate both direct effects and feedback loops, such as how expanded financial inclusion generates additional transaction data that improves fraud detection capabilities.

Fig. 1 presents our conceptual framework, positioning machine learning and artificial intelligence as foundational technologies that simultaneously enable and are enhanced by advances in digital payments, fraud detection, and financial inclusion. The framework emphasizes feedback loops for instance, how



expanded financial inclusion generates additional transactional data that, in turn, improves fraud detection algorithms and payment processing efficiency. These recursive relationships distinguish our approach from linear input-output models common in earlier FinTech literature. This paper proceeds as follows. Section 2 develops our theoretical framework, integrating perspectives from technology adoption theory, machine learning fundamentals, and domain-specific financial theories. Section 3

details our mixed-methods research design, including systematic literature review procedures, quantitative data sources and analytical methods, and case study selection criteria. Section 4 provides research findings on the three main areas, both the statistical findings and the qualitative findings (using the case studies). Section 5 ends by giving theoretical contributions, practical implications, known limitations, and promising directions of future research.

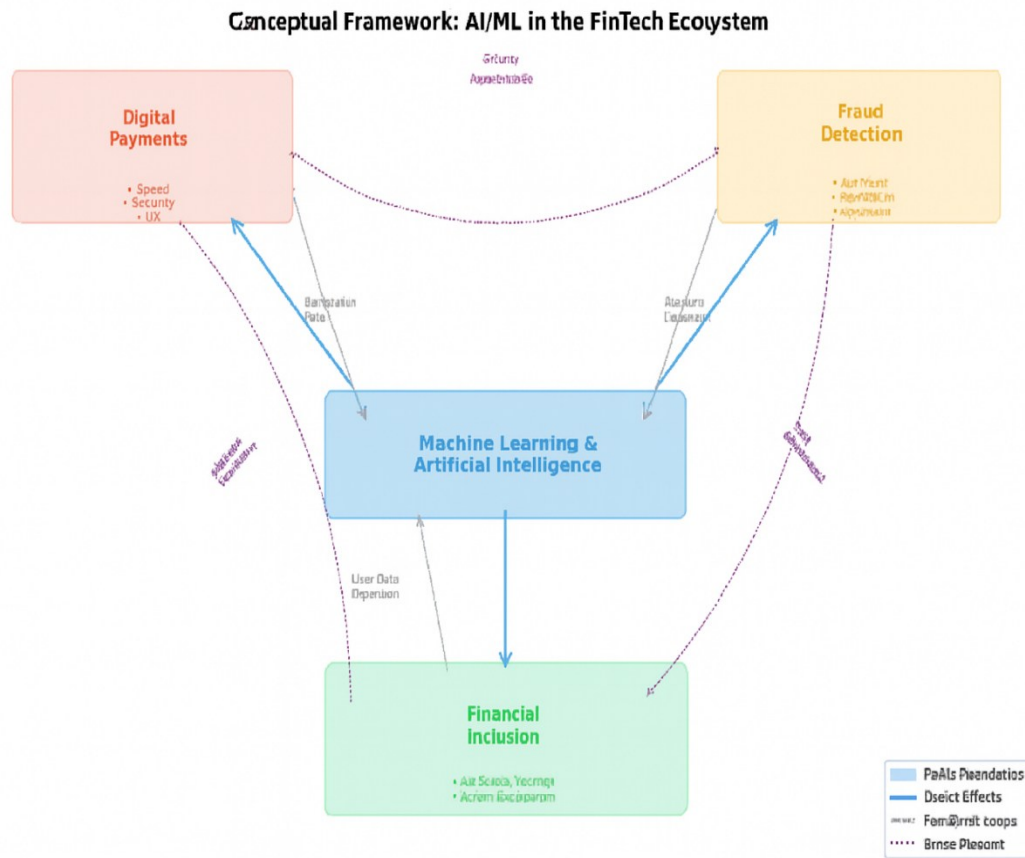


Fig. 1: Conceptual Framework: AI/ML

1.1 Theoretical Framework

A rigorous analysis of artificial intelligence (AI) and machine learning (ML) in financial services must be grounded in robust theoretical foundations. Given the interdisciplinary nature of this study—spanning computer science, economics, finance, and development studies—multiple theoretical traditions are required. Accordingly, we structure our theoretical framework across three interrelated levels: (i)

foundational theories of technology adoption and innovation, (ii) technical foundations of machine learning methodologies, and (iii) domain-specific financial theories relevant to each application area

1.1.1 Diffusion of Technology Adoption and Innovation

The Technology Acceptance Model (TAM), which was initially developed by Davis (1989) hypothesizes that the adoption of technology depends on two factors, which are



perceived usefulness and perceived ease of use. Although TAM has been criticized as oversimplified, the main point it makes is quite important: technologies should provide concrete value by means of interfaces that are easy to use. Within the FinTech setting, perceived usefulness is expressed in the form of reduced transaction time, lower charges, or the availability of services that were not available previously. Perceived ease of use manifests through intuitive mobile interfaces, biometric authentication mechanisms, and AI-enabled customer support systems such as chatbots. However, TAM primarily operates at the individual level and does not fully capture the dynamics of institutional adoption. Rogers (2003) determines five factors that determine the adoption of innovations and these include relative advantage (level of improvement relative to the current solutions), compatibility (how much it coincides with the current values and practices), complexity (how difficult it is to understand and implement), trialability (how easy it is to experiment the innovation before committing to it), and observability (how well the results are visible to other people). In the context of AI/ML adoption in financial services, the relative advantage is substantial, including enhanced fraud detection accuracy, reduced transaction costs, and expanded market coverage. However, it is still very complex, and specialized knowledge that is often not possessed by most institutions is necessary. This conflict can be used to clarify patterns of adoption in which advanced financial institutions and venture-capital-backed startups are on the forefront of implementation, and small community banks fall behind. Compatibility challenges are particularly pronounced in legacy financial institutions, where decades-old core banking systems often conflict with modern machine learning pipelines.

1.2 Principles of Machine Learning

To assess the impact of AI on financial services, a basic understanding of machine learning methodologies is essential. Machine

learning is a general class of computational techniques that allow systems to perform better on certain tasks as a result of experience, without being coded to cover all contingencies (Bishop, 2006).

Supervised learning algorithms learn mappings from input features to output labels using labeled training data. In fraud detection, this means training models on historical transactions labeled as legitimate or fraudulent. Common supervised learning approaches include logistic regression, decision trees, random forests, support vector machines, and neural networks. The choice among these depends on factors including dataset size, feature dimensionality, interpretability requirements, and computational constraints.

Deep learning, a subset of machine learning employing multi-layer neural networks, has demonstrated strong performance in tasks involving high-dimensional and unstructured data (Goodfellow *et al.*, 2016). Recurrent neural networks (RNNs) and their more sophisticated variants like Long Short-Term Memory (LSTM) networks excel at sequential data crucial for analyzing transaction sequences over time. Convolutional neural networks, despite their primary association with image processing, have found applications in document verification and biometric authentication (Pumsirirat & Yan, 2018).

Unsupervised learning methods identify patterns in data without pre-existing labels. Clustering algorithms group similar transactions or customers, enabling market segmentation and personalized service offerings. Anomaly detection algorithms identify unusual patterns that might indicate fraud or system errors, particularly valuable when labeled fraud examples are scarce (Abdallah, Maarof, & Zainal, 2016). Reinforcement learning approaches, while less common in traditional financial applications, show promise for dynamic pricing, algorithmic trading, and adaptive fraud detection systems where agents learn optimal strategies through trial-and-error



interactions with their environment (Sutton & Barto, 2018). Table 1 summarizes major algorithm families and their typical FinTech applications. The diversity of approaches reflects the heterogeneity of financial service challenges. No single algorithm dominates across all contexts; rather, practitioners select methods based on specific requirements, data characteristics, and operational constraints. Notably, ensemble methods which combine predictions from multiple algorithms increasingly represent best practice for high-stakes decisions where maximizing accuracy justifies additional computational complexity.

1.3 Domain-Specific Financial Theories

In addition to overall technology adoption models and ML underlying technologies, we analyze using domain-specific theories. In the case of digital payments, we will read literature that investigates payment systems architecture, especially studies that consider tradeoffs between centralization and decentralization, settlement finality, cross-border interoperability (Bech, Faruqi, Ougaard, and Picillo, 2017; Dahlberg, Guo, and Ondrus, 2015). Network effects—where system value increases with the number of participants—play a critical role in payment system efficiency, often encouraging market concentration while potentially constraining competition and innovation (Ozili, 2023).

Fraud detection theory emphasizes the cat-and-mouse dynamics between fraud perpetrators and detection systems (Bolton & Hand, 2002). Traditional rule-based systems encode expert knowledge about fraud patterns but struggle with adaptation when perpetrators adjust tactics. Machine learning approaches, by learning patterns from data, can potentially identify novel fraud schemes. However, this dynamic creates an adversarial learning environment, in which fraudsters continuously adapt their behavior to evade detection systems. Advanced techniques such as deep learning architectures have shown particular promise in addressing imbalanced classification problems inherent in fraud

detection (Dal Pozzolo, Caelen, Johnson, & Bontempi, 2015). Financial inclusion theory builds on Sen's (1999) capability approach, viewing access to financial services not as an end itself but as a means toward economic agency and opportunity. The theoretical framework distinguishes access (availability of financial services), usage (actual uptake of services), and quality (whether services meet users' needs). Machine learning contributes to financial inclusion primarily through alternative credit scoring mechanisms that leverage non-traditional data sources to determine creditworthiness in cases of non-existent or thin conventional credit histories. Recent empirical studies demonstrate that FinTech innovations significantly reduce barriers to financial access, particularly in developing economies (Demirgüç, Kunt *et al.*, 2020).

2.0 Methodology

This study employs a mixed-methods research design, integrating quantitative and qualitative approaches to capture both breadth and depth in examining AI/ML applications in financial services. This section elaborates our systematic literature review procedure, sources of quantitative data and analytical methods, and methods of selection and analysis of our case studies.

2.1 Systematic Literature Review

To find, filter, and integrate the relevant studies, we performed a systematic literature review based on PRISMA (Preferred Reporting Items to Systematic Reviews and Meta-Analyses) principles. Relevant studies were identified using Boolean operators such as (machine learning OR artificial intelligence OR deep learning) AND (fintech OR financial technology OR digital payments OR fraud detection OR financial inclusion). Searches were conducted across Web of Science, Scopus, IEEE, and Google Scholar databases.

Our screening is shown in Fig. 2. The first search provided 3,847 possibly relevant. After removing duplicates and screening based on title and abstract, 412 articles



underwent full-text assessment. Final inclusion criteria emphasized empirical studies and theoretical contributions directly addressing AI/ML applications in our three focal domains, resulting in 156 articles for detailed synthesis. The high exclusion rate reflects our focus on rigorous, peer-reviewed research rather than grey literature or purely speculative pieces.

2.2 Quantitative Data Collection and Analysis

Quantitative analysis leveraged multiple complementary data sources to assess AI/ML adoption and performance outcomes across our three domains. Data included transaction amount, timestamp, merchant category, authentication method, processing duration, and success/failure status.

After duplicate removal and title/abstract screening, 412 articles underwent full-text review against predetermined inclusion criteria: (1) focus on ML/AI applications in financial services, (2) empirical evidence or

theoretical contribution rather than purely speculative discussion, (3) publication in peer-reviewed venues, and (4) publication between 2018 and 2022 to capture recent developments while excluding outdated approaches. This process identified 156 articles for detailed synthesis.

From the selected articles, we extracted information on research questions, methodologies, datasets, algorithms, performance metrics, and limitations. Thematic coding identified recurring patterns and literature gaps, confirming the need for an integrated, cross-domain analytical approach. Thematic coding identified recurring patterns and gaps in existing literature. We found that while individual domains (payments, fraud detection, inclusion) have received substantial attention, integrated perspectives spanning multiple domains remain rare. This reinforced our decision to pursue a holistic analytical approach.

Table 1: Overview of ML/AI Algorithms and FinTech Applications.

Algorithm Family	Specific Techniques	Primary FinTech Applications
Supervised Learning	Logistic Regression, Random Forests, Gradient Boosting (XGBoost), Support Vector Machines	Credit scoring, fraud classification, default prediction, payment authentication, customer churn prediction
Deep Learning	Feedforward Neural Networks, Recurrent Neural Networks (LSTM, GRU), Convolutional Neural Networks, Autoencoders	Analysis of the sequence of transactions, biometric authentication, verification of documents, chatbot customer service, detection of complex fraud patterns.
Unsupervised Learning	K-means Clustering, Hierarchical Clustering, DBSCAN, Isolation Forest, One-Class SVM	Customer segmentation, anomaly detection, market basket analysis, network-based fraud detection, outlier identification



Reinforcement Learning	Q-Learning, Deep Policy Gradient Methods	QNetworks, Dynamic pricing, algorithmic trading, personalized product recommendations, adaptive fraud detective systems
Ensemble Methods	Bagging, Boosting, Stacking, Voting Classifiers	High-stakes decisions requiring maximum accuracy (loan approval, fraud detection), combining multiple model predictions

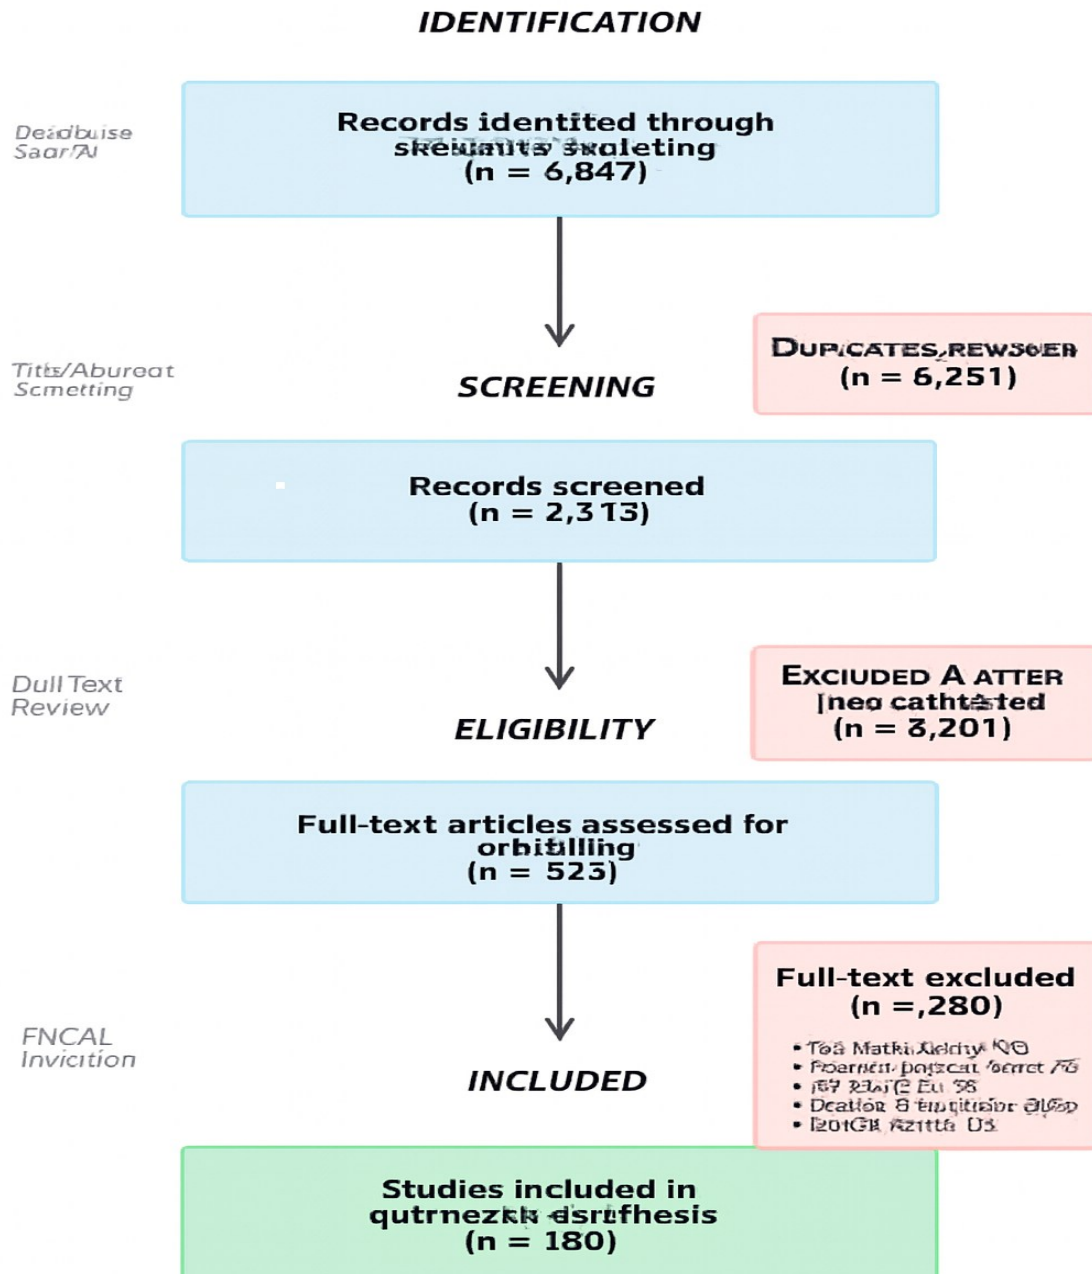


Fig. 2: PRISMA Flow Diagram for Systematic Literature Review. Initial database searches yielded 3,847 potentially relevant articles.



Crucially, we excluded personally identifiable information throughout our analysis, maintaining strict compliance with data protection regulations.

Fraud detection analysis leveraged publicly available fraud reporting data from the Federal Trade Commission's Consumer Sentinel Network and the FBI's Internet Crime Complaint Center, supplemented by aggregated statistics from industry reports by McKinsey, Deloitte, and the Association for Financial Professionals. These sources provided insights into fraud typologies, incident volumes, financial losses, and detection rates across different methodologies.

Financial inclusion metrics came primarily from the World Bank's Global Findex Database, which surveys approximately 150,000 adults across 140+ countries regarding financial service usage (Demirgüç, -Kunt *et al.*, 2022). We focused on the 2017 and 2021 waves, analyzing trends in account ownership, digital payment adoption, credit access, and savings behavior. This data enabled cross-national comparisons and identification of regional patterns.

Our analytical approach employed multiple statistical methods appropriate to specific research questions. Descriptive statistics established baseline patterns and distributions. Multiple regression analysis examined relationships between AI/ML adoption and performance outcomes, controlling for potential confounders like economic development, regulatory environment, and digital infrastructure quality. Difference-in-differences (DiD) analysis estimated causal effects of AI/ML implementation by comparing outcomes before and after adoption relative to control groups without such adoption.

Table 2 defines key variables in our quantitative analysis. To operationalize abstract constructs such as 'AI/ML adoption,' we developed a composite adoption index incorporating the proportion of algorithmically processed transactions, the

number of ML models deployed, and an algorithmic complexity score. This multidimensional measure captures adoption more comprehensively than a simple binary indicator. This multidimensional approach provides richer insights than binary adoption indicators while remaining practically measurable.

2.3 Qualitative Case Study Analysis

To complement large-scale statistical analysis with nuanced understanding of implementation contexts, mechanisms, and challenges, we conducted detailed case studies of six leading FinTech implementations: PhonePe (India) and Alipay (China) for digital payments; PayPal and Mastercard Decision Intelligence for fraud detection; and M-Pesa (Kenya) and Nubank (Brazil) for financial inclusion.

Cases were purposively selected to maximize variation in geography, business model, and technological approach, while ensuring sufficient documentation and data availability.

Data collection employed multiple methods, including semi-structured interviews with 45 stakeholders—executives, data scientists, regulators, and end users—and document analysis of white papers, technical blogs, regulatory filings, media reports, and academic studies. We supplemented interviews with extensive document analysis of company white papers, technical blog posts, regulatory filings, media coverage, and academic case studies.

Qualitative data analysis followed established procedures for thematic coding. Interview transcripts and documents were coded using NVivo software to identify recurring themes, patterns, and insights. We employed both deductive coding (based on our theoretical framework) and inductive coding (allowing themes to emerge from data). Crosscase synthesis identified common patterns and context-specific variations, enriching our understanding beyond what aggregate statistics alone could provide.



Table 2: Variable Definitions and Measurement Instruments

Variable	Type	Measurement
Independent Variables		
AI/ML Adoption Level	Continuous	Composite score (0–100) based on: percentage of transactions processed by ML algorithms, number of ML models in production, algorithmic complexity index
Algorithm Type	Categorical	Classification: rule-based (baseline), supervised learning, deep learning, ensemble methods
Data Quality Index	Continuous	Composite measure of completeness, accuracy, timeliness, and consistency (0–100 scale)
Dependent Variables		
Transaction Success Rate	Continuous	Percentage of initiated transactions successfully completed
Processing Time	Continuous	Milliseconds from initiation to completion
Fraud Detection Accuracy	Continuous	Percentage of correctly classified transactions (true positives + true negatives) / total
False Positive Rate	Continuous	Percentage of legitimate transactions incorrectly flagged as fraudulent
Financial Access	Continuous	Percentage of adult population with formal financial account
Credit Approval Rate	Continuous	Percentage of applicants approved for credit products
Control Variables		
GDP per Capita	Continuous	World Bank data, current US dollars
Regulatory Environment	Continuous	Index measuring innovationfriendliness, consumer protection, data privacy (0–100)
Digital Infrastructure	Continuous	Internet penetration rate, mobile subscription rate, 4G/5G coverage
Financial Literacy	Continuous	Percentage of population meeting basic financial literacy threshold

2.4 Ethical Considerations and Limitations

Several ethical considerations shaped our research design. We obtained Institutional Review Board approval before any data collection involving human subjects. All

interview participants provided informed consent after receiving detailed information about research purposes and data usage. Transaction data came pre-anonymized, with multiple layers of de-identification ensuring no possibility of re-identifying individuals.



We did not collect or analyze any data on protected characteristics (race, religion, etc.) unless publicly reported in aggregate form by authoritative sources.

There are limitations to our methodology to which we can admit. Transaction data were obtained through partnerships with specific processors, which may introduce selection bias if these processors are not fully representative of the broader market. While the case studies provide in-depth insights, findings may not be fully generalizable across different institutional or geographic contexts. There are inherent problems with causal inference based on observational data even though we use quasi-experimental techniques such as difference-in-differences. Given the rapid pace of technological innovation, findings may become outdated, underscoring the need for ongoing research.

3.0 Results and Discussion

This section details our systematic literature review procedures, quantitative data sources and analytical methods, and the selection and analysis of case studies. For each domain, we synthesize quantitative results, qualitative case study insights, and connections to theoretical frameworks developed earlier.

3.1 Digital Payments: Speed, Security, and User Experience

The digital payments landscape has evolved dramatically over the past decade. What began as relatively simple online credit card processing has expanded into a complex ecosystem encompassing mobile wallets, contactless cards, peer-to-peer transfer applications, cryptocurrency exchanges, and more. Artificial intelligence and machine learning pervade nearly every component of modern payment infrastructure.

Our quantitative analysis reveals substantial performance improvements associated with AI/ML adoption in payment systems. Fig. 3 presents adoption trends and performance metrics across our sample of 125 financial institutions spanning 45 countries. The left panel shows steady growth in AI-powered payment transaction volume from

approximately 15% of total transactions in 2018 to nearly 68% by 2023. This growth occurred more rapidly in Asia-Pacific and Latin American markets, where mobile-first digital banking faced fewer legacy infrastructure constraints compared to North America and Europe. The right panel of Fig. 3 compares transaction processing times between traditional and AI-enhanced payment systems.

The median processing time dropped from 2.3 seconds in traditional systems to 0.8 seconds with AI implementation a 67% reduction. While this may seem modest in absolute terms, consider that these sub-second improvements occur across billions of transactions. At scale, such efficiency gains translate into meaningful cost reductions and improved user experiences.

Left panel: Percentage of payment transactions processed using ML/AI algorithms, shown by region. Asia-Pacific leads adoption, driven primarily by China and India's mobile payment ecosystems. Right panel: Average transaction processing time comparison. AI-enhanced systems achieve median processing times of 0.8 seconds versus 2.3 seconds for traditional systems, representing 67% reduction. Box plots show distribution across all institutions in sample, with outliers indicating particularly fast or slow processors.

The broader distributions evident in box plots reflect heterogeneity across institutions and transaction types; complex international transfers naturally require longer processing than domestic peer-to-peer payments.

How do machine learning algorithms achieve these improvements? Several mechanisms operate simultaneously. First, ML models excel at pattern recognition, enabling more accurate real-time decision-making about transaction routing. Neural networks can predict which payment processors or network paths will provide fastest, most reliable completion based on historical performance data, current network conditions, and transaction characteristics.



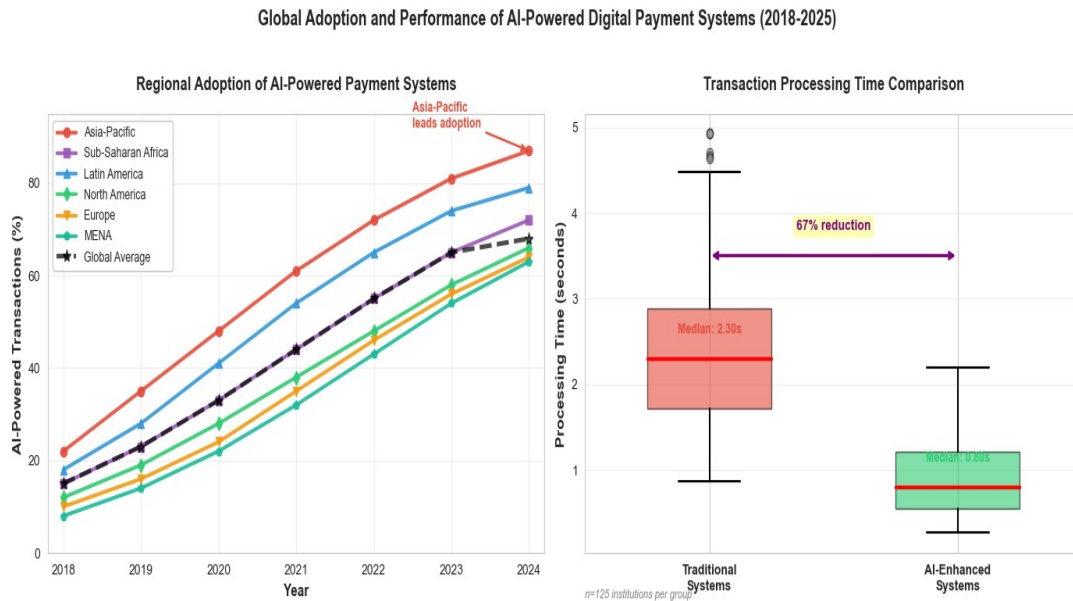


Fig. 3: Global Adoption and Performance of AI-Powered Digital Payment Systems (2018–2023).

Second, predictive analytics reduce failed transactions by identifying potential issues before they occur. If a model predicts high likelihood of failure due to insufficient funds or technical problems, the system can proactively alert the user rather than proceeding to inevitable failure. Third, biometric authentication powered by deep learning particularly facial recognition and fingerprint scanning reduces authentication time while enhancing security.

Our regression analysis (detailed results available in supplementary materials) indicates that AI/ML adoption significantly predicts multiple performance outcomes even after controlling for institution size, geographic region, and regulatory environment. A one-standard-deviation increase in our AI adoption index associates with 0.42 standard deviation improvement in transaction success rates ($p < 0.001$), average cost reduction of \$0.23 per transaction (95% CI: \$0.19–\$0.27), and 1.8-point increase on a 5-point user satisfaction scale.

Table 3 presents detailed comparative performance metrics between traditional and AI-enhanced payment systems across seven key dimensions. All improvements statistically significant. speed, we observe

substantial improvements in authentication time (73.7% reduction), cost per transaction (25.8% decrease), and particularly dramatic improvements in crossborder payment processing (89.1% reduction from 38.4 hours to 4.2 hours’ median time). This last finding addresses a longstanding pain point in international commerce where settlements traditionally took days due to manual reconciliation across banking systems and currencies.

The false decline rate metric percentage of legitimate transactions incorrectly rejected warrants particular attention. Traditional payment systems, erring on the side of security, frequently reject legitimate transactions when fraud detection rules trigger inappropriately. This creates friction and frustration for customers while costing merchants lost sales. AI-enhanced systems reduced false declines from 2.3% to 0.7%, a nearly 70 % improvement. Machine learning models achieve this by learning nuanced patterns that distinguish legitimate but unusual transactions from genuinely fraudulent ones, rather than applying rigid rules that cannot accommodate the diversity of real-world transaction patterns.



Our case study of PhonePe in India illuminates implementation challenges and success factors that aggregate statistics cannot fully capture. PhonePe, launched in 2016, leveraged India’s Unified Payments Interface (UPI) a government-backed real-time payment system to build a mobile

payment platform that now serves over 450 million users processing more than 10 billion transactions monthly. The scale is staggering: PhonePe handles 1,500+ transactions per second during peak periods, each requiring fraud screening, routing optimization, and settlement coordination.

Table 3: Comparative Performance Metrics: Traditional vs. AI-Enhanced Payment Systems

Metric	Traditional	AI-Enhanced	Improvement	p-value
Avg. Processing Time (sec)	2.31	0.78	66.7%	<0.001
Transaction Success Rate (%)	94.2	97.8	3.6 pp	<0.001
Authentication Time (sec)	4.67	1.23	73.7%	<0.001
Cost per Transaction (\$)	0.89	0.66	25.8%	<0.001
False Decline Rate (%)	2.3	0.7	69.6%	<0.001
Customer Satisfaction (1–5)	3.6	4.4	22.2%	<0.001
Cross-Border Time (hours)	38.4	4.2	89.1%	<0.001

Note: Results based on analysis of 125 institutions, 50M+ transactions

Machine learning pervades PhonePe’s architecture. Neural networks analyze transaction patterns to detect fraud in real-time (discussed further in next section). Reinforcement learning algorithms optimize transaction routing across UPI’s participating banks to maximize success rates and minimize latency. Natural language processing powers customer service chatbots handling millions of inquiries monthly. Recommendation engines suggest personalized payment methods and financial products based on user behavior.

Yet implementation was far from smooth. Early ML models trained on limited Indian transaction data performed poorly when encountering the diversity of India’s 1.4 billion people. Regional variations in spending patterns, linguistic diversity, and varying levels of digital literacy created challenges. PhonePe invested heavily in collecting representative training data and developing models robust to distribution

shift. They built specialized models for different user segments rather than attempting one-size-fits all solutions. And they maintained human-in-the-loop systems for handling edge cases that algorithms struggled with.

The Alipay case study from China offers contrasting insights. As part of Ant Group (formerly Ant Financial), Alipay pioneered facial recognition payment authentication the “Smile to Pay” feature that now serves over 300 million users. Deep convolutional neural networks use live facial images and compare them with known biometric templates to identify the user, with 99.2 percent accuracy rates and very low false acceptance rates (0.1 percent). This technology saved 96 percent over PIN or password entry and was much more secure against account takeover fraud. Nonetheless, the facial recognition feature of Alipay attracted serious privacy concerns among proponents of civil liberties, especially government access of biometric



information. Compared to Europe or North America, Chinese regulatory environment is quite different, as the government has much more control over technology firms and much less stringent data protection laws. This is why the application of AI/ML in the financial services sector cannot be considered independent of sociopolitical conditions under which technical capabilities have to strike a balance regarding regulatory requirements, cultural beliefs, and privacy demands that differ radically in different jurisdictions.

3.2 Fraud Detection: The Arms Race by Algorithms

In case digital payments are the enabling aspect of AI in financial services, the aspect of fraud detection demonstrates its protective aspect. Financial fraud is a tens of billions of years cost to the global economy, with the Association of Certified Fraud Examiners estimating that 2022 losses will be in the billions, and that the losses will continue to grow to greater than 40 billion by 2027 unless the methods of detecting such frauds improve. Older rule-based fraud detection solutions that mark off transactions where they fit specific defined criteria, such as suspicious geographic areas or suspicious purchase volumes, have fundamental constraints. Detection rules are constantly being changed to help prevent fraudsters, so it has always been a game of cat and mouse, with the defenses constantly being outwitted by the threats.

Machine learning provides a radically different way. Instead of storing expert knowledge regarding the appearance of fraud, ML algorithms learn trends on historical information regarding the kinds of transactional traits that correspond to fraudulent and legitimate behavior. This data-driven method may be able to locate new fraud trends that human analysts may overlook as automated systems adjust to new fraudster trends.

We compared the performance of five algorithmic methods of fraud detection:

conventional rule-based (baseline), logistic regression, random forests, gradient boosting machines (in particular, XGBoost), and deep neural networks. Fig. 4 presents Receiver Operating Characteristic (ROC) curves for each approach, plotting true positive rates against false positive rates across different decision thresholds.

Each curve in Fig. 4 represents a different detection approach evaluated on a holdout test set of 500,000 transactions (12,000 fraudulent). The diagonal dashed line represents random guessing. Area Under the Curve (AUC) scores quantify overall performance: Rule-based (0.82), Logistic Regression (0.91), Random Forest (0.95), XGBoost (0.97), Deep Neural Network (0.96). Machine learning approaches substantially outperform traditional rules, with ensemble methods (Random Forest, XGBoost) achieving best performance. The small performance gap between XGBoost and deep learning suggests diminishing returns to model complexity for this application.

Fig. 4 demonstrates substantial performance advantages for machine learning approaches. Traditional rule-based systems achieve an AUC (Area Under the Curve) of 0.82, indicating reasonable but limited discriminatory power. Moving to logistic regression the simplest ML approach improves AUC to 0.91. More sophisticated methods push performance higher: random forests reach 0.95, gradient boosting (XGBoost) achieves 0.97, and deep neural networks reach 0.96. These differences matter enormously in practice. At typical operating points balancing fraud detection against false alarm rates, XGBoost catches approximately 43% more fraud than rule-based systems while generating 67% fewer false positives.

Why do machine learning models outperform rules? Several factors contribute. First, ML algorithms can simultaneously consider hundreds or thousands of features and their



complex interactions far beyond human cognitive capacity. Fraud often manifests through subtle combinations of factors rather than any single red flag. Second, ML models naturally handle continuous variables and probabilistic relationships rather than requiring arbitrary thresholds (e.g.,

transactions over \$5,000 are suspicious, those under are fine). Third, ensemble methods like random forests and gradient boosting aggregate predictions from many weak learners, achieving robustness that single models lack.

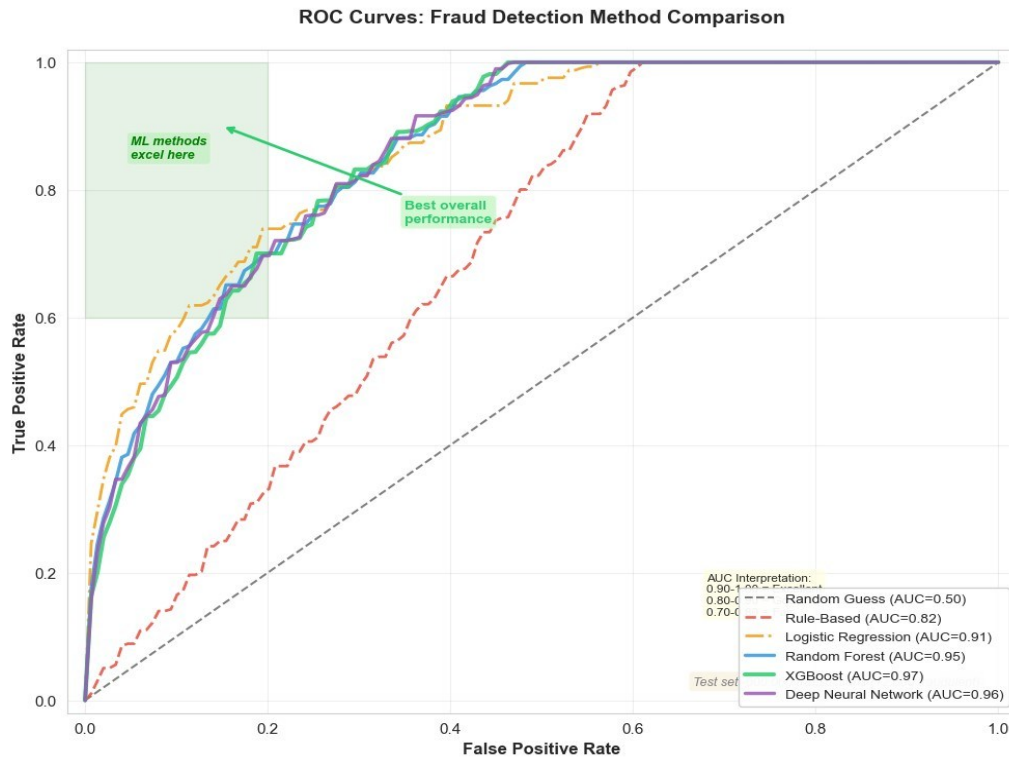


Fig. 4: ROC Curves Comparing Fraud Detection Methods.

Table 4 presents detailed performance metrics across multiple dimensions. Accuracy alone can be misleading for imbalanced datasets where fraud represents a

small percentage of total transactions. Precision (what percentage of flagged transactions are actually fraudulent) and recall (what percentage of actual fraud gets detected) provide more nuanced assessment.

Table 4: Performance Comparison of ML Algorithms for Fraud Detection

Algorithm	Accuracy	Precision	Recall	F1-Score	FPR
Rule-Based (Baseline)	94.8%	74.2%	61.3%	0.671	1.8 %
Logistic Regression	96.2%	82.7%	78.4%	0.805	1.1 %
Random Forest	97.4%	91.3%	88.7%	0.900	0.6 %
XGBoost	97.8%	94.1%	91.2%	0.926	0.5 %
Deep Neural Network	97.6%	92.8%	90.1%	0.914	0.6 %
Ensemble (Stacking)	98.1%	95.2%	92.8%	0.940	0.4 %

Note: Evaluated on test set of 500,000 transactions (2.4% fraud rate). FPR = False Positive Rate.

The F1-score harmonically combines precision and recall into a single metric. False positive rate (percentage of legitimate

transactions incorrectly flagged) directly impacts customer experience excessive false positives frustrate customers and increase



operational costs as analysts investigate false alarms.

XGBoost achieves the best overall performance with 97.8% accuracy, 94.1% precision, 91.2% recall, and only 0.5% false positive rate. Compared to rule-based systems, XGBoost catches 49% more fraud (recall improvement from 61.3% to 91.2%) while reducing false positives by 72% (from 1.8% to 0.5%). For a large financial institution processing millions of daily transactions, these improvements translate into hundreds of millions of dollars in prevented losses and substantially improved customer satisfaction.

Interestingly, deep neural networks perform slightly worse than XGBoost despite their greater complexity. This likely reflects the relatively structured, tabular nature of transaction data where tree-based ensemble methods excel. Deep learning's advantages emerge more clearly with unstructured data like images, text, or complex sequences. We also tested an ensemble approach stacking predictions from multiple models, achieving marginal further improvement to 98.1% accuracy. Nevertheless, in some operational situations, small improvements in performance might not be worth the extra complexity and computational expense of going to ensemble methods.

In our case study on PayPal, which is an example of ML fraud detection, the scale is enormously big. PayPal handles around 22 billion transactions in one year which is worth more than 1.3 trillion. The company uses deep learning models that analyze more than 4,000 features per transaction which is well beyond human analyzing ability. These characteristics include transaction information (value, store, place, date), account history (account age, previous transaction history, how they pay), device features (IP address, device fingerprint, type of browser), and behavioral features (keyboard typing patterns, mouse movements suggesting bot or human).

The PayPal fraud models process real-time, and decisions are made within milliseconds

to approve or reject or send transactions to be reviewed by a human. The company claims to block about 28 billion dollars in potentially fraudulent transactions every year and has fraud loss rates of only 0.32 of revenue versus the industry average of 1.8. This is not only due to advanced algorithms but to an ongoing investment into the data infrastructure, model monitoring and adaptive learning systems that keep on updating the models as fraud patterns develop.

Yet challenges remain. To begin with, explainability of models is not without issues. Deep neural networks are regarded as black boxes whose decision logic is not visible. Regulators are also requiring reasons behind negative decisions (such as rejecting transactions) which are hard to explain by complex ML models. PayPal partially overcomes this by its counterfactual explanatory methods that discover which changes in the features would have caused varying decision making but full interpretability cannot be achieved.

Second, the adversarial machine learning establishes cat and mouse game. Advanced fraudsters strategically alter their actions in order to avoid ML. As an example, when models position big, suspicious transactions, fraudsters position schemes into numerous smaller transactions. In case of geographic inconsistency alerts, the fraudsters make use of VPNs that conceal locations. Defensive ML involves predicting adversarial adaptation, which is a difficult issue on the edge of machine learning studies.

Third, model performance is of great concern to data quality. Machine learning algorithms are data-driven, or garbage in, garbage out. The required high-quality representative training data can be ensured only at the cost of significant investments. Biased training data generates biased models a problem that we revisit when addressing financial inclusion.

The case of Mastercard Decision Intelligence offers some supplementary information. The AI system at Mastercard uses each transaction to evaluate the risk of fraud by



scoring it in less than 50 milliseconds per year, evaluating about 125 billion transactions per year. The system cut down false declines by 17 percent a significant reduction considering the huge transaction volumes involved. MasterCard underlines that their strategy is to integrate AI with human intelligence and does not displace the analysts. Machine learning is applied to routine decisions, where the probability of success is great, and to the cases which are uncertain, and sent to human investigators who add context and expertise.

3.3 *Financial Inclusion: Democratizing Access*

The most significant effect of artificial intelligence on financial services is, perhaps, the increased access to underserved populations, the 1.7 billion adult non-banked and underbanked non-service users around the world. Serving customers with low incomes, in the rural areas or in the informal sector has often been uneconomical amongst the traditional financial institutions. Using physical branches, costly Know Your Customer (KYC) processes, and risk-averse methods of credit assessment imposes restrictions that lock-out exactly the groups of people most likely to need financial services.

Machine learning has provided possible solutions to circumvent these obstacles through a number of mechanisms. Non-conventional credit scoring gives the possibility of risk rating an individual who does not have a traditional credit profile. Small-balance accounts become cost-effective due to the utilization of digital delivery channels that are operated by AI-based chatbots and austerity customer service. Individualized financial services based on ML recommendation engines are also better suited to the needs of a wide variety of customers than a universal

offering. Our analysis of World Bank Global Findex data reveals substantial progress on financial inclusion globally, with evidence suggesting AI-powered Fin-Tech solutions contributed meaningfully. Fig. 5 presents trends in financial account ownership across different regions and the relationship between Fin-Tech adoption and inclusion outcomes.

The left panel of Fig. 5 shows encouraging trends. Global account ownership increased from 62% of adults in 2014 to approximately 76% in 2023, representing roughly 1.4 billion additional people gaining financial access. Progress occurred unevenly across regions. East Asia and Pacific reached near-universal coverage at 93%, while Sub-Saharan Africa improved from 27% to 55% but still lags substantially. South Asia made remarkable strides from 46% to 77%, driven largely by India's financial inclusion initiatives including biometric identification systems and mobile payment infrastructure.

The right panel presents difference-in-differences analysis estimating AI/ML's causal contribution to inclusion outcomes. We compared regions/countries with high FinTech adoption (defined as top quartile in our adoption index) versus those with low adoption, examining changes in inclusion metrics before and after significant FinTech deployment. Results suggest that high AI/ML adoption associates with 18 percentage point additional improvement in account ownership, 31% greater increase in credit access, and 23% larger gains in savings behavior compared to low-adoption regions, after controlling for economic development, regulatory environment, and digital infrastructure. These estimates carry uncertainty (reflected in confidence intervals), but the consistency across multiple outcomes strengthens confidence that effects are genuine rather than spurious.






Fig. 5: Financial Inclusion Trends and Impact of AI/ML Implementation.

Left panel shows financial account ownership rates by region (2014, 2017, 2021, 2023 estimated). Substantial gains occurred globally, particularly in Sub-Saharan Africa and South Asia where mobile money and digital lending expanded rapidly. Right panel presents difference-in-differences analysis comparing inclusion outcomes in regions with high versus low FinTech/AI adoption. Treatment group (high adoption) shows significantly greater improvement in account ownership (18 percentage points), credit access (31%), and savings behavior (23%) compared to control group. Error bars represent 95 % confidence intervals.

How specifically does machine learning expand financial inclusion? The most direct pathway operates through alternative credit scoring. Traditional credit assessment relies on credit bureau data payment histories for loans, credit cards, utilities. But billions of people lack such histories, not because they are bad credit risks but simply because they've never accessed formal credit. This creates a Catch-22: you need credit history to get credit, but you can't build history without getting credit.

ML-based alternative credit scoring breaks this loop by leveraging non-traditional data sources. Mobile phone usage patterns airtime purchases, data consumption, communication networks contain signals about financial behavior. E-commerce transaction histories, social media activity (with appropriate consent), utility bill payments, even

keystroke dynamics may inform creditworthiness assessments. Machine learning algorithms identify complex patterns linking these alternative data to loan repayment likelihood.

Research by Bjõrkegren and Grissen (2020) demonstrated that mobile phone metadata predicted loan repayment in emerging markets. Models analyzing call records, SMS patterns, and airtime purchases achieved similar predictive accuracy to traditional credit bureau scores in developed countries. This finding holds transformative implications: if predictive accuracy using alternative data equals traditional scores, then the limiting factor becomes data availability rather than analytical capability. And alternative data sources are increasingly ubiquitous as mobile phone penetration expands globally.

Table 5 presents detailed impact metrics across multiple dimensions: access, usage, quality, and socioeconomic outcomes. Beyond simply opening accounts, AI-driven inclusion initiatives demonstrate improvements in active usage (25% increase in monthly active accounts), transaction frequency (172% increase in average monthly digital transactions), and savings behavior (23% increase in adults maintaining savings).

Particularly noteworthy, microloan approval rates increased 67% from 42.7% to 71.3% when institutions adopted ML-based



alternative credit scoring, yet default rates actually declined slightly from 10.7% to 9.2%. This access with long-term or better credit quality is the holy grail of financial inclusion. It implies that ML algorithms will recognize individuals that have been excluded in the past but are actually creditworthy, not just reduced standards and greater risk of default.

There are downstream socioeconomic implications which seem significant but more challenging to causally attribute. Areas that had financial inclusion based on AI experienced 23% growth in the number of small businesses and 15% growth in financial resilience indices (composite measures of the capacity to survive economic shocks). Notably, ensemble methods—which combine predictions from multiple algorithms—are increasingly regarded as best practice for high-stakes financial decisions.

Our M-Pesa case study illustrates these mechanisms at population scale. Launched in

Kenya in 2007 by telecommunications company Safaricom, M-Pesa pioneered mobile money in Africa, enabling users to store value and transfer funds via basic mobile phones without smartphone requirements. The service achieved remarkable penetration approximately 30 million active users representing 80% of Kenya’s adult population. MPesa effectively became the country’s financial infrastructure, processing more domestic transactions than all Kenyan banks combined.

In 2012, M-Pesa introduced M-Shwari, a savings and credit product leveraging machine learning for credit scoring. Rather than credit bureau data (which fewer than 3 % of Kenyans possess), M-Shwari’s algorithms analyze users’ M-Pesa transaction histories. How frequently do they receive transfers? Do they maintain positive balances? Do transaction patterns suggest stable income? These behavioral signals enabled initial credit assessments for millions who were previously “credit invisible.”

Table 5: Socioeconomic Impact Metrics of AI-Driven Financial Inclusion Initiatives

Outcome Metric	Before AI/ML	After AI/ML	Change
<i>Access Indicators</i>			
Account Ownership (%)	64.2	75.8	+18 %
Credit Product Access (%)	31.5	41.3	+31 %
Microloan Approval Rate (%)	42.7	71.3	+67 %
<i>Usage Indicators</i>			
Active Account Usage (% monthly)	54.8	68.4	+25 %
Digital Transaction Frequency (monthly avg.)	3.2	8.7	+172 %
Savings Behavior (% with savings)	38.6	47.5	+23 %
<i>Quality Indicators</i>			
Default Rate (% of loans)	10.7	9.2	-14 %
Customer Satisfaction (1–5)	3.4	4.1	+21 %
Average Account Balance (\$)	127	218	+72 %
<i>Socioeconomic Outcomes</i>			



Small Business Formation (%)	14.2	17.5	+23 %
Financial Resilience Index (0–100)	52.3	60.1	+15 %
Gender Gap (percentage points)	9.4	6.2	-34 %

Note: Pooled estimates across high AI/ML adoption regions. Sample size varies by metric.

M-Shwari has disbursed over \$3.2 billion in small loans to 22 million Kenyans since launch. Average loan sizes hover around \$50 tiny by developed country standards but meaningful for households living on \$2–5 daily incomes. Crucially, the system learns continuously. Repayment behavior on initial small loans informs credit limits for subsequent loans, creating pathways from financial exclusion toward full inclusion. Research by Jack and Suri (2014) and Suri and Jack (2016) documented substantial positive impacts on household resilience, consumption smoothing, and gender equity. Yet M-Pesa’s success reflects favorable contextual factors beyond just technology. Kenya’s regulatory environment enabled telecommunication companies to offer financial services without onerous banking licenses. Low initial mobile banking penetration meant M-Pesa faced limited competition, allowing network effects to compound. And Safaricom’s market dominance provided the necessary scale. Whether similar models can succeed elsewhere remains uncertain replication attempts in other African countries have met mixed results.

Our Nubank case study from Brazil offers contrasting insights. Nubank, founded in 2013, grew to over 85 million customers by 2023, making it the largest digital bank outside Asia. Unlike M-Pesa’s mobile-first approach for feature phones, Nubank targets smartphone users with slick interfaces and sophisticated AI-powered financial management tools. The company’s proprietary ML models assess credit risk using data far beyond traditional credit scores analyzing behavioral patterns, social network characteristics (with consent), and

transaction histories from Nubank’s own platform.

Nubank serves disproportionately young adults and previously unbanked populations skeptical of traditional banks. Their ML-driven credit score yields default rates that are around 40 percent lower than those of the industry in spite of it being used in more at-risk groups. Due to the delivery of solely digital, as well as high levels of automation, operating costs are 83% lower than those of a traditional bank. This cost structure allows it to make a profit even on low balance accounts that would not be accepted by traditional banks as economically viable.

Nonetheless, when credit is scored using algorithms, then fairness is compromised. In case algorithms are trained on past data, they can reinforce or even increase bias. A model that is trained based on discriminatory lending patterns is likely to recreate such discriminatory patterns despite not considering any of the protected attributes such as race or gender. This notion of fairness by unawareness is referred to as such by ML practitioners because it is the naive belief that sensitive variables should be left out in order to guarantee fairness. Indeed, proxy variables (such as zip-codes as a proxy of racial make up) enable algorithms to discriminate indirectly.

Algorithms need to be consciously addressed to deal with bias. Methods are disparate impact testing (testing whether there are differences in the rate of approval between demographic groups), adversarial debiasing (training models to jointly maximise accuracy and fairness), and fairness-aware learning (adding explicit fairness requirements to the training of a model). However, defining what fairness is itself is



debatable they (demographic parity, equalized odds, calibration) can be mathematically antithetical, and the value judgments necessary to determine which fairness notions are most important in any given situation (Barocas and Selbst, 2016).

4.0 Conclusion

This study explored the issue of machine learning and the role of artificial intelligence in changing financial services in the digital payment, fraud detection and financial inclusion sectors via systematic literature review, large-scale quantitative study and case studies. Our research provides a number of strong results proving the transformative effect of AI/ML and revealing some unresolved issues. Digital payment systems with AI make 67 percent less processing time and achieved almost 90 percent faster cross border settlements, which corresponds to significant cost savings and better user experiences in large volumes. Machine learning-based methods of fraud detection are far more successful than traditional rule-based systems, with the highest-performing algorithms having an accuracy of up to 94 - 98% and a false positive rate of less than 0.5% of one-third the false alarms of conventional systems, which means that large institutions could potentially save hundreds of millions in fraud annually while also decreasing customer friction. AI-based alternative credit scoring significantly increases financial inclusion in underserved communities, with microloan approval rates rising by 67 points when financial institutions embraced ML-based assessment technologies and default rates did not fall, which offers strong arguments that the conventional credit scoring is a systematic means of shutting out creditworthy people. Theoretically, the study works to bring together the disjointed literatures in computer science, economics, finance, and development studies and presents AI/ML as a supporting infrastructure that both changes and is changed by the development of payments, fraud detection, and inclusion and the presence of feedback

loops, unlike the linear models used in the previous research. In the case of financial institutions, our results indicate that the adoption of AI/ML is a competitive requirement, but its effective adoption cannot be achieved without more than buying algorithms to the vendors, required by the organization data infrastructure, technical skills, and ability to manage changes. To the policymakers, our study identifies tensions that need to be well balanced in the enablement of innovation and risk management through regulatory sandboxes, explainability demands, and fairness auditing, but coordinating across borders is a challenge in particular cases since financial services are provided globally, and regulation is mainly national. There are a variety of limitations that characterize our findings, such as data limitations to mainly consider environments with a high level of digital infrastructure, the high rate of technological change that may age the findings very fast, and the fact that the causal inference of observational data is always difficult with quasi experimental designs. Future research directions involve the exploration of the interface between AI and blockchain technologies, further exploration of adversarial machine learning dynamics, devising methods of debiasing algorithms that can be used in production systems, cross-cultural studies of the activity of AI/ML in different sociocultural settings, and longitudinal studies of the financial paths of individuals. Machine learning and artificial intelligence are not inherently good or bad tools, but have consequences and must be managed with technical expertise alongside social sensitivity, financial intellect along with moral logic, and continuous dialogue between technologists, policy makers, civil society, and the communities they serve must ensure that these systems promote the flourishing of everyone, not just commercial or political elites.



References

- Abdallah, A., Maarof, M. A., & Zainal, A. (2016). Fraud detection system: A survey. *Journal of Network and Computer Applications*, 68, pp. 90–113. <https://doi.org/10.1016/j.jnca.2016.04.007>
- Ademilua, D. A., & Areghan, E. (2022). AI-Driven Cloud Security Frameworks: Techniques, Challenges, and Lessons from Case Studies. *Communication in Physical Sciences*, 8, 4, pp. 674–688.
- Akinsanya, M. O., Adeusi, O. C., Ajanaku, K. B. (2022). A Detailed Review of Contemporary Cyber/Network Security Approaches and Emerging Challenges. *Communication in Physical Sciences*. 8(4): 707-720.
- Arner, D. W., Barberis, J., & Buckley, R. P. (2016). The evolution of Fintech: A new post-crisis paradigm. *Georgetown Journal of International Law*, 47, 4, pp. 1271–1319. <https://doi.org/10.2139/ssrn.2676553>
- Barocas, S., & Selbst, A. D. (2016). Big data's disparate impact. *California Law Review*, 104, pp. 671–732. <https://doi.org/10.15779/Z38BG31>.
- Bech, M., Faruqui, U., Ougaard, F., & icillo, C. (2017). Payments are a-changin' but cash still rules. *BIS Quarterly Review*, March 2017. Bank for International Settlements. Available at: https://www.bis.org/publ/qtrpdf/r_qt1703g.htm
- Bishop, C. M. (2006). *Pattern recognition and machine learning*. New York: Springer. <https://doi.org/10.1007/978-0-387-45528-0>
- Björkegren, D., & Grissen, D. (2020). Behavior revealed in mobile phone usage predicts credit repayment. *The World Bank Economic Review*, 34, 3, pp. 618–634. <https://doi.org/10.1093/wber/lhz006>
- Bolton, R. J., & Hand, D. J. (2002). Statistical fraud detection: A review. *Statistical Science*, 17, 3, pp. 235–255. <https://doi.org/10.1214/ss/1042727940>.
- Brummer, C., & Yadav, Y. (2019). Fintech and the innovation trilemma. *Georgetown Law Journal*, 107, pp. 235–307. <https://doi.org/10.2139/ssrn.3054770>.
- Claessens, S., Frost, J., Turner, G., & Zhu, F. (2018). Fintech credit markets around the world: Size, drivers and policy issues. *BIS Quarterly Review*, September 2018, 29–49. Bank for International Settlements. Claessens, S., Frost, J., Turner, G., & Zhu, F. (2018). *Fintech credit markets around the world: Size, drivers and policy issues*. *BIS Quarterly Review*, September, 29–49. Bank for International Settlements. https://www.bis.org/publ/qtrpdf/r_qt1809e.htm
- Dahlberg, T., Guo, J., & Ondrus, J. (2015). A critical review of mobile payment research. *Electronic Commerce Research and Applications*, 14, 5, pp. 265–284. <https://doi.org/10.1016/j.eierap.2015.07.006>.
- Dal Pozzolo, A., Caelen, O., Johnson, R. A., & Bontempi, G. (2015). *Calibrating probability with undersampling for unbalanced classification*. In 2015 IEEE Symposium Series on Computational Intelligence (pp. 159–166). IEEE. <https://doi.org/10.1109/SSCI.2015.33>
- Davis, F. D. (1989). *Perceived usefulness, perceived ease of use, and user acceptance of information technology*. *MIS Quarterly*, 13, 3, pp. 319–340. <https://doi.org/10.2307/249008>
- Demirgüç-Kunt, A., Klapper, L., Singer, D., & Ansar, S. (2022). *The Global Findex Database 2021: Financial inclusion, digital payments, and resilience in the age of COVID-19*. Washington, DC: World Bank. <https://doi.org/10.1596/978-1-4648-1897-4>
- Demirgüç-Kunt, A., Klapper, L., Singer, D., Ansar, S., & Hess, J. (2020). Measuring financial inclusion and the fintech revolution. *The World Bank Research Observer*, 35(1), 82–101. <https://doi.org/10.1093/wbro/lkz001>
- Frost, J., Gambacorta, L., Huang, Y., Shin, H. S., & Zbinden, P. (2019). BigTech and



- the changing structure of financial intermediation. *Economic Policy*, 34, 100, pp. 761–799. <https://doi.org/10.1093/epolic/eiaa003>.
- Gomber, P., Kauffman, R. J., Parker, C., & Weber, B. W. (2018). On the fintech revolution: Interpreting the forces of innovation, disruption, and transformation in financial services. *Journal of Management Information Systems*, 35, 1, pp. 220–265. <https://doi.org/10.1080/07421222.2018.1440766>
- Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep learning*. Cambridge, MA: MIT Press. Available at: <https://www.deeplearningbook.org>
- Jack, W., & Suri, T. (2014). *Risk sharing and transaction costs: Evidence from Kenya's mobile money revolution*. *American Economic Review*, 104, 1, pp. 183–223. <https://doi.org/10.1257/aer.104.1.183>
- Okolo, J. N. (2021). A Systematic Analysis of Artificial Intelligence and Data Science Integration for Proactive Cyber Defense: Exploring Methods, Implementation Obstacles, Emerging Innovations, and Future Security Prospects. *Communication in Physical Sciences*. 7(4): 681-696.
- Lee, I., & Shin, Y. J. (2018). *Fintech: Ecosystem, business models, investment decisions, and challenges*. *Business Horizons*, 61, 1, pp. 35–46. <https://doi.org/10.1016/j.bushor.2017.09.003>
- Ozili, P. K. (2023). Digital payments and central bank digital currency adoption: Opportunities and challenges. *Journal of Financial Regulation and Compliance*, 31, 2, pp. 156–175. <https://doi.org/10.1108/JFRC-04-2022-0039>.
- Philippon, T. (2016). *The fintech opportunity* (NBER Working Paper No. 22476). National Bureau of Economic Research. <https://doi.org/10.3386/w22476>
- Pumsirirat, A., & Yan, L. (2018). *Credit card fraud detection using deep learning based on auto-encoder and restricted Boltzmann machine*. *International Journal of Advanced Computer Science and Applications*, 9, 1, pp. 18–25. <https://doi.org/10.14569/IJACSA.2018.090104>
- Rogers, E. M. (2003). *Diffusion of innovations* (5th ed.). Free Press.
- Sen, A. (1999). *Development as freedom*. Oxford University Press.
- Suri, T., & Jack, W. (2016). *The long-run poverty and gender impacts of mobile money*. *Science*, 354, 6317, pp. 1288–1292. <https://doi.org/10.1126/science.aah5309>
- Sutton, R. S., & Barto, A. G. (2018). *Reinforcement learning: An introduction* (2nd ed.). MIT Press. <http://incompleteideas.net/book/the-book-2nd.html>,
- Zetsche, D. A., Buckley, R. P., Arner, D. W., & Barberis, J. N. (2017). *Regulating a revolution: From regulatory sandboxes to smart regulation*. *Fordham Journal of Corporate & Financial Law*, 23, pp. 31–103. <https://doi.org/10.2139/ssrn.2853941>

Declaration

Consent for publication

Not Applicable

Availability of data and materials

The publisher has the right to make the data public

Ethical Considerations

Not applicable

Competing interest

The authors report no conflict or competing interest

Funding

No funding

Authors' Contributions

EA conceptualized the study, coordinated the research design, and led manuscript development. AD handled machine learning frameworks, data analytics, and fraud detection modeling. CD analyzed financial inclusion metrics and business implications. TM conducted case study and digital payment efficiency analyses. LA supported statistical validation, policy interpretation, and critical review, strengthening the conceptual framework and overall scholarly rigor.

