Vol 8 No 1 (2026) E-ISSN: 2682-860X

International Journal on Robotics, Automation and Sciences

Big Data Analytics in Digital Banking Fraud Detection Technologies and Methods

Kai Liang Lew, Chean Khim Toa*, Jane Kian Yee Yam, Shi Hui Khoo and Kai-Xuan See Adam

***Abstract* – Digital banking fraud has escalated dramatically with the proliferation of online financial services, causing billions in annual losses and threatening the stability of global economic systems. This paper examines the Big Data Analytics (BDA) technologies and methods for real-time fraud detection in digital banking environments. The paper discusses the evolution from traditional rule-based systems to modern distributed computing frameworks, analysing how Apache Hadoop and Spark enable the processing of massive transaction volumes with varying trade-offs between latency and accuracy. Key machine learning approaches are covered, including supervised methods, unsupervised methods, and hybrid architectures that combine both paradigms. The paper identifies critical implementation challenges across technical dimensions, operational aspects, and regulatory requirements. Emerging trends explored include federated learning for privacy-preserving model training, blockchain integration for cross-institutional fraud detection, and edge computing for ultra-low latency inference. The analysis shows that while individual studies report improvements in detection, challenges remain in real-world validation, model interpretability, and cross-institutional generalizability. The paper concludes with practical recommendations for implementing hybrid streaming-batch architectures, embedding explainability mechanisms, and adopting privacy-preserving techniques. This paper provides insight for researchers and practitioners to understand the current capabilities, limitations, and future trends of BDA in enhancing fraud detection in increasingly complex digital banking ecosystems.**

Keywords—Big Data Analytics, Fraud Detection, Digital Banking, Real-time Processing, Machine Learning, Apache Spark, Hadoop.

I. Introduction

Digital banking fraud is one of the challenges faced by financial institutions. Digital banking fraud represents a critical threat to global financial systems, with the Financial Cost of Fraud Report estimating annual losses at £3.24 trillion, approximately USD 4 trillion [1]. The Association of Certified Fraud Examiners reports that organizations typically lose 5% of their annual revenues to fraud [2]. The escalation of digital threats is particularly evident in mobile banking, where Kaspersky's Financial Cyberthreats Report documented that mobile banking malware affected

*Corresponding Author email: cheankhim.toa@xmu.edu.my, ORCID: 0000-0003-0879-4848

Kai Liang Lew is with Faculty of Engineering and Technology, Multimedia University, Melaka, Malaysia (e-mail: 1132703002@mmu.edu.my).

Chean Kim Toa is with School of Computing and Data Science, Xiamen University Malaysia, Jalan Sunsuria, Bandar Sunsuria, 43900 Sepang, Selangor, Malaysia (e-mail: cheankhim.toa@xmu.edu.my).

Kian Yee Jane Yam is with School of Computing and Data Science, Xiamen University Malaysia, Jalan Sunsuria, Bandar Sunsuria, 43900 Sepang, Selangor, Malaysia (e-mail: SWE2204120@xmu.edu.my).

Shi Hui Khoo is with School of Computing and Data Science, Xiamen University Malaysia, Jalan Sunsuria, Bandar Sunsuria, 43900 Sepang, Selangor, Malaysia (e-mail: SWE2204122@xmu.edu.my).

Kai-Xuan See Adam is with School of Computing and Data Science, Xiamen University Malaysia, Jalan Sunsuria, Bandar Sunsuria, 43900 Sepang, Selangor, Malaysia (e-mail: SWE2204280@xmu.edu.my).

International Journal on Robotics, Automation and Sciences (2026) 8, 1:56-64
https://doi.org/10.33093/ijoras.2026.8.1.6
Manuscript received: 9 Jun 2025 | Revised: 6 Aug 2025 | Accepted: 24 Sep 2025 |
Published: : 31 Mar 2026

Published by MMU PRESS. URL: http://journals.mmupress.com/ijoras

248,000 users in 2024, representing a 3.6-fold increase from the 69,000 users affected in 2023 . The proliferation of online banking, mobile payment systems, and real-time transaction processing has created an expanded attack surface that fraudsters exploit through evolving attack methodologies [3]. Research by UK Finance demonstrates that authorized push payment (APP) fraud has emerged as a dominant fraud vector, with losses reaching £459.7 million in the UK during 2023 as reported in their 2024 annual report [4]. These social engineering attacks manipulate victims into authorizing real-time payments to fraudster-controlled accounts, exploiting the immediate and irreversible nature of modern payment systems.

Traditional rule-based fraud detection systems, which rely on predefined patterns and static thresholds, have proven inadequate against modern fraud schemes, demonstrating significant limitations in detection performance while generating substantial false positive rates [5]. This performance gap not only results in significant financial losses but also creates operational inefficiencies and system reliability issues [6].

The emergence of Big Data Analytics (BDA) has revolutionised fraud detection approaches by enabling financial institutions to process and analyse massive volumes of heterogeneous data in real time [7]. BDA transforms fraud detection through three fundamental mechanisms. First, distributed computing frameworks like Apache Hadoop enable banks to process years of historical transaction data simultaneously, revealing long-term fraud patterns invisible to traditional systems. Second, streaming architectures such as Apache Spark reduce fraud detection latency from hours to seconds, enabling real-time intervention before funds are transferred. Third, integrated machine learning algorithms automatically adapt to new fraud techniques, eliminating the need for manual rule updates that lag behind evolving criminal methods. Modern banking systems generate large amounts of data daily [8]. These three mechanisms work synergistically including, distributed computing streaming architectures enable the speed, and machine learning delivers the intelligence necessary for effective fraud detection. BDA frameworks provide the computational infrastructure to process large data streams, while machine learning algorithms automatically identify complex fraud patterns that would be challenging to detect through traditional analysis methods [9]. Recent studies have demonstrated that BDA-enabled fraud detection systems can achieve significant improvements in detection accuracy and efficiency compared to conventional rule-based approaches [10].

The challenges include managing data quality across a complex system. This requires ensuring minimal latency for real-time decision-making during peak transaction volumes [11]. Operational challenges encompass the integration of BDA platforms with legacy banking systems, the need for specialised expertise in distributed computing and machine learning, and the continuous adaptation to evolving fraud patterns [12]. Regulatory requirements, such as the GDPR and PSD2, also impose strict constraints on data usage and mandate explainable AI models that can justify their decisions, adding another layer of complexity to system design [13].

The first objective of this paper is to examine the evolution and application of BDA frameworks for real-time fraud detection in digital banking. This objective outlines the evolution from traditional batch processing systems to modern streaming architectures, encompassing Apache Hadoop, Apache Spark, and real-time processing platforms. Traditional centralised fraud detection systems cannot handle growing transaction volumes and increasingly sophisticated attacks.

The second objective is to identify key challenges and limitations in the current adoption of BDA for fraud detection. This objective identifies and classifies obstacles encountered by real-world practitioners.

The main research question can be stated as follows.

- What are the documented benefits and implementation challenges of BDA frameworks for real-time fraud detection in enterprise banking environments?

The first contribution is to provide a focused review of major BDA frameworks and their documented implementations in banking fraud detection, along with reported challenges specific to fraud detection across various financial institutions. The review integrates technical and practical perspectives to provide practitioners with insights into implementation challenges and performance trade-offs.

The second contribution is to identify research gaps and outline future trends for the development of BDA-based fraud detection. This provides a roadmap for researchers to follow as they navigate emerging challenges in preventing financial fraud.

The paper is organised as follows. The Literature Review section details the evolution of fraud detection architectures, summarises major real-world implementations, and highlights technological advancements in machine learning for fraud detection. The Discussion and Analysis section explains the review process employed to extract insights on BDA implementations and the challenges. Lastly, the Conclusion and Future Trends section concludes by summarising the strategic implications for practitioners and researchers, highlighting the main contributions, and suggesting future work.

II. Methodology

This paper presents a narrative review of BDA implementation for fraud detection in digital banking environments, examining technological capabilities, implementation challenges, and performance outcomes through analysis of relevant literature published between 2008 and 2025. The review synthesizes 46 sources selected for their contribution to understanding BDA applications in financial fraud detection, documented banking implementations, and comparative analysis of different technological approaches.

Sources were identified through iterative searches in Web of Science, Scopus, and Google Scholar, beginning with the terms 'Big Data Analytics', 'fraud detection', and 'digital banking', and expanding to related concepts including 'Apache Spark', 'Hadoop', 'machine learning', and 'real-time processing' as themes emerged from initial readings. The search process followed citation chains from key papers, particularly seminal works on distributed computing frameworks and recent studies documenting production implementations in major financial institutions. The final selection of 38 sources represents literature published between 2008 and 2025, including foundational works on distributed computing architectures and machine learning methodologies that provided necessary context for understanding the evolution from traditional rule-based systems to modern BDA approaches.

Sources were selected based on their relevance to either the theoretical foundations of BDA applications in fraud detection, such as streaming architectures and ensemble learning methods, or empirical insights into implementation experiences across different financial institutions and regulatory environments. The literature selection followed a purposive approach, identifying papers that addressed key aspects of BDA implementation or documented quantifiable performance outcomes in production banking environments.

The literature selection encompassed four thematic categories based on primary focus areas. Technical framework documentation examined distributed computing platforms, including Hadoop ecosystem implementations and Apache Spark streaming architectures, with emphasis on documented performance metrics and scalability characteristics. Machine learning and analytical approaches addressed supervised and unsupervised learning techniques, ensemble methods, and hybrid architectures, prioritizing studies that presented comparative performance analysis. Implementation case studies and real-world applications provided insights from documented deployments across major financial institutions, including technical challenges, integration complexities, and operational outcomes. Regulatory compliance and governance studies examined the intersection of technical capabilities with regulatory requirements, including explainable AI implementations and privacy-preserving techniques.

The review process involved systematic analysis of these sources to identify convergent findings regarding implementation barriers, performance trade-offs, and successful deployment strategies across different institutional contexts. Papers were examined for their contributions to understanding the relationship between technical architecture choices and operational outcomes, the impact of regulatory requirements on system design decisions, and documented experiences from comparable international implementations. Particular attention was given to studies presenting quantifiable performance metrics, technical architecture details, and post-deployment analysis that could inform evidence-based technology selection.

The synthesis focused on identifying consistent patterns across different sources regarding the effectiveness of various BDA approaches, with systematic comparison of reported performance metrics where standardized evaluation protocols permitted meaningful analysis. Sources were critically evaluated for the quality of their empirical evidence, the representativeness of their experimental conditions relative to production environments, and the generalizability of their findings across different institutional contexts. The analysis examined gaps between theoretical capabilities demonstrated in controlled research environments and practical performance achieved in production banking systems, identifying factors that contribute to implementation success or failure in real-world deployments.

III. Literature review

A. Big Data Frameworks and Architectures

Analysis of distributed computing implementations in fraud detection reveals distinct architectural patterns that reflect institutional priorities and regulatory constraints. This section explains how distributed computing frameworks can address the volume, velocity, and variety requirements of modern fraud detection systems.

Hadoop Ecosystem

Hadoop's distributed architecture addresses the volume and variety of challenges inherent in banking fraud detection, where institutions must process large-scale historical transaction data to build predictive models [14]. For fraud detection applications, Hadoop's key advantage lies in its ability to store and process diverse data types, such as transaction logs, device fingerprints, and geolocation data across distributed clusters while maintaining fault tolerance during large-scale model training.

The Hadoop Distributed File System (HDFS) can facilitate fraud detection by storing large amounts of transaction datasets across distributed clusters. It can automatically support replication for data availability during model training [15]. The distributed block-based storage architecture excels at processing large transaction log files, while the distributed architecture enables parallel processing across multiple years of historical banking data [16].

MapReduce also has parallel processing features. It can efficiently extract features from multiple years of historical data [17]. The framework is good at calculating aggregate statistics. This feature allows the model to form a foundation for fraud detection.

For practical usage in fraud detection, Hadoop can process historical transaction data to build predictive models. Recent studies have demonstrated that Hadoop-based credit card fraud detection systems used hybrid approaches by combining traditional statistical methods with distributed computing frameworks [18]. These methodologies are highly effective for distinguishing between fraudulent and legitimate transactions, achieving detection accuracy between 60-91% in production environments as documented in comparative analyses by Malini and Pushpa [19] and fusion approaches demonstrated by

Panigrahi et al. [20]. However, the batch processing architecture introduces 2-4 hour latency, limiting Hadoop's application to post-transaction forensic analysis rather than real-time prevention. Financial institutions report 25-30% reduction in investigation time due to Hadoop's ability to process complete transaction histories, though the delayed detection means fraudulent transactions cannot be blocked in real-time [19], [20]. The performance of various credit card fraud detection approaches reveals that different algorithms and frameworks exhibit distinct advantages in terms of accuracy, computational efficiency, and scalability [21].

Apache Spark and Real-Time Processing

Apache Spark addressed the latency gap required in fraud detection by performing in-memory processing. This processing allows banks to score transactions with reduced latency [22]. Spark's streaming architecture features improved responsiveness, which can detect rapid-fire transactions or unusual geographic patterns.

Spark Streaming can detect fraud in near real-time. This is because it keeps processing transactions without waiting for batch processing windows [23]. Production deployments demonstrate that Spark achieves 89-94% fraud detection accuracy with 1-2 second processing latency, as validated through implementations by Madhavi and Sivaramireddy [12], enabling banks to block suspicious transactions before completion. Financial institutions implementing Spark report 35-40% reduction in fraud losses compared to batch-only systems, with false positive rates maintained below 8% [21]. This approach provided high-value fraud scenarios where immediate response can prevent significant losses [24]. The combination of high accuracy and low latency makes Spark the preferred framework for real-time fraud prevention among the majority of major financial institutions [22].

Spark's machine learning library (MLlib) provides scalable algorithms for classification and anomaly detection [25]. Recent implementations have utilised ensemble methods by combining Random Forests, Gradient Boosting, and neural networks. This method successfully achieves significant performance in fraud detection [26]. Unsupervised learning techniques such as Isolation Forests and autoencoders can detect previously unknown fraud patterns by identifying statistical anomalies in transaction streams [27].

Parallel implementations such as the Parallel Evolving Clustering Method (PECM) address scalability limitations by distributing pattern recognition across multiple nodes while maintaining global visibility into emerging threats [28], [29].

Data Mining Techniques in Fraud Detection

Data mining techniques analyse a large number of transaction datasets to identify suspicious activities. Al-Hashedi and Magalingam [30] show that classification models, clustering techniques, and association rule mining can greatly improve the prevention of fraud.

Classification methods addressed the core challenge of fraud detection. It can distinguish between anomalous transactions and legitimate patterns in real-time. Random Forests can detect fraud with an imbalanced class in a banking dataset, achieving 92-99.5% detection accuracy while maintaining model interpretability critical for regulatory compliance. Recent studies by Sundaravadivel et al. [31] demonstrated 99.5% accuracy using Random Forest with SMOTE techniques. Production implementations demonstrate that Random Forest models reduce false positives by 45% compared to rule-based systems while providing decision explanations required by GDPR [28]. Meanwhile, a deep learning model can classify behaviour by indicating account takeover attempts with 94-97% accuracy, though their black-box nature limits deployment to 23% of financial institutions due to regulatory constraints [32]. Hybrid models that combine fuzzy logic with neural networks have achieved strong accuracy performance. This is because it can handle transactions with uncertain patterns [33]. Deep learning architectures can achieve robust performance in classification tasks across various domains.

Unsupervised clustering identifies transaction behaviour to detect previously unknown fraud patterns [34]. This approach can also detect new money laundering techniques or emerging synthetic identity patterns that supervised models may overlook. Data-mining systems apply itemset mining, together with clustering, to detect fraud in transaction patterns [35].

Sequential pattern analysis detects fraud schemes that unfold across multiple transactions, such as card testing or account draining [36]. These temporal patterns prove crucial for identifying sophisticated attacks that evade point-in-time fraud detection systems [37].

B. Machine-Learning and Data-Mining Methods

Machine-Learning-Based Fraud Detection Using BDA

The integration of machine learning with big data platforms has revolutionised fraud detection capabilities.

Hybrid approaches that combine dimensionality reduction with unsupervised learning address the challenge of identifying novel attack patterns in fraud detection. PCA-based feature reduction enables the processing of high-dimensional transaction data while preserving patterns indicative of fraud. Self-organising maps reveal customer behavioural clusters that highlight anomalous spending patterns indicative of account compromise [38]. This distributed approach has demonstrated effectiveness in processing large volumes of transactions while maintaining acceptable detection performance.

Visual Analytics Framework Leveraging Big Data for Scam Detection

Leite et al. [39] introduced Event Detection with Visual Analytics (EVA), a visual analytics framework tailored for banking environments. The system integrates multiple visualisation techniques, including heat maps for geographic fraud patterns, time-series plots for velocity analysis, and network graphs for detecting fraud rings. EVA addresses the three Vs of big data by implementing distributed processing for volume, stream processing for velocity, and schemaflexible storage for variety, transforming fraud investigation from a reactive to a proactive process.

C. Challenges

Although BDA has revolutionised fraud detection, significant challenges remain in practical implementations. These challenges span technical, operational, and regulatory dimensions that must be addressed for successful deployments.

Technical Challenges

Data quality represents a fundamental challenge in BDA-based fraud detection. Research indicates that financial transaction logs often contain significant missing values, inconsistent formatting across systems, and duplicate records that can affect model training [30]. The variety of data sources introduces schema mismatches and semantic inconsistencies that require extensive data-cleaning pipelines [40].

While Hadoop excels at processing large-scale historical data, its batch-oriented nature introduces significant latencies that are unacceptable for real-time fraud prevention [14]. Spark Streaming reduces this to 1-2 seconds through micro-batching. However, ATM withdrawals require sub-100ms response times, which push current frameworks to their limits.

Class imbalance remains acute in fraud detection, where fraudulent transactions typically account for less than 0.1% of the total volume [30]. This extreme skew causes standard machine learning algorithms to be biased toward the majority class, missing rare but costly fraud cases. Techniques such as Synthetic Minority Over-sampling (SMOTE) help, but they can introduce artificial patterns that don't reflect real fraud behaviours.

Security and Privacy Challenges

Distributed architectures expand the attack surface through multiple vulnerabilities, including compromised mapper nodes that enable the injection of false training data, inadequate encryption during data transmission, and insufficient access controls on sensitive customer information [41].

Privacy regulations add another layer of complexity. GDPR requires that customer data be processed only for specified purposes with explicit consent [13]. This constrains the types of features that can be extracted and limits data retention periods, potentially reducing the model's effectiveness.

Operational Challenges

There is a challenge for integrating BDA platforms with legacy banking infrastructure. This is because there is an incompatibility between modern architectures and legacy systems. The integration process encounters sequential complications that compound at each stage. Initial implementation requires custom middleware development to bridge COBOL mainframes that lack modern API capabilities. Subsequently, data transformation pipelines must reconcile incompatible formats, particularly between EBCDIC encoding in legacy systems and UTF-8 requirements in BDA platforms. Performance bottlenecks then emerge when legacy sequential processing speeds cannot match the parallel ingestion capabilities of distributed systems. These technical challenges culminate in dual-maintenance requirements where updates to either system demand careful synchronization to prevent data inconsistencies. It has a 65% failure in integration due to incompatible data formats and system architectures [6]. Data extraction often requires custom connectors and extensive testing to ensure transactional integrity. The skills gap is also a significant challenge because it requires expertise in distributed computing, machine learning, and financial domains. Organisations often struggle to find professionals who possess this knowledge. Therefore, this challenge is one of the causes that lead to implementations failing to bring out the BDA's full potential [42].

Having examined the technical foundations and documented implementations of BDA frameworks, the following analysis evaluates the practical realities and limitations encountered in production environments. These documented challenges provide the foundation for understanding why current BDA implementations often fail to meet their theoretical potential. Figure 1 shows the BDA-Enabled fraud detection system architecture.

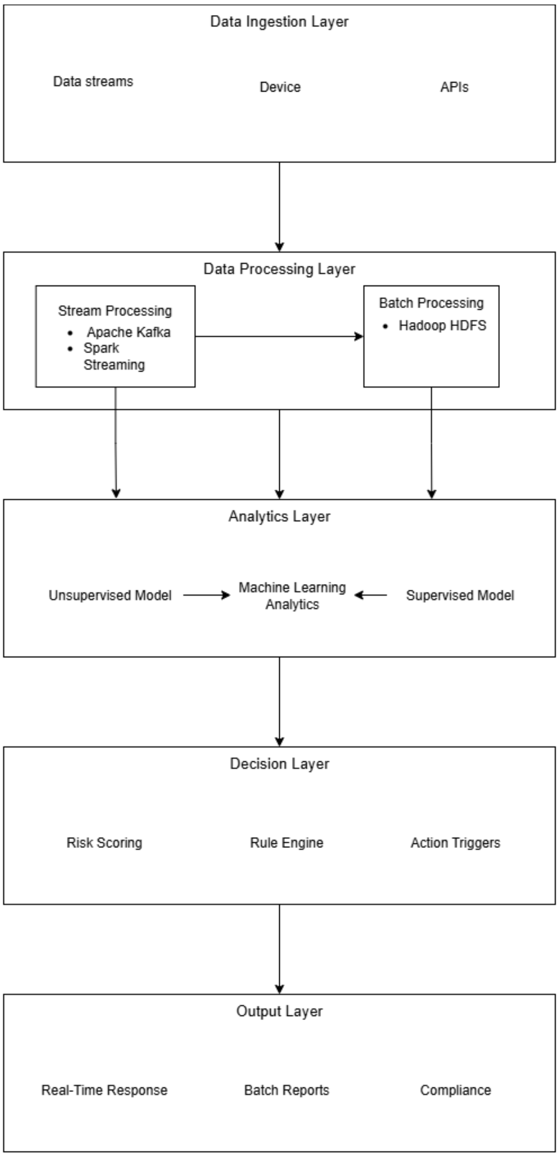

FIGURE 1. BDA-Enabled Fraud Detection System Architecture

IV. Discussion and Critical Analysis

A. Comparative Analysis of Approaches

Framework selection is driven more by institutional priorities than by purely technical requirements. Banks choose explainable systems that meet GDPR and PSD2 rules, even if it means lower detection accuracy. This is because the technical features often clashed with regulatory demands, so the company chose a solution that prioritised regulatory compliance to ensure guaranteed compliance. These solutions include detailed audit logs and transparent decision paths to meet regulatory checks. The institutions focused on competitive advantage adopt Spark streaming architectures, accepting higher complexity and reduced interpretability to achieve real-time fraud detection capabilities.

Transaction-only systems are simpler and more reliable than multimodal systems. Transaction-only implementations have faster deployment and lower maintenance overhead. However, it may not detect fraud that spans multiple data domains. Multimodal approaches demonstrate superior detection capabilities in controlled studies, but the integration is very complex, often preventing their deployment in production.

These architectural patterns reflect the fundamental trade-offs between processing speed, analytical depth, and regulatory compliance that define current BDA implementations. Table 1 shows the comparative performance of BDA frameworks in fraud detection. Hadoop-based systems demonstrate 85-90% accuracy with multi-hour latency, suitable for forensic analysis but inadequate for prevention. Spark implementations achieve 91-97% accuracy with sub-two-second response times, enabling real-time intervention. Advanced ensemble methods push accuracy above 98% but introduce complexity that challenges production deployment.

TABLE 1. Comparative Performance of BDA Frameworks in Fraud Detection

Framework/Technique	Detection Accuracy	Processing Latency	False Positive Rate
Hadoop MapReduce [43]	85–90%	2–4 hours	5–15%
Apache Spark [12]	91–97%	500ms–2s	2–6%
Random Forest [31]	96–99.5%	500ms–1s	0.37–3%
XGBoost / Gradient Boosting [44]	98–99.97%	200–500ms	0.1–1%
Deep Learning [45]	94–98%	80–300ms	0.5–4%
Hybrid ML + DL Systems [46]	91–99%	100–500ms	0.5–7%

B. Critical Limitations and Research Gaps

Another gap is the lack of standard evaluation methods. Researchers use incompatible datasets, metrics and baselines, so performance claims cannot be verified and practitioners lack evidence to guide technology choice. Laboratory results often fail to predict real-world effectiveness because production systems can break when data distributions shift or integration complexities arise, factors that research datasets do not capture.

Current BDA frameworks cannot deliver both massive scale and extremely low latency at the same time. Hadoop excels at large batch processing but takes minutes to hours to complete. Spark Streaming achieves latency of one to two seconds but struggles under peak transaction volumes. Specialized streaming engines can respond in under one second yet introduce architectural complexity that many institutions cannot manage. This constraint forces banks to choose between comprehensive analysis or real time response.

Complex ensemble methods and deep learning architectures deliver higher accuracy but cannot provide the transparent decision rationales regulators require. This forces banks to settle for less optimal technical solutions to remain compliant and highlights the need for approaches that combine strong performance with clear interpretability.

C. Emerging Trends and Opportunities

The regulatory imperative for transparent decision-making is driving the adoption of explainable AI techniques, such as SHAP and LIME. Early implementations demonstrate that explainability tools can be integrated into Spark-based scoring pipelines without significant performance degradation. This trend represents a convergence of regulatory requirements and technical capability, making explainable AI the highest priority area for both research and industry investment.

Privacy-preserving federated learning addresses both the data generalizability challenge and regulatory privacy requirements. Initial implementations demonstrate promise in enabling collaborative fraud detection across institutions without requiring the sharing of data. However, significant practical barriers limit near-term viability in banking environments. Communication overhead between participating institutions creates latency bottlenecks that contradict real-time fraud detection requirements, with synchronization delays typically exceeding 5-10 seconds per model update cycle. Regulatory frameworks lack clear guidance on cross-institutional model sharing, creating legal uncertainty that risk-averse banking institutions avoid. Technical challenges include ensuring model convergence when participant institutions have vastly different data distributions and transaction volumes, with smaller banks contributing insufficient data to meaningfully influence model training while larger institutions risk diluting their competitive advantages through knowledge sharing. This approach offers the most realistic path toward improving industry-wide fraud detection.

The deployment of lightweight fraud detection models directly onto payment devices represents a paradigm shift toward prevention rather than detection. Recent advances in edge-based neural network implementations demonstrate the feasibility of deploying sophisticated classification models on resource-constrained devices [36]. Early implementations in payment card chips demonstrate feasibility for simple rule-based models, but extending this approach to sophisticated machine learning remains technically challenging. Success in this area could eliminate the latency constraints that currently limit real-time fraud prevention. IoT-enabled devices with embedded processing capabilities are increasingly being integrated into financial transaction infrastructure, enabling distributed fraud detection at the point of transaction [37].

V. Conclusions and Future Trends

This review reveals significant gaps between BDA theoretical capabilities and practical fraud detection implementations, with regulatory compliance requirements often superseding technical performance considerations in real-world deployments. The absence of standardised evaluation protocols makes performance claims across studies essentially meaningless, preventing evidence-based technology adoption and impeding scientific progress in the field.

The paper demonstrates that current BDA implementations face fundamental trade-offs between detection accuracy and regulatory compliance, real-time response and analytical depth, and system performance and integration complexity. These gaps can be addressed with both technical advances and coordinated efforts from each sector to align the requirements.

A. Critical Research Gaps

The field's most pressing limitation is the absence of standard benchmarks and evaluation protocols. Without standardised datasets and metrics, the study's performance claims remain unverifiable, preventing practitioners from making informed technology selections and researchers from building upon previous work. Academic research relies heavily on synthetic datasets that fail to capture the complexity of the production environment. The gap between laboratory performance and real-world effectiveness remains largely unmeasured, creating substantial implementation risk for financial institutions. Current regulatory frameworks fundamentally conflict with the machine learning techniques that achieve the highest fraud detection performance, forcing banks to accept suboptimal technical solutions to maintain compliance. Critical research gaps remain largely unaddressed in current literature. Adversarial attacks on federated fraud detection systems represent a significant vulnerability, where malicious participants could poison shared models through carefully crafted fraudulent transaction data, potentially compromising fraud detection capabilities across entire banking consortiums. The literature lacks comprehensive analysis of adversarial robustness in distributed fraud detection architectures, despite the high-stakes nature of financial applications. Additionally, future quantum computing capabilities will threaten current cryptographic frameworks used in BDA systems, requiring proactive research into quantum-resistant encryption methods, as quantum computers capable of breaking current encryption are expected to emerge around 2030. The intersection of privacy-preserving techniques and adversarial robustness remains unexplored, creating potential security vulnerabilities in systems designed to protect customer privacy while maintaining fraud detection effectiveness.

B. Future Research Trending

The field requires coordinated development of standard benchmark datasets, evaluation metrics, and testing protocols. This foundational work must precede advances in algorithms or architectures, as current research lacks the evaluation infrastructure necessary for meaningful progress. Federated learning approaches offer the most promising path toward addressing data generalizability challenges while satisfying regulatory privacy requirements. Research should focus on communication efficiency, security against adversarial participants, and performance validation across heterogeneous banking environments. Developing explainable AI techniques that maintain detection performance while providing regulatory-compliant transparency represents a crucial bridge between technical capability and compliance requirements.

Based on current technological trajectories and regulatory developments, several testable hypotheses emerge for future research validation. First, hybrid edge-cloud architectures will achieve sub-50ms fraud detection latency while maintaining 90-93% accuracy within three years as specialized AI processors become cost-effective for payment infrastructure.. Second, quantum-resistant encryption methods will become mandatory for financial institutions by 2028, requiring complete redesign of current BDA security frameworks. Third, regulatory convergence toward explainable AI requirements will drive adoption of inherently interpretable models, potentially limiting deep learning applications in fraud detection despite superior performance. Fourth, cross-border regulatory harmonization will enable federated learning implementations among major financial institutions by 2030, though initial deployments will likely focus on low-risk applications such as synthetic identity detection rather than real-time transaction scoring

C. Practical Recommendation

Banks should prioritise explainability from system design inception. Institutions should evaluate BDA platforms based on integration complexity and regulatory compliance requirements rather than relying solely on performance metrics, given the current absence of reliable performance benchmarks and the primacy of regulatory constraints. Financial institutions should collaborate on developing shared evaluation standards and anonymised benchmark datasets before investing in proprietary algorithm development, as the evaluation infrastructure represents a greater barrier to progress than algorithmic limitations.

Successful BDA implementation requires a systematic four-phase approach over twelve months. Phase 1 (months 1-3) involves comprehensive gap analysis between current capabilities and BDA requirements, assessing infrastructure compatibility and regulatory constraints. Phase 2 (months 4-6) deploys proof-of-concept implementations using Spark

Streaming with transaction-only data, establishing baseline metrics while developing compliant explainability mechanisms. Phase 3 (months 7-9) addresses integration challenges through robust data quality pipelines and parallel running with existing systems, validating performance against production data distributions. Phase 4 (months 10-12) implements gradual production deployment, beginning with low-risk transactions while establishing continuous monitoring and feedback loops for iterative refinement.

Acknowledgment

We want to thank peer reviewers for helping to review the paper.

Funding Statement

There is no funding agencies supporting the research work.

Author contributions

Kian Yee Jane Yam: Writing – Original Draft Preparation, Review & Editing;

Shi Hui Khoo: Writing – Original Draft Preparation, Review & Editing;

Kai-Xuan See Adam: Writing – Original Draft Preparation, Review & Editing;

Kai Liang Lew: Project Administration, Validation, Writing – Review & Editing;

Chean Khim Toa: Project Administration, Supervision, Writing – Review & Editing.

Conflict of interests

No conflict of interests were disclosed.

Ethics statements

Ethical approval was not applicable to this research since it did not involve human participants, animals, or sensitive data.

References

[1] J. Gee and M. Button, “The Financial Cost of Fraud 2018: The Latest Data from Around the World”, *Crowe Clark Whitehill*, 2018.
URL: https://www.crowe.com/uk/croweuk/insights/financial-cost-of-fraud-2018 (Accessed 6 August 2025)

[2] ACFE, “Occupational Fraud 2024: A Report to the Nations,” *Association of Certified Fraud Examiners (ACFE)*, 2024.
URL:https://legacy.acfe.com/report-to-the-nations/2024/ (Accessed 6 August 2025)

[3] Kaspersky, “Financial Cyberthreats Report 2024,” *SecureList by Kaspersky*, 2024.
URL:https://securelist.com/financial-threat-report-2024/115966/ (Accessed 6 August 2025)

[4] UK Finance, “Annual Fraud Report 2025,” *UK Finance Limited trading*, 2025.
URL:https://www.ukfinance.org.uk/policy-and-guidance/reports-and-publications/annual-fraud-report-2025 (Accessed 8 August 2025)

[5] A. Ali, S.A. Razak, S.H. Othman, T.A.E. Eisa, A. Al-Dhaqm, M. Nasser, T. Elhassan, H. Elshafie and A. Saif, "Financial Fraud Detection Based on Machine Learning: A Systematic Literature Review," *Applied Sciences*, vol. 12, no. 19, pp. 9637, 2022.
DOI: https://doi.org/10.3390/app12199637

[6] M.U. Hassan, M.H. Rehmani and J. Chen, "Anomaly Detection in Blockchain Networks: A Comprehensive Survey," *IEEE Communications Surveys & Tutorials*, vol. 25, no. 1, pp. 289-318, 2023.
DOI: https://doi.org/10.1109/COMST.2022.3205643

[7] C.P. Chen and C. Zhang, "Data-intensive applications, challenges, techniques and technologies: A survey on Big Data," *Information Sciences*, vol. 275, pp. 314-347, 2014.
DOI: https://doi.org/10.1016/j.ins.2014.01.015

[8] R. Vinayakumar, M. Alazab, K.P. Soman, P. Poornachandran, A. Al-Nemrat and S. Venkatraman, "Deep Learning Approach for Intelligent Intrusion Detection System," *IEEE Access*, vol. 7, pp. 41525-41550, 2019.
DOI: https://doi.org/10.1109/ACCESS.2019.2895334

[9] Y. Sahin and E. Duman, "Detecting credit card fraud by ANN and logistic regression," *2011 International Symposium on Innovations in Intelligent Systems and Applications*, pp. 315-319, 2011.
DOI: https://doi.org/10.1109/INISTA.2011.5946108

[10] F. Almalki and M. Masud, "Financial Fraud Detection Using Explainable AI and Stacking Ensemble Methods," *arXiv*, 2025.
DOI: https://doi.org/10.48550/arXiv.2505.10050

[11] R. P., S.V. E., C. Anilkumar, P. Thilakaveni and U. Moorthy, "Big Data Analytics and Implementation Challenges of Machine Learning in Big data," *Applied and Computational Engineering*, vol. 2, no. 1, pp. 484-489, 2023.
DOI: https://doi.org/10.54254/2755-2721/2/20220584

[12] A. Madhavi and T. Sivaramireddy, "Real-Time Credit Card Fraud Detection Using Spark Framework," *Algorithms for Intelligent Systems*, pp. 287-298, 2021.
DOI: https://doi.org/10.1007/978-981-33-4046-6_28

[13] S. Fritz-Morgenthal, B. Hein and J. Papenbrock, "Financial Risk Management and Explainable, Trustworthy, Responsible AI," *Frontiers in Artificial Intelligence*, vol. 5, 2022.
DOI: https://doi.org/10.3389/frai.2022.779799

[14] J. Dean and S. Ghemawat, "MapReduce," *Communications of the ACM*, vol. 51, no. 1, pp. 107-113, 2008.
DOI: https://doi.org/10.1145/1327452.1327492

[15] K. Shvachko, H. Kuang, S. Radia and R. Chansler, "The Hadoop Distributed File System," *2010 IEEE 26th Symposium on Mass Storage Systems and Technologies (MSST)*, pp. 1-10, 2010.
DOI: https://doi.org/10.1109/MSST.2010.5496972

[16] T. White, “Hadoop: the definitive guide, Fourth edition,” *O’Reilly*, 2015.

[17] A. Ghoting, R. Krishnamurthy, E. Pednault, B. Reinwald, V. Sindhwani, S. Tatikonda, Y. Tian and S. Vaithyanathan, "SystemML: Declarative machine learning on MapReduce," *2011 IEEE 27th International Conference on Data Engineering*, 2011.
DOI: https://doi.org/10.1109/ICDE.2011.5767930

[18] M. Zareapoor and P. Shamsolmoali, "Application of Credit Card Fraud Detection: Based on Bagging Ensemble Classifier," *Procedia Computer Science*, vol. 48, pp. 679-685, 2015.
DOI: https://doi.org/10.1016/j.procs.2015.04.201

[19] N. Malini and M. Pushpa, "Analysis on credit card fraud identification techniques based on KNN and outlier detection," *2017 Third International Conference on Advances in Electrical, Electronics, Information, Communication and Bio-Informatics (AEEICB)*, pp. 255-258, 2017.
DOI: https://doi.org/10.1109/AEEICB.2017.7972424

[20] S. Panigrahi, A. Kundu, S. Sural and A. Majumdar, "Credit card fraud detection: A fusion approach using Dempster–Shafer theory and Bayesian learning," *Information Fusion*, vol. 10, no. 4, pp. 354-363, 2009.
DOI: https://doi.org/10.1016/j.inffus.2008.04.001

[21] Suman and D. Kumar, “Performance Analysis of Various Credit Card Fraud Detection Approaches: A Review,” *International Journal of Advanced Research in Science and Engineering*, vol. 5, no. 9, pp. 120–126, 2016.
URL:https://www.ijarse.com/images/fullpdf/1473254973_506ijarse.pdf

[22] M. Zaharia, M. Chowdhury, M. J. Franklin, S. Shenker and I. Stoica, “Spark: cluster computing with working sets,” *Proceedings of the 2nd USENIX Conference on Hot Topics in Cloud Computing*, p. 10, 2010.
URL: https://dl.acm.org/doi/10.5555/1863103.1863113

[23] M. Zaharia, T. Das, H. Li, T. Hunter, S. Shenker and I. Stoica, "Discretized streams," *Proceedings of the Twenty-Fourth ACM Symposium on Operating Systems Principles*, pp. 423-438, 2013.
DOI: https://doi.org/10.1145/2517349.2522737

[24] K. Dabas, A. Dubey, and "Real-Time Fraud Detection using Apache Kafka, Apache Spark, and PySpark MLlib," *International Journal Of Scientific Research In Engineering And Management*, vol. 9, no. 05, pp. 1-9, 2025.
DOI: https://doi.org/10.55041/IJSREM47661

[25] X. Meng, J. Bradley, B. Yavuz, E. Sparks, S. Venkataraman, D. Liu, J. Freeman, D. Tsai, M. Amde, S. Owen, D. Xin, R. Xin, M.J. Franklin, R. Zadeh, M. Zaharia and A. Talwalkar, "MLlib: Machine Learning in Apache Spark," *arXiv*, 2015.
DOI: https://doi.org/10.48550/arXiv.1505.06807

[26] T. Chen and C. Guestrin, "XGBoost," *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 785-794, 2016.
DOI: https://doi.org/10.1145/2939672.2939785

[27] M. Almansoori and M. Telek, "Anomaly Detection using combination of Autoencoder and Isolation Forest," *1st Workshop on Intelligent Infocommunication Networks, Systems and Services*, pp. 25-30, 2023.
DOI: https://doi.org/10.3311/WINS2023-005

[28] R.D. Baruah and P. Angelov, "DEC: Dynamically Evolving Clustering and Its Application to Structure Identification of Evolving Fuzzy Models," *IEEE Transactions on Cybernetics*, vol. 44, no. 9, pp. 1619-1631, 2014.
DOI: https://doi.org/10.1109/TCYB.2013.2291234

[29] R.D. Baruah and P. Angelov, "Evolving local means method for clustering of streaming data," *2012 IEEE International Conference on Fuzzy Systems*, pp. 1-8, 2012.
DOI: https://doi.org/10.1109/FUZZ-IEEE.2012.6251366

[30] L.H. Aros, L.X.B. Molano, F. Gutierrez-Portela, J.J.M. Hernandez and M.S.R. Barrero, "Financial fraud detection through the application of machine learning techniques: a literature review," *Humanities and Social Sciences Communications*, vol. 11, no. 1, 2024.
DOI: https://doi.org/10.1057/s41599-024-03606-0

[31] P. Sundaravadivel, R.A. Isaac, D. Elangovan, D. KrishnaRaj, V.V.L. Rahul and R. Raja, "Optimizing credit card fraud detection with random forests and SMOTE," *Scientific Reports*, vol. 15, no. 1, 2025.
DOI: https://doi.org/10.1038/s41598-025-00873-y

[32] L. Breiman, "Random Forests," *Machine Learning*, vol. 45, no. 1, pp. 5-32, 2001.
DOI: https://doi.org/10.1023/A:1010933404324

[33] N.S. Halvaiee and M.K. Akbari, "A novel model for credit card fraud detection using Artificial Immune Systems," *Applied Soft Computing*, vol. 24, pp. 40-49, 2014.
DOI: https://doi.org/10.1016/j.asoc.2014.06.042

[34] M. Ester, H. -P. Kriegel, and X. Xu, “A Density-Based Algorithm for Discovering Clusters in Large Spatial Databases with Noise,” *Second International Conference on Knowledge Discovery and Data Mining*, pp. 226–231, 1996,.
URL: https://cdn.aaai.org/KDD/1996/KDD96-037.pdf

[35] N. Carneiro, G. Figueira and M. Costa, "A data mining based system for credit-card fraud detection in e-tail," *Decision Support Systems*, vol. 95, pp. 91-101, 2017.
DOI: https://doi.org/10.1016/j.dss.2017.01.002

[36] J. Pei, J. Han, B. Mortazavi-Asl, J. Wang, H. Pinto, Q. Chen, U. Dayal and M. Hsu, "Mining sequential patterns by pattern-growth: the PrefixSpan approach," *IEEE Transactions on Knowledge and Data Engineering*, vol. 16, no. 11, pp. 1424-1440, 2004.
DOI: https://doi.org/10.1109/TKDE.2004.77

[37] B. Krawczyk and M. Woźniak, "Diversity measures for one-class classifier ensembles," *Neurocomputing*, vol. 126, pp. 36-44, 2014.
DOI: https://doi.org/10.1016/j.neucom.2013.01.053

[38] T. Kohonen, "The self-organizing map," *Neurocomputing*, vol. 21, no. 1-3, pp. 1-6, 1998.
DOI: https://doi.org/10.1016/S0925-2312(98)00030-7

[39] R.A. Leite, T. Gschwandtner, S. Miksch, S. Kriglstein, M. Pohl, E. Gstrein and J. Kuntner, "EVA: Visual Analytics to Identify Fraudulent Events," *IEEE Transactions on Visualization and Computer Graphics*, vol. 24, no. 1, pp. 330-339, 2018.
DOI: https://doi.org/10.1109/TVCG.2017.2744758

[40] S. Kaisler, F. Armour, J.A. Espinosa and W. Money, "Big Data: Issues and Challenges Moving Forward," *2013 46th Hawaii International Conference on System Sciences*, 2013.
DOI: https://doi.org/10.1109/HICSS.2013.645

[41] P. Jain, M. Gyanchandani and N. Khare, "Big data privacy: a technological perspective and review," *Journal of Big Data*, vol. 3, no. 1, 2016.
DOI: https://doi.org/10.1186/s40537-016-0059-y

[42] U. Sivarajah, M.M. Kamal, Z. Irani and V. Weerakkody, "Critical analysis of Big Data challenges and analytical methods," *Journal of Business Research*, vol. 70, pp. 263-286, 2017.
DOI: https://doi.org/10.1016/j.jbusres.2016.08.001

[43] E. Hormozi, M.K. Akbari, H. Hormozi and M.S. Javan, "Accuracy evaluation of a credit card fraud detection system on Hadoop MapReduce," *The 5th Conference on Information and Knowledge Technology*, pp. 35-39, 2013.
DOI: https://doi.org/10.1109/IKT.2013.6620034

[44] J. Liu, S. Zhang and H. Fan, "A two-stage hybrid credit risk prediction model based on XGBoost and graph-based deep neural network," *Expert Systems with Applications*, vol. 195, pp. 116624, 2022.
DOI: https://doi.org/10.1016/j.eswa.2022.116624

[45] T.T. Nguyen, H. Tahir, M. Abdelrazek and A. Babar, "Deep Learning Methods for Credit Card Fraud Detection," *arXiv*, 2020.
DOI: https://doi.org/10.48550/arXiv.2012.03754

[46] I.D. Mienye and T.G. Swart, “A Hybrid Deep Learning Approach with Generative Adversarial Network for Credit Card Fraud Detection,” *Technologies*, vol. 12, no. 10, p. 186, 2024.
DOI: https://doi.org/10.3390/technologies12100186