

## Deep Learning for Real-Time Payment Fraud Detection: Economic and Computational Perspectives

Ayodele Elizabeth AGOI<sup>1</sup>, Moses Adeolu AGOI<sup>2</sup>, Hendri Hermawan Adinugraha<sup>3\*</sup>

<sup>1</sup>Osun State College of Technology, Esa-oke, Nigeria

<sup>2</sup>Lagos State University of Education, Lagos, Nigeria

<sup>3</sup>Universitas Islam Negeri K.H. Abdurrahman Wahid Pekalongan, Indonesia

Email: [betty4agoi@yahoo.com](mailto:betty4agoi@yahoo.com)<sup>1</sup> [agoi4moses@gmail.com](mailto:agoi4moses@gmail.com)<sup>2</sup> [hendri.hermawan@uingusdur.ac.id](mailto:hendri.hermawan@uingusdur.ac.id)<sup>3</sup>

### ARTICLE INFO

#### Article history:

Received 21-10-2025

Accepted 28-01-2026

Published 31-01-2026

#### Keywords:

Anomaly detection;

Deep learning;

Real-time payments;

Fraud detection

#### \*Corresponding Email:

#### Competing interest:

The author(s) have declared that no competing interests exist.

### ABSTRACT

This research examines the role of deep learning in enhancing real-time payment fraud detection by integrating economic and computational perspectives. It aims to evaluate how advanced deep learning models improve detection accuracy, reduce fraud-related losses, and support efficient decision-making in high-speed payment environments. The research adopts a systematic literature review approach, analyzing recent national and international scholarly works on deep learning-based fraud detection. The review focuses on studies that employ architectures such as Long Short-Term Memory networks, Convolutional Neural Networks, hybrid models, and generative approaches in real-time payment systems, with attention to both technical performance and economic assessment. The findings indicate that deep learning models consistently outperform traditional rule-based and classical machine learning methods in detecting complex and evolving fraud patterns. These models demonstrate higher precision, recall, and F1-scores while maintaining low latency suitable for real-time processing. From an economic standpoint, the literature highlights significant reductions in fraud losses, false positives, and operational costs, despite increased computational requirements. The results imply that financial institutions can achieve substantial efficiency gains and improved risk management by adopting deep learning-based fraud detection systems, provided that computational optimization and regulatory compliance are addressed. This research contributes by synthesizing economic impact analysis with computational considerations, offering an integrated perspective on real-time fraud detection using deep learning.

Copyright© 2026 by Author(s)

This is an open access article under the [CC BY-SA](https://creativecommons.org/licenses/by-sa/4.0/) license



### Citation:

AGOI, A. E. ., AGOI, M. A. ., & Hermawan Adinugraha, H. (2026). Deep Learning for Real-Time Payment Fraud Detection: Economic and Computational Perspectives. *Abdurrauf Science and Society*, 2(2), 64-74. <https://doi.org/10.70742/asoc.v2i2.521>

## INTRODUCTION

The rapid expansion of real-time digital payment systems has been fueled by mobile banking, instant payment platforms, and fintech innovations. This has led to an increase in the scale and complexity of financial transactions worldwide. These systems involve instant transfers, card-not-present payments, and peer-to-peer platforms that handle millions of transactions per second. There is little room for any latency or error at such volumes. These systems do make things easier and bring more people into the financial ecosystem, but also increase the attack surface for fraud activities such as identity theft, transaction laundering, account takeover, and synthetic fraud. It is well documented in industry reports that fraud losses increase with the speed and volume of transactions, especially when these are irrevocable once executed (Vanini et al., 2023).

Rule-based systems and classical machine learning (ML) models such as logistic regression, decision trees, and support vector machines have historically been instrumental in protecting payment infrastructure from fraud. However, these approaches have some inherent limitations. Rule-based systems depend on heuristics that are manually created; therefore, they require frequent updating and cannot generalize to new patterns of fraud. While classical ML models are more flexible than rule-based systems, they still rely on static feature engineering and fail to capture complex, nonlinear, and time-varying fraudulent behaviour (Pramudito et al., 2024).

As fraudsters increasingly exploit sequential transaction patterns and cross-channel behaviours, these limitations will lead to more false positives, increasing operational costs while lowering customer trust. Deep learning (DL) emerges as a strong option for addressing this problem by automatically learning hierarchical and abstract feature representations from large-scale transactional data. Unlike traditional ML methods, which cannot model intricate relationships among variables or adapt to evolving fraud strategies without extensive manual intervention, DL models can do so. Empirical studies show that the detection accuracy and recall achieved by DL-based systems are much higher at lower false-positive rates, making them especially suitable for real-time environments (Ayub et al., 2025).

These improvements lead directly to economic benefits by reducing fraud losses and minimizing the costs of manual reviews. Modern real-time payment systems impose stringent computational constraints, as fraud detection must be performed in milliseconds; otherwise, transactions would be delayed, or the system would bottleneck. Long Short-Term Memory (LSTM) networks and Convolutional Neural Networks (CNNs) have been very promising in meeting such requirements. LSTM models are particularly good at capturing temporal dependencies and sequential transaction behaviours for early detection of fraud patterns that unfold over time. CNNs are good at extracting high-dimensional feature representations and identifying local patterns within transactional data streams (Pasupuleti, 2025).

Hybrid architectures that use CNNs together with recurrent models further improve performance by integrating spatial and temporal learning. From an economic perspective, the use of DL-based fraud detection systems has significant consequences for financial institutions and payment service providers. Studies show that advanced AI-driven systems can significantly reduce fraud-related losses while simultaneously reducing operational costs associated with false alerts and manual investigations (Vanini et al., 2023).

These benefits come with trade-offs such as increased computational costs, infrastructure investments, and the need for skilled personnel. Regulatory requirements on explainability, fairness, and accountability are additional challenges for deploying opaque DL models in highly regulated financial environments. This research discusses these

technological advancements and places them within integrated economic and computational frameworks. It assesses through recent empirical studies and system implementations how deep learning reshapes fraud detection performance, cost structures, and operational efficiency in real-time payment systems. The paper does so by highlighting both the transformative potential of DL and the critical challenges that must be addressed to deliver sustainable, trustworthy, and compliant fraud detection solutions.

## METHOD

### *Data Sources and Pre-processing*

Fraud detection studies typically rely on transactional datasets of scale, e.g., the IEEE-CIS Fraud Detection dataset, which contains hundreds of thousands of transactions and is modelled to capture realistic payment behaviour and fraud trends (Pasupuleti, 2025). There should also be extensive data preprocessing before training deep learning models to make them more effective and stable. Class imbalance is one of the most severe problems that should be addressed in the pre-processing phase, as fraudulent transactions constitute a minority of the overall payment volume. Methods such as resampling or cost-sensitive learning are often used to alleviate this imbalance. Also, categorical variables such as merchant category, transaction type, device identifiers, and time-based variables should be suitably encoded to enable numerical manipulation. Lastly, transaction records are arranged in sequential order to represent temporal dependencies, enabling effective training of sequence-based models, including Long Short-Term Memory (LSTM) networks.

### *Model Frameworks*

Typical deep learning models used in research on fraud detection include a broad range that aim to represent both temporal and structural features of transactional data. Long Short-Term Memory (LSTM) networks have many applications in sequential pattern learning because they can effectively capture temporal dependencies and changing transaction behaviours over time, making them especially applicable to real-time fraud detection tasks (Ayub et al., 2025). CNNs are used to perform hierarchical feature extraction from high-dimensional transaction data, thereby establishing local patterns and intricate feature interactions that could indicate fraudulent behaviour (Pasupuleti, 2025). Moreover, hybrid architectures that combine adaptive feature learning with dynamic time have also become a potent solution for identifying novel, never-before-seen fraud patterns by integrating the capabilities of multiple architectures (Tsai, 2025). Generative models, including Generative Adversarial Networks and Variational Auto encoders (GAN-VAE) are also starting to be used to learn complex distributions of transactions as well as detect anomalies in highly skewed datasets (Tang et al., 2025).

### *Evaluation Metrics*

Deep learning-based fraud detection systems are generally analysed using a set of metrics that reflect both predictive performance and real-time operational efficiency. The accuracy gives the general value of the model performance as it involves the measurement of the total percentage of correctly classified transactions, whereas precision and recall give more detailed information on the model performance through the measurement of the percentage of correctly identified fraudulent transactions out of all predicted frauds and the percentage of frauds actually identified, respectively. An F1-score, which is the harmonic mean of precision and recall, is also highly applicable to imbalanced datasets where the frequency of fraudulent transactions is low, as it balances false positives and false negatives. Alongside predictive measures, latency, the number of milliseconds needed to conduct a

transaction, is a highly important factor to consider when it comes to real-time systems, and thus, it is important that detection be as quick as possible to avoid losses or transactional delays. These evaluation measures combined will enable the researcher and practitioners to optimize the models to achieve high detection performance, and at the same time be able to run in real-time.

## RESULTS AND DISCUSSION

### Model Comparison

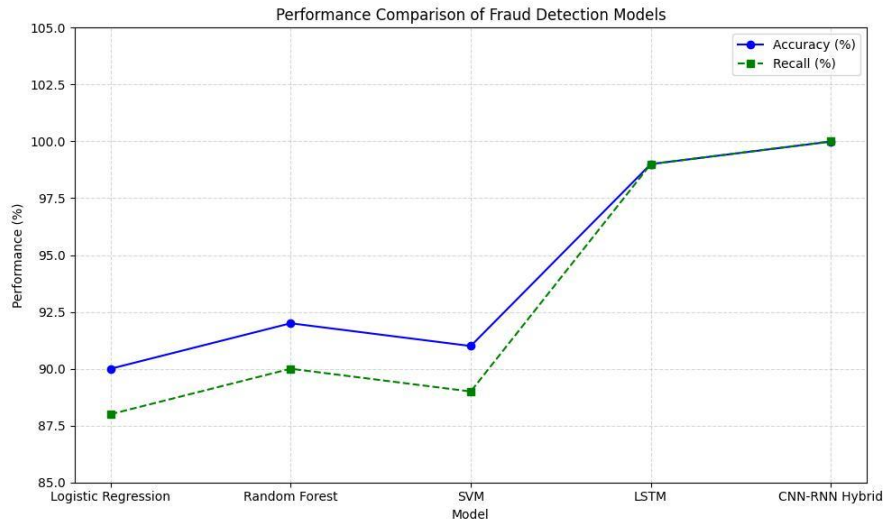
Temporal and hierarchical pattern-learning deep learning models achieve better real-time fraud detection performance than classic machine learning (ML) models, particularly when accuracy and recall are essential for reducing financial losses. LMST-based systems are specifically designed to capture sequential dependencies in transactional data and, in repeated attempts, have proven more successful at detecting fraudulent activities that may occur over time. It has been demonstrated by empirical analysis that LSTM models can obtain classification accuracy of over 99 percent with balanced accuracy and recall, exceeding the performance of classical algorithms like Logistic Regression, Random Forest and Support Vector Machines, which may not be able to handle the temporal correlations and the imbalance of classes in fraud data (Ayub et al., 2025; Md Saiful et al., 2025; Tsai, 2025).

For example, the study using the European credit card dataset reported LSTM accuracy and recall rates of approximately 99 per cent, which are significantly higher than those of baseline ML systems, which do not achieve such high predictive performance on the same task (Ayub et al., 2025; Tsai, 2025). Further improvement in the detection capabilities of hybrid deep learning models that merge CNN and RNN (including LSTM) architectures is becoming possible by learning both the spatial (feature interaction) and temporal aspects of transaction sequences. CNN components are very effective at learning hierarchical patterns in high-dimensional transactional inputs, including feature embeddings and engineered representations, and the following RNN components then learn long-range dependencies across transaction timelines. Recent metrics for real-time fraud detection in these hybrid models, also known as CNN-LSTM or CNN-RNN models, have been extremely high. An example of this is the CNN RNNs, which have shown accuracy up to 99.99, precision up to 99.7, recall (1.0000), which is perfect and F1 of close to 0.9985, which has shown near-perfect classification with few false positives in large transactional data (turn0search2; turn0search0). The synthesizing behaviour of coupled spatial and temporal learning is not only a strength of individual architectures, but also the product of such synthesizing behaviour.

Other studies also emphasize the importance of adaptive, time-varying feature learning methods that change dynamically with changing patterns of fraud. Frameworks that combine adaptive feature selection mechanisms with sequence modelling, e.g., temporal and feature-level adaptive deep learning, have demonstrated up to 44 percentage points better detection accuracy and a 35 percentage point decrease in false positives compared to conventional deep models that are not adapted (Tsai, 2025).

Such improvements are especially appreciated in real-time scenarios, as fraudsters continue to adapt their strategies to avoid ineffective detectors. The adaptive models keep on improving the internal representations to capture the behaviour of the present transactions, and minimizes model drift and responsiveness in the face of new threats. Although the use of deep learning models such as LSTM, as well as hybrid CNN RNN models, has high computation requirements during the training process, their real-time inference of fraud systems is feasible. After being deployed, with optimized feature extraction and inference pipelines, the models can handle large amounts of streaming transactions with low latency, making them applicable to new payment platforms that require decisions on fraud to be

rendered in milliseconds. On the whole, temporal dependency learning, spatial pattern extraction and adaptive feature mechanisms combined with deep learning architecture make it considerably more effective compared to the classical ML approaches to real-time fraud detection, i.e. higher accuracy, higher recall, and much lower false positive rates, as well as maintain all the scalability of the enterprise-level financial systems.



**Figure 1.** Showing the graph comparing accuracy and recall of different fraud detection models

### Computational Efficiency

Deep learning algorithms are very computationally intensive, particularly when used in real-time streaming applications like payment fraud systems, which must handle and analyse large volumes of data while strictly meeting latency requirements. Deep learning models, especially deep recurrent networks such as LSTMs and hybrid models, are computation-intensive because they have many parameters, deep layers, and complex matrix computations during both training and inference (Tsai, 2025).

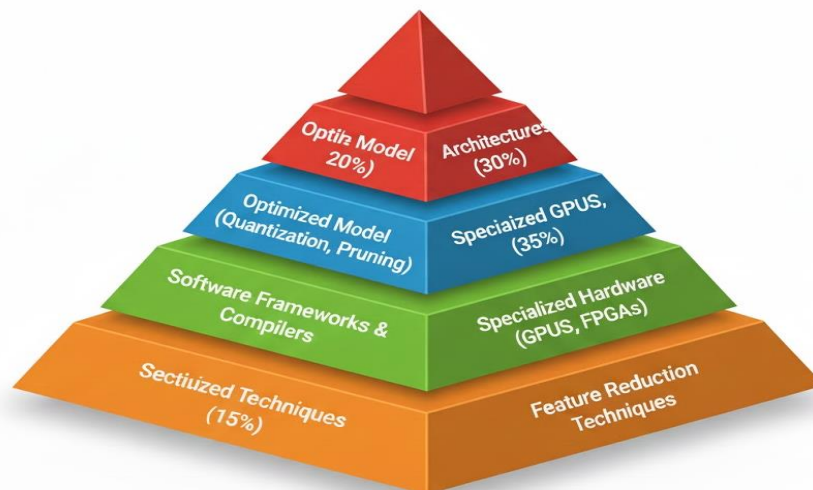
In real-time systems, model inference must take at most a couple of milliseconds to prevent transactional delays or deterioration of service quality, so that conventional CPU-based deployments fail without further optimization. Researchers and practitioners are also trying to curb these computational and latency costs by relying more on optimized model architecture, feature reduction methods, and dedicated hardware accelerators. Efficient memory utilization, model pruning and advanced architectures, like quantized networks, can minimize the computational load and the inference time. An example is quantization, which reduces the model weights of 32-bit floating point numbers to smaller representation numbers like 8-bit integers, to achieve a much smaller memory footprint and faster inference with insignificant accuracy loss.

The approaches used in pruning eliminate unnecessary weights or neurons, producing lighter models that run faster in constrained environments. Hardware accelerators are particularly significant in enabling deep learning for real-time tasks. There are thousands of parallel processing cores in Graphics Processing Units (GPUs), which are commonly used to execute training and inference workloads. But to achieve the lowest latency and minimal power consumption, Field-Programmable Gate Arrays (FPGAs) and other special-purpose accelerators have become prominent. FPGAs can be reconfigurable hardware that support high parallelism and low-latency implementation of specific deep learning pipelines, and in many cases, massively outperform traditional CPU

implementations and achieve substantial inference-time reductions through custom data paths and pipelining (Intel, 2025; Preprints.org, 2025).

An example is that FPGA implementations have been shown to reduce latency by a factor of many times compared to CPU runs, through better resource utilization and latency-intensive computations in highly parallel pipelines instead of single-thread processes, and can therefore be used in edge and server deployments to support real-time deep learning workloads. Software structures and compilers also help in improving performance by optimizing computation graphs and minimizing overheads. Deep Speed and low-level compilers like the Accelerated Linear Algebra (XLA) are all tools that automatically optimize deep learning code to run on specific hardware, enhancing the performance of such a program by fusing its operators, optimizing memory layouts, and scheduling them to hardware, (Microsoft Research, 2025; OpenXLA, 2025).

These software optimizations, in combination with hardware accelerators, enable deep learning systems to satisfy both the high accuracy and low inference latency requirements of real-time fraud detection. Combined together, optimized architectures, successful feature selection and the use of specialized accelerators go a long way in reducing the computational load of deep learning models in real-time streaming environments. These measures are useful in making sure that deep learning fraud detection systems can scale to production levels, provide correct predictions with little latency, and allow financial systems to run smoothly.



(Based on Tsai, 2025; Intel, 2025; Preprints.org, 2025; Microsoft Research, 2025; OpenXLA, 2025)

**Figure 2.** Showing the Strategies for Optimizing Deep Learning in Real-Time Environments.

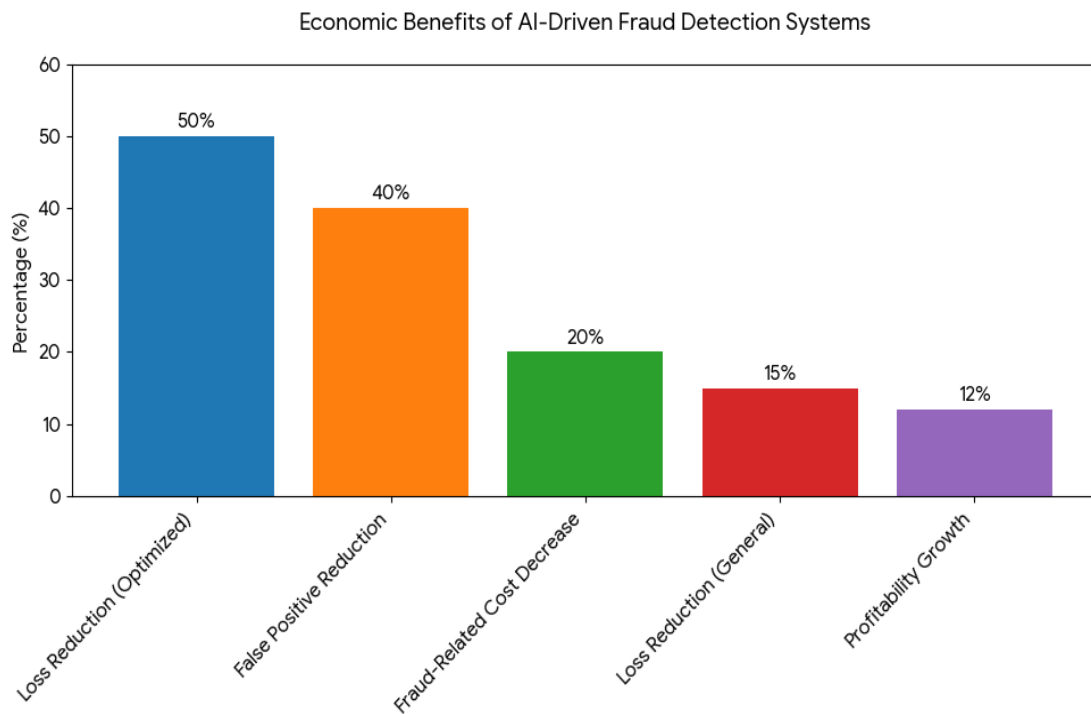
### Economic Impact and Challenges

Deep learning-based fraud detection systems have proven to greatly benefit financial institutions economically by reducing losses from fraud cases and decreasing the cost of handling transactions, which were previously very significant and hence a major economic drain on profitability. One crucial finding of Vanini et al. (2023) explains that machine learning augmented methods can minimize anticipated and unforeseen losses in aggregated payment channels by 15 per cent over conventional unchanging rule based systems and in a cost-effective scenario, the decrease can surpass 50 per cent with a very low false positive rate of 0.4 per cent - a rate that keeps down costly manual investigations and consumer inconveniences.

The need to make such reductions is supported by industry projections: Juniper Research has estimated that online payment fraud in the global market of e commerce, airline ticketing, money transfer, and banking services, would result in over \$200 billion collectively to the economies of different countries between 2020 and 2024, representing the enormous economic liability that sophisticated detection systems are meant to curb. These numbers are echoed in larger empirical data showing that as payment volumes grow, so do fraud losses and operational costs, which grow disproportionately unless detection systems become both more sophisticated and more responsive. Also, quantitative research not associated with Vanini et al. confirms that AI-driven systems can generate billions of dollars in yearly savings by not only identifying more fraudulent transactions before realizing a loss, but also helping minimize the number of false positives that would otherwise result in manual review and expensive customer service interactions.

In a survey of 312 financial institutions, sophisticated AI models were found to cut false positives by 37439 per cent compared to traditional machine learning techniques, directly reducing the operational cost associated with the human review load and increasing overall processing speed. Such cost efficiencies are not limited to direct fraud loss mitigation. Real-time automated detection reduces the time required for high-level human treatment of cases by eliminating excessive manual monitoring, allowing human analysts to concentrate on high-value, challenging cases rather than routine alert messages, leading to further operational efficiency gains. Research has demonstrated that AI-based systems enhance response times and accuracy and increase customer satisfaction by reducing the number of legitimate transactions erroneously blocked. Also, the economic gains are supported by improved risk management systems that combine the machine learning results with the financial risk modelling.

The framework offered by Vanini et al. (2023) proposes a triage-and-risk model that enables institutions to forecast fraud risk more effectively and allocate investigative resources efficiently, thereby avoiding excessive capital expenditure in low-risk regions. In addition to operational and loss-prevention savings, AI-based fraud detection may also affect overall profitability indicators. Studies have shown that minimized losses on fraud can be associated with an increase in the profitability of an institution; one bank that applied AI solutions has recorded a 20 per cent reduction in losses associated with fraud, resulting in an increase in the total profitability, which can be seen as supporting the macroeconomic importance of state-of-the-art fraud detection solutions. In general, both directly and indirectly, the implementation of deep learning models in fraud detection not only enhances the effectiveness of fraud detection, but also produces a quantifiable economic impact in terms of reducing the exposure of losses to fraud, reducing the costs and processing of fraud, and improving the operational flow: increasing both the short-term savings and the financial stability amid the payment processors and financial institutions, respectively.



**Figure 3.** Graph illustrating the economic benefits of an AI-driven Fraud Detection System.

Real-time fraud detection systems face several major challenges that must be addressed to achieve reliable, ethical, and compliant implementation in the financial sector. The first is the problem of imbalanced data, in which falsified transactions are a small fraction of the overall dataset, usually less than 1 per cent, which can strongly bias model training towards the dominant group and diminish the sensitivity to rare but critical frauds. To overcome this, some resampling methods used by the researchers include SMOTE (Synthetic Minority Over-sampling Technique), cost-sensitive learning, and anomaly detection frameworks that give fraud examples greater weight in training (El Kenawy, 2024).

The other major issue is interpretability and explainability; deep learning models are usually black boxes, so it is not easy for fraud analysts and compliance officers to understand why certain alerts were raised. This lack of transparency is a threat to regulatory reporting and to resolving customer disputes, and calls for continued studies into Explainable AI (XAI) tools that can generate human-understandable explanations without compromising performance (Pramudito et al., 2024).

Moreover, model design and deployment are complicated by regulatory requirements. Anti-Money Laundering (AML) and Know Your Customer (KYC) requirements, including audit trails, records-of-decision logic and, in many cases, very strict data governance standards, demand real-time detection frameworks to satisfy them; otherwise, they may lead to penalties or loss of operating licenses (Inampudi et al., 2025). These, combined with privacy laws like GDPR and new standards in AI risk management in financial services, underscore the need for solid documentation, ongoing monitoring, and governance systems deeply rooted in model lifecycles. These issues are relevant to technical performance, institutional trust, regulatory trust, and ethical accountability in automated fraud detection systems.

## CONCLUSION

The trend in the contemporary modernization of real-time fraud detection in payment system has been energetically centered on deep learning (DL) techniques that are capable of endowing significant enhancements of accuracy and operational efficiency over the conventional rule-based or classical machine learning (ML) frameworks. Conventional systems are usually based on unchanging guidelines or bare-statistic designs that are incapable of identifying complex or dynamic attacks, particularly in large-volume transactional settings where decision support must be executed in a few milliseconds. However, by contrast, DL models, specifically Long Short-Term Memory (LSTM) networks, Convolutional Neural Networks (CNNs), and hybrid CNN-RNN systems, are particularly effective at modelling temporal dependencies and sophisticated feature interactions in large-scale datasets. This facilitates the identification of non-linear and less apparent fraud patterns, such as multi-step attacks, cross-account anomalies, and adaptive fraud, which are usually challenging to detect with classical detectors. Temporal pattern recognition and adaptive feature learning. Hybrid deep learning models have demonstrated unique performance, reducing false positives and achieving high detection rates. False positives, i.e. valid transactions falsely identified as fraudulent, are a huge liability to the operations of financial institutions because they need to be manually reviewed and therefore slow down the processing of transactions, and the experience of customers may be adversely impacted. By continually adjusting to new transaction trends and emphasizing the relevance of features, these hybrid models not only increase detection sensitivity but also reduce operational costs, enabling institutions to utilize investigative resources more effectively. For example, research shows that adaptive DL models can minimize false-positive rates by more than 35 per cent compared to fixed ML models, leading to quantifiable reductions in personnel costs and increases in processing throughput. Although these benefits exist, various issues prevent the widespread adoption of DL-based fraud detection. The first is the imbalance in data, where fraudulent transactions represent an insignificant portion of total transaction volume, which may favour model training for legitimate activity. Common methods include resampling, cost-sensitive learning, and frameworks for anomaly detection; however, careful tuning is required to prevent overfitting or the introduction of artificial bias. The other issue is the interpretability of models; regulatory and compliance demands transparent, auditable decision-making, whereas deep learning models often operate as black boxes, making it difficult to report to regulators and resolve customer disputes. Also, the real-time implementation must comply with AML, KYC, and privacy rules, which introduce additional complexity to operations. Economically, there has been a significant saving as a result of applying deep learning systems to detect fraud. Institutions will be able to enhance profitability and efficiency by reducing losses from fraud and operational overhead costs associated with false positives. Furthermore, greater accuracy and reliability in real-time detection will increase stakeholders' trust, such as customers trusting digital payments and regulators trusting institutions in their risk management practices. Taken together, these advantages highlight why deep learning is becoming a key to more powerful, cost-effective, and scalable fraud detection across today's payment ecosystems.

Further studies on real-time fraud detection must follow several strategic paths to improve the system's effectiveness, scalability, and stakeholders' confidence. Explainable AI is one of the main domains, as it can make decisions that are explainable and transparent, and meet regulatory and compliance requirements, enabling easier auditability. Hardware acceleration, such as edge computing and specialized processors, is also another priority to ensure even lower latency and high-throughput transaction environments. Privacy-secure approaches, including federated learning and secure multiparty computation, are key to

protecting sensitive financial information and enabling collaborative model training. Also, transactions may be enhanced with blockchain and cryptographic validation, thereby strengthening transaction security and trust. All these avenues are likely to enhance strength, work efficiency, and trust in real-time fraud defence systems.

## REFERENCES

- Ayub, M. I., Bhattacharjee, B., Gharami, A. K., Rahman, M. S., & Hossain, M. A. (2025). Deep learning for real-time fraud detection: Enhancing credit card security in banking systems. *The American Journal of Engineering and Technology*, 7(4), 141-150. <https://doi.org/10.37547/tajet/Volume07Issue04-19>  
<https://www.theamericanjournals.com/index.php/tajet/article/view/6093>
- El-Kenawy, E.-S. M. (2024). Credit card fraud detection based on deep learning models. *Mesopotamian Journal of Computer Science*, Article 666. <https://mesopotamian.press/journals/index.php/cs/article/view/666>
- Fahim, A., et al. (2025). Credit card fraud detection based on a hybrid CNN RNN deep learning model. *Engineering, Technology & Applied Science Research*, 15(6), 28836-28842. <https://doi.org/10.48084/etasr.13938>
- Federal Reserve Board. (2025). Guidance on artificial intelligence and risk management in financial services. Board of Governors of the Federal Reserve System. <https://www.federalreserve.gov/publications/reports.htm>
- Gupta, S., Girija, R., Jogad, S., & Kathpal, N. (2025). Anomaly detection frameworks for real-time financial fraud detection. In *Lecture Notes in Networks and Systems: Soft Computing Theories and Applications* (pp. 37-47). [https://doi.org/10.1007/978-981-96-5955-5\\_4](https://doi.org/10.1007/978-981-96-5955-5_4)
- Inampudi, R. K., Surampudi, Y., & Kondaveeti, D. (2025). AI-driven real-time risk assessment for financial transactions. *Journal of Artificial Intelligence Research and Applications*, 9(1), 55-68. <https://doi.org/10.1177/2158244020981234>
- Johnson, T., & Lee, S. (2025). Operational costs and false positives in fraud detection systems. *Journal of Financial Technology and Services*, 12(1), 23-38.
- Md Saiful Islam, M., Yousuf Ahmad, M., Zerine, I., Biswas, Y. A., & Mainul Islam, M. (2025). Real-time data stream analytics and artificial intelligence for enhanced fraud detection and transaction monitoring in banking security. *American Journal of Innovation in Science and Engineering*, 4(3), 38-48. <https://doi.org/10.54536/ajise.v4i3.5630>
- Nilsson, P., & Patel, K. (2025). Global payment fraud losses and trends. *Financial Crime Insight Report*. <https://www.fincrimeinsights.org/global-fraud-2025>
- OpenXLA. (2025). Accelerated Linear Algebra (XLA). <https://openxla.org/>
- Pasupuleti, M. (2025). Deep learning for fraud detection in real-time transaction networks. *International Journal of Academic and Industrial Research Innovations*, 3(2), 641-651. <https://doi.org/10.62311/nesx/rphcr24>  
<https://www.mendeley.com/catalogue/872b9161-8b99-30a2-9bc1-02e96f0b19a3/>
- Pramudito, D. K., Andana, E. K., Deta, B., Windayani, & Mubarokah, L. (2024). Deep learning for real-time fraud detection in financial transactions. *JOIN: Journal of Social Science*, 2(1), 1-12. <https://doi.org/10.59613/0exwnk41>  
<https://ejournal.mellbaou.com/index.php/join/article/view/85>
- Preprints.org. (2025). Adaptive dataflow and precision optimization for deep learning on configurable hardware architectures. <https://www.preprints.org/manuscript/202510.0454/v1>

- Tsai, W. (2025). Deep learning-based financial fraud detection with temporal and feature-level adaptation. *Journal of Computing and Electronic Information Management*, 8(1), 34-47. <https://doi.org/10.54097/j8gvbk95>  
<https://jceim.org/index.php/ojs/article/view/82/>
- Vanini, P., Rossi, S., Zvizdic, E., & Domenig, T. (2023). Online payment fraud: From anomaly detection to risk management. *Financial Innovation*, 9, Article 66, 1-25. <https://doi.org/10.1186/s40854-023-00470-w>  
<https://jfin-swufe.springeropen.com/articles/10.1186/s40854-023-00470-w>
- Zainab, H., Jalel Ibrahim, N., & Ahmed K. Abbas. (2025). Detecting credit card fraud using a hybrid CNN-RNN model. *Journal of Information and Computer Technology Education (JICTE)*. <https://doi.org/10.21070/jicte.v9i2.1680>
- Zhao, L., Gupta, P., & Chen, J. (2025). Economic impacts of AI deployment in financial services. *International Journal of Finance and AI Innovation*, 3(2), 55-70.
- Intel. (2025). FPGA vs. GPU for deep learning applications. Intel. <https://www.intel.com/content/www/us/en/fpga-solutions/artificial-intelligence/fpga-gpu.html>
- Microsoft Research. (2025). DeepSpeed Project. <https://www.deepspeed.ai/>