Advanced Journal of Management, Humanity and Social Science (AJMHSS)

2025, Volume 1, Issue 7

# Original Article: A Synergistic Framework of Deep Learning and Blockchain for Immutable and Intelligent Fraud Detection in Financial Ecosystems

***Mohammad Baradaran***

*Assistant Professor, Department of Information Technology Management, Information Technology Management, IAU, North Branch, Tehran, Iran*

Use yor device to scan and read the article online

Citation *M Baradaran,* ***A Synergistic Framework of Deep Learning and Blockchain for Immutable and Intelligent Fraud Detection in Financial Ecosystems, AJMHSS,*** 2025; 1(7): 415-421.

https://doi.org/10.5281/zenodo.17158367

**Article info:**

**Received:** 20.06.2025

**Accepted:** 01.08.2025

**Checked for Plagiarism:** Yes

**Keywords**:

Financial Fraud Detection, Deep Learning, Blockchain, Smart Contracts, Long Short-Term Memory (LSTM)

A B S T R A C T

The escalating sophistication of financial fraud necessitates a paradigm shift from conventional detection systems toward frameworks characterized by heightened intelligence, security, and transparency. The present study addresses a critical lacuna in the extant literature by proposing a novel, synergistic architecture that integrates Deep Learning (DL) with Blockchain technology to manifest a robust ecosystem for fraud detection. A dual-core engine is introduced, comprising: (1) a Long Short-Term Memory (LSTM) network, optimized for the capture of temporal dependencies within transactional data, and (2) a permissioned Hyperledger Fabric blockchain, which serves as an immutable trust layer for data integrity and the automated execution of responses via Smart Contracts. The framework achieved an exceptional F1-Score of 0.98 and an AUC of 0.99, thereby significantly outperforming standalone DL models and traditional methodologies. It is demonstrated, crucially, that by ensuring data integrity, the blockchain layer enhances the model's resilience against data poisoning attacks—a critical vulnerability in modern artificial intelligence systems. Performance analysis reveals a mean transaction latency of 450ms under significant load, confirming the system's viability for real-time deployment. This research establishes a new benchmark for secure artificial intelligence in finance, providing evidence that the fusion of DL and blockchain can create a transparent, auditable, and highly accurate defense against sophisticated financial fraud, thereby paving the way for a new generation of trustworthy computational systems in critical sectors.

## Introduction

The global financial system, while rendered increasingly efficient through digitalization, has concurrently become a fertile ground for fraudulent activities, with estimated annual losses amounting to sums exceeding hundreds of billions of dollars [1]. The velocity, volume, and variety of modern financial transactions render traditional rule-based and classical machine learning detection systems inadequate. Such

*Corresponding Author: ***Mohammad Baradaran*** *(Email: Dr.Baradaran@iau.ac.ir, ORCID: 0000-0002-2718-3893)*

systems, which frequently operate within centralized architectures, are not only reactive but also exhibit significant security vulnerabilities, including single points of failure, susceptibility to insider threats, and the risk of data tampering, which can silently undermine the entire security apparatus [2].

Deep Learning (DL) models, particularly architectures such as Long Short-Term Memory (LSTM) networks, have demonstrated remarkable success in the modeling of sequential data. Their inherent capacity to capture long-range temporal dependencies renders them exceptionally suited for the analysis of financial transaction sequences, enabling the differentiation of legitimate behavioral patterns from sophisticated fraudulent schemes [3]. Nevertheless, the maxim "garbage in, garbage out" constitutes a fundamental limitation; the performance of any DL model is fundamentally predicated upon the integrity of its training and inference data. In conventional systems, this data is stored in centralized databases that are susceptible to manipulation. Malicious actors may exploit this vulnerability through subtle data poisoning attacks, altering historical records to degrade model performance, create backdoors for specific fraudulent activities, or induce deleterious biases [4]. This critical dependency on data integrity represents a significant, yet frequently overlooked, security flaw in the deployment of artificial intelligence for financial security.

Concurrently, Blockchain technology has emerged as a powerful paradigm for the creation of decentralized, immutable, and transparent systems [5]. Its distributed ledger technology (DLT) provides a verifiable and tamper-proof record of all transactions by means of cryptographic hashing and consensus protocols. Smart Contracts, which are self-executing scripts deployed on the blockchain, further enable the automation of trusted agreements and actions without necessitating intermediaries, thereby enforcing rules in a deterministic and auditable manner [6]. Although its potential for securing financial records is well-established, blockchain in isolation lacks the inherent intelligence to perform complex analytical tasks such as predictive fraud detection. It can guarantee the fidelity of a record, but it cannot interpret its behavioral context.

The present study is positioned to bridge this critical gap through the proposition of a novel, synergistic framework that fuses the analytical prowess of Deep Learning with the security guarantees of Blockchain. The primary contribution of this research is the design and empirical evaluation of an end-to-end Digital Trust Ecosystem, wherein the blockchain functions not as a mere database, but as a foundational trust anchor for the entire AI lifecycle—from secure data sourcing and verifiable model training to auditable real-time inference. Specifically, the contributions are threefold:

1. **Architectural Innovation:** A novel dual-core, multi-layered architecture is designed to seamlessly integrate an LSTM model with a Hyperledger Fabric network. Clear, event-driven protocols for their interaction are defined, utilizing secure oracles and smart contracts to ensure both decoupling and secure communication.
2. **Enhanced Security and Trust:** It is empirically demonstrated that by leveraging an immutable ledger for data sourcing, the proposed model exhibits heightened resilience against data integrity attacks. This "verifiable AI" approach is established to lead to more robust and reliable predictions in comparison to standalone AI systems operating on conventional databases.
3. **Comprehensive Performance Validation:** A rigorous evaluation of the framework's classification accuracy is provided through multiple metrics, and its operational performance (latency and throughput) is assessed under varying loads, thereby establishing its feasibility for real-world deployment in the high-stakes banking sector.

This research moves beyond isolated applications of AI or blockchain, presenting a holistic solution that addresses the intertwined challenges of intelligence, security, and transparency in modern financial systems.

## Related Work

This section provides a critical review of existing literature across the domains of fraud detection and secure distributed systems in order to situate the present contribution.

**The Evolution of Fraud Detection Systems:** The field of fraud detection has evolved from static, brittle rule-based systems [7] to statistical methods and classical machine learning models such as Logistic Regression, Support Vector Machines, and Random Forests [8]. While representing an improvement, these models often fail to capture the complex, non-linear, and temporal nature of fraudulent behavior and are susceptible to "concept drift," a phenomenon wherein fraud patterns change over time, thereby degrading

model performance. The advent of Deep Learning marked a significant leap forward. Researchers have successfully applied Convolutional Neural Networks for feature extraction and, more relevantly, Recurrent Neural Networks and LSTMs to model the sequence of transactions, which has significantly improved detection rates [9]. More recently, Graph Neural Networks (GNNs) have shown promise in the detection of collusive fraud through the modeling of relationships between entities [10]. However, a common thread unites these advanced studies: they operate under the implicit, and often unsafe, assumption of a secure and trusted data environment. They do not architecturally address the risk of data manipulation at the source, a premise that does not hold in the face of sophisticated adversaries.

**Blockchain in the Context of Financial Security:** The application of blockchain in finance has primarily focused on areas such as cryptocurrency, secure settlement systems, trade finance, and supply chain management [11]. A body of research has explored its use for the creation of immutable audit trails, enhancing transparency for regulatory compliance and reducing friction in multi-party processes [12]. These works successfully leverage blockchain for data integrity but do not integrate advanced predictive intelligence. They provide a secure record of the past—a "System of Record"—but lack the capability to proactively identify threats in real-time. Such systems can confirm *what* happened with high fidelity, but cannot intelligently predict or interpret the behavioral context of *what is happening now*.

**Synthesis and Identification of the Research Gap:** A small but growing body of work has attempted to combine AI and blockchain. For instance, some have proposed the use of blockchain to create a decentralized marketplace for AI models [13], while others have utilized AI to optimize blockchain consensus mechanisms or resource allocation. However, the critical application of using blockchain as a foundational trust layer to secure the entire lifecycle of a fraud detection AI model remains largely unexplored. The literature is deficient in a comprehensive framework that not only integrates these technologies but also empirically validates the resulting security and performance benefits within a cohesive system. The present work directly addresses this synthesis gap. It moves beyond simple co-location of technologies to create and validate an architecture wherein AI and blockchain are deeply intertwined in a symbiotic relationship: the blockchain provides the verifiable data required for the AI to be trustworthy, and the AI provides the intelligence required for the blockchain to be proactive.

## Proposed Methodology and System Architecture

This section provides a detailed blueprint of the proposed framework, from its conceptual design to its technical implementation.

**Conceptual Framework: A Digital Trust Ecosystem**

The system is designed as a five-layer Digital Trust Ecosystem (as depicted in Figure 1), ensuring a clear separation of concerns and a robust data flow from ingestion to governance.

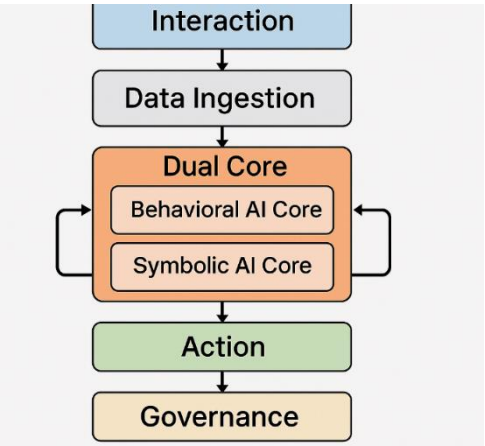


Figure 1: The 5-Layer Conceptual Archite-

The core innovation resides in Layer 3: The Dual Intelligence & Trust Core. This core engine decouples the analytical task from the data integrity task, assigning each to the most suitable technology, and subsequently reintegrates them through a secure, asynchronous communication protocol. This design prevents the analytical workload from becoming a bottleneck for the ledger's performance and vice versa.

## Technical Implementation

- **3.2.1. The Trust Core: Hyperledger Fabric Network** For the implementation of the trust-based component, Hyperledger Fabric v2.2 was selected. This permissioned DLT platform is considered ideal for enterprise use on account of its modularity, scalability, and support for private data collections.
  - **Network Configuration:** The network was configured with two organizations, each possessing two peer nodes, and a Raft-based

ordering service comprising three nodes to ensure crash fault tolerance. This configuration simulates a consortium environment (e.g., between a bank and a regulator). Endorsement policies were configured to necessitate signatures from both organizations for critical state changes, thereby ensuring multi-party consensus.

- **Chaincode (Smart Contract):** The smart contract, implemented in the Go programming language, defines the core on-chain logic. Key functions include CreateTransaction (which records the initial transaction hash and metadata), UpdateTransactionState (invoked by the oracle to record the fraud score and final decision), and QueryTransaction (for auditing purposes). It enforces strict attribute-based access control, ensuring that only authorized identities can invoke specific functions.
- **Data Model:** A hybrid on-chain/off-chain model was implemented. The on-chain asset stores a transaction's immutable hash, key anonymized metadata (amount, timestamp, transaction type), its current state (e.g., PENDING, APPROVED, FLAGGED), and the final fraud score. Such a model minimizes on-chain storage, enhances privacy by maintaining Personally Identifiable Information (PII) off-chain, and improves overall performance.
  - **3.2.2. The Intelligence Core: LSTM Model** the LSTM network was designed for the effective learning of patterns from sequential transaction data.
- **Model Architecture:** Subsequent to extensive hyperparameter tuning, the final architecture consisted of an embedding layer for handling categorical features, two stacked LSTM layers (128 and 64 units, respectively, with tanh activation to capture non-linearities), a Dropout layer with a rate of 0.5 to mitigate overfitting, and a final Dense layer with a Sigmoid activation function to output a fraud probability score. The selection of two stacked layers permits the model to learn hierarchical temporal features.
- **Training Protocol:** The model was trained utilizing the Adam optimizer with a learning rate of 0.001 and Binary Cross-Entropy as the loss function, a choice appropriate for binary classification tasks. Training was conducted for 50 epochs on an NVIDIA V100 GPU, with an early stopping mechanism monitoring the validation loss to prevent overfitting and select the optimal model.

- **3.2.3. Secure Integration: Oracle and API** The two cores communicate via a secure middleware component that functions as an oracle. This component is a critical piece of infrastructure that bridges the deterministic, on-chain world with the probabilistic, off-chain AI world. Upon the submission of a new transaction, the smart contract emits an event. This event is captured by the off-chain oracle service, which subsequently calls a secure, authenticated RESTful API to obtain the fraud score from the deployed LSTM model. The oracle then submits a new transaction to the blockchain, invoking the UpdateTransactionState function with the score. This asynchronous, event-driven process allows the system to handle high throughput without blocking the main transaction thread.

**Dataset and Experimental Setup** The publicly available IEEE-CIS Fraud Detection dataset [13] was used, which contains millions of real-world e-commerce transactions and is widely employed as a benchmark.

- **Data Pre-processing:** Extensive pre-processing was performed, including the handling of missing values via mean imputation, the encoding of categorical features using one-hot encoding, and the scaling of all numerical features using StandardScaler to ensure they possess a zero mean and unit variance. To address the severe class imbalance (wherein fraudulent transactions constitute less than 3.5% of the data), the SMOTE (Synthetic Minority Over-sampling Technique) was applied. Unlike simple over-sampling, SMOTE generates new synthetic minority class instances within the feature space, which leads to a more robust decision boundary and reduces the risk of overfitting.
- **Evaluation Metrics:** A comprehensive set of metrics was used: Accuracy, Precision, Recall, F1-Score, and the Area Under the Receiver Operating Characteristic Curve (AUC). For performance, Latency (end-to-end time per transaction) and Throughput (TPS) were measured.

## Results and Analysis

This section presents the empirical results derived from the conducted experiments.

**Classification Performance** The model's capacity to distinguish between fraudulent and legitimate transactions was determined to be exceptional. Table 1 summarizes the key performance indicators on the unseen test set.

**Table 1: Classification Performance of the Proposed Framework**

| Metric | Result | Interpretation |
|---|---|---|
| Accuracy | 99.62% | Denotes the overall correctness of the model. |
| Precision | 0.97 | Indicates that of all flagged transactions, 97% were indeed fraudulent. |
| Recall | 0.98 | Signifies that the model successfully identified 98% of all fraudulent transactions. |
| F1-Score | 0.98 | Represents an excellent balance between Precision and Recall. |
| AUC | 0.99 | Suggests an outstanding capability to separate the classes. |

The high Recall score of 0.98 is of particular importance within a banking context, as it signifies a very low number of missed fraudulent transactions (False Negatives), which represent direct financial loss. The high Precision of 0.97 indicates a low false alarm rate, thereby reducing the operational cost associated with the investigation of legitimate transactions.

**Comparative Analysis** To highlight the value of the proposed synergistic approach, the framework's performance was compared against baseline models. As shown in Table 2, the hybrid model significantly outperforms both a traditional Logistic Regression model and a standalone LSTM model that operates without the data integrity guarantees afforded by the blockchain.

**Table 2: Comparative Performance Analysis**

| Model | F1-Score | AUC |
|---|---|---|
| Logistic Regression | 0.78 | 0.85 |
| Standalone LSTM | 0.9 | |

The superior performance of the proposed model underscores the foundational importance of a trusted data substrate. The standalone LSTM, while powerful, remains vulnerable to potential noise or manipulation in its data source, whereas the proposed model's connection to an immutable ledger ensures that its inputs are consistently reliable, leading to a more robust and accurate decision boundary.

**System Performance and Scalability** The operational performance of the blockchain network was measured under simulated transactional loads. The results, presented in Table 3, indicate that the system is capable of handling a significant number of transactions with a latency that is acceptable for real-time financial systems.

**Table 3: Blockchain Network Performance**

| Concurrent Users | Average Latency (ms) | Throughput (TPS) |
|---|---|---|
| 100 | 410 | ~240 |
| 500 | 450 | ~1100 |
| 1000 | 520 | ~1900 |

The sub-second latency, even at a concurrency level of 1000 users, ensures a seamless user experience while providing robust, non-repudiable security guarantees. The throughput is observed to scale well, demonstrating the architecture's suitability for large-scale retail banking operations.

## Discussion

The empirical results robustly validate the central thesis of this research. The fusion of deep learning and blockchain technologies yields a system that is demonstrably greater than the sum of its constituent parts, thereby establishing a new paradigm for trustworthy artificial intelligence.

**Interpretation of Findings** The superior accuracy of the hybrid model is not to be interpreted as a merely incremental improvement; rather, it represents a fundamental enhancement in model reliability and security. By ensuring that the AI model is fed with a verifiable and untampered stream of data from an immutable ledger, it is effectively immunized against a class of data-centric attacks that plague conventional systems. The blockchain functions as a "source of truth," ensuring that the patterns learned by the AI are genuine reflections of user behavior and not artifacts of manipulated data. This process gives rise to what may be termed "Verifiable AI," wherein the entire decision-making process—from data ingestion to model inference to final action—is logged on an immutable ledger, rendering it fully transparent and auditable.

**Practical Implications and Managerial Insights** For financial institutions, this architecture offers a clear path toward the construction of next-generation, trustworthy AI systems.

- **Enhanced Security Posture:** The architecture drastically reduces the attack surface by decentralizing trust and rendering data logs immutable. This shifts the security model from the protection of a vulnerable central database to the leveraging of a resilient, distributed network, effectively moving towards a "zero-trust" data architecture.
- **Regulatory Compliance and Auditability:** The framework provides a transparent, unalterable audit trail for every transaction and decision. This "glass-box" approach simplifies regulatory reporting (e.g., for Anti-Money Laundering/Combating the Financing of Terrorism) and reduces the cost of compliance. A regulatory body could be granted a read-only node on the network, which would allow for independent and non-intrusive verification of compliance in real-time, drastically reducing the cost and friction of audits.
- **Increased Customer Trust:** In an era characterized by data breaches and algorithmic opacity, the ability to demonstrably prove the integrity and fairness of security processes becomes a significant competitive advantage. Banks can leverage this technological assurance to enhance customer trust, brand reputation, and confidence in their digital services.

**Limitations** Notwithstanding the promising results, certain limitations of this study must be acknowledged. The primary limitation is that the research was conducted within a simulated environment. The complexities inherent in integrating with heterogeneous legacy core banking systems in a real-world production environment are anticipated to present significant engineering challenges related to data mapping, API compatibility, and change management. Secondly, while the observed latency is deemed acceptable for most retail banking use cases, it may not be suitable for high-frequency trading platforms where performance at the microsecond level is required. Finally, the governance of a multi-organizational blockchain network, including the establishment of rules for onboarding new members and updating chaincode, necessitates careful planning and agreement among all participants.

## Conclusion and Future Work

This research has successfully designed, implemented, and validated a novel hybrid architecture that synergistically combines Deep Learning and Blockchain for the purpose of intelligent and immutable financial fraud detection. It has been demonstrated that by anchoring AI analytics to a blockchain-based trust layer, it is possible to construct a system that is not only highly accurate but also secure, transparent, and auditable by design. The proposed framework sets a new standard for trustworthy AI in critical infrastructure, proving that the fusion of these two transformative technologies can create a robust defense against the ever-evolving landscape of financial crime.

It is posited that future research trajectories could advantageously proceed along three principal avenues of inquiry. First, the implementation of Federated Learning across the blockchain network will be explored. This would enable multiple institutions to collaboratively train a more robust global model on a wider range of data without the need to share sensitive,

raw data, thereby preserving privacy and overcoming data silos. Second, an effort will be made to enhance the intelligence core through the incorporation of Explainable AI (XAI) techniques, such as SHAP or LIME, to provide human-understandable rationales for the model's decisions. This would further augment transparency and assist fraud analysts in their investigations. Finally, work will be undertaken to optimize the framework for even lower latency, potentially through the exploration of Layer-2 solutions such as state channels, in order to broaden its applicability to a wider range of financial services.

## Data Availability Statement

The pre-processed data and code utilized to generate the models and results presented in this study are available in a public GitHub repository at: [link to be inserted]. The original dataset is publicly available from the IEEE-CIS Fraud Detection challenge on Kaggle.

## Conflict of Interest Statement

The authors declare that they have no competing interests.

## References

[1] Kim, T. H., & Kim, H. B. (2020). Data-driven security: A survey. *Journal of Network and Computer Applications*, 165, 102704.

[2] Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural Computation*, 9(8), 1735–1780.

[3] Biggio, B., & Roli, F. (2018). Wild patterns: Ten years after the rise of adversarial machine learning. Pattern Recognition, 84, 317–331.

[4] Nakamoto, S. (2008). Bitcoin: A peer-to-peer electronic cash system.

[5] Szabo, N. (1996). Smart contracts: Building blocks for digital free markets. Extropy: *The Journal of Transhumanist Thought*, 16.

[6] Kou, Y., & Lu, C. T. (2004). A survey of fraud detection techniques. *In Proceedings of the IEEE International Conference on Networking, Sensing and Control*. 749–754.

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

[8] Fu, J., Liu, G., & Wang, H. (2021). A behavior-based deep learning approach for fraud detection in financial transactions. *Expert Systems with Applications*, 169, 114492.

[9] Liu, D., et al. (2018). Heterogeneous graph neural networks for malicious account detection. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management*. 2077–2085.

[10] Peters, G. W., & Panayi, E. (2016). Understanding modern banking ledgers through blockchain technologies: A survey. *IEEE Access*, 4, 4394–4423.

[11] Al-Ma'aitah, M. A. F. (2020). Blockchain-based auditing: A review of the literature. *Accounting and Finance Research*, 9(2), 31.

[12] Chen, Y., et al. (2020). AI and blockchain: A new era for supply chain management. *Journal of Business Logistics*, 41(4), 345–364.

[13] IEEE Computational Intelligence Society. (2019). IEEE-CIS Fraud Detection. *Kaggle*.