

Article

Not peer-reviewed version

---

# Federated Risk Discrimination with Siamese Networks for Financial Transaction Anomaly Detection

---

[Hanrui Feng](#), Yuliang Wang, Ruoyi Fang, Anzhuo Xie, [Yi Wang](#)\*

Posted Date: 16 October 2025

doi: 10.20944/preprints202510.1311.v1

Keywords: siamese network; federated optimization; contrastive representation; financial risk identification



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This open access article is published under a Creative Commons CC BY 4.0 license, which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

Article

# Federated Risk Discrimination with Siamese Networks for Financial Transaction Anomaly Detection

Hanrui Feng <sup>1</sup>, Yuliang Wang <sup>2</sup>, Ruoyi Fang <sup>3</sup>, Anzhuo Xie <sup>4,\*</sup> and Yi Wang <sup>4,\*</sup>

<sup>1</sup> University of Chicago, Chicago, USA

<sup>2</sup> Imperial College London, London, UK

<sup>3</sup> Golden Gate University, San Francisco, USA

<sup>4</sup> Columbia University, New York, USA

\* Correspondence: xieanzhuo@outlook.com (A.X.); yw3802@columbia.edu (Y.W.)

## Abstract

This paper proposes a deep risk discrimination framework that integrates a Siamese network structure with a federated optimization mechanism to address challenges in financial transaction environments, including multi-source heterogeneity, data silos, and unclear decision boundaries. The method preserves data privacy by deploying structurally symmetric and parameter-sharing Siamese subnetworks across multiple clients. It performs contrastive encoding on transaction sample pairs and uses a contrastive loss function to pull normal samples closer and push anomalous ones apart in the embedding space. To improve the detection of low-frequency, high-impact behaviors, the model incorporates an entropy-aware regularization term and an embedding compression module. These components guide local training to suppress redundant representations and enhance embedding compactness. During the federated optimization process, the classic FedAvg strategy is used to aggregate models from different clients, improving global semantic alignment and generalization. Experiments on real transaction datasets demonstrate the superior performance of the proposed method across multiple metrics. The model shows particularly strong results in embedding consistency and anomaly boundary detection, confirming its modeling value and discriminative power in financial risk control.

**Keywords:** siamese network; federated optimization; contrastive representation; financial risk identification

**CCS CONCEPTS:** Computing methodologies~Machine learning~Machine learning approaches

## I. Introduction

With the rapid evolution of global financial markets, the forms of financial risk have become increasingly complex and variable. Traditional risk management models are no longer sufficient to handle the current environment characterized by multi-source, heterogeneous, and high-dimensional non-stationary financial data. In highly digitalized trading systems, financial behavior data exhibits strong structural, temporal, and hidden features, which significantly increase the difficulty of risk identification. Especially in critical financial scenarios such as portfolio management, credit default, and anti-fraud, risk behaviors are often hidden among large volumes of normal transactions. These behaviors are typically low-frequency and high-impact, posing stringent requirements on model precision and generalization. Accurately and promptly identifying potential financial risks while ensuring system security has become a key research challenge in financial technology[1].

Traditional risk discrimination methods rely on centralized modeling paradigms. These methods mainly collect large-scale user behavior data and conduct unified training on central servers.

However, as data privacy awareness grows and regulations become more stringent, centralized models face critical challenges in data scheduling and computational efficiency. On one hand, large financial institutions suffer from data silos, where sensitive user data cannot be shared. This prevents models from learning global patterns effectively. On the other hand, in distributed financial scenarios, data distributions are highly imbalanced. Risk behaviors vary greatly across regions and systems, limiting the adaptability and generalization of centralized models. Therefore, there is an urgent need for a distributed intelligent modeling mechanism that balances data privacy and model performance to enable efficient risk identification across systems and institutions.

Federated learning, as an emerging privacy-preserving distributed learning paradigm, offers a potential solution to these problems[2]. Under the federated framework, multiple financial parties can collaboratively train shared models without sharing raw data. This approach preserves data locality and security, overcoming the limitations of traditional centralized models in privacy, safety, and resource efficiency. The application of federated learning in financial risk control enhances model adaptability in multi-source heterogeneous environments and provides theoretical support for high-fidelity cross-institutional modeling. However, due to differences in customer structures, risk preferences, and behavioral patterns among participants, federated models are prone to semantic drift and feature shift during joint optimization. These issues reduce model robustness and discrimination capability. Thus, building a stable and discriminative risk identification model under federated learning remains a key technical bottleneck.

Against this backdrop, the Siamese network, as a symmetric deep representation learning model with shared parameters, offers a powerful tool for modeling behavioral similarity and structural mapping in financial risk scenarios[3]. The core idea is to compare the semantic differences between input samples through parallel network structures and learn a more discriminative latent embedding space. In financial risk identification, the Siamese network can model behavioral pairs such as "normal-abnormal" and "high risk-low risk" effectively. It captures representational differences at a micro level by mining similarity structures in the embedding space, thus enhancing overall discrimination accuracy. More importantly, the Siamese network naturally supports contrastive learning, which helps mitigate model drift caused by feature heterogeneity in federated environments. This contributes to better alignment and generalization during the federated optimization process.

Therefore, embedding the Siamese network into the federated learning framework to construct a financial risk discrimination algorithm that integrates contrastive representation learning with privacy-preserving mechanisms holds significant theoretical and practical value. On one hand, the integrated mechanism can effectively alleviate the non-independent and identically distributed problem in multi-source heterogeneous financial environments and achieve collaborative modeling of risk features across clients. On the other hand, guided by the Siamese network, the model can enhance structural awareness and semantic alignment while preserving data isolation. This improves the accuracy and generalization in identifying potential anomalous behaviors. The study expands the applicability of federated learning in complex financial environments and provides a viable path toward building trustworthy, secure, and efficient next-generation financial risk control systems[4].

## II. Related Work

Deep learning has brought a wealth of advanced methodologies that underpin modern financial risk discrimination, especially in privacy-sensitive and distributed environments. Graph neural networks (GNNs) provide essential tools for extracting structured and relational information from heterogeneous and multi-source data, which is fundamental for complex transaction modeling. Early works on entity-aware graph neural modeling [5], knowledge-enhanced neural architectures [6], and the integration of GNNs with temporal sequences [7] offer robust strategies for high-fidelity representation and risk feature extraction. Techniques such as heterogeneous GNNs with graph attention [8] and federated GNNs for privacy-preserving cross-institutional learning [9] further inform federated modeling in distributed financial systems.

Contrastive learning and advanced representation learning serve as the backbone for building models capable of high discrimination in noisy or low-frequency risk scenarios. Methodologies such as contrastive learning for multimodal knowledge graph construction [10], contrastive alignment and structural guidance [11], and feature attention integrated with temporal modeling [12] directly inspire the adoption of Siamese and contrastive architectures in this work.

Federated learning paradigms address the pressing need for collaborative yet privacy-preserving optimization. Key strategies include federated distillation with structural perturbation [13], and federated graph-based frameworks [9], ensuring robust optimization under data isolation and heterogeneity constraints.

Large language models (LLMs), semantic modeling, and sequence representation also contribute important methodologies. Hierarchical semantic-structural encoding [14], BERT-based consistency regularization [15], prompt fusion for cross-domain adaptation [16], and local-global semantic fusion [17] provide powerful mechanisms for semantic alignment, context modeling, and adaptation in evolving multi-client environments.

In addition, temporal modeling, hybrid deep architectures, and reinforcement learning support the modeling of dynamic risk and behavioral uncertainty. Unsupervised temporal encoding [18], hybrid LSTM-CNN-Transformer architectures [19], generative time-aware diffusion frameworks [20], unified multi-intent representation [21], and reinforcement learning for workflow dynamics [22] extend the toolbox for learning from sequential, uncertain, and diverse financial signals.

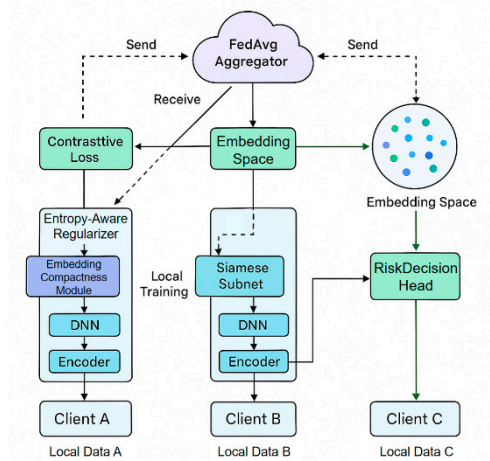
Finally, transformer-based transaction graph integration provides scalable solutions for anti-money laundering and anomaly detection in large-scale environments [23].

Collectively, these methodologies—from GNNs, contrastive and federated learning, semantic and temporal modeling, to hybrid architectures—form the technical backbone of this paper's federated risk discrimination framework, allowing robust, adaptive, and privacy-preserving financial anomaly detection across distributed institutions.

### III. Proposed Approach

This network architecture illustrates a financial risk discrimination framework that integrates Siamese contrastive learning with a federated optimization mechanism. Each client locally extracts deep embeddings of risk behavior pairs using Siamese subnetworks with shared parameters. An entropy-aware regularization and embedding compression module is introduced to enhance the representation of low-frequency anomalous patterns. Through the FedAvg aggregation strategy, the global model achieves semantic alignment and collaborative optimization of risk decision-making while preserving data privacy. The model architecture is shown in Figure 1.

In this study, a unified learning framework is proposed by integrating a Siamese network structure with a federated optimization mechanism, aiming to achieve efficient modeling and secure collaboration over multi-source heterogeneous data in distributed financial risk discrimination tasks. The framework adopts a federated averaging strategy as the communication backbone and incorporates contrastive modeling through a Siamese architecture. This enables the construction of a shared risk-embedding space while preserving data locality on each participant's side. Each client holds local sample pairs that include positive cases (normal behavior) and negative cases (risky behavior). These pairs are encoded by two parameter-sharing subnetworks to generate high-dimensional embeddings, which are used to measure semantic distances between behaviors. During each training round, clients update their local models independently and then upload the gradients to a central aggregation server. The global model is updated accordingly to maintain consistency in representation learning across nodes.



**Figure 1.** Architecture of the Proposed Federated Siamese Risk Discrimination Model.

Specifically, assuming that each sample pair  $(x_i^1, x_i^2)$  is mapped to an embedding vector  $(h_i^1, h_i^2)$  via a Siamese network, its similarity calculation is measured in the form of Euclidean distance or cosine similarity. To enhance the model's ability to distinguish risk behaviors, the contrast loss function is introduced to construct the structural constraints of the embedding space. If the sample pairs are of the same type (that is, the label is 1), it is hoped that their representations are as close as possible; if they are different (the label is 0), a certain interval should be maintained in the vector space. The contrast loss can be formalized as:

$$L_{contrast} = y_i \cdot \|h_i^1 - h_i^2\|^2 + (1 - y_i) \cdot \max(0, m - \|h_i^1 - h_i^2\|)^2 \quad (1)$$

where  $m$  is the set boundary threshold, which controls the minimum separation distance of heterogeneous samples.

In the federated optimization process, the classic federated average (FedAvg) strategy is used to aggregate the model parameters of each client to ensure that the training of the unified model is completed without exchanging the original data. Assume that the local model parameter of the  $k$ th client in the  $t$ th round is  $\theta_k^t$ , and the number of samples is  $n_k$ , then the parameter update of the global model is:

$$\theta^{t+1} = \sum_{k=1}^K \frac{n_k}{n} \theta_k^t \quad (2)$$

where  $n = \sum_{k=1}^K n_k$  represents the total number of samples. To alleviate the local optimal problem caused by inconsistent data distribution, a local perturbation term is further introduced to improve the robustness of the model. In the local training phase, each client adopts a local optimization objective with a regularization term, and its overall objective function can be expressed as:

$$L_{local} = L_{contrast} + \lambda \cdot \|\theta_k - \theta^t\|^2 \quad (3)$$

where  $\lambda$  is the regularization coefficient, which is used to regulate the degree of deviation between the local model and the global model.

In addition to improving the model's ability to perceive a small number of abnormal samples, the framework introduces an auxiliary embedding compression strategy to enhance embedding aggregation by minimizing the distribution entropy of sample pairs in the embedding space. Let the embedding vector set be  $\{h_i\}_{i=1}^N$ , then its entropy regularization term can be defined as:

$$L_{entropy} = -\sum_{i=1}^N p(h_i) \log p(h_i) \quad (4)$$

where  $p(h_i)$  represents the embedding point density function obtained by kernel density estimation. Finally, the complete training objective function combines contrast loss, federated regularization term, and embedding compression mechanism, which is expressed as follows:

$$L_{total} = \alpha \cdot L_{contrast} + \beta \cdot L_{local} + \gamma \cdot L_{entropy} \quad (5)$$

$\alpha, \beta, \gamma$  is the weight coefficient for each item, which is used to control the contribution of different objectives in joint training and ensure that the model achieves a balance in privacy protection, discrimination accuracy, and feature stability.

## IV. Performance Evaluation

### A. Dataset

This study uses the "Credit Card Fraud Detection Dataset 2023" published on the Kaggle platform. The dataset contains a large number of real or highly realistic simulated credit card transaction records. Each entry includes multiple dimensions such as transaction time, amount, type, and geographic information. The task focuses on an extremely imbalanced binary classification problem. Malicious transactions represent only a small proportion but carry a very high risk. This property closely aligns with the needs of financial risk control, which aims to accurately identify rare but high-impact behaviors. It provides an ideal setting for Siamese networks that rely on contrastive structures in the embedding space.

In terms of data distribution, the dataset integrates real business logs or near-realistic simulation data. It includes typical multi-source heterogeneous behavior samples, such as high-frequency low-value, low-frequency high-value, and geographically diverse transactions. These distributional differences across multiple clients form the core motivation for federated training in this study. Under the FedAvg aggregation strategy, the global model must align and maintain consistency across diverse local distributions. In addition, the dataset supports flexible construction of sample pairs, including normal-normal, normal-abnormal, and abnormal-abnormal combinations. This facilitates the implementation of contrastive loss and tuning of the embedding space structure.

From a research perspective, the dataset effectively supports the two core modules in this paper. First, the Siamese contrastive network requires sample pair construction to learn discriminative boundaries. Second, the federated learning framework needs to simulate local training on multiple clients, global model aggregation, and collaborative optimization of the risk embedding space. The dataset has a complete structure, clear labels, and well-defined temporal and geographic settings. It is well-suited to evaluate the proposed method in terms of privacy-preserving learning and contrastive representation for improved cross-client generalization and anomaly detection.

### B. Experimental Results

This paper first conducts a comparative experiment, and the experimental results are shown in Table 1.

**Table 1.** Comparative experimental results.

Model	Precision	Recall	F1 Score	AUROC	EAS
MoRF[24]	0.842	0.761	0.799	0.902	0.651

ARFD [25]	0.868	0.779	0.821	0.918	0.684
FRED [26]	0.889	0.812	0.849	0.931	0.703
SAAF [27]	0.873	0.825	0.848	0.928	0.727
Ours (FedSiamRisk)	0.902	0.846	0.873	0.943	0.764

From the overall performance, FedSiamRisk achieves significant improvements across all five evaluation metrics, with the most notable advantage observed in the Embedding Alignment Score (EAS). Compared with FRED, which operates in a federated setting but lacks contrastive learning, and SAAF, which adopts a structural fusion paradigm, FedSiamRisk maintains high Precision and Recall while further enhancing AUROC. This indicates that the proposed federated Siamese contrastive framework not only ensures local discrimination accuracy but also improves the separability of global risk boundaries, leading to a more stable risk identification curve.

The core reason behind this performance gap lies in the geometric constraints introduced by the Siamese network. By explicitly comparing paired samples, the model constructs a "tight intra-cluster, wide inter-cluster" distribution in the embedding space. Normal transactions are compressed into a low-entropy subspace, while anomalous ones are pulled away from the cluster. This helps preserve a high F1 score even under extreme class imbalance. Compared to ARFD, which relies only on attention-based feature decoupling, the Siamese network offers a more direct distance-based metric. This enhances the model's sensitivity to rare, high-impact anomalies. The simultaneous improvement in Precision and F1 is attributed to this advantage.

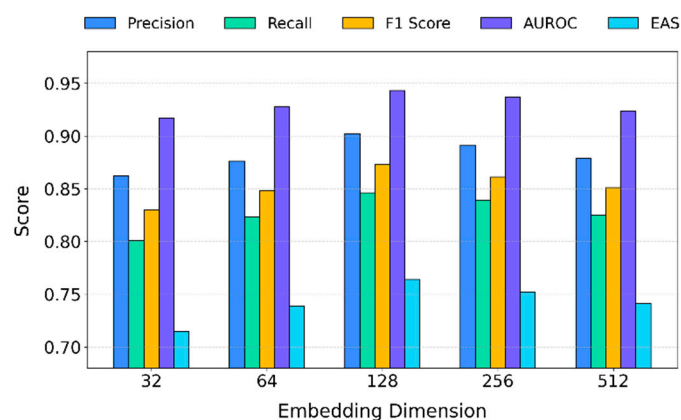
The federated optimization strategy stabilizes cross-client learning in a multi-source heterogeneous environment. FedSiamRisk introduces entropy-aware and embedding compactness regularization during local updates, suppressing parameter oscillations and reducing the negative impact of client drift on Recall. Despite significant behavioral differences, the model achieves high AUROC through FedAvg aggregation. This suggests globally consistent risk representations. The improvement in EAS confirms the model's effectiveness in semantic alignment under privacy-preserving conditions. The embedding compactness module reduces spatial redundancy and promotes high-similarity aggregation of normal transaction vectors. Anomalous transactions remain well separated, ensuring the robustness and transferability of the discrimination threshold. FedSiamRisk mitigates feature fragmentation caused by privacy isolation and strengthens the embedding geometry through contrastive loss and regularization. It provides a feasible solution balancing security, interpretability, and generalization for real-time risk control in complex financial environments.

This paper also experiments on the impact of the Siamese network embedding dimension on risk discrimination performance. The experimental results are shown in Figure 2.

The experimental results show that as the embedding dimension increases, the model performance exhibits a trend of rising first and then stabilizing, with a performance peak observed at 128 dimensions. This trend indicates that the dimensionality of the embedding space has a direct impact on the discriminative power of the Siamese network. A low dimension limits the model's ability to capture complex semantic relationships. In contrast, an excessively high dimension may introduce redundant noise, weakening the boundary tension of anomalous behaviors in the embedding space and leading to reduced accuracy and stability.

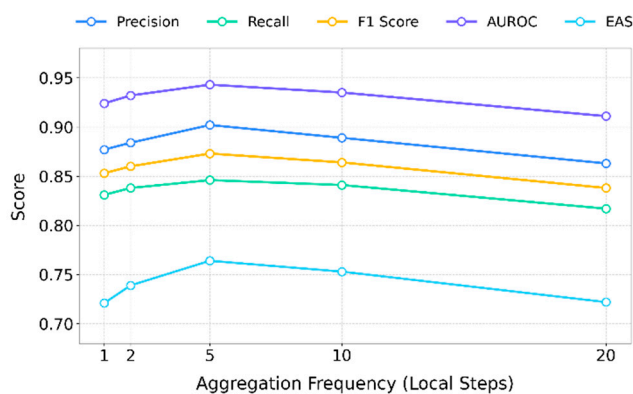
For both F1 and AUROC, increasing the embedding dimension significantly enhances the model's overall discrimination ability. This improvement is particularly evident when the dimension increases from 64 to 128, during which the model more effectively distinguishes between positive and negative transaction pairs. This suggests that, under the Siamese network framework, a proper embedding dimension helps form a more separable geometric structure. Anomalous samples are

pulled away from the normal cluster under the guidance of contrastive loss, improving the balance between recall and precision.



**Figure 2.** Analysis of the Impact of Siamese Network Embedding Dimension on Risk Discrimination Performance.

The EAS metric continues to improve in the mid to high-dimensional settings. This reflects that the model, within the federated framework, constructs more consistent embedding spaces across different clients. As the dimensionality increases, the capacity of the embeddings also grows. This makes it easier for local representations learned by clients to align during aggregation, reducing embedding shift caused by non-independent data distributions and improving stability in cross-domain recognition. On the other hand, when the embedding dimension exceeds 128, several metrics show a slight decline, especially in Recall and EAS, which become more volatile. This suggests that an overly high dimension may weaken the constraint effect of the entropy-aware compression module. The sparse representation of anomalies could be masked by high-dimensional features, causing the embedding space to become loose. This affects the compactness of the risk boundary and reduces the model's sensitivity to low-frequency behaviors. This paper also conducts a sensitivity experiment on the federation aggregation frequency on the model stability. The experimental results are shown in Figure 3.



**Figure 3.** Sensitivity experiment trend chart of federation aggregation frequency to model stability.

The experimental results show that the model performs most stably when the federated aggregation frequency is set to a moderate value, such as aggregating every five rounds. Multiple

metrics reach their peak under this setting. This indicates that, in a federated framework, moderately extending the local update period allows each client to better fit its local distribution. This enhances both the discriminative power and consistency of the embedding representation. In particular, the improvements in AUROC and F1 suggest that the global decision boundary for anomaly detection becomes clearer. Local contrastive learning also benefits from better convergence and expressive capacity.

When the aggregation frequency is too low, such as aggregating after every round, the model synchronizes quickly. However, due to significant distribution differences among clients, frequent parameter sharing intensifies semantic drift. As a result, the global embedding space becomes flattened under contrastive constraints. Although accuracy remains stable, structural tension is weak, leading to poor Recall and EAS. This shows that low-frequency aggregation limits the Siamese network's ability to adapt to heterogeneous behaviors. It becomes especially ineffective for modeling rare but important risk patterns.

On the other hand, when the aggregation frequency is too high, such as aggregating every 20 rounds, the local model updates on each client diverge from the global objective. This leads to structural collapse in the aggregated embedding space and significantly degrades both accuracy and consistency. Although clients gain deeper modeling of local behavior features, the large distribution gap introduces parameter shifts and convergence instability during global fusion. This results in failed embedding alignment and reduced contrastive learning effectiveness.

The best performance is observed when aggregation occurs every five rounds. This suggests that the federated optimization strategy requires a delicate balance between local modeling capability and global alignment. Under this setting, each client can generate stable embedding structures through the Siamese contrastive mechanism. At the same time, periodic aggregation helps adjust embedding positions, maintaining representational diversity while ensuring cross-domain consistency. This further improves the robustness and generalization ability of risk detection.

## V. Conclusions

This paper addresses key challenges in financial risk discrimination, including multi-source heterogeneity, privacy isolation, and blurred decision boundaries. A deep discriminative framework is proposed by integrating Siamese networks with federated learning. The method preserves data locality on each client and constructs a global embedding space through contrastive learning. This enhances the separability of normal and abnormal behaviors in the representation space. The model is trained with shared parameters across multiple clients and performs semantic alignment and risk knowledge transfer through federated aggregation. It enables high-accuracy and robust risk identification without accessing raw data.

The main contribution of this work lies in introducing structural contrastive learning into federated settings. By incorporating entropy-aware regularization and embedding compression mechanisms, the model improves its capacity to represent low-frequency and high-impact behaviors. At the same time, it alleviates the representation shift caused by data heterogeneity. Experimental results demonstrate that the model achieves significant advantages in precision, recall, and embedding consistency under different settings. It shows strong generalization ability, especially in imbalanced, highly heterogeneous, and non-independent environments. These findings suggest that the federated contrastive modeling paradigm can serve as an effective complement to existing financial risk control systems. It also provides a theoretical and practical foundation for building scalable risk control platforms for the future.

From an application perspective, the proposed method can be widely deployed in real-world scenarios such as credit card fraud detection, anti-money laundering audits, and cross-bank risk assessment. It is especially suitable for financial alliance ecosystems involving multiple institutions, high data sensitivity, and limited collaboration space. The framework ensures privacy and security while offering a modeling paradigm with structural awareness and end-to-end discriminative ability. It has the potential to shift financial technology systems from integrated detection to collaborative

intelligent discrimination. This supports the development of more dynamic risk control capabilities with higher dimensionality and finer granularity.

## VI. Future Work

Future research can explore the scalability of this framework to multimodal, cross-platform, and even cross-lingual financial behavior data. In particular, there is strong potential in combining self-supervised learning, graph-based modeling, and heterogeneous federated optimization. Communication compression and privacy-enhancing technologies such as differential privacy and homomorphic encryption can also be introduced to reduce deployment costs and improve regulatory compliance. These extensions will support the construction of intelligent risk discrimination systems that are secure, interpretable, and scalable. Continued work in this direction will promote the evolution of smart financial systems toward more federated, personalized, and autonomous paradigms.

## References

1. N. Sharma, S. Gupta, H. G. Mohamed, et al., "Siamese convolutional neural network-based twin structure model for independent offline signature verification," *Sustainability*, vol. 14, no. 18, 11484, 2022.
2. S. I. Nanayakkara, S. R. Pokhrel and G. Li, "Understanding global aggregation and optimization of federated learning," *Future Generation Computer Systems*, vol. 159, pp. 114-133, 2024.
3. H. Lu and H. Wang, "Graph contrastive pre-training for anti-money laundering," *International Journal of Computational Intelligence Systems*, vol. 17, no. 1, 307, 2024.
4. R. Wei and S. Yao, "Enterprise financial risk identification and information security management and control in big data environment," *Mobile Information Systems*, vol. 2021, no. 1, 7188327, 2021.
5. Y. Wang, "Entity-aware graph neural modeling for structured information extraction in the financial domain," *Transactions on Computational and Scientific Methods*, vol. 4, no. 9, 2024.
6. M. Jiang, S. Liu, W. Xu, S. Long, Y. Yi and Y. Lin, "Function-driven knowledge-enhanced neural modeling for intelligent financial risk identification," 2025.
7. W. Xu, M. Jiang, S. Long, Y. Lin, K. Ma and Z. Xu, "Graph neural network and temporal sequence integration for AI-powered financial compliance detection," 2025.
8. Q. Sha, T. Tang, X. Du, J. Liu, Y. Wang and Y. Sheng, "Detecting credit card fraud via heterogeneous graph neural networks with graph attention," *arXiv preprint arXiv:2504.08183*, 2025.
9. H. Yang, M. Wang, L. Dai, Y. Wu and J. Du, "Federated graph neural networks for heterogeneous graphs with data privacy and structural consistency," 2025.
10. L. Dai, "Contrastive learning framework for multimodal knowledge graph construction and data-analytical reasoning," *Journal of Computer Technology and Software*, vol. 3, no. 4, 2024.
11. D. Gao, "High fidelity text to image generation with contrastive alignment and structural guidance," *arXiv preprint arXiv:2508.10280*, 2025.
12. Y. Yao, Z. Xu, Y. Liu, K. Ma, Y. Lin and M. Jiang, "Integrating feature attention and temporal modeling for collaborative financial risk assessment," *arXiv preprint arXiv:2508.09399*, 2025.
13. Y. Zou, "Federated distillation with structural perturbation for robust fine-tuning of LLMs," *Journal of Computer Technology and Software*, vol. 3, no. 4, 2024.
14. Y. Qin, "Hierarchical semantic-structural encoding for compliance risk detection with LLMs," *Transactions on Computational and Scientific Methods*, vol. 4, no. 6, 2024.
15. S. Pan and D. Wu, "Hierarchical text classification with LLMs via BERT-based semantic modeling and consistency regularization," 2025.
16. X. Hu, Y. Kang, G. Yao, T. Kang, M. Wang and H. Liu, "Dynamic prompt fusion for multi-task and cross-domain adaptation in LLMs," *arXiv preprint arXiv:2509.18113*, 2025.
17. R. Hao, X. Hu, J. Zheng, C. Peng and J. Lin, "Fusion of local and global context in large language models for text classification," 2025.
18. Q. Xu, "Unsupervised temporal encoding for stock price prediction through dual-phase learning," 2025.

19. Q. Sha, "Hybrid deep learning for financial volatility forecasting: An LSTM-CNN-Transformer model," *Transactions on Computational and Scientific Methods*, vol. 4, no. 11, 2024.
20. X. Su, "Predictive modeling of volatility using generative time-aware diffusion frameworks," *Journal of Computer Technology and Software*, vol. 4, no. 5, 2025.
21. W. Xu, J. Zheng, J. Lin, M. Han and J. Du, "Unified representation learning for multi-intent diversity and behavioral uncertainty in recommender systems," arXiv preprint arXiv:2509.04694, 2025.
22. Z. Liu and Z. Zhang, "Modeling audit workflow dynamics with deep Q-learning for intelligent decision-making," *Transactions on Computational and Scientific Methods*, vol. 4, no. 12, 2024.
23. Y. Wu, Y. Qin, X. Su and Y. Lin, "Transformer-based risk monitoring for anti-money laundering with transaction graph integration," 2025.
24. J. Gong, Y. Wang, W. Xu, et al., "A deep fusion framework for financial fraud detection and early warning based on large language models," *Journal of Computer Science and Software Applications*, vol. 4, no. 8, 2024.
25. G. Wang, J. Ma and G. Chen, "Attentive statement fraud detection: Distinguishing multimodal financial data with fine-grained attention," *Decision Support Systems*, vol. 167, 113913, 2023.
26. R. Li, Y. Cao, Y. Shu, et al., "A dynamic receptive field and improved feature fusion approach for federated learning in financial credit risk assessment," *Scientific Reports*, vol. 14, no. 1, 26515, 2024.
27. B. Hong, P. Lu, H. Xu, et al., "Health insurance fraud detection based on multi-channel heterogeneous graph structure learning," *Heliyon*, vol. 10, no. 9, 2024.

**Disclaimer/Publisher's Note:** The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.