

Article

Not peer-reviewed version

Multi-Hop Relational Modeling for Credit Fraud Detection via Graph Neural Networks

[Kewei Cao](#) , Yinghao Zhao , Hejing Chen , [Xinyi Liang](#) , Yihan Zheng , Sumeng Huang *

Posted Date: 31 December 2025

doi: 10.20944/preprints202512.2728.v1

Keywords: graph learning; credit risk identification; relationship modeling; financial networks



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

Multi-Hop Relational Modeling for Credit Fraud Detection via Graph Neural Networks

Kewei Cao ¹, Yinghao Zhao ², Hejing Chen ³, Xinyi Liang ¹, Yihan Zheng ¹ and Sumeng Huang ^{4,*}

¹ Columbia University, New York, USA

² Pace University, New York, USA

³ Walsh College, Troy Michigan, USA

⁴ Georgia Institute of Technology, Atlanta, USA

* Correspondence: shuang610@gatech.edu

Abstract

Credit fraud detection is a central problem in financial risk management. It is challenging due to complex relationships among transaction entities, strong coordination of fraudulent activities, and high levels of behavioral camouflage at the individual level. Traditional approaches mainly rely on features from independent samples. They often fail to represent multi-entity interaction structures effectively. From a relational modeling perspective, this study investigates a graph neural network-based method for credit fraud detection. The method represents transaction entities and their interactions as a unified graph. It performs information propagation and neighborhood aggregation in graph space. It jointly models node attributes and structural context. It therefore learns more discriminative risk representations. During modeling, the network captures both local interaction patterns and multi-hop relational information. This makes anomalous patterns embedded in complex transaction networks more explicit. Based on the proposed method, this study establishes a standardized comparative evaluation on a public credit fraud dataset. It systematically compares the method with multiple existing models. The overall results confirm the effectiveness of graph-structured modeling for credit fraud detection. The findings show that, compared with detection methods that ignore or weaken relational information, graph neural network models achieve clear advantages in stability and consistency of risk identification. They better reflect the propagation characteristics of fraudulent behavior within network structures. These results support modeling credit fraud detection as a relation-aware graph learning problem. This formulation enhances risk characterization in complex financial scenarios. It also provides a structured modeling perspective for building intelligent risk control systems.

Keywords: graph learning; credit risk identification; relationship modeling; financial networks

I. Introduction

The widespread adoption of digital payments and online credit has enabled financial services to scale and operate in real time. It has also made fraud more covert and more organized. Traditional credit fraud detection often relies on static rules or anomalous signals from a single transaction. In internet platforms, cross-channel payments, and multi-party ecosystems, fraudsters can coordinate across multiple accounts, devices, merchants, and interaction paths [1]. They dilute single-point risk signals and create behavior traces that appear normal. At the same time, rapid business iteration continuously reshapes transaction patterns and customer profiles. Fraud strategies evolve accordingly. Methods built only on historical patterns can therefore lag and become ineffective. Reliable risk identification under high concurrency, strict timeliness, and complex associations remains a central challenge in credit risk control.

In terms of data structure, credit fraud identification is not simply a tabular classification problem. It is a network-oriented risk perception task with strong dependencies. Transactions,

accounts, devices, addresses, merchants, and behavioral sequences are naturally linked through multiple entities and multiple relation types. Risk does not always appear as isolated samples. It often emerges as local clusters, chain-like propagation, or implicit interactions across groups. In practice, many fraud activities depend on relationships. Examples include shared devices, reused collection routes, circular fund transfers, or intermediary accounts that conceal the final beneficiary. These behaviors form interaction graphs with clear structural characteristics. If such a structure is ignored and features are extracted only from individual records, the model will struggle to capture key signals such as group coordination and path dependence. This limits the ability to detect complex fraud patterns [2].

Against this background, graph neural networks offer a modeling paradigm that matches the nature of the problem. These methods aggregate and propagate contextual information over graph structures. The model can use both individual attributes and relational signals. It can learn more discriminative risk representations from neighborhood relations, interaction intensity, and connection patterns [3]. Graph neural networks can also capture higher-order relations. Multi-hop message passing enables the model to sense latent long-range associations. It can therefore identify risk behaviors that look normal at the surface but are structurally abnormal. In financial applications, heterogeneous graphs are common. Entity types differ, and relation semantics are diverse. Graph neural networks can jointly model multiple node types and multiple edge types within a unified framework. This provides flexibility for representing complex risk control scenarios [4]. Overall, reframing credit fraud detection as relationship-driven graph learning supports a more complete understanding of risk formation mechanisms. It also lays theoretical and practical foundations for building more robust and reliable intelligent risk control systems.

II. Related Work

Recent risk identification research increasingly emphasizes learning under realistic constraints such as sample scarcity, class imbalance, and evolving patterns. Meta-learning has been explored to improve adaptability when fraud patterns shift and labeled positives are limited, providing a mechanism to quickly adjust decision boundaries with few informative examples [5]. Self-supervised learning similarly reduces dependence on labels by extracting transferable representations from large volumes of unlabeled records, which is particularly useful when fraud labels are noisy or heavily imbalanced [6]. In time-dependent risk settings, deep attention models have been used to capture long-range temporal dependencies and improve forecasting stability, indicating that attention mechanisms can effectively model non-stationary dynamics in complex sequences [7].

A parallel line of work focuses on explicitly modeling structure rather than treating instances as independent. Graph neural networks provide a natural way to represent entities and interactions and to perform neighborhood aggregation for relational reasoning. Graph-integrated risk monitoring further combines sequence modeling with transaction graph structure, demonstrating that relational context can strengthen detection signals beyond isolated features [8]. Beyond standard graph learning, dynamic spatiotemporal causal graph neural networks incorporate temporal evolution and causal dependencies to improve representation fidelity under changing relational patterns, suggesting that causal structure can regularize message passing and enhance robustness [9]. These approaches directly motivate formulating credit fraud detection as a relation-aware learning problem where coordinated behaviors emerge as subgraph-level patterns rather than single-record outliers. Interpretability and traceability are also becoming central, especially when model decisions must be explainable and auditable. Knowledge-augmented agent frameworks and explainable decision-making pipelines leverage structured knowledge and reasoning components to provide clearer rationales for predictions [10]. More generally, explainable representation learning aims to produce feature spaces that support fine-grained attribution and decision transparency, which can be adapted to graph-based fraud detection by connecting risk scores to salient neighbors, paths, or interaction motifs [11]. Semantic knowledge graph frameworks provide another transferable methodology for

structuring entity relations and supporting evidence-based reasoning over connected data, which complements GNN message passing with explicit relational semantics [12].

Finally, robustness-oriented learning for complex systems has inspired techniques that can strengthen graph-based fraud models under distribution shift and adversarial conditions. Modular task decomposition and dynamic collaboration highlight that decomposing complex objectives into coordinated subtasks can improve stability and controllability in multi-component learning pipelines [13]. Parameter-efficient adaptation with structural priors and modular adapters provides a way to update models with constrained parameter changes, which can be useful when fraud patterns drift and frequent full retraining is impractical [14]. Multi-scale LoRA extends this idea by enabling adaptation at different granularities while keeping updates lightweight [15]. At the training objective level, contrastive knowledge transfer and robust optimization emphasize aligning representations across conditions and improving generalization under perturbations [16]. Reinforcement learning has also been used for adaptive optimization and scheduling in complex pipelines, illustrating a general strategy for learning policies that remain effective under changing workloads and constraints [17]. Moreover, graph-based modeling has been applied to infer high-level satisfaction states from relational data, supporting the general claim that graph representations can capture latent outcomes that depend on multi-entity interactions rather than independent samples [18]. Together, these methods motivate combining relation-aware graph learning with adaptation-friendly training and robustness mechanisms to improve the stability and consistency of credit fraud detection.

III. Proposed Framework

In terms of overall modeling, this paper represents the credit fraud identification problem as a node risk discrimination task on a graph structure. First, a graph representation is constructed containing multiple types of entities and their interactions, where each node corresponds to a risk subject, and each edge characterizes the interaction or association between subjects. Let the graph be:

$$G = (V, E) \quad (1)$$

where V represents the set of nodes and E represents the set of edges. Each node $v \in V$ is associated with an initial feature vector, denoted as:

$$H^{(0)} \in R^{|V| \times d} \quad (2)$$

This representation serves as the basis for the fundamental attribute information of the entity and as the starting point for subsequent relationship modeling and information dissemination. By explicitly introducing a graph structure, the model can incorporate the connections between entities into a unified risk representation space, beyond the feature level. This paper also presents the overall model architecture, as shown in Figure 1.

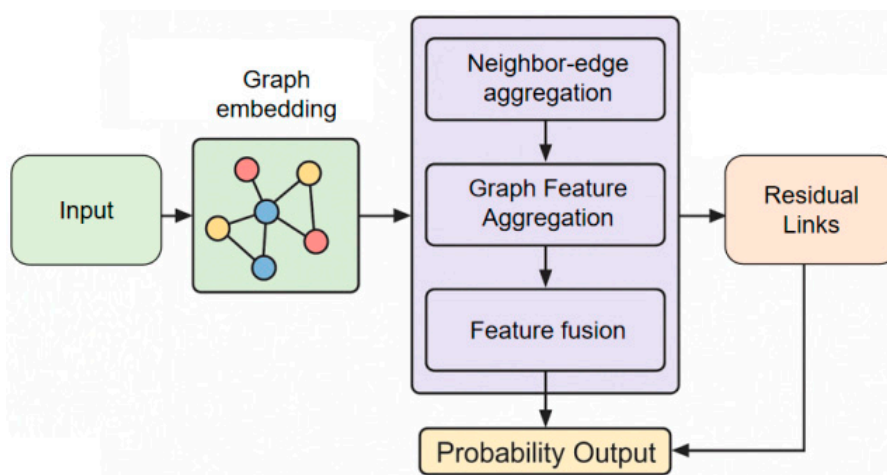


Figure 1. Overall model architecture.

During the representation learning stage, the model applies a graph neural network framework that updates node features via neighborhood aggregation. Each node's new representation is derived by applying a function that fuses its own attributes with aggregated signals from its neighbors. This approach has been successfully applied in the context of anomaly detection for financial data, where attention-based neighborhood aggregation enhances the model's ability to distinguish between normal and anomalous patterns [19]. To further strengthen relational modeling, our method applies attention mechanisms during the aggregation process, allowing the model to adaptively weight neighboring nodes—a strategy that has proven effective in previous deep learning frameworks for anomaly detection [20].

In handling heterogeneous graphs characteristic of transaction networks, the method further incorporates advanced message passing techniques, as applied by Xie & Chang [21], enabling the integration of both node-level and relational context for more expressive representations. Accordingly, the update rule in our model is designed to jointly encode local features and multi-type relational information, and is formally expressed as:

$$H^{(l+1)} = \sigma(\tilde{A}H^{(l)}W^{(l)}) \quad (3)$$

where \tilde{A} is the normalized adjacency matrix, $W^{(l)}$ is the learnable parameter matrix of the l -th layer, and $\sigma(\cdot)$ is the nonlinear mapping function. This process enables risk information to propagate layer by layer in the graph structure, allowing nodes to perceive the behavioral patterns of their directly and indirectly related entities, thereby forming a potential representation of the fused structural context.

To enhance the model's ability to characterize complex relational structures, a multi-layer propagation mechanism is further introduced, capturing higher-order dependencies by stacking several graph convolutional layers. After L layers of updates, the final representation of the nodes is obtained:

$$Z \in H^L \quad (4)$$

This representation integrates information from multi-hop neighborhoods, which helps identify fraud patterns that rely on cooperative relationships or path structures. Compared to shallow modeling methods that only use local neighborhood information, multi-layer propagation can mitigate the impact of single-interaction noise, making risk features more stable and coherent in the structural space.

In the risk discrimination stage, the model constructs a probability output based on the final representation of the nodes to characterize the likelihood of the subject engaging in fraudulent behavior. Specifically, a linear mapping and normalization function are used to map the representation to a risk score:

$$y = \text{Softmax}(ZW_o) \quad (5)$$

where W_o represents the output layer parameters. This design enables the model to complete the end-to-end reasoning process from relationship modeling to risk scoring within a unified framework. By jointly encoding structural information and node attributes into the representation space, the method can theoretically more fully reflect the correlation and systemic characteristics of credit fraud, providing a consistent modeling foundation for subsequent risk control and decision support.

IV. Experimental Analysis

A. Dataset

The dataset used in this study is the public Credit Card Fraud Detection dataset, which is widely adopted in credit fraud identification research. It consists of credit card transaction records from European cardholders. It is available from open source platforms for machine learning and anomaly detection tasks. The dataset contains a large volume of transaction behavior information generated in real payment environments. It includes several hundred thousand records. The fields include

transaction time, transaction amount, and anonymized transaction features. To protect privacy, the transaction information has been preprocessed, such as by principal component analysis. The original attributes are transformed into numerical features. Key information relevant to fraud pattern analysis is retained. The labels clearly distinguish legitimate transactions from fraudulent behavior. This provides supervised semantic annotations for credit fraud identification.

A prominent characteristic of this dataset is severe class imbalance. Fraudulent samples account for only a very small proportion of all transactions. This setting is representative of real-world credit fraud scenarios. Since fraud is inherently rare in the overall transaction population, the minority class constitutes only a tiny fraction of the total volume. This imbalanced distribution imposes stricter requirements on learning strategies. In addition, fields such as transaction time and amount support studies of temporal dynamics and monetary scale. These attributes are directly useful for building graph structures and defining node features in graph neural networks. Therefore, this open dataset is not only a standard benchmark for classical credit fraud detection. It also provides a rich basis for mining complex relational patterns.

Although the dataset is originally provided in a flat tabular format, graph neural network modeling typically requires graph construction based on relations among entities. Entities such as transaction participants, for example, accounts and merchants, device identifiers, and time windows can be treated as nodes. Transaction activities can be defined as edges. In this way, raw transaction logs can be transformed into a multi-entity relational graph. This representation reflects the relational nature of financial activities. It also provides richer contextual information for deep graph models. It helps capture risk propagation patterns across nodes. The relational graph built from this dataset contains both attribute information and behavioral links. This supports the modeling and discrimination of complex fraudulent behaviors with graph neural networks.

B. Experimental Results

This article first presents the results of the comparative experiments, as shown in Table 1.

Table 1. Comparative experimental results.

Method	Acc	Recall	AUC	F1-Score
CATCHM [22]	0.641	0.628	0.662	0.635
NUS [23]	0.654	0.641	0.671	0.646
RHSOFS [24]	0.662	0.649	0.678	0.653
NAG [25]	0.667	0.653	0.683	0.659
CTCN [26]	0.673	0.661	0.689	0.665
Ours	0.689	0.675	0.703	0.679

Results show clear stratification across methods on credit fraud detection, indicating that modeling paradigm strongly determines risk identification quality: traditional or weakly structured models struggle to capture coordinated and concealed fraud within complex transaction relations and are more sensitive to local noise, while relational/graph-aware methods are more stable because independent features alone cannot represent the underlying risk structure. Performance generally improves as models leverage more transaction connectivity, reflecting that fraud is networked and often propagates along account interactions and transaction paths, so explicitly modeling structure helps capture anomalous signals and improves robustness. Our method achieves consistent advantages across metrics by jointly modeling node attributes and relational structure to learn more discriminative risk representations under multi-account coordination and path obfuscation, reinforcing the view that fraud detection is fundamentally a structured relational learning problem with strong practical potential in large financial networks. Finally, because GNN depth controls the neighborhood range for aggregation and thus local vs. higher-order relationship modeling, we conduct a sensitivity analysis across layer settings, with results reported in Figure 2.

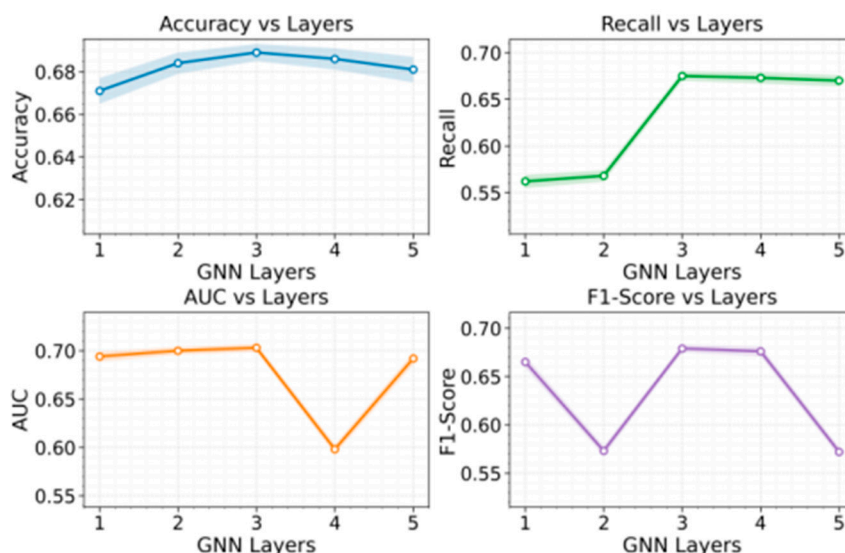


Figure 2. The impact of the number of layers in a graph neural network on experimental results.

The depth sensitivity curves indicate that the model exhibits a clear and controllable change in how it captures structural fraud patterns under different propagation depths. With a shallow depth, node representations rely more on local neighborhoods. The model can capture short-range anomaly cues induced by direct interactions. However, coverage is limited for cross-account paths and covert coordinated behavior. As the depth increases, the aggregation range expands. The model can incorporate risk signals from multi-hop relations into a shared representation space. This makes it more sensitive to graph-structured risks such as group-based fraud and path-based transfers. The trend supports the need to formulate fraud detection as a relational learning task.

The accuracy trend suggests that a moderate propagation depth better balances local details and global context. The model can maintain stable discrimination under the diverse patterns of legitimate transactions. If the depth is too shallow, fraud patterns that depend on higher-order relations are more likely to be missed. This reflects insufficient modeling of complex structures. When the depth continues to increase, accuracy shows signs of decline. This suggests effects from over-smoothing or noise diffusion. In credit transaction graphs, legitimate behavior is denser and more frequent. Excessive depth can mix large amounts of non-risk neighborhood information into node representations. This weakens the decision boundary for key anomalous structures.

The recall curve shows a sharp increase followed by a plateau. This matches the dependence of fraud detection on relational diffusion. Multi-hop aggregation brought by deeper propagation makes hidden fraud associations along paths more observable. This is especially important when fraud samples are rare and behavioral camouflage is strong. Cross-node propagation helps compensate for the limitations of single-point features. After the depth reaches a certain level, the gain from adding more layers diminishes. This indicates that most critical fraud relations have already entered the effective receptive field. Further depth mainly introduces redundant neighborhood information rather than additional evidence.

The variations in AUC and F1 are more pronounced. This indicates that depth changes affect not only whether fraud can be captured, but also the trade-off between ranking ability and precise coverage. AUC drops markedly at a deeper setting. This suggests that overly deep propagation may dilute risk representations with non-risk neighborhoods. It reduces overall separability and ranking stability. AUC then rises again, which implies a nonlinear relationship between structural signal absorption and noise accumulation across depths. The fluctuation of F1 shows that, under class imbalance, depth adjustment changes the structural distribution of false positives and false negatives. Moderate depth is more likely to yield a relatively balanced decision state. Depth that is too shallow or too deep can bias the model toward one type of error. This provides direct modeling evidence for selecting an appropriate propagation depth in credit fraud scenarios.

V. Conclusion

This study addresses real-world characteristics of credit fraud detection, including complex relational structures, strong behavioral concealment, and highly coordinated risk patterns. It conducts a systematic investigation of graph neural network-based modeling. It elevates the conventional risk identification problem from an independent sample-centered setting to a structure-driven graph learning task. By representing transaction entities and their interactions as a unified graph, the model can exploit relational information among multiple parties beyond feature-level signals. This design aligns more closely with the intrinsic mechanisms of risk formation and diffusion in real financial systems. The analysis indicates that relational modeling plays an irreplaceable role in improving the completeness of risk characterization and the consistency of discrimination. It offers a new methodological perspective for intelligent risk control in complex credit scenarios.

From a methodological perspective, graph neural networks propagate and aggregate neighborhood information layer by layer. This enables the model to perceive local interactions and higher-order relations at multiple scales. It shows stronger adaptability to typical risk patterns, including organized group fraud, chain-like transfers, and implicit coordination. The model no longer relies only on explicit anomalous features. It learns latent risk representations from structural context. This shifts fraud detection from static judgment to relation-aware perception and structural reasoning. The shift supports more stable risk identification. It also provides a theoretical basis for building more robust risk control systems.

From an application perspective, graph neural network-based credit fraud detection is important for risk management in financial institutions. On the one hand, relation-aware models can capture cross-account and cross-channel risk associations earlier. This provides richer support for real-time interception and early warning. On the other hand, structured risk representations can improve interpretability. They make risk decisions easier to align with business rules, compliance requirements, and audit processes. The approach also has strong transfer potential across related domains, including payment risk control, credit approval, and anti-money laundering.

Looking ahead, credit fraud detection will remain in environments with continuous distribution shifts and evolving attack strategies. Several questions require further study. One is how to achieve continual learning under dynamic graph structures. Another is how to maintain stable structural representations under noisy and incomplete relations. A third is how to support cross-scenario risk modeling while satisfying privacy and compliance constraints. As graph learning methods become more integrated with financial business systems, graph neural network-based risk identification is expected to play a key role in larger-scale and more complex financial networks. It can provide long-term support for building secure, robust, and intelligent financial infrastructure.

References

1. S. Xiang, M. Zhu, D. Cheng et al., "Semi-Supervised Credit Card Fraud Detection via Attribute-Driven Graph Representation," Proceedings of the AAAI Conference on Artificial Intelligence, vol. 37, no. 12, pp. 14557–14565, 2023.
2. Y. Duan, G. Zhang, S. Wang et al., "CAT-GNN: Enhancing Credit Card Fraud Detection via Causal Temporal Graph Neural Networks," arXiv preprint, arXiv:2402.14708, 2024.
3. A. Ali, S. Abd Razak, S. H. Othman et al., "Financial Fraud Detection Based on Machine Learning: A Systematic Literature Review," Applied Sciences, vol. 12, no. 19, p. 9637, 2022.
4. Y. Zou and D. Cheng, "Effective High-Order Graph Representation Learning for Credit Card Fraud Detection," arXiv preprint, arXiv:2503.01556, 2025.
5. H. Fan, Y. Yi, W. Xu, Y. Wu, S. Long and Y. Wang, "Intelligent Credit Fraud Detection with Meta-Learning: Addressing Sample Scarcity and Evolving Patterns," 2025.
6. J. Lai, A. Xie, H. Feng, Y. Wang and R. Fang, "Self-Supervised Learning for Financial Statement Fraud Detection with Limited and Imbalanced Data," 2025.

7. Q. Xu, W. Xu, X. Su, K. Ma, W. Sun and Y. Qin, "Enhancing Systemic Risk Forecasting with Deep Attention Models in Financial Time Series," Proceedings of the 2025 2nd International Conference on Digital Economy, Blockchain and Artificial Intelligence, pp. 340–344, 2025.
8. Y. Wu, Y. Qin, X. Su and Y. Lin, "Transformer-Based Risk Monitoring for Anti-Money Laundering with Transaction Graph Integration," Proceedings of the 2025 2nd International Conference on Digital Economy, Blockchain and Artificial Intelligence, pp. 388–393, 2025.
9. Q. Gan, R. Ying, D. Li, Y. Wang, Q. Liu and J. Li, "Dynamic Spatiotemporal Causal Graph Neural Networks for Corporate Revenue Forecasting," 2025.
10. Q. Zhang, Y. Wang, C. Hua, Y. Huang and N. Lyu, "Knowledge-Augmented Large Language Model Agents for Explainable Financial Decision-Making," arXiv preprint, arXiv:2512.09440, 2025.
11. Y. Xing, M. Wang, Y. Deng, H. Liu and Y. Zi, "Explainable Representation Learning in Large Language Models for Fine-Grained Sentiment and Opinion Classification," 2025.
12. L. Yan, Q. Wang and C. Liu, "Semantic Knowledge Graph Framework for Intelligent Threat Identification in IoT," 2025.
13. S. Pan and D. Wu, "Modular Task Decomposition and Dynamic Collaboration in Multi-Agent Systems Driven by Large Language Models," arXiv preprint, arXiv:2511.01149, 2025.
14. Y. Wang, D. Wu, F. Liu, Z. Qiu and C. Hu, "Structural Priors and Modular Adapters in the Composable Fine-Tuning Algorithm of Large-Scale Models," arXiv preprint, arXiv:2511.03981, 2025.
15. H. Zhang, L. Zhu, C. Peng, J. Zheng, J. Lin and R. Bao, "Intelligent Recommendation Systems Using Multi-Scale LoRA Fine-Tuning and Large Language Models," 2025.
16. J. Zheng, H. Zhang, X. Yan, R. Hao and C. Peng, "Contrastive Knowledge Transfer and Robust Optimization for Secure Alignment of Large Language Models," arXiv preprint, arXiv:2510.27077, 2025.
17. K. Gao, Y. Hu, C. Nie and W. Li, "Deep Q-Learning-Based Intelligent Scheduling for ETL Optimization in Heterogeneous Data Environments," arXiv preprint, arXiv:2512.13060, 2025.
18. R. Liu, R. Zhang and S. Wang, "Graph Neural Networks for User Satisfaction Classification in Human-Computer Interaction," arXiv preprint, arXiv:2511.04166, 2025.
19. J. Li, Q. Gan, Z. Liu, C. Chiang, R. Ying and C. Chen, "An Improved Attention-Based LSTM Neural Network for Intelligent Anomaly Detection in Financial Statements," 2025.
20. H. Wang, C. Nie and C. Chiang, "Attention-Driven Deep Learning Framework for Intelligent Anomaly Detection in ETL Processes," 2025.
21. A. Xie and W. C. Chang, "Deep Learning Approach for Clinical Risk Identification Using Transformer Modeling of Heterogeneous EHR Data," arXiv preprint, arXiv:2511.04158, 2025.
22. R. Van Belle, B. Baesens and J. De Weerd, "CATCHM: A Novel Network-Based Credit Card Fraud Detection Method Using Node Representation Learning," Decision Support Systems, vol. 164, p. 113866, 2023.
23. H. Zhu, M. C. Zhou, G. Liu et al., "NUS: Noisy-Sample-Removed Undersampling Scheme for Imbalanced Classification and Application to Credit Card Fraud Detection," IEEE Transactions on Computational Social Systems, vol. 11, no. 2, pp. 1793–1804, 2023.
24. B. K. Padhi, S. Chakravarty, B. Naik et al., "RHSOFS: Feature Selection Using the Rock Hyrax Swarm Optimization Algorithm for Credit Card Fraud Detection System," Sensors, vol. 22, no. 23, p. 9321, 2022.
25. K. Ghosh Dastidar, J. Jurgovsky, W. Siblini et al., "NAG: Neural Feature Aggregation Framework for Credit Card Fraud Detection," Knowledge and Information Systems, vol. 64, no. 3, pp. 831–858, 2022.
26. X. Zhao and S. Guan, "CTCN: A Novel Credit Card Fraud Detection Method Based on Conditional Tabular Generative Adversarial Networks and Temporal Convolutional Network," PeerJ Computer Science, vol. 9, p. e1634, 2023.

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.