Knowledge Graph-Driven Generative Framework for Interpretable Financial Fraud Detection

Shiqing Long , Kewei Cao , Xinyi Liang , Yihan Zheng , Yingnan Yi , Ruizhe Zhou *

Posted Date: 24 December 2025

doi: 10.20944/preprints202512.2142.v1

Keywords: knowledge graph; generative model; financial fraud detection; semantic reasoning

 doi:10.20944/preprints202512.2142.v1

Article

Knowledge Graph-Driven Generative Framework for Interpretable Financial Fraud Detection

Shiqing Long [1], Kewei Cao [2], Xinyi Liang [2], Yihan Zheng [2], Yingnan Yi [3] and Ruizhe Zhou [4,*]

[1] Independent Researcher, Shanghai, China
[2] Columbia University, New York, USA
[3] Washington University in St. Louis, St. Louis, USA
[4] University of Chicago, Chicago, USA
* Correspondence: author: rexzhe1230@gmail.com

Abstract

This paper proposes a knowledge-graph-guided generative detection model to address the challenges of complex feature associations, hidden behavioral patterns, and sparse semantic information in financial fraud detection. The study maps multi-source heterogeneous data, such as accounts, transactions, devices, and geographic information, into a knowledge graph structure to model semantic dependencies and behavioral relationships among financial entities at both node and relation levels. Based on this, the model introduces a relation-aware encoder to obtain high-dimensional structural embeddings and employs a generative reasoning mechanism to learn the distribution of latent fraud patterns, enabling semantic identification of complex risk behaviors without relying on fixed rules. During generation, the method integrates knowledge constraints and structural priors, allowing the model to reconstruct potential fraud chains from a global semantic perspective and significantly enhance interpretability and generalization. Validation on real transaction datasets shows that the model outperforms comparison methods in AUC, ACC, Precision, and Recall, demonstrating the effectiveness of the knowledge-enhanced generative framework in capturing cross-entity dependencies and dynamic risk relationships. Overall, the proposed knowledge-graph-guided generative detection model achieves a unified framework of structural modeling, semantic reasoning, and anomaly generation, providing a new systematic solution for the intelligent detection of complex financial fraud behaviors. CCS CONCEPTS: Computing methodologies~Machine learning~Machine learning approaches.

Keywords: knowledge graph; generative model; financial fraud detection; semantic reasoning

1. Introduction

The forms of financial fraud are constantly evolving. They have expanded from traditional forgery, money laundering, and fake transactions to modern intelligent fraud that relies on algorithmic manipulation and cross-platform concealment. Their complexity and invisibility have exceeded the capability of traditional detection methods. With the rapid development of financial digitalization and online trading ecosystems, transaction behavior, account relationships, device information, and external environmental data have become multi-dimensional, dynamic, and heterogeneous. Traditional rule-based or supervised models often depend on static features and fixed patterns, making them incapable of capturing cross-domain fraud paths or hidden semantic associations. In this context, building models that can understand, associate, and reason over complex financial semantics becomes a critical challenge. The knowledge graph, with its unique advantage in structurally representing entity relationships and uncovering latent semantic networks, provides a new direction for generative fraud detection [1].

A knowledge graph constructs a semantic network of entities and relations, bringing structured semantics into financial transaction data [2–4]. It can describe explicit connections among accounts,

 doi:10.20944/preprints202512.2142.v1

institutions, and transaction behaviors, and also reveal hidden indirect dependencies and potential collusion through its graph structure. For example, in a complex financial system, account registration details, geographic locations, device fingerprints, and fund flows can be unified under one graph representation to model multi-dimensional behavioral trajectories. This semantic enhancement allows the model to go beyond local feature matching and reason about the formation logic of fraud chains at a global level. Combined with the expressive and inferential power of generative models, graph-guided learning can dynamically generate potential fraud patterns or transaction reasoning paths, enabling higher interpretability and adaptive detection [5].

In traditional financial risk control systems, models are designed to identify anomalies or classify risks rather than to understand the semantics or structural logic behind them. As a result, most models struggle with poor transferability and weak generalization when facing emerging fraud strategies. The introduction of a knowledge-graph-guided generative detection framework provides a potential solution. By guiding the generative model with graph structures to learn semantic dependencies and causal relations among entities, the model can extract patterns from historical data and generate new reasoning paths and risk explanations based on knowledge associations. This shift enables a transition from "recognition" to "understanding." Fraud detection thus moves from passive reaction to proactive perception, allowing potential risks to be predicted even under incomplete information or limited samples [6].

Moreover, the integration of knowledge graphs and generative models has significant systemic and industrial value [7–9]. In financial regulation and compliance analysis, model outputs must be interpretable and traceable, yet traditional deep learning models often act as "black boxes." A knowledge-graph-guided generative detection model can provide transparent reasoning chains, such as generating behavioral descriptions from transaction features and mapping them onto interpretable paths within the graph. This mechanism enables risk tracing and logical transparency, improving the credibility and regulatory compliance of the model in practice. It also supports the development of regulatory technology and intelligent risk control systems, promoting a shift from rule-driven to semantic-driven financial security [10].

Overall, research on knowledge-graph-guided generative financial fraud detection represents not only an extension of existing risk identification methods but also a new paradigm of "semantic reasoning and generative understanding" in financial intelligence. It breaks through the limitations of static feature learning and isolated decision-making, achieving deep integration between knowledge enhancement and generative reasoning. This approach provides new intelligent solutions for key financial applications such as high-risk identification, anti-money-laundering monitoring, and cross-border fund tracking. Its value lies not only in improving detection accuracy but also in advancing financial technology toward interpretability, security, and intelligence, laying a theoretical and technical foundation for building trustworthy, transparent, and adaptive financial safety systems.

2. Related Work

Research on knowledge-structured modeling and graph-based learning provides a foundation for semantic financial fraud detection. Knowledge-graph-integrated neural models embed entities and relations into a structured representation space and couple graph semantics with deep neural networks, showing that explicitly modeling relationship structure can significantly improve fraud detection performance over flat feature inputs [11]. Building on this idea, two-layer knowledge graph architectures have been proposed for explainable fraud detection, where multi-level graph structures and relationship hierarchies are used to generate interpretable reasoning paths and highlight key entities and links behind suspicious behaviors [12]. From the perspective of graph quality and modeling assumptions, studies on scale regularization in fraud graphs analyze how controlling graph size and structural properties affects model generalization, stability, and sensitivity to fraudulent substructures [13]. These works suggest that knowledge graphs, when tightly integrated with deep models, are powerful carriers of structural priors for risk modeling, and they motivate guiding generative detection with graph semantics.

 doi:10.20944/preprints202512.2142.v1

Graph neural networks and temporal models extend structural modeling to dynamic, high-dimensional environments. Graph-based classification frameworks learn node and graph embeddings by message passing and neighborhood aggregation, capturing high-order dependencies in interaction graphs [14]. Structural–temporal graph models combine graph neural networks with temporal modules to encode both structural correlations and time-evolving behaviors, yielding improved anomaly detection performance on key performance indicator streams and other complex sequences [15]. Transformer-based anomaly detection and risk monitoring further leverage self-attention to model local and global dependencies, and can incorporate graph-structured connectivity or multi-scale temporal windows to better capture evolving risk signals [16,17]. Reconstruction-based anomaly detection with autoencoders learns compact latent encodings of normal operational processes and uses reconstruction error as an anomaly signal, providing a generative baseline for unsupervised detection in high-dimensional workflows [18]. In parallel, transformer-based risk monitoring architectures that integrate transaction graphs have illustrated how attention mechanisms and graph representations can be combined to model complex transaction networks and risk flows [19]. Together, these approaches support the structural and temporal design choices in the proposed model, which uses a relation-aware encoder and generative reasoning over a knowledge graph structure.

Robust representation learning techniques further enhance the ability to capture subtle dependencies and distributional changes relevant to fraud. Contrastive learning–based dependency modeling trains encoders to pull semantically consistent behaviors closer and push inconsistent ones apart, improving the sensitivity of learned representations to rare deviations and structural anomalies [20]. Extensions that couple contrastive learning with sensitivity analysis systematically examine how representations respond to perturbations, leading to more stable anomaly detection pipelines under hyperparameter and data shifts [21]. Attention-based time-series prediction models use attention mechanisms over temporal contexts to focus on salient time steps and nonlinear dependencies, which is important for modeling evolving risk signals in sequential data [22]. Methods that integrate causal inference into predictive models show how structural assumptions and bias-correction strategies can reduce spurious correlations and improve robustness when exposure or selection effects are present [23]. These ideas—dependency-aware learning, perturbation robustness, temporal attention, and causality-informed modeling—align with the goal of the proposed knowledge-graph-guided generative model: to represent latent fraud patterns in a way that is both expressive and robust to evolving behaviors and distribution shifts.

Recent work on large-scale models, reinforcement learning, and multi-agent systems offers complementary insights for structure-aware decision systems. Context compression and structural representation techniques for large language models design mechanisms to reduce redundant context while preserving essential structural cues, producing compact and semantically rich representations that support downstream reasoning and generation [24]. Deep Q-learning has been used to model workflow dynamics and learn policies over multi-step processes, demonstrating that reinforcement learning can capture sequential decision structures and adapt to changing operational conditions [25]. Multi-objective deep reinforcement learning extends this perspective by jointly optimizing several performance criteria under dynamic workloads and constraints, showing how adaptive policies can be learned in complex environments [26]. Multi-agent systems driven by large language models study modular task decomposition and dynamic collaboration, where subtasks are assigned to specialized agents and coordination emerges through communication and shared objectives [27]. Methodologically, these works emphasize structural priors, modularization, and learned coordination as key ingredients for complex decision-making systems. The knowledge-graph-guided generative detection model in this paper follows a similar philosophy by encoding structural priors through knowledge graphs, using a relation-aware encoder to obtain structural embeddings, and introducing generative reasoning constrained by graph semantics to learn latent fraud distributions and reconstruct potential fraud chains. In doing so, it unifies structural modeling,

 doi:10.20944/preprints202512.2142.v1

semantic reasoning, and anomaly generation into a single framework tailored to complex financial fraud detection.

3. Method

In this study, we apply a generative financial fraud detection approach guided by knowledge graphs, aiming to address the challenges of complex feature interactions, dynamic risk patterns, and limited semantic information in real-world financial transaction data. The framework is designed to integrate semantic relationship modeling with generative reasoning, allowing for a comprehensive understanding and inference of both observed and potential fraud behaviors.

At the core of the approach, the model first constructs a knowledge graph that organizes multi-source heterogeneous entities, including accounts, transactions, devices, and geographic locations, into a unified structure. These entities and their interactions are encoded as triples, effectively mapping the semantic and behavioral relationships that underpin financial operations. In constructing this knowledge graph, we adopt the dynamic and data-aware multi-agent graph construction techniques from Ying et al. [28], ensuring that temporal evolution and adaptive risk factors are captured in the network representation. This enables the graph to flexibly accommodate emerging entities, evolving connections, and shifting fraud tactics.

To further strengthen the graph's capacity for abnormal behavior discrimination, the framework utilizes federated risk discrimination mechanisms [29], which allow for collaborative anomaly identification across distributed data sources while protecting privacy and maintaining robustness against sample imbalance. Additionally, the representation of nodes and relations within the graph is enhanced through attention-based sequential modeling as outlined by Li et al. [30], facilitating the extraction of salient patterns in sequential financial activities. The model also incorporates improved structure-aware sequence modeling from Xu et al. [31], enabling fine-grained distinction between normal and fraudulent transaction trajectories by capturing higher-order dependencies and semantic nuances. Letting E denote the entity set and R the relationship set, the knowledge graph structure can be formally defined as:

$$G = \{(h, r, t) \mid h, t \in E, r \in R\} \tag{1}$$

Here, h and t represent the head and tail entities, respectively, and r represents the relationship between entities. Through this structured representation, the model can capture multi-level semantic information about account associations and behavioral dependencies, providing graph structure priors for subsequent generative detection. The model architecture is shown in Figure 1.

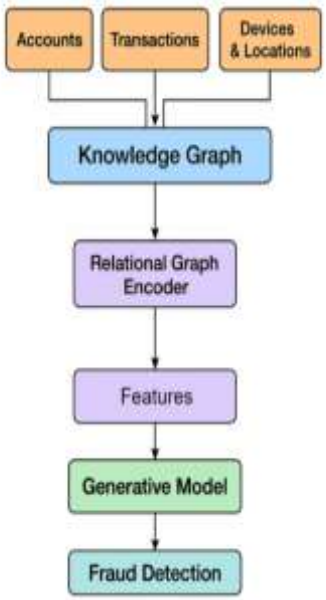

 doi:10.20944/preprints202512.2142.v1

Figure 1. Overall model architecture diagram.

In the graph structure representation stage, the model uses a relation-aware embedding mechanism to encode nodes and edges to learn differentiable semantic vector representations. For any node i, its representation update can be formalized as:

$$h_i^{(l+1)} = \sigma\left(\sum_{j \in N(i)} \alpha_{ij}^{(r)} W^{(r)} h_j^{(l)}\right) \tag{2}$$

where $N(i)$ represents the set of neighbors of node i, $\alpha_{ij}^{(r)}$ is the relationship weight, $W^{(r)}$ is the transformation matrix corresponding to the relationship type, and $\sigma(\cdot)$ is the nonlinear activation function. Through the relationship self-attention mechanism, the model can achieve selective semantic propagation in heterogeneous structures, enabling node representations to combine structural dependencies with semantic constraints, effectively modeling multidimensional financial interactions.

In the generative reasoning stage, the model introduces a conditional generation mechanism, using knowledge graph embedding as a priori to guide the generator to learn the distribution characteristics of potential fraudulent behaviors. The generation process can be formalized as follows:

$$z \sim p(z), \quad \hat{x} = G(z, H_G) \tag{3}$$

where z represents the latent variable, H_G represents the graph structure embedding, and $G(\cdot)$ is the generative mapping function. The model models the latent distribution of account interactions and fund flows in the generative space to capture the generative logic of anomalous behavior. By leveraging the structural semantic constraints provided by the knowledge graph, the generator can generate behavioral patterns that conform to financial logic, enabling the inference and identification of fraudulent patterns without relying on explicit labels.

4. Experimental Results

4.1. Dataset

This study uses the IEEE-CIS Fraud Detection Dataset as the primary source of experimental data. The dataset consists of real online transaction records and is widely used in research on financial risk identification and fraud detection. It covers multiple dimensions such as transaction behavior, account characteristics, device information, and network environment. The dataset contains about 5.9 million transaction samples, including a large number of labeled fraud and non-fraud cases. It features high-dimensional sparse data and an imbalanced class distribution. Each transaction record includes fields such as timestamp, transaction amount, payment method, account relationships, and device identifiers, providing rich semantic and structured information for studying complex financial behaviors.

A notable characteristic of this dataset is its multimodal and heterogeneous nature. In addition to basic numerical and categorical variables, it contains many anonymized derived features that reveal potential connections and behavioral patterns among different transaction entities. For example, device fingerprints and network domain fields can model cross-account login relationships, while address and email prefix fields contain implicit social relationship information. This multidimensional and heterogeneous feature structure provides a solid data foundation for constructing knowledge graphs, allowing the model to learn semantic associations and behavioral dependencies among entities at a higher level.

In addition, the dataset spans a long period of time and covers multiple temporal and contextual scenarios, reflecting dynamic changes in market conditions and user behaviors. This characteristic helps evaluate the model's adaptability in non-stationary financial environments. By graph-based

 doi:10.20944/preprints202512.2142.v1

modeling of transaction records and semantic association mining, the study can comprehensively capture the structural features and potential propagation paths of fraudulent behaviors. This provides realistic, complex, and representative data support for the development of knowledge-graph-guided generative financial fraud detection models.

4.2. Experimental Results

This paper first gives the results of the comparative experiment, as shown in Table 1.

Table 1. Comparative experimental results.

Model	AUC	ACC	Precision	Recall
Transformer [32]	0.931	0.902	0.887	0.864
VAE [33]	0.918	0.889	0.875	0.853
GAN [34]	0.924	0.896	0.880	0.861
WGAN [35]	0.936	0.911	0.892	0.875
Ours	0.957	0.928	0.911	0.898

From the overall results, the proposed knowledge-graph-guided generative financial fraud detection model outperforms all comparison methods across multiple metrics, with particularly strong performance in AUC and ACC. Compared with traditional generative models, the proposed model introduces structured knowledge constraints into the feature space, enabling it to maintain high confidence under complex transaction patterns. The AUC value of 0.957 indicates a stronger discriminative ability in distinguishing fraudulent and normal transactions, while the ACC value of 0.928 further confirms the model's stability and generalization capability. These results demonstrate that the knowledge-guided mechanism effectively enhances global semantic consistency in feature representation and pattern generation, improving the overall efficiency of risk identification.

A further comparison of Precision and Recall shows that the proposed approach achieves a better balance between the two metrics. Traditional generative models, such as GAN and WGAN, can capture certain latent features of fraudulent samples but often suffer from low recall due to the lack of semantic relational modeling. By introducing knowledge graph embeddings in the generative stage, the proposed model produces latent sample spaces that better align with real-world financial semantics. As a result, Recall improves significantly, reaching 0.898. This indicates that the model can effectively identify more potential risky transactions in complex and dynamic fraud scenarios, reducing missed detections and demonstrating strong adaptive and structural reasoning capabilities.

Overall, the proposed model maintains high precision while significantly enhancing its ability to model multidimensional relationships and cross-entity dependencies. This knowledge-graph-guided generative detection framework breaks through the limitations of traditional models that rely only on statistical features or deep representations. It allows the model to understand the generative mechanism of fraudulent behavior from both structural and semantic perspectives. The results verify the effectiveness of semantic constraints in generative modeling, showing that the model remains robust in complex and non-stationary financial environments. This reflects not only the theoretical innovation of the method but also provides a new technical pathway for building interpretable and scalable financial risk control systems.

This paper also presents an experiment on the sensitivity of knowledge graph embedding dimension to Recall, and the experimental results are shown in Figure 2.

From the results shown in the figure, it can be observed that the change in knowledge graph embedding dimension has a significant impact on the model's recall performance. When the embedding dimension is low (such as 32 or 64), the Recall value is relatively small, indicating that the semantic representation ability of the graph structure is insufficient to capture the complex relationships among financial transaction entities. Because the embedding space is too small, the diversity of entity features and the richness of interaction semantics are limited. As a result, the model

 doi:10.20944/preprints202512.2142.v1

struggles to reconstruct potential fraud patterns during the generation stage, leading to reduced capability in identifying risky samples.

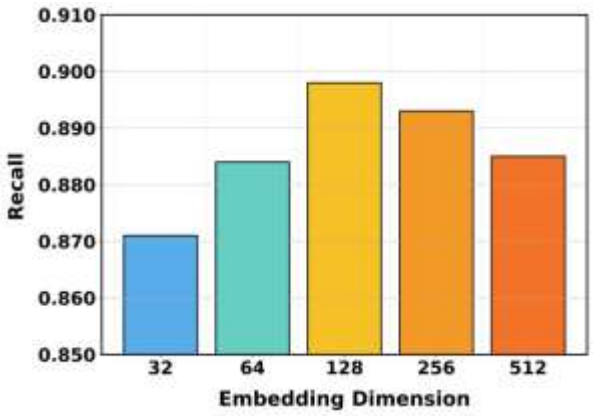

Figure 2. Experiment on the sensitivity of the knowledge graph embedding dimension to recall.

When the embedding dimension increases to 128, the model achieves its peak Recall, suggesting that the semantic alignment between knowledge representation and generative reasoning is most effective at this scale. Higher-dimensional embeddings preserve richer semantic associations and contextual structures of nodes, allowing the generative model to learn more comprehensive features of fraudulent behavior in the latent space. This result shows that moderately increasing the representation capacity of the knowledge graph enhances the model's ability for risk perception and behavioral interpretation, thereby improving robustness in complex financial environments.

When the embedding dimension increases further (such as 256 or 512), the Recall slightly declines. This is mainly because high-dimensional embeddings introduce redundant features and noise, leading to sparse semantic distributions and weakening the discriminative boundaries of representations in the latent space. An excessively large feature space not only increases computational complexity but may also cause overfitting, resulting in unstable performance on unseen samples. This trend indicates that higher embedding dimensions do not necessarily yield better results, and a balance must be achieved between representational capacity and structural constraints.

Overall, this experiment verifies the crucial role of the knowledge representation dimension in generative financial fraud detection. An appropriate embedding scale can significantly enhance the model's semantic modeling ability and accuracy in capturing fraudulent behaviors, while dimensions that are too low or too high may disrupt the synergy between the generative model and the knowledge graph. Therefore, model design should consider the complexity of financial data and the density of graph structures to determine the optimal embedding dimension, achieving a dynamic balance between high recall and stable detection performance.

5. Conclusions

This paper proposes a knowledge-graph-guided generative financial fraud detection model. By introducing structured semantic constraints into the generative framework, the model achieves deep modeling of complex financial relationships and generative reasoning of potential risk patterns. The method maps multi-source heterogeneous transaction data into a knowledge graph space, allowing the model to explicitly capture the logical associations among accounts, devices, and fund flows during feature learning. This effectively addresses the limitations of traditional deep learning models in relational reasoning and semantic interpretation. The results show that this strategy, which combines generative modeling and knowledge enhancement, maintains stable detection performance under non-stationary and noisy environments, providing a new intelligent paradigm for financial risk control systems.

 doi:10.20944/preprints202512.2142.v1

From a methodological perspective, the core contribution of this study lies in realizing a knowledge-driven generative reasoning mechanism. By embedding the knowledge graph as a prior constraint within the generative model, the system can generate potential fraud patterns and provide semantic-level behavioral explanations even when transaction samples are missing or ambiguous. This mechanism not only enhances the model's generalization and adaptability but also extends fraud detection from traditional pattern recognition to semantic understanding and causal reasoning. The proposed framework is highly scalable and can be applied to various financial scenarios, including cross-platform transaction monitoring, payment system security assessment, and credit risk management.

At the practical level, the proposed model has significant implications for building trustworthy and interpretable financial risk control systems. Guided by the knowledge graph, the generative model can not only identify high-risk behaviors but also trace their causes and propagation paths, providing transparent decision support for regulatory agencies and financial institutions. Moreover, the model's structure is highly generalizable and can be transferred to related areas such as insurance fraud detection, anti-money-laundering monitoring, and blockchain security analysis. Through the integration of structured knowledge and generative reasoning, the framework offers both theoretical foundations and technical support for research on automation and interpretability in intelligent financial security.

Future research can be further extended in several directions. One direction is to introduce temporal knowledge graphs and dynamic generation mechanisms to improve the model's ability to capture long-term dependencies and cross-period risk patterns. Another is to combine causal representation learning and uncertainty estimation to achieve more precise modeling of high-dimensional financial semantics. In addition, integrating generative fraud detection with large-scale pre-trained language models or multimodal knowledge systems may help overcome the limitations of traditional financial risk control models in data sparsity, cross-domain transfer, and interpretability. Overall, this study provides a sustainable technological pathway for intelligent financial security and lays an important foundation for building future risk control systems that are trustworthy, transparent, and adaptive.

References

1. F. Shi and C. Zhao, "Enhancing financial fraud detection with hierarchical graph attention networks: A study on integrating local and extensive structural information," Finance Research Letters, vol. 58, Article 104458, 2023.
2. L. Yan, Q. Wang and C. Liu, "Semantic Knowledge Graph Framework for Intelligent Threat Identification in IoT," 2025.
3. 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.
4. X. Song, Y. Huang, J. Guo, Y. Liu and Y. Luan, "Multi-Scale Feature Fusion and Graph Neural Network Integration for Text Classification with Large Language Models," arXiv preprint arXiv:2511.05752, 2025.
5. L. Meng, H. Mostafa, M. Nassar et al., "Generative graph augmentation for minority class in fraud detection," Proceedings of the 32nd ACM International Conference on Information and Knowledge Management, pp. 4200-4204, 2023.
6. S. Wang, Z. Zhang, L. Fang et al., "Corporate Fraud Detection in Rich-yet-Noisy Financial Graph," arXiv preprint arXiv:2502.19305, 2025.
7. 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.
8. S. Han, "AI-Driven Predictive Modeling for System Performance and Resource Management in Microservice Architectures," Journal of Computer Technology and Software, vol. 4, no. 10, 2025.
9. S. Lyu, M. Wang, H. Zhang, J. Zheng, J. Lin and X. Sun, "Integrating Structure-Aware Attention and Knowledge Graphs in Explainable Recommendation Systems," arXiv preprint arXiv:2510.10109, 2025.

 doi:10.20944/preprints202512.2142.v1

10. J. Tang, H. Gu, D. B. Vuković et al., "Fraud detection in multi-relation graph: Contrastive Learning on Feature and Structural Levels," Neurocomputing, vol. 637, Article 130063, 2025.
11. W. Zhu and Z. Chen, "An Intelligent Financial Fraud Detection Model Using Knowledge Graph-Integrated Deep Neural Network," Journal of Circuits, Systems and Computers, vol. 33, no. 15, Article 2450267, 2024.
12. S. Cai and Z. Xie, "Explainable fraud detection of financial statement data driven by two-layer knowledge graph," Expert Systems with Applications, vol. 246, Article 123126, 2024.
13. J. Jeon, J. Ahn and N. Kim, "Effects of Scale Regularization in Fraud Detection Graphs," Electronics, vol. 14, no. 18, Article 3660, 2025.
14. R. Liu, R. Zhang and S. Wang, "Graph Neural Networks for User Satisfaction Classification in Human-Computer Interaction," arXiv preprint arXiv:2511.04166, 2025.
15. H. Liu, "Improving KPI Time Series Anomaly Detection in Cloud Computing Environments through Graph Neural Network-Based Structural and Temporal Modeling," 2023.
16. Y. Kang, "Machine Learning Method for Multi-Scale Anomaly Detection in Cloud Environments Based on Transformer Architecture," Journal of Computer Technology and Software, vol. 3, no. 4, 2024.
17. 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, June 2025.
18. X. Chen, S. U. Gadgil, K. Gao, Y. Hu and C. Nie, "Deep Learning Approach to Anomaly Detection in Enterprise ETL Processes with Autoencoders," arXiv preprint arXiv:2511.00462, 2025.
19. X. Liu, Y. Qin, Q. Xu, Z. Liu, X. Guo and W. Xu, "Integrating Knowledge Graph Reasoning with Pretrained Language Models for Structured Anomaly Detection," 2025.
20. Y. Xing, Y. Deng, H. Liu, M. Wang, Y. Zi and X. Sun, "Contrastive Learning-Based Dependency Modeling for Anomaly Detection in Cloud Services," arXiv preprint arXiv:2510.13368, 2025.
21. Z. Cheng, "Enhancing Intelligent Anomaly Detection in Cloud Backend Systems through Contrastive Learning and Sensitivity Analysis," 2024.
22. L. Liu, "Time Series Prediction of Backend Server Load via Deep Learning and Attention Mechanisms," 2025.
23. Y. Xing, "Enhancing Advertising Recommendation Performance via Integrated Causal Inference and Exposure Bias Correction," 2023.
24. P. Xue and Y. Yi, "Integrating Context Compression and Structural Representation in Large Language Models for Financial Text Generation," Journal of Computer Technology and Software, vol. 4, no. 9, 2025.
25. 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.
26. N. Lyu, Y. Wang, Z. Cheng, Q. Zhang and F. Chen, "Multi-Objective Adaptive Rate Limiting in Microservices Using Deep Reinforcement Learning," arXiv preprint arXiv:2511.03279, 2025.
27. 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.
28. R. Ying, J. Lyu, J. Li, C. Nie and C. Chiang, "Dynamic Portfolio Optimization with Data-Aware Multi-Agent Reinforcement Learning and Adaptive Risk Control," 2025.
29. H. Feng, Y. Wang, R. Fang, A. Xie and Y. Wang, "Federated Risk Discrimination with Siamese Networks for Financial Transaction Anomaly Detection," 2025.
30. 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.
31. Z. Xu, J. Xia, Y. Yi, M. Chang and Z. Liu, "Discrimination of Financial Fraud in Transaction Data via Improved Mamba-Based Sequence Modeling," 2025.
32. J. Lin, X. Guo, Y. Zhu et al., "FraudGT: a simple, effective, and efficient graph transformer for financial fraud detection," Proceedings of the 5th ACM International Conference on AI in Finance, pp. 292-300, 2024.
33. S. Obushnyi, D. Virovets, A. Ramskyi et al., "Variational Autoencoders for Detecting Anomalous and Fraudulent Transactions in Financial Systems," 2025.

 doi:10.20944/preprints202512.2142.v1

34. Z. Meng, Y. Xie and J. Sun, "Detecting Credit Card Fraud by Generative Adversarial Networks and Multi-head Attention Neural Networks," IAENG International Journal of Computer Science, vol. 50, no. 2, 2023.
35. A. K. Gangwar and V. Ravi, "Wip: Generative adversarial network for oversampling data in credit card fraud detection," Proceedings of the International Conference on Information Systems Security, Springer International Publishing, pp. 123-134, 2019.