

Article

Not peer-reviewed version

Graph Neural Network and Temporal Sequence Integration for AI-Powered Financial Compliance Detection

[Weiyao Xu](#) , Mohan Jiang , Shiqing Long , [Yuxiu Lin](#) , Kunyuan Ma , [Zhen Xu](#) *

Posted Date: 26 August 2025

doi: 10.20944/preprints202508.1902.v1

Keywords: trajectory graph modeling; graph neural network; time series modeling; compliance detection; trading behaviors



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

Graph Neural Network and Temporal Sequence Integration for AI-Powered Financial Compliance Detection

Weiyao Xu ¹, Mohan Jiang ², Shiqing Long ³, Yuxiu Lin ⁴, Kunyuan Ma ⁵ and Zhen Xu ^{6,*}

- ¹ Fordham University, New York, USA
- ² New York University, New York, USA
- ³ Independent Researcher, Shanghai, China
- ⁴ Columbia University, New York, USA
- ⁵ New York University, New York, USA
- ⁶ Independent Researcher, Shanghai, China
- * Correspondence: xuzhen8309@hotmail.com

Abstract

Paper proposes a compliance anomaly detection algorithm for financial transaction data, addressing the challenges of behavioral complexity, structural dependence, and strong temporal characteristics. The method constructs a dynamic trajectory graph with multi-type edges based on account interactions, providing a unified representation of structural relationships and temporal order in financial transactions. A graph neural network is employed to embed node relationships and extract high-order behavioral structure features. A gated recurrent neural network is then used to model the temporal dynamics of trajectory sequences and capture the evolution patterns of transaction behaviors. In the classification stage, a joint optimization strategy combines contrastive learning with an anomaly scoring function to enhance intra-class cohesion and inter-class separation, improving the model's ability to identify abnormal behaviors. To evaluate the effectiveness of the proposed method, experiments are conducted across various sensitivity dimensions, including time window length, trajectory length, node dropout rate, anomaly ratio, and class imbalance. The results show that the method achieves superior detection performance and strong adaptability under diverse perturbation conditions, demonstrating the potential of structure-temporal joint modeling for financial compliance anomaly detection tasks. CCS CONCEPTS: Computing methodologies~Machine learning~Machine learning approaches.

Keywords: trajectory graph modeling; graph neural network; time series modeling; compliance detection; trading behaviors

1. Introduction

In modern financial markets, compliance regulation plays a critical role in preventing systemic risks, maintaining market order, and protecting investor interests [1]. As financial instruments grow more complex and business models become increasingly diverse, traditional compliance detection mechanisms that rely on rule libraries, expert knowledge, or static indicators show clear limitations. These methods often suffer from inefficiency, significant delays, and poor adaptability to emerging behavioral patterns. Against this backdrop, the rapid development of financial technology has driven the rise of data-driven intelligent compliance detection approaches. These methods offer new directions for identifying complex behaviors and achieving dynamic risk control. A key challenge lies in how to effectively capture the evolving characteristics of trading behaviors, which is essential for building efficient anomaly detection algorithms [2].

Trading behavior trajectories, defined as sequences of operations by a trading entity over time, contain rich semantic relationships, temporal dependencies, and interactive features. Unlike isolated and static transaction records, trajectory data reveals evolving behavioral patterns and latent intentions. This makes it a crucial source of evidence for detecting sophisticated violations such as masked operations or coordinated activities across multiple accounts. However, due to their strong temporal nature, high-dimensional sparsity, and heterogeneous structure, traditional rule-based or statistical feature extraction methods struggle to capture their deep structural information. As a result, these methods often fail to detect subtle or distribution-shifted anomalies. A unified modeling framework that integrates trajectory structure, behavioral context, and semantic dynamics is urgently needed to improve the depth and accuracy of compliance detection [3].

Graph structures offer a natural form for modeling complex behavioral relationships [4–7]. It has been demonstrated that such representations are highly effective in modeling transaction paths, fund flows, and account interaction patterns, and their applicability has been successfully extended to domains such as large language models (LLMs) [8–10], recommendation systems [11–13], and distributed computing architectures [14–16]. In financial transaction scenarios, trading trajectories can be abstracted as dynamic graphs with multiple nodes, edge types, and temporal information. By introducing graph modeling mechanisms, it is possible to preserve the structural integrity of behavioral sequences while uncovering abnormal substructures and potential risk propagation paths through node interactions. With the help of structure-aware techniques such as graph neural networks, it becomes feasible to embed behavioral patterns and dynamically capture context dependencies [17–19]. This enhances the model's capability to perceive complex compliance anomalies. Furthermore, graph structures support multi-level information aggregation, which lays a solid foundation for integrating diverse risk signals [20].

In addition, compliance anomalies often exhibit strong asymmetry and concealment. These anomalies are rarely identifiable through single indicators or isolated actions. Instead, they require comprehensive modeling of temporal continuity, structural jumps, and contextual irregularities within behavioral sequences. Effective detection algorithms should be capable of multi-modal fusion, temporal modeling, and structural alignment. They must separate weak abnormal signals from a large volume of normal trajectories. This demands that models understand global deviations in account behaviors while also detecting local anomalies at a fine-grained scale. The integration of graph representation learning with behavioral trajectory modeling supports both local behavioral shifts and the overall rationality of behavior paths. This provides a novel technical path for financial compliance anomaly detection.

In summary, constructing graph-based representations around trading behavior trajectories and integrating deep learning techniques such as graph neural networks has become a major direction in financial compliance anomaly detection. This direction expands the modeling capacity for behavioral data from a theoretical perspective and offers more intelligent and efficient tools for risk identification in practice. In-depth research on this topic contributes not only to improving the intelligence and real-time capabilities of financial supervision but also provides a general modeling framework and technical reference for other anomaly detection tasks centered on sequential trajectories.

2. Relevant Literature

Recent advances in deep learning have catalyzed substantial progress in the field of anomaly detection, with a growing emphasis on unified models that integrate temporal, structural, and multi-modal information. Central to this evolution are hybrid neural architectures, which blend recurrent and transformer components to capture both local sequential dependencies and global contextual signals. These designs enable the effective modeling of long-term behavioral evolution and rich feature interactions, thereby enhancing the detection of subtle and evolving anomalies in complex data streams [21].

To further improve model robustness in the absence of dense labels, unsupervised learning frameworks—particularly those based on contrastive objectives—have gained traction. By leveraging

temporal augmentations and representation similarity learning, these approaches facilitate the discovery of rare or previously unseen anomalous behaviors, reducing reliance on manual annotation and domain-specific heuristics [22]. In parallel, ensemble and data balancing techniques have been employed to address severe class imbalance, ensuring that deep models remain sensitive to minority anomaly classes while maintaining high generalization [23].

Graph neural networks have become an essential methodological foundation for encoding intricate relational structures. Advances in heterogeneous GNNs and attention-based mechanisms allow for the flexible representation of multi-type nodes and edges, supporting the identification of complex interactions and structural irregularities that may signify anomalous events [24]. These methods have shown particular promise in learning high-order connectivity patterns and in adapting to dynamic, rapidly changing environments—further supported by architectures tailored for high-frequency data and hierarchical feature extraction [25], as well as probabilistic generative models such as mixture density networks, which provide calibrated anomaly scores through uncertainty estimation [26]. Recent innovations in neural backbone design, including separable convolutions and attention layers, further contribute to efficient and discriminative feature learning in structured sequences [27].

Extending beyond pointwise detection, a new class of unsupervised and diffusion-based anomaly detection techniques harnesses structure-aware denoising to learn robust latent representations that respect both temporal continuity and relational dependencies [28]. Temporal graph neural networks represent a natural progression in this space, offering direct modeling of structural evolution across sequential observations and facilitating time-respecting analysis of dynamic behavioral trajectories [29]. The broader application of deep graph modeling supports risk identification and dependency reasoning in structured data by uncovering latent operational patterns through high-fidelity representations [30]. Similarly, heterogeneous network learning has proven effective in revealing subtle, multi-hop associations within complex relational graphs, highlighting the power of graph-based abstractions for anomaly detection [31].

The methodological landscape has also been enriched by advances in sequence modeling, notably through self-attention-based architectures such as transformers. These models are capable of capturing long-range dependencies and effectively integrating multi-modal features, setting new benchmarks in both predictive accuracy and representational expressiveness [32–34]. Complementing these architectures, causal and invariant representation learning methods have been developed to improve model robustness under shifting distributions, enabling reliable anomaly detection even in the presence of changing or adversarial behavior [35,36]. Generative models, particularly those inspired by diffusion processes, contribute by providing a principled mechanism for modeling sequence normality and detecting deviations through probabilistic reconstruction [37].

In parallel, representation learning research continues to emphasize the importance of effective data augmentation and contrastive learning strategies. Techniques such as relational masking, feature perturbation, and temporal shuffling have been demonstrated to improve the generalizability and invariance of learned representations, further supporting robust anomaly detection in complex environments [38]. Probabilistic graphical models and variational inference have emerged as vital tools for handling class imbalance and uncertainty, offering principled solutions for rare-event detection and calibrated risk assessment [39]. Advances in few-shot learning and dual-loss optimization, often implemented within transformer architectures, have also enhanced the adaptability of deep models to low-data regimes and shifting patterns [40]. Moreover, deep contextual modeling has enabled more nuanced semantic understanding of sequential and textual data, broadening the scope of anomaly detection techniques [41].

Finally, a broader methodological perspective is informed by foundational work in risk forecasting, early-warning systems, and planning under uncertainty. The strategies and frameworks developed in these areas reinforce the value of anticipatory modeling and structured risk quantification, principles that directly inform the design and evaluation of contemporary anomaly detection pipelines [42,43]. Taken together, these methodological advances provide a comprehensive

toolkit for structure–temporal joint modeling, supporting the development of more intelligent, robust, and adaptive solutions for anomaly detection in complex sequential and relational data contexts.

3. Proposed Methodology

This paper proposes a financial compliance anomaly detection algorithm that integrates transaction behavior trajectory graphs. The overall method consists of trajectory graph construction, graph representation learning, behavior modeling, and anomaly scoring. First, the continuous transaction behavior is abstracted into a multi-node multi-edge trajectory graph, where each node represents an account or transaction event, and each edge represents the fund flow relationship between accounts or the chronological behavior connection. The model architecture is shown in Figure 1.

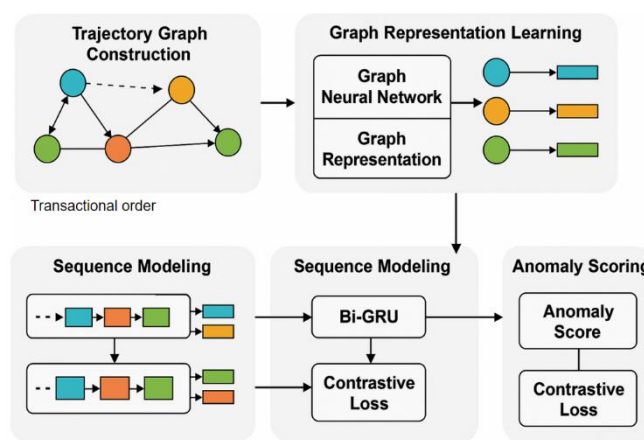


Figure 1. Overall model architecture.

To characterize the intensity and timing characteristics of transactions, the time decay function is introduced to define the edge weight as:

$$w_{ij} = \exp(-\lambda \cdot \Delta t_{ij})$$

Δt_{ij} is the time interval between nodes i and j , and λ is the time decay factor, which is used to control the impact of historical behavior. This graph structure retains the coupling characteristics of transaction activities in time and structure, providing a basic expression for subsequent modeling.

In the graph representation learning stage, a structure-aware embedding method based on a graph neural network is used to encode the representation of nodes. Specifically, the graph attention mechanism is used to aggregate neighbor node information at each layer, and the representation update process is as follows:

$$h_t = GRU(z_t, h_{t-1})$$

z_t is the node embedding representation corresponding to time t , and h_t is the hidden state at the current moment, which is used to capture the dependency of the trajectory over time. This process realizes the fusion of structural representation and time series modeling, effectively improving the model's ability to model behavior changes and sequence shifts.

Finally, based on the representation vector of the entire trajectory, an anomaly scoring function is designed to measure the degree of compliance deviation of the current behavior sequence. The final encoding result is mapped to the anomaly score space, and the scoring function is defined as:

$$s = \|h_{traj} - \hat{h}_{nom}\|_2^2$$

Where h_{traj} represents the representation vector of the current trajectory, and \hat{h}_{norm} is the aggregate representation of the baseline normal behavior. In addition, the contrast constraint loss is introduced to enhance the distinction between classes, and the loss function is:

$$L = L_{rec} + \alpha \cdot L_{contrast}$$

Where L_{rec} is the reconstruction error, $L_{contrast}$ is the contrast loss term, and α is the balance coefficient. Through joint optimization, the model can accurately judge the compliance of financial transactions based on structural perception and temporal dependency.

4. Experimental Dataset

This study uses the publicly available financial transaction dataset, the Elliptic Dataset, as the primary data source for the compliance anomaly detection task. The dataset consists of real transaction data from the Bitcoin blockchain. It includes over 200,000 nodes and more than 200,000 transaction edges. Each node is labeled as a normal transaction, an illicit transaction, or an unknown transaction. The dataset is widely used in tasks such as anti-money laundering and financial behavior identification. Each node represents a transaction event and includes a timestamp, a feature vector, and label information. It supports the design of tasks that combine graph structure modeling with time series modeling.

The Elliptic Dataset is notable for its natural graph structure and strong temporal characteristics. These features make it highly suitable for building anomaly detection models based on trajectory graphs. Each transaction can be viewed as a node in the graph. Transactions are connected by directed edges that represent fund transfers, forming a transaction network. The node features include statistical attributes, aggregated indicators, and on-chain behavioral descriptions. These features are structured and semantically rich, providing reliable inputs for both graph neural networks and temporal modeling approaches.

To better adapt the data to the proposed modeling framework, the preprocessing stage divides the transactions into a sequence of sliding time windows. The edge structure of the graph is retained to construct the transaction trajectory graph. Based on this setup, continuous transaction trajectories are extracted and transformed into dynamic graph representations. This unified data format supports both behavioral modeling and compliance classification. Due to its realism, complexity, and wide acceptance, the dataset is frequently used in financial risk control and regulatory technology research.

5. Results and Analysis

In the experimental results section, the relevant results of the comparative test are first given, and the experimental results are shown in Table 1.

Table 1. Comparative experimental results.

Method	AUC-ROC	F1-Score	Accuracy
STGNN[44]	0.951	0.842	0.911
MDST-GNN[45]	0.963	0.859	0.925
AMNet[46]	0.940	0.830	0.898
GraphCAD[47]	0.918	0.795	0.875
Ours	0.973	0.875	0.937

As shown in the table, the proposed method achieves significantly better performance in terms of AUC-ROC compared to other baseline models. This indicates stronger discrimination capability in distinguishing between normal transactions and compliance anomalies. In financial transaction trajectories, anomalous behavior often involves multi-hop transfers and concealed paths. Traditional models typically fail to capture such latent structures, leading to blurred decision boundaries. By modeling trajectory graphs and incorporating structural awareness, the proposed method enhances sensitivity to higher-order graph patterns and improves overall classification performance.

In terms of F1-Score, the model also outperforms others. This demonstrates its ability to maintain high recall while preserving precision, which is essential for real-world regulatory applications that demand low miss rates. Given that compliance anomalies are often rare and diverse, F1 serves as a valuable metric for evaluating model performance on minority classes. The proposed method integrates both structural dependencies among transaction nodes and the evolution of behavioral sequences. This allows unified modeling of suspicious behaviors across accounts and time spans, thereby improving intra-class cohesion and inter-class separation.

For Accuracy, although all models perform at a high level, the proposed method still shows an advantage while maintaining prediction stability. This suggests that the model can accurately identify violations while effectively avoiding false positives on normal transactions. This improvement is attributed to the dynamic modeling of heterogeneous edge weights and temporal decay during the graph representation learning phase, which enhances the semantic distinction between different types of transaction behavior.

The proposed framework models transactional behavior through structured trajectory graphs, integrating graph neural networks with temporal encoding mechanisms. This design markedly enhances the capability to identify anomalous patterns within complex transaction pathways. In contrast to existing graph-based anomaly detection approaches, the method delivers richer structural graph representations and achieves a more effective fusion of behavioral semantics with compliance deviations in an end-to-end manner. As a result, it significantly advances both the practicality and robustness of financial compliance anomaly detection systems. This paper also gives the impact of changes in node loss ratio on model recognition accuracy, and the experimental results are shown in Figure 2.




Figure 2. The impact of node loss ratio changes on model recognition accuracy.

As shown in the experimental results in Figure 2, both the accuracy and F1-score of the model exhibit a clear downward trend as the node dropout ratio increases. This indicates that node information is highly sensitive for modeling financial behavior trajectories. In particular, when the dropout ratio increases from 0% to 50%, the accuracy drops from 0.937 to approximately 0.802, and the F1-score decreases from 87.5% to around 70%. This shows that node loss significantly weakens the structural integrity of trajectories and impairs the model's ability to distinguish abnormal behaviors.

This observation reinforces the structural dependence of the proposed trajectory graph modeling framework. In transaction trajectories, nodes encode essential temporal and behavioral

relationship information. When nodes are missing, transaction chains may be interrupted and contextual cues lost, limiting the ability of graph neural networks to capture global dependencies. This weakens the expressive capacity of behavioral representations and degrades anomaly detection accuracy. The disruption illustrates how even localized perturbations in financial graph data can undermine the overall stability of the model. From the trend in F1-score, the impact on recall is particularly notable. When the node dropout exceeds 30%, the F1-score declines more rapidly. This indicates that the model becomes more likely to miss anomalous behaviors under high dropout conditions. The result suggests that the model still faces challenges in robustness when dealing with incomplete graph structures. It is necessary to enhance its resistance to noise through mechanisms such as graph completion or robustness-aware learning strategies.

This paper also gives the impact of changes in the proportion of abnormal samples on the robustness of the model, and the experimental results are shown in Figure 3.




Figure 3. The impact of changes in the proportion of abnormal samples on model robustness.

As shown in Figure 3, the proportion of anomalous samples has a clear impact on model robustness, especially in terms of F1-score, which exhibits a nonlinear trend. When the anomaly ratio is low (between 1% and 10%), the model performance improves as the proportion increases. The F1-score rises from 84.1% to 87.5%. This indicates that a moderate increase in anomalous samples helps the model learn more representative patterns of abnormal behavior, enhancing its recognition ability under imbalanced data conditions.

However, when the anomaly proportion continues to grow beyond 20%, the F1-score starts to decline. The performance drops sharply to 82.3% and 78.5% when the anomaly ratio reaches 40% and 50%, respectively. This suggests that an excessive number of anomalies interferes with the model's ability to learn sparse and rare abnormal patterns. The model becomes more prone to misclassifying normal samples, which weakens the class boundary. This also reflects a decrease in structural stability when the model encounters extremely skewed data distributions.

The results highlight an important observation: model robustness in compliance detection depends not only on the presence of anomalies but also on the quality and proportion of those anomalies. When the ratio is appropriate, the model is better able to construct clear decision boundaries for anomalies. Once the anomaly proportion becomes too high, the model loses a stable understanding of “normal” behavior, leading to performance degradation. This underlines the importance of dynamically controlling anomaly proportions during training to preserve class discriminability.

This paper also presents the recognition stability analysis under the change of transaction trajectory length, and the experimental results are shown in Figure 4.




Figure 4. Analysis of recognition stability under changing transaction trajectory length.

As shown in the experimental results in Figure 4, the length of transaction trajectories has a significant impact on the model's recognition performance. As the trajectory length increases from 2 to 8, the model's accuracy rises markedly from 0.821 to 0.889. This indicates that shorter trajectories do not provide sufficient contextual information, which limits the model's ability to capture behavioral patterns. Medium-length trajectories effectively cover the core stages of behavioral evolution. They help the model extract key structural and sequential features, thereby enhancing its classification capability.

However, when the trajectory length is further increased to 14, accuracy begins to decline. This suggests that excessively long trajectories may introduce redundant or noisy information. Such input can blur the model's focus and reduce its generalization ability. This result indicates a nonlinear trade-off between precision and information density in trajectory modeling. The model should avoid overload and preserve its focus on essential behavioral signals when constructing transaction sequence representations.

On the other hand, the trend of recall differs from that of accuracy. Recall increases steadily with longer trajectories, rising from 78.5% to 88.1%. This suggests that longer sequences help the model detect potential anomalies more effectively and improve its ability to retrieve rare abnormal samples. This may be due to the presence of richer behavioral clues in longer trajectories, which make the model more sensitive when identifying edge cases, thereby improving coverage in anomaly detection.

Overall, the experiment confirms the critical role of trajectory length in financial behavior modeling. It also demonstrates the flexibility and adaptability of the proposed method in structure-sequence fusion modeling. By carefully controlling trajectory length, it is possible to achieve a dynamic balance between precision and recall. This provides a more stable and reliable technical path for anomaly detection in complex financial scenarios.

This paper also gives the impact of changing the time window length on the model performance, and the experimental results are shown in Figure 5.

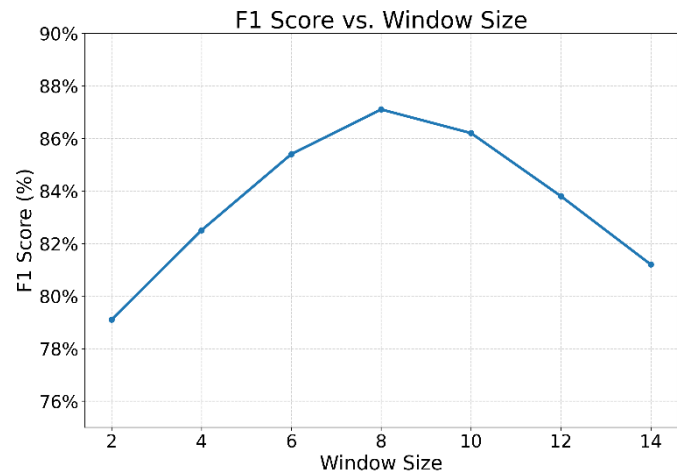


Figure 5. The impact of changing the time window length on model performance.

As shown in Figure 5, the length of the time window has a significant effect on the model's F1-score. As the window increases from 2 to 8, the F1-score rises steadily, reaching a peak of approximately 87%. This indicates that within a suitable period, the model can more effectively capture temporal patterns and abnormal fluctuations in financial transactions. This improves both the accuracy and consistency of anomaly detection. When the window is too short, the information is insufficient to construct a complete semantic context, limiting the model's expressive capacity.

When the time window exceeds 8, the model's performance begins to decline. At a window size of 14, the F1-score drops to 82.5%. This suggests that an overly long window introduces redundant or irrelevant historical behaviors. Such noise interferes with the model's ability to identify anomaly signals. Long periods may also blur the local characteristics of events, making it difficult for the model to focus on critical risk points. This weakens the clarity and stability of decision boundaries.

The experimental results show that the time window plays a dual role in constructing transaction trajectories. It functions as both an information filter and a modeling perspective controller. An appropriate window length ensures sufficient feature coverage while avoiding the negative effects of noise on graph learning and temporal modeling. It is a key hyperparameter that significantly influences model performance. The results form a clear inverted U-shaped curve concerning window size.

In conclusion, this experiment further confirms the high sensitivity of the proposed structure-temporal joint modeling mechanism to time granularity. By adjusting the window length, the model's performance can be dynamically optimized for different scenarios. In real-world financial supervision systems, the window should be flexibly set according to the periodicity and density of the target anomalies. This enhances the model's generalization ability and robustness under complex behavioral patterns.

This paper also gives an analysis of the model generalization ability under the condition of imbalanced positive and negative samples, and the experimental results are shown in Figure 6.




Figure 6. Analysis of model generalization ability under the condition of imbalanced positive and negative samples.

Figure 6 presents the generalization performance of the model under different positive-to-negative sample ratios in terms of F1-score. As the class balance improves, the model performance steadily increases. When the ratio is adjusted from 1:9 to 5:5, the F1-score rises from 80.1% to 87.5%. This indicates that a balanced data distribution helps the model better distinguish between normal and abnormal behaviors, significantly enhancing the stability of classification boundaries. The trend shows that the model struggles in highly imbalanced scenarios, while proper ratio adjustment improves its representational capacity and generalization ability.

It is worth noting that when the proportion of positive samples continues to increase (such as 6:4 or 7:3), the F1-score drops significantly. This suggests that too many positive samples weaken the

model's ability to learn from negative samples, which represent normal behavior. The resulting performance degradation is a sign of overfitting. Even in anomaly detection tasks, excessive emphasis on positive samples can disrupt the model's understanding of the real data distribution, reducing detection stability. Therefore, it is important to maintain a relative balance between classes during training to avoid structural bias caused by class imbalance.

From the perspective of generalization, these results reveal the nonlinear impact of data distribution on model performance. In compliance anomaly detection, anomalies are naturally rare and difficult to collect. Thus, building a high-quality and moderately balanced training set is a key step in improving model robustness. By integrating graph structures and behavioral sequence features, the proposed method can effectively learn discriminative representations of transaction trajectories under moderately balanced conditions, showing strong structural adaptability.

6. Conclusions

This paper addresses the problem of anomaly detection in financial compliance scenarios by proposing an algorithm that integrates transaction behavior trajectory graphs. The method constructs structured trajectory graphs to represent dynamic interactions and temporal evolution between accounts. This enhances the model's ability to perceive the integrity of behavior chains and underlying structural patterns. At the modeling level, a unified framework is introduced by combining graph neural networks with sequence modeling mechanisms. This allows the model to achieve stronger expressive capacity and classification accuracy when dealing with complex and heterogeneous data structures. As a result, the model demonstrates superior stability and robustness under various interference and variation conditions.

Through experimental analysis across multiple dimensions of data sensitivity and environmental changes, the proposed method shows strong adaptability to scenarios such as node dropout, class imbalance, and varying window lengths. These results indicate that the deep integration of structure-aware and sequence-based modeling effectively addresses structural disruptions and dynamic fluctuations in financial transaction data. It improves the detection of highly concealed and deceptive anomalous behaviors. At the same time, the use of trajectory graphs provides a unified input representation space for behavior modeling, offering good task transferability and extensibility.

This study provides new methodological insights for the development of financial regulatory technologies. It shows strong practical value and theoretical relevance in real-world applications involving highly complex and multi-dimensional data. Compared to traditional rule-based or single-modality approaches, the proposed method captures latent risk patterns behind behavioral structures more effectively. It offers key technical support for building precise supervision, intelligent auditing, and risk warning systems.

7. Future Work

Future research can be extended in two directions. First, it is possible to explore mechanisms such as heterogeneous graph modeling and cross-graph contrastive learning to strengthen the modeling of potential relationships across multi-source transaction data. Second, the model can be embedded into large-scale real-time regulatory systems. By incorporating incremental learning and graph stream updating strategies, the model can improve its responsiveness and adaptability in dynamic financial environments. Continuous optimization of model architecture and algorithm design may further promote the application of intelligent compliance technologies in finance, government, taxation, and other key domains.

8. USE OF AI

I used AI tools exclusively to assist with grammar and wording; all the core concepts, analysis, and writing were my original work.

References

1. B. I. Adekunle, E. C. Chukwuma-Eke, E. D. Balogun *et al.*, "Developing a digital operations dashboard for real-time financial compliance monitoring in multinational corporations," *International Journal of Scientific Research in Computer Science, Engineering and Information Technology*, vol. 9, no. 3, pp. 728-746, 2023.
2. K. Vasista, "Regulatory compliance and supervision of artificial intelligence, machine learning and also possible effects on financial institutions," *International Journal of Innovative Research in Computer and Communication Engineering*, 2021.
3. Balakrishnan, "Leveraging artificial intelligence for enhancing regulatory compliance in the financial sector," *International Journal of Computer Trends and Technology*, 2024.
4. N. Jiang, W. Zhu, X. Han, W. Huang, and Y. Sun, "Joint graph convolution and sequential modeling for scalable network traffic estimation," *arXiv:2505.07674*, 2025.
5. D. Gao, "Graph neural recognition of malicious user patterns in cloud systems via attention optimization," *Transactions on Computational and Scientific Methods*, vol. 4, no. 12, 2024.
6. M. Gong, "Modeling microservice access patterns with multi-head attention and service semantics," *Journal of Computer Technology and Software*, vol. 4, no. 6, 2025.
7. Y. Sun, R. Meng, R. Zhang, Q. Wu, and H. Wang, "A deep Q-network approach to intelligent cache management in dynamic backend environments," 2025.
8. Q. Wu, "Task-aware structural reconfiguration for parameter-efficient fine-tuning of LLMs," *Journal of Computer Technology and Software*, vol. 3, no. 6, 2024.
9. Y. Ma, G. Cai, F. Guo, Z. Fang, and X. Wang, "Knowledge-informed policy structuring for multi-agent collaboration using language models," *Journal of Computer Science and Software Applications*, vol. 5, no. 5, 2025.
10. Y. Xing, "Bootstrapped structural prompting for analogical reasoning in pretrained language models," *Transactions on Computational and Scientific Methods*, vol. 4, no. 11, 2024.
11. L. Zhu, W. Cui, Y. Xing, and Y. Wang, "Collaborative optimization in federated recommendation: Integrating user interests and differential privacy," *Journal of Computer Technology and Software*, vol. 3, no. 8, 2024.
12. R. Wang, Y. Luo, X. Li, Z. Zhang, J. Hu, and W. Liu, "A hybrid recommendation approach integrating matrix decomposition and deep neural networks for enhanced accuracy and generalization," *Proceedings of the 2025 5th International Conference on Neural Networks, Information and Communication Engineering*, pp. 1778-1782, 2025.
13. Y. Xing, Y. Wang, and L. Zhu, "Sequential recommendation via time-aware and multi-channel convolutional user modeling," *Transactions on Computational and Scientific Methods*, vol. 5, no. 5, 2025.
14. Y. Wang, T. Tang, Z. Fang, Y. Deng, and Y. Duan, "Intelligent task scheduling for microservices via A3C-based reinforcement learning," *arXiv:2505.00299*, 2025.
15. Y. Zou, N. Qi, Y. Deng, Z. Xue, M. Gong, and W. Zhang, "Autonomous resource management in microservice systems via reinforcement learning," *arXiv:2507.12879*, 2025.
16. X. Sun, Y. Yao, X. Wang, P. Li, and X. Li, "AI-driven health monitoring of distributed computing architecture: Insights from XGBoost and SHAP," *Proceedings of the 2024 4th International Conference on Communication Technology and Information Technology*, pp. 480-484, 2024.
17. J. Liu, "Global temporal attention-driven transformer model for video anomaly detection," *Proceedings of the 2025 5th International Conference on Artificial Intelligence and Industrial Technology Applications*, pp. 1909-1913, 2025.

18. J. Zhan, "Single-device human activity recognition based on spatiotemporal feature learning networks," *Transactions on Computational and Scientific Methods*, vol. 5, no. 3, 2025.
19. T. Yang, Y. Cheng, Y. Ren, Y. Lou, M. Wei, and H. Xin, "A deep learning framework for sequence mining with bidirectional LSTM and multi-scale attention," *arXiv:2504.15223*, 2025.
20. A. Adeniran, A. O. Abhulimen, A. N. Obiki-Osafiele *et al.*, "Strategic risk management in financial institutions: Ensuring robust regulatory compliance," *Finance & Accounting Research Journal*, vol. 6, no. 8, pp. 1582-1596, 2024.
21. P. Feng, "Hybrid BiLSTM-transformer model for identifying fraudulent transactions in financial systems," *Journal of Computer Science and Software Applications*, vol. 5, no. 3, 2025.
22. X. Li, Y. Peng, X. Sun, Y. Duan, Z. Fang, and T. Tang, "Unsupervised detection of fraudulent transactions in e-commerce using contrastive learning," *Proceedings of the 2025 4th International Symposium on Computer Applications and Information Technology*, pp. 1663-1667, 2025.
23. Y. Wang, "A data balancing and ensemble learning approach for credit card fraud detection," *Proceedings of the 2025 4th International Symposium on Computer Applications and Information Technology*, pp. 386-390, 2025.
24. 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:2504.08183*, 2025.
25. Q. Bao, J. Wang, H. Gong, Y. Zhang, X. Guo, and H. Feng, "A deep learning approach to anomaly detection in high-frequency trading data," *Proceedings of the 2025 4th International Symposium on Computer Applications and Information Technology*, pp. 287-291, 2025.
26. L. Dai, W. Zhu, X. Quan, R. Meng, S. Chai, and Y. Wang, "Deep probabilistic modeling of user behavior for anomaly detection via mixture density networks," *arXiv:2505.08220*, 2025.
27. X. Du, "Audit fraud detection via EfficiencyNet with separable convolution and self-attention," *Transactions on Computational and Scientific Methods*, vol. 5, no. 2, 2025.
28. H. Xin and R. Pan, "Unsupervised anomaly detection in structured data using structure-aware diffusion mechanisms," *Journal of Computer Science and Software Applications*, vol. 5, no. 5, 2025.
29. Q. R. Xu, "Capturing structural evolution in financial markets with graph neural time series models," 2025.
30. D. Gao, "Deep graph modeling for performance risk detection in structured data queries," *Journal of Computer Technology and Software*, vol. 4, no. 5, 2025.
31. Z. Liu and Z. Zhang, "Graph-based discovery of implicit corporate relationships using heterogeneous network learning," *Journal of Computer Technology and Software*, vol. 3, no. 7, 2024.
32. Y. Yao, "Stock price prediction using an improved transformer model: Capturing temporal dependencies and multi-dimensional features," *Journal of Computer Science and Software Applications*, vol. 5, no. 2, 2025.
33. Y. Wang, "Stock prediction with improved feedforward neural networks and multimodal fusion," *Journal of Computer Technology and Software*, vol. 4, no. 1, 2025.
34. Liu, "Multimodal data-driven factor models for stock market forecasting," *Journal of Computer Technology and Software*, vol. 4, no. 2, 2025.
35. Y. Sheng, "Market return prediction via variational causal representation learning," *Journal of Computer Technology and Software*, vol. 3, no. 8, 2024.
36. Y. Wang, Q. Sha, H. Feng, and Q. Bao, "Target-oriented causal representation learning for robust cross-market return prediction," *Journal of Computer Science and Software Applications*, vol. 5, no. 5, 2025.
37. X. Su, "Predictive modeling of volatility using generative time-aware diffusion frameworks," *Journal of Computer Technology and Software*, vol. 4, no. 5, 2025.
38. M. Wei, H. Xin, Y. Qi, Y. Xing, Y. Ren, and T. Yang, "Analyzing data augmentation techniques for contrastive learning in recommender models," 2025.

39. Y. Lou, J. Liu, Y. Sheng, J. Wang, Y. Zhang, and Y. Ren, "Addressing class imbalance with probabilistic graphical models and variational inference," *Proceedings of the 2025 5th International Conference on Artificial Intelligence and Industrial Technology Applications*, pp. 1238-1242, 2025.
40. X. Han, Y. Sun, W. Huang, H. Zheng, and J. Du, "Towards robust few-shot text classification using transformer architectures and dual loss strategies," *arXiv:2505.06145*, 2025.
41. Y. Qin, "Deep contextual risk classification in financial policy documents using transformer architecture," *Journal of Computer Technology and Software*, vol. 3, no. 8, 2024.
42. P. Efunniyi, A. O. Abhulimen, A. N. Obiki-Osafiele *et al.*, "Strengthening corporate governance and financial compliance: Enhancing accountability and transparency," *Finance & Accounting Research Journal*, vol. 6, no. 8, pp. 1597-1616, 2024.
43. E. C. Chukwuma-Eke, O. Y. Ogunsola, and N. J. Isibor, "A conceptual approach to cost forecasting and financial planning in complex oil and gas projects," *International Journal of Multidisciplinary Research and Growth Evaluation*, vol. 3, no. 1, pp. 819-833, 2022.
44. Yin, A. Qiu, L. Fang *et al.*, "STGNN: A novel spatial-temporal graph neural network for predicting complicated business process performance under multi-event parallelism," *Expert Systems with Applications*, 2025, Art. no. 128391.
45. S. Chen, Y. Liu, Q. Zhang *et al.*, "Multi-distance spatial-temporal graph neural network for anomaly detection in blockchain transactions," *Advanced Intelligent Systems*, 2025, Art. no. 2400898.
46. W. Lin and Y. Liu, "Deep learning-based attention mechanism algorithm for blockchain credit default prediction," *International Journal of Advanced Computer Science & Applications*, vol. 16, no. 2, 2025.
47. F. Liu, G. Guo, Y. Ye *et al.*, "GraphCAD: Leveraging graph neural networks for accuracy prediction handling crosstalk-affected delays," *Proceedings of the 2025 International Symposium on Physical Design*, pp. 125-133, 2025.

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.