

Article

Not peer-reviewed version

---

# Self-Supervised Learning for Financial Statement Fraud Detection with Limited and Imbalanced Data

---

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

Posted Date: 10 November 2025

doi: 10.20944/preprints202511.0452.v1

Keywords: financial fraud detection; self-supervised learning; feature representation; sensitivity analysis



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

# Self-Supervised Learning for Financial Statement Fraud Detection with Limited and Imbalanced Data

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

<sup>1</sup> Babson College, Wellesley, USA

<sup>2</sup> Columbia University, New York, USA

<sup>3</sup> University of Chicago, Chicago, USA

<sup>4</sup> Golden Gate University, San Francisco, USA

\* Correspondence: fangruoyi@gmail.com

## Abstract

This study addresses the challenges of scarce fraudulent samples, complex data distributions, and the limited adaptability of traditional methods in financial statement fraud detection by proposing a self-supervised learning algorithm. The approach first standardizes multidimensional financial indicators to mitigate scale differences, then employs an encoder to construct latent representations that capture high-order nonlinear relationships among indicators. A reconstruction task is introduced as an auxiliary signal, where a decoder approximates the input and minimizes reconstruction error to enhance the fidelity of representations. In parallel, a classification module distinguishes normal from fraudulent statements, with the model jointly optimizing reconstruction and classification losses to improve both feature completeness and discriminative ability. Experiments on a public financial fraud dataset show that the proposed method significantly outperforms existing baselines on Precision, Recall, F1-Score, and AUC, with particular strength in minority class recognition under imbalanced and limited data. Additional sensitivity experiments demonstrate that the method remains stable and robust across variations in optimizer type and imbalance ratios, confirming its effectiveness in complex financial environments. Overall, the algorithm provides an efficient and reliable pathway for fraud detection and exhibits distinctive advantages in accuracy and adaptability.

**Keywords:** financial fraud detection; self-supervised learning; feature representation; sensitivity analysis

## 1. Introduction

In the context of rapid global economic development and the increasing complexity of capital markets, financial statements serve not only as a key medium for corporate information disclosure but also as the core basis for decision-making by investors, regulators, and the public. However, financial fraud remains persistent and difficult to eliminate. It undermines transparency and credibility in markets and creates potential risks for economic operations. Traditional audit and regulatory approaches often face delays and limitations when dealing with large-scale enterprises, complex information, and hidden fraudulent practices [1]. They struggle to provide effective early warning and prevention. Therefore, it is essential to explore methods that can automatically learn and deeply capture latent features in financial data. Such approaches have important theoretical and practical value for maintaining market order and protecting investment security.

In recent years, advances in artificial intelligence have offered new perspectives and technical pathways for detecting financial fraud. Unlike traditional rule-based or experience-driven approaches, self-supervised learning can model features directly from the structure and semantics of the data without requiring extensive labeled samples [2]. This property is well aligned with the challenges of financial reporting, where fraudulent samples are scarce and costly to obtain, while

normal samples are abundant. Self-supervised learning has been widely applied across multiple domains [3–5], including large language models (LLMs) [6–9], computer vision (CV) [10–12], and system architecture [13–15]. In LLMs, it serves as the foundational training paradigm that enables models to learn linguistic representations from massive unlabeled text corpora through tasks such as masked token prediction and next-word generation [16–21]. In computer vision, self-supervised techniques leverage contrastive and generative objectives to extract semantic and structural information from unlabeled images, achieving performance comparable to supervised methods. In system architecture and intelligent optimization, self-supervised learning facilitates anomaly detection, workload prediction, and adaptive control by capturing latent patterns from operational data without manual annotation. These advances collectively demonstrate the versatility and scalability of self-supervised learning as a unifying framework for representation learning across heterogeneous data modalities and domains. By constructing auxiliary tasks, self-supervised learning enhances representation ability, enabling the detection of complex relationships and hidden patterns in financial data. With the continuous expansion of data types and dimensions in corporate reporting, this method not only improves the automation of fraud detection but also supports large-scale applications.

At the macro level, research on fraud detection is closely related to financial stability and the efficient allocation of resources. When fraudulent behavior is not effectively identified, capital may be misallocated to seemingly strong but deceptive firms. This misallocation can create risks and even trigger systemic crises. In an era of globalization and cross-border capital flows, financial fraud in one country or region can produce chain reactions and erode market confidence more broadly. Thus, applying new intelligent algorithms to improve the scientific and forward-looking nature of fraud detection has become a pressing issue in financial supervision and corporate governance. Modeling the multi-level structure of financial statements through self-supervised learning provides fresh momentum for this field.

At the micro level, fraud detection directly concerns corporate sustainability and the interests of stakeholders. False financial information misleads investment decisions, undermines shareholder value, accumulates debt risks, and damages corporate reputation. Once exposed, fraud results not only in legal penalties and economic losses but also in long-term brand crises. For investors and the public, fraudulent activity signals the collapse of risk assessment and trust mechanisms. The use of self-supervised learning in systematic fraud detection can improve the quality of financial disclosure. It also encourages firms to strengthen compliance and risk management [22].

In summary, research on fraud detection in financial statements using self-supervised learning follows the broader trend of artificial intelligence in finance. It responds to urgent market demands for efficient, intelligent, and scalable detection methods. From stabilizing markets to supporting corporate governance and rebuilding public trust, this direction carries far-reaching significance [23]. It is not only a technical exploration but also a critical measure to ensure fairness and transparency in the economy. Continued research and application in this area will improve the scientific basis and effectiveness of fraud detection and provide solid support for the healthy development of capital markets.

## 2. Related Work

Recent advances in financial fraud detection have seen the widespread adoption of data-driven and intelligent modeling approaches, with a strong emphasis on self-supervised learning, representation learning, and graph-based analysis. Early work on dynamic process modeling, such as using deep Q-learning to optimize audit workflows, highlights how reinforcement learning methods can capture sequential dependencies and adaptively improve decision-making in complex transactional environments[24]. This reinforcement learning perspective lays a solid foundation for integrating more sophisticated neural architectures in downstream tasks.

Building on this foundation, knowledge-enhanced neural modeling has been introduced to address the need for robust feature representation in financial risk identification[25]. By embedding

expert-guided constraints within neural networks, these methods not only enhance the interpretability of the extracted features but also improve the generalization of fraud detection systems. This evolution naturally connects to the emerging use of self-supervised transfer learning, where shared encoders facilitate effective knowledge transfer across domains, making models more resilient to new data distributions and reducing reliance on extensive labeled datasets[26]. The application of self-supervised learning, in particular, bridges the gap between feature completeness and data scarcity, which is a common challenge in financial statement analysis.

As these modeling paradigms mature, there is an increasing focus on advanced sequence architectures. For example, the adoption of improved Mamba-based sequence models offers enhanced capacity for capturing long-range and nonlinear dependencies in financial transaction data, further boosting anomaly and fraud detection performance in real-world scenarios[27]. To complement sequential modeling, causal-aware regression methods have been explored to strengthen interpretability and stability in time series forecasting, leveraging structured attention to reveal cause-effect relationships within noisy and temporally complex financial data[28].

The integration of graph-based learning and temporal modeling has proven particularly fruitful in this field. The combination of graph neural networks with temporal sequence integration offers a comprehensive framework for compliance detection and anomaly discovery, capturing both relational structures and temporal dynamics within financial systems[29]. Contrastive learning methods applied to multimodal knowledge graphs further enrich latent representations, supporting cross-modal association mining and making models more adept at identifying complex fraudulent behavior[30]. When combined with unsupervised graph neural networks and Transformer models, these approaches can robustly uncover subtle anomalies, even when explicit labels are limited[31].

Such methodological advancements are complemented by the use of temporal-semantic graph attention, which introduces an adaptive focus mechanism, enabling models to attend to critical structures and context for risk pattern recognition[32]. Causal inference and structure-aware graph attention are increasingly employed to improve the resilience and interpretability of extracted features in environments characterized by noise, uncertainty, and rapid change[33]. In the same vein, unsupervised contrastive learning provides a robust basis for anomaly detection, even within heterogeneous transactional systems, by maximizing the distinction between normal and abnormal patterns[34].

Further extending these capabilities, recent work on graph neural time series models allows the joint modeling of market structure evolution and temporal dependencies, supporting more accurate identification of emerging risks and shifts in complex financial networks[35]. Hierarchical semantic-structural encoding also contributes to improved regulatory compliance and risk classification, by integrating multi-level contextual and structural cues for deeper understanding of policy-related texts[36].

As time-series forecasting continues to evolve, the use of granularity-aware attention networks provides models with adaptive temporal focus, enhancing predictive accuracy in volatile or rapidly changing financial environments[37]. These techniques complement graph-based discovery methods, which are effective for uncovering implicit relationships in large-scale corporate networks, thus informing risk assessment and early warning systems[38]. Graph neural networks equipped with graph attention have also demonstrated remarkable effectiveness for tasks such as credit card fraud detection, further illustrating the value of deep structural modeling[39].

Feature-rich sequential models—such as those leveraging structured text factors and dynamic time windows—can efficiently extract the nuanced behaviors that drive asset returns, supporting both real-time monitoring and historical analysis of financial trends[40]. Deep contextual risk classification methods, particularly those based on transformer architectures, enable the extraction of layered semantic structures, which are valuable for analyzing complex policy documents and compliance requirements[41]. Finally, dual-phase learning and unsupervised temporal encoding strategies enable robust modeling of sequential data in environments where labeled anomalies are

rare, supporting effective fraud detection and risk assessment across a wide range of financial settings[42].

Collectively, these advances form a cohesive methodology landscape that underpins state-of-the-art financial fraud detection and risk analysis. Each methodological component—whether it be self-supervised learning, sequential modeling, graph-based reasoning, or contrastive representation learning—contributes to a comprehensive and resilient approach, ensuring models remain adaptive and reliable in increasingly complex and dynamic financial environments.

### 3. Method

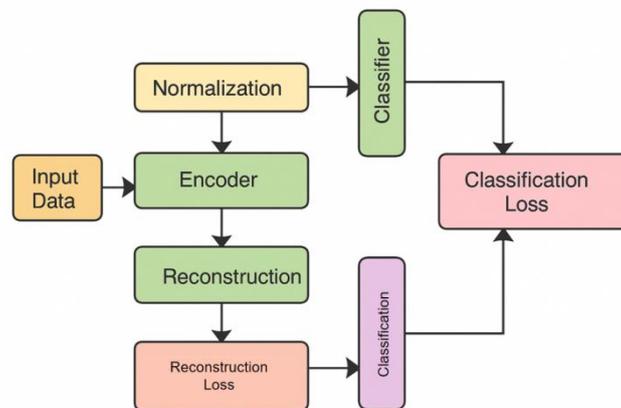
The core idea of the method design in this study is to learn the latent representation of financial statements through a self-supervisory mechanism and transform it into a discriminative task that can be used to distinguish normal from abnormal. First, the financial statement input matrix is set as:

$$X = \{x_1, x_2, \dots, x_n\}, \quad x_i \in R^d \quad (1)$$

where  $n$  represents the number of samples and  $d$  represents the dimension of the financial indicator. To enhance the adaptability of the model to different scales and distributions, the input data is preprocessed by standardization:

$$x_i = \frac{x_i - \mu}{\sigma} \quad (2)$$

Here,  $\mu$  and  $\sigma$  represent the sample mean and standard deviation, respectively. This process can eliminate the impact of inconsistent dimensions of different financial indicators and lay the foundation for subsequent representation learning. The model architecture is shown in Figure 1.



**Figure 1.** The model architecture is shown in Figure 1.

In the representation learning stage, the encoder function  $f_{\theta}(x_i)$  is introduced to map the input to the latent space representation:

$$h_i = f_{\theta}(x_i), \quad h_i \in R^k \quad (3)$$

where  $k$  is the dimension of the latent space. The encoder captures high-order relationships between financial indicators through multi-layer nonlinear mapping, resulting in  $h_i$  that takes into account both global financial features and local fine-grained information. To further enhance the discriminability of features, a reconstruction task is introduced as an auxiliary signal, generating an approximate input through the decoder  $g_{\theta}(\cdot)$ :

$$\tilde{x}_i = g_\phi(h_i) \quad (4)$$

Also, minimize the difference between the original input and the reconstructed result:

$$L_{rec} = \frac{1}{n} \sum_{i=1}^n \|x_i - \tilde{x}_i\|_2^2 \quad (5)$$

In the discrimination stage, the potential representation  $h_i$  is used to input the classifier  $C(\cdot)$ , and the output sample is the probability distribution of normal or abnormal:

$$y_i = C(h_i), \quad \hat{y}_i \in [0,1] \quad (6)$$

When  $\hat{y}_i$  is close to 1, it indicates normality, and when it is close to 0, it indicates abnormality. The objective function of the discrimination process is defined in the form of cross-entropy as:

$$L_{cls} = -\frac{1}{n} \sum_{i=1}^n [y_i \log(\hat{y}_i) + (1 - y_i) \log(1 - \hat{y}_i)] \quad (7)$$

where A represents the true label. The final optimization objective is composed of reconstruction loss and discrimination loss to ensure that the feature representation has both fidelity and discriminative power:

$$L = \lambda L_{rec} + (1 - \lambda) L_{cls} \quad (8)$$

where  $\lambda \in [0,1]$  is a trade-off parameter. This optimization process ensures that the model can learn stable and discriminative feature representations based on a large number of normal financial statements, even with a limited number of fraud samples, to effectively detect financial fraud.

## 4. Experimental Results

### 4.1. Dataset

In this study, the dataset used is the Enron Financial Statement Fraud Dataset. This dataset is widely applied in research on financial fraud detection. It contains a large number of corporate financial statements and related disclosure information across different years. It includes both normal and fraudulent labels. The features provided cover financial ratios, cash flows, profitability indicators, debt structures, and other dimensions. These features reflect corporate operations and financial status comprehensively and provide a solid data foundation for fraud detection tasks.

The dataset used in this study is of moderate scale, comprising a substantial number of normal corporate financial statement samples alongside a smaller, yet carefully labeled, subset of fraudulent cases. This inherent class imbalance mirrors the real-world rarity of financial fraud, providing a realistic benchmark for evaluating model robustness under complex and skewed data distributions. Furthermore, the dataset encompasses firms from various industries and time periods, enhancing its representativeness and supporting a degree of cross-industry generalization in model evaluation and application. The choice of this dataset ensures reproducibility of the study and aligns with existing research in the field. It also provides a reliable reference for comparing and improving methods. Modeling and validation on this dataset allow for an effective evaluation of the practical value of the proposed approach in fraud detection. It also supplies data support for advancing intelligent auditing and risk management.

### 4.2. Experimental Results

This paper first gives the comparative experimental results, which are shown in Table 1.

**Table 1.** Comparative experimental results.

Model	Precision	Recall	F1-Score	AUC
<b>Realnet [43]</b>	0.842	0.816	0.829	0.892
<b>Mambaad [44]</b>	0.857	0.823	0.839	0.901
<b>Real-iad [45]</b>	0.871	0.835	0.852	0.914
<b>Promptad [46]</b>	0.886	0.848	0.867	0.926
<b>Adaclip [47]</b>	0.893	0.856	0.874	0.934
<b>Ours</b>	0.917	0.881	0.899	0.951

From the experimental results, it can be seen that the selected baseline models show certain differences in overall performance. However, the proposed method achieves the best performance on all metrics. In particular, the Precision reaches 0.917, which is clearly higher than traditional methods. This indicates that the model can effectively reduce false positives when distinguishing between normal and fraudulent financial statements and maintain high accuracy. This is important in real financial contexts, as false positives not only waste resources but may also affect the normal operation and reputation of legitimate firms.

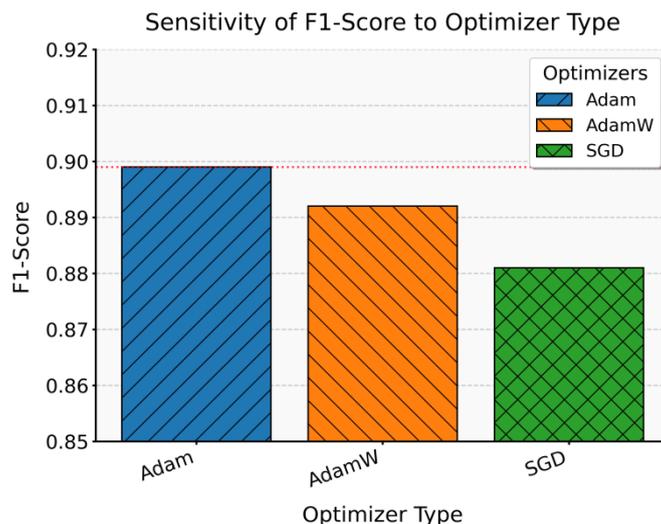
In terms of Recall, the model also shows strong advantages, reaching 0.881, which is higher than other baselines. This means the method can capture more real fraud cases when fraudulent samples are limited and reduce missed detections. For financial fraud detection, improvements in Recall are critical because any missed fraudulent case can lead to severe economic and social risks. The results demonstrate that combining self-supervised feature learning with classification tasks significantly enhances the ability of the model to capture potential abnormal patterns.

In terms of the balanced performance metric F1-Score, the proposed method once again surpasses the comparison models, achieving a value of 0.899. This indicates that the model attains both high precision and strong recall, reflecting an effective balance between accuracy and completeness. Unlike other approaches, it maintains high predictive accuracy without compromising detection coverage. Such an equilibrium is particularly consistent with the requirements of financial risk management, where both comprehensiveness and accuracy are critical. By maintaining this balance, the model exhibits enhanced robustness and reliability when operating under complex and uncertain conditions.

Finally, regarding overall discriminative ability, the AUC of the proposed method reaches 0.951, which is significantly higher than other baselines. This metric reflects the model's ability to distinguish cases under different thresholds. It shows that the model performs well not only under specific conditions but also adapts effectively to multiple detection standards. Considering the practical background of financial fraud detection, this result suggests that the proposed method maintains superior generalization across different firms, industries, and data distributions. It provides stronger technical support for intelligent financial fraud detection.

This paper also presents an experiment on the sensitivity of optimizer type to F1-Score, and the experimental results are shown in Figure 2.

From the figure, it can be observed that the model's F1-Score shows some fluctuation under different optimizers. Taking Adam as an example, the result is closest to the baseline value of 0.899. This indicates that under Adam, the model achieves the most stable discriminative performance and balances Precision and Recall effectively. It suggests that in financial statement fraud detection, Adam shows better adaptability to self-supervised feature learning and classification tasks, which helps improve overall detection performance.



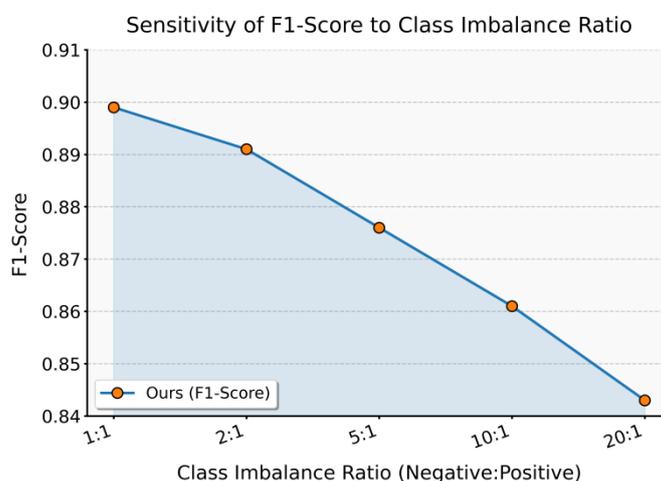
**Figure 2.** Sensitivity experiment of optimizer type to F1-Score.

The F1-Score of AdamW is slightly lower than Adam, with a value of 0.892. This difference shows that the introduction of weight decay helps reduce overfitting but also reduces flexibility in identifying minority classes. As a result, Recall may be affected. This finding suggests that when fraudulent samples are scarce in financial data, overly strong regularization may restrict the model's ability to capture abnormal patterns.

The result of SGD is the lowest, only 0.881, which is far below Adam and AdamW. This shows that basic stochastic gradient descent is not sufficient for complex financial data representation. Financial statements contain many nonlinear relationships and feature interactions. However, the limitations of SGD in convergence speed and global optimization make it difficult for the model to learn effective decision boundaries within limited iterations. This result highlights the critical role of optimizer choice in fraud detection tasks.

Overall, this sensitivity experiment confirms the significant impact of optimizer type on model performance. In highly imbalanced financial data environments, the choice of optimizer directly determines the trade-off between accuracy and coverage. The results emphasize that when designing fraud detection models, optimizers with stronger adaptability and robustness should be prioritized to ensure stability and discriminative ability across different scenarios.

This paper also gives the impact of class imbalance on experimental results, and the experimental results are shown in Figure 3.



**Figure 3.** The impact of class imbalance on experimental results.

From the experimental results, it can be seen that as the class imbalance ratio increases, the model's F1-Score shows a clear downward trend. Under the balanced condition of 1:1, the model achieves the best performance and balances Precision and Recall effectively. This indicates that when the sample distribution is balanced, the proposed method can fully demonstrate its representation and classification ability. The result highlights the direct impact of data distribution on model performance.

When the imbalance ratio reaches 2:1 or 5:1, the F1-Score drops slightly to 0.891 and 0.876, respectively. This shows that as fraudulent samples become more scarce, the model's sensitivity to the minority class is affected. Although the overall performance remains at a relatively high level, the bias caused by imbalance begins to appear. This observation is consistent with real financial scenarios, where fraudulent samples are often far fewer than normal samples.

Under the more extreme condition of 10:1, the F1-Score further decreases to 0.861. This means that the recall rate in detecting fraudulent samples is strongly impacted, and the number of missed cases increases. For financial statement fraud detection, this decline is an important warning. If not addressed, the model may overlook many critical fraud cases in real applications, reducing its value for risk warning.

When the imbalance ratio reaches 20:1, the F1-Score falls to 0.843, showing the most severe degradation. This indicates that in extremely imbalanced settings, the discriminative ability of the model is significantly weakened. The overall trend shows that class imbalance systematically reduces the robustness and generalization of the model. Therefore, techniques such as data sampling, loss function adjustment, or generative methods will be crucial to mitigating imbalance and improving the practical value of fraud detection systems.

## 5. Conclusions

This study focuses on financial statement fraud detection based on self-supervised learning. It proposes an algorithmic framework that maintains strong discriminative ability even with limited labeled samples. By introducing a dual-task structure of reconstruction and classification, the method improves both the stability and the discriminative power of feature representations. As a result, it shows superior detection performance in complex financial environments. Compared with traditional statistical approaches and supervised learning methods that depend on large labeled datasets, this study achieves high accuracy while adapting better to the scarcity of fraudulent samples. It provides a new direction for building efficient and intelligent financial risk management systems.

The proposed method has theoretical value and also shows practical significance. Fraud detection is a crucial element in the healthy functioning of capital markets. Its outcomes directly influence investor confidence and the effectiveness of regulatory governance. The findings demonstrate that self-supervised learning can establish robust detection mechanisms in heterogeneous and multi-source data environments. This helps mitigate problems of information asymmetry and resource misallocation. More importantly, the method provides a reference pathway for intelligent auditing and automated risk control. It adds new momentum to the integration of financial technology and corporate governance.

Future research can extend in several directions. One direction is to integrate cross-modal information and external knowledge graphs to enhance reasoning and discrimination across multidimensional data [48]. Another direction is to validate and optimize the model on larger and more complex real-world datasets to improve generalization and robustness. In addition, strengthening interpretability while maintaining performance is critical for applying the method in auditing and regulatory contexts. In conclusion, this study not only provides an innovative solution for financial fraud detection but also lays a solid foundation for the future development of intelligent financial analysis.

## References

1. Y. Zhang and B. Duan, "Accounting data anomaly detection and prediction based on self-supervised learning," *Frontiers in Applied Mathematics and Statistics*, vol. 11, 1628652.
2. K. Reynisson, M. Schreyer and D. Borth, "GraphGuard: Contrastive self-supervised learning for credit-card fraud detection in multi-relational dynamic graphs," *arXiv preprint arXiv:2407.12440*, 2024.
3. Y. Zou, "Hierarchical large language model agents for multi-scale planning in dynamic environments," *Transactions on Computational and Scientific Methods*, vol. 4, no. 2, 2024.
4. Y. Wang, "Structured compression of large language models with sensitivity-aware pruning mechanisms," *Journal of Computer Technology and Software*, vol. 3, no. 9, 2024.
5. Q. Wang, X. Zhang and X. Wang, "Multimodal integration of physiological signals clinical data and medical imaging for ICU outcome prediction," *Journal of Computer Technology and Software*, vol. 4, no. 8, 2025.
6. S. Wang, S. Han, Z. Cheng, M. Wang and Y. Li, "Federated fine-tuning of large language models with privacy preservation and cross-domain semantic alignment," 2025.
7. R. Wang, Y. Chen, M. Liu, G. Liu, B. Zhu and W. Zhang, "Efficient large language model fine-tuning with joint structural pruning and parameter sharing," 2025.
8. Z. Xue, "Dynamic structured gating for parameter-efficient alignment of large pretrained models," *Transactions on Computational and Scientific Methods*, vol. 4, no. 3, 2024.
9. W. Zhu, "Fast adaptation pipeline for LLMs through structured gradient approximation," *Journal of Computer Technology and Software*, vol. 3, no. 6, 2024.
10. J. Hu, B. Zhang, T. Xu, H. Yang and M. Gao, "Structure-aware temporal modeling for chronic disease progression prediction," *arXiv preprint arXiv:2508.14942*, 2025.
11. D. Gao, "High fidelity text to image generation with contrastive alignment and structural guidance," *arXiv preprint arXiv:2508.10280*, 2025.
12. Y. Zi and X. Deng, "Joint modeling of medical images and clinical text for early diabetes risk detection," *Journal of Computer Technology and Software*, vol. 4, no. 7, 2025.
13. X. Zhang, X. Wang and X. Wang, "A reinforcement learning-driven task scheduling algorithm for multi-tenant distributed systems," *arXiv preprint arXiv:2508.08525*, 2025.
14. C. Hu, Z. Cheng, D. Wu, Y. Wang, F. Liu and Z. Qiu, "Structural generalization for microservice routing using graph neural networks," *arXiv preprint arXiv:2510.15210*, 2025.
15. H. Liu, Y. Kang and Y. Liu, "Privacy-preserving and communication-efficient federated learning for cloud-scale distributed intelligence," 2025.
16. S. Pan and D. Wu, "Trustworthy summarization via uncertainty quantification and risk awareness in large language models," 2025.
17. X. Hu, Y. Kang, G. Yao, T. Kang, M. Wang and H. Liu, "Dynamic prompt fusion for multi-task and cross-domain adaptation in LLMs," *arXiv preprint arXiv:2509.18113*, 2025.
18. X. Quan, "Structured path guidance for logical coherence in large language model generation," *Journal of Computer Technology and Software*, vol. 3, no. 3, 2024.
19. R. Zhang, "Privacy-oriented text generation in LLMs via selective fine-tuning and semantic attention masks," *Journal of Computer Technology and Software*, vol. 4, no. 8, 2025.
20. R. Zhang, L. Lian, Z. Qi and G. Liu, "Semantic and structural analysis of implicit biases in large language models: An interpretable approach," *arXiv preprint arXiv:2508.06155*, 2025.
21. X. Song, Y. Liu, Y. Luan, J. Guo and X. Guo, "Controllable abstraction in summary generation for large language models via prompt engineering," *arXiv preprint arXiv:2510.15436*, 2025.
22. S. Visbeek, E. Acar and F. den Hengst, "Explainable fraud detection with deep symbolic classification," *Proceedings of the World Conference on Explainable Artificial Intelligence*, Cham: Springer Nature Switzerland, pp. 350-373, 2024.

23. M. Tayebi and S. El Kafhali, "Generative modeling for imbalanced credit card fraud transaction detection," *Journal of Cybersecurity and Privacy*, vol. 5, no. 1, pp. 9, 2025.
24. 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.
25. M. Jiang, S. Liu, W. Xu, S. Long, Y. Yi and Y. Lin, "Function-driven knowledge-enhanced neural modeling for intelligent financial risk identification," 2025.
26. Y. Zhou, "Self-supervised transfer learning with shared encoders for cross-domain cloud optimization," 2025.
27. 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.
28. C. Liu, Q. Wang, L. Song and X. Hu, "Causal-aware time series regression for IoT energy consumption using structured attention and LSTM networks," 2025.
29. W. Xu, M. Jiang, S. Long, Y. Lin, K. Ma and Z. Xu, "Graph neural network and temporal sequence integration for AI-powered financial compliance detection," 2025.
30. L. Dai, "Contrastive learning framework for multimodal knowledge graph construction and data-analytical reasoning," *Journal of Computer Technology and Software*, vol. 3, no. 4, 2024.
31. Y. Zi, M. Gong, Z. Xue, Y. Zou, N. Qi and Y. Deng, "Graph neural network and transformer integration for unsupervised system anomaly discovery," arXiv preprint arXiv:2508.09401, 2025.
32. H. Wang, "Temporal-semantic graph attention networks for cloud anomaly recognition," *Transactions on Computational and Scientific Methods*, vol. 4, no. 4, 2024.
33. L. Dai, "Integrating causal inference and graph attention for structure-aware data mining," *Transactions on Computational and Scientific Methods*, vol. 4, no. 4, 2024.
34. W. Cui, "Unsupervised contrastive learning for anomaly detection in heterogeneous backend system," *Transactions on Computational and Scientific Methods*, vol. 4, no. 7, 2024.
35. Q. R. Xu, "Capturing structural evolution in financial markets with graph neural time series models," 2025.
36. Y. Qin, "Hierarchical semantic-structural encoding for compliance risk detection with LLMs," *Transactions on Computational and Scientific Methods*, vol. 4, no. 6, 2024.
37. X. Su, "Deep forecasting of stock prices via granularity-aware attention networks," *Journal of Computer Technology and Software*, vol. 3, no. 7, 2024.
38. 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.
39. Q. Sha, T. Tang, X. Du, J. Liu, Y. Wang and Y. Sheng, "Detecting credit card fraud via heterogeneous graph neural networks with graph attention," arXiv preprint arXiv:2504.08183, 2025.
40. X. Su, "Forecasting asset returns with structured text factors and dynamic time windows," *Transactions on Computational and Scientific Methods*, vol. 4, no. 6, 2024.
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. Q. Xu, "Unsupervised temporal encoding for stock price prediction through dual-phase learning," 2025.
43. X. Zhang, M. Xu and X. Zhou, "Realnet: A feature selection network with realistic synthetic anomaly for anomaly detection," *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 16699-16708, 2024.
44. H. He, Y. Bai, J. Zhang, et al., "Mambaad: Exploring state space models for multi-class unsupervised anomaly detection," *Advances in Neural Information Processing Systems*, vol. 37, pp. 71162-71187, 2024.
45. C. Wang, W. Zhu, B. B. Gao, et al., "Real-iad: A real-world multi-view dataset for benchmarking versatile industrial anomaly detection," *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 22883-22892, 2024.
46. Y. Wang, "Gradient-guided adversarial sample construction for robustness evaluation in language model inference," *Transactions on Computational and Scientific Methods*, vol. 4, no. 7, 2024.

47. Y. Cao, J. Zhang, L. Frittoli, et al., "Adaclip: Adapting clip with hybrid learnable prompts for zero-shot anomaly detection," *Proceedings of the European Conference on Computer Vision*, Cham: Springer Nature Switzerland, pp. 55-72, 2024.
48. L. Lian, "Semantic and factual alignment for trustworthy large language model outputs," *Journal of Computer Technology and Software*, vol. 3, no. 9, 2024.

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