

Article

Not peer-reviewed version

Enhancing Systemic Risk Forecasting with Deep Attention Models in Financial Time Series

[Qinggong \(Raisy\) Xu](#) , [Weiyao Xu](#) , Xin Su , Kunyuan Ma , Wenying Sun , Ying Qin *

Posted Date: 9 July 2025

doi: 10.20944/preprints202507.0737.v1

Keywords: deep learning; multi-scale convolution; attention mechanism; systemic risk



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

Enhancing Systemic Risk Forecasting with Deep Attention Models in Financial Time Series

Qingqing Xu ¹, Weiyao Xu ², Xin Su ³, Kunyuan Ma ⁴, Wenying Sun ⁵ and Ying Qin ^{6,*}

¹ Washington University in St. Louis, St. Louis, USA

² Fordham University, New York, USA

³ University of Chicago, Chicago, USA

⁴ New York University, New York, USA

⁵ Southern Methodist University, Dallas, USA

⁶ University of Illinois at Urbana-Champaign, Champaign, USA

* Correspondence: author: yingqinuiuc@gmail.com

Abstract

This paper proposes a deep learning model for systemic risk identification and prediction in financial markets, integrating multi-scale convolutional structures, attention mechanisms, and residual connections. The model is designed to capture dynamic features at different temporal resolutions, enhance focus on key risk factors, and maintain training stability in complex financial data environments. Using the US Systemic Risk Dataset released by the Federal Reserve, the model transforms the task into a binary classification problem based on crisis-related indicators. All time series data are standardized, interpolated, and structured through sliding windows to preserve temporal continuity. Experimental comparisons show that the proposed model outperforms traditional methods such as MLP, CNN, Transformer, and BERT across accuracy, precision, and recall. Further analysis reveals the effectiveness of multi-scale and attention mechanisms in modeling complex dependencies and detecting early risk signals. Sensitivity experiments on convolutional kernel size demonstrate that integrating cross-feature information significantly improves performance, with 3x3 and 5x3 kernels showing optimal results. The training process exhibits fast convergence and low loss stability, confirming the model's robustness and efficiency. These findings validate the technical advantages of the proposed architecture in addressing the challenges of systemic risk modeling in high-dimensional and heterogeneous financial time series.

Keywords: deep learning; multi-scale convolution; attention mechanism; systemic risk

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

1. Introduction

In the context of the growing global economic integration and the increasing interconnectedness of financial markets, the identification and prediction of systemic financial risks have emerged as central concerns within the realm of financial research and practice. Systemic risks are characterized by their sudden onset, wide contagion, and deep destructiveness [1]. Once triggered, they often have far-reaching impacts on the entire financial system and the real economy. Especially in the current environment of frequent extreme economic events and heightened market volatility, early identification and accurate forecasting of systemic risk are critical to ensuring the stable operation of financial markets and preventing the spread of financial crises [2]. Traditional risk management methods struggle to capture the nonlinear relationships and potential interaction mechanisms within complex financial systems. There is an urgent need for more efficient and intelligent technologies to meet this challenge.

In recent years, with the rapid development of big data and artificial intelligence, financial technology has become a key force in enhancing financial risk management. In particular, deep learning has achieved breakthroughs in image recognition, natural language processing, and time series analysis [3]. Its powerful feature extraction and nonlinear modeling capabilities offer new tools for addressing the high dimensionality [4], nonlinearity, and time-variability of financial risk prediction. However, financial markets are far more complex than many traditional tasks. The data are multi-source, heterogeneous, highly time-dependent [5], and updated at high frequency [6]. Designing a model that can fully extract deep structural information from data while maintaining strong generalization ability has become a key challenge in applying deep learning to systemic risk forecasting [7].

Against this backdrop, the introduction of multi-scale modeling and attention mechanisms offers more interpretable and flexible solutions for risk identification. Multi-scale structures can extract useful information from different time scales or feature dimensions, capturing the layered and asynchronous nature of risk transmission in financial systems. Attention mechanisms dynamically assign weights to different features or time segments, highlighting key variables and suppressing noise [8]. This enhances the model's sensitivity to risk factors. The integration of multi-scale analysis and attention mechanisms allows for better modeling of the multi-level and multi-path transmission of risks in financial markets. It enables the model to learn early-warning patterns of systemic risk directly from data [9].

In addition, residual network structures introduce cross-layer connections that effectively alleviate gradient vanishing and performance degradation in deep neural networks [10]. This makes deeper feature learning possible. For financial markets—characterized by complex signals and numerous latent factors—deep networks help to capture more abstract and essential risk patterns [11]. The structural advantages of residual networks, combined with multi-scale attention mechanisms, form an efficient, stable, and scalable deep learning framework. This framework not only enables early identification of systemic risk but also ensures strong generalization and interpretability. It holds significant theoretical and practical value for financial regulation, risk warning, and policy-making. Systemic risk remains a major challenge in financial research and requires breakthroughs through advanced modeling approaches. The integration of multi-scale attention mechanisms and residual networks for identifying and forecasting systemic risk in financial markets addresses the urgent need for accurate and dynamic risk management. This study contributes to more scientific and forward-looking systemic risk prevention. It also provides theoretical support and technical assurance for building intelligent, data-driven financial regulation systems.

2. Method

This study proposes a deep learning framework that integrates multi-scale convolutional structures, attention mechanisms, and residual connections to enable high-precision identification and prediction of systemic risks in financial markets. The framework is designed to process multi-dimensional financial time series by constructing parallel convolutional channels that operate at different temporal scales, thereby capturing both short-term volatility and long-term trend dependencies. The preprocessing step involves standardizing the time series data and segmenting it into overlapping windows to maintain temporal continuity. These windows are then passed through multi-scale convolutional layers, where each channel applies filters of varying kernel sizes. This design facilitates the extraction of fine-grained local features and broader contextual patterns simultaneously—an approach inspired by Gunathilaka et al. [12], who demonstrated the efficacy of hierarchical convolutional modeling in detecting complex anomalies in financial transaction data.

To enhance the model's capacity for dynamic feature selection, an attention mechanism is incorporated after multi-scale feature extraction. This module adaptively recalibrates the significance of extracted features, allowing the model to prioritize critical systemic risk indicators, drawing upon strategies similar to those employed in their early warning system for financial fraud [13].

Furthermore, to stabilize gradient flow and mitigate information degradation across layers, residual connections are embedded throughout the network. This design choice is motivated by the success of residual structures in deep temporal models, as highlighted by Sheng [14], particularly in maintaining performance when modeling long-term dependencies in financial sequences.

The complete architecture of the proposed model is illustrated in Figure 1, with the initial phase involving window-based segmentation and convolution operations structured as follows:

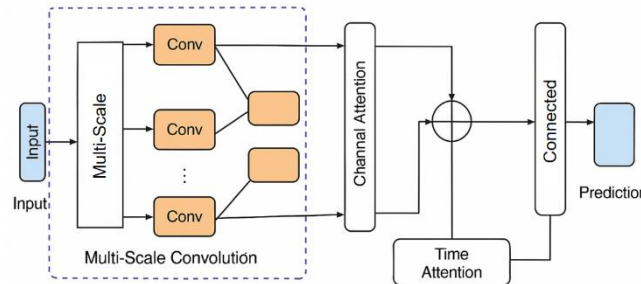


Figure 1. Overall model architecture diagram.

The model receives multi-dimensional financial time series as input and extracts deep feature representations at multiple temporal scales in parallel through multi-scale convolutional modules. This parallel structure allows the network to capture both localized variations and broader temporal dynamics, which aligns with prior approaches that emphasize adaptive feature extraction across granularities in structured data contexts [15,16]. The resulting feature tensors are then processed through a sequence of attention mechanisms, including channel attention and temporal attention modules. These components dynamically reweight the importance of different channels and time steps, effectively emphasizing signals most indicative of systemic risk. Such attention-based weighting strategies have been applied successfully in related work dealing with anomaly detection in high-dimensional behavioral data streams [17]. To ensure information integrity across network depth, residual connections are incorporated, followed by fully connected layers that integrate multi-layered features into a final predictive output. The use of residual pathways in conjunction with dense representations supports the extraction of robust, causally aligned predictive features in volatile and cross-market environments, as seen in recent studies on return prediction under systemic uncertainty [18].

Assume that the input financial sequence is $X \in \mathbb{R}^{T \times D}$, where T is the time step and D is the feature dimension. After multi-scale convolution, multiple scale feature maps $\{F_s\}_{s=1}^S$ are obtained, which can be expressed as:

$$F_s = \text{Conv}_{k_s}(X), \quad s = 1, 2, \dots, S$$

In order to enhance the model's ability to identify key risk features, a dual fusion strategy of channel attention mechanism and time attention mechanism is introduced. In the channel dimension, global average pooling is used to obtain the importance distribution of each channel, and the attention weight $\alpha \in \mathbb{R}^D$ is generated through nonlinear transformation to weight the original features. In the time dimension, a weighting mechanism based on self-attention is constructed to make the contribution of different time segments to the output learnable. Let the feature sequence be $H \in \mathbb{R}^{T \times D}$, and its weighted result is:

$$\tilde{H}_t = \alpha \otimes H_t, \quad t = 1, 2, \dots, T$$

$$A = \text{Softmax}\left(\frac{QK^T}{\sqrt{d_k}}\right), \tilde{Z} = A \cdot V$$

where Q, K, V is the query, key and value matrix obtained by linear transformation, A represents the attention score matrix, and \tilde{Z} is the time-weighted representation. In terms of the overall structure, the residual connection mechanism is used to improve the stability of deep network training. The output $F(x)$ of each layer is added to its input x to obtain the residual block output:

$$y = F(x) + x$$

Finally, after multiple layers of stacking and fully connected layer mapping, the model outputs a systemic risk prediction value or classification result. In order to ensure the prediction performance and model robustness, the loss function uses a combination of mean square error and regularization terms to establish an objective function for the predicted value \hat{y} and the true value y :

$$L = \frac{1}{N} \sum_{i=1}^N (\hat{y}_i - y_i)^2 + \lambda \|\theta\|_2^2$$

where λ is the regularization coefficient and θ is the model parameter set. Through end-to-end training, the model can automatically learn the risk signal characteristics at different time scales and enhance the ability to capture potential risk outbreaks. This method not only integrates multiple advanced mechanisms in structure, but also has strong scalability and interpretability in mathematical modeling, and is suitable for systemic risk modeling tasks in complex financial data environments.

3. Experiment

A. Datasets

This study uses a real-world financial dataset—the US Systemic Risk Dataset—published by the Federal Reserve. The dataset covers major market data of financial institutions and macroeconomic indicators in the United States from 2000 to the present. It includes multiple dimensions such as bank stock prices, interest rate spreads, volatility indices, credit spreads, and leverage ratios. For example, interest rate spreads capture the difference between short-term and long-term government bond yields, indicating market expectations about future economic conditions. Credit spreads reflect the risk premium demanded by investors and serve as a proxy for perceived credit risk in the financial system. Leverage ratios measure the financial fragility of institutions by comparing their debt to equity levels.

The data are primarily at daily and monthly frequencies, with long time spans, high dimensionality, and heterogeneity. These characteristics make it well-suited for deep modeling tasks on systemic risk. In total, the dataset comprises over 500,000 time-series records, with approximately 10% of samples labeled as high risk (positive class) and 90% as low risk (negative class), indicating a class imbalance that must be considered during model training.

To construct the classification task, this study defines systemic risk events as the prediction target. Positive and negative samples are labeled based on historical crises, such as the 2008 financial crisis and the 2020 market turmoil during the COVID-19 outbreak. Specifically, if systemic risk indicators—such as SRISK or CATFIN—exceed historical quantile thresholds within a future time window, the current sample is labeled as high risk (positive class). Otherwise, it is labeled as low risk (negative class). This transforms the problem into a binary classification task. The labeling method incorporates financial stability metrics, allowing the model to learn early warning signals of risk during training.

In the data preprocessing stage, all time series features are standardized. Missing values are filled using linear interpolation over time. A sliding window approach is then applied to construct samples [19]. Each sample consists of a continuous sequence of multi-dimensional features and serves as input to the model. This processing method ensures temporal continuity and captures the dynamic evolution of risk, providing a solid data foundation for systemic risk identification.

B. Experimental Results

First, this paper conducted a comparative test, and the experimental results are shown in Table 1.

Table 1. Comparative experimental results.

| Method | Acc | Precision | Recall |
|-----------------|-------|-----------|--------|
| MLP[20] | 0.764 | 0.728 | 0.702 |
| CNN[21] | 0.781 | 0.749 | 0.719 |
| Transformer[22] | 0.805 | 0.771 | 0.746 |
| Bert[23] | 0.812 | 0.778 | 0.753 |
| Ours | 0.843 | 0.813 | 0.794 |

As shown in the comparison results in Table 1, the proposed model outperforms all baseline methods in terms of three key metrics: accuracy (Acc), precision, and recall. Traditional models such as MLP and CNN perform relatively poorly on financial time series data. This is mainly due to their limited ability to model long-term dependencies. Their accuracy scores are 0.764 and 0.781, respectively. The recall scores are only 0.702 and 0.719. These models fail to effectively capture the complex features of systemic financial risk.

In contrast, Transformer and BERT show improved performance. This is attributed to the advantage of self-attention mechanisms in modeling sequential features. These models can extract some critical risk signals. BERT achieves a precision of 0.778 and a recall of 0.753, outperforming the previous shallow models. This indicates that introducing pre-trained language modeling helps enhance the model's ability to understand dependencies between features.

The proposed multi-scale attention residual network achieves the best overall performance. It reaches an accuracy of 0.843, a precision of 0.813, and a recall of 0.794. These results demonstrate that the multi-scale convolutional structure can capture risk features at different temporal resolutions. The attention mechanism improves the identification of key variables. The residual connections enhance network depth and stability. Overall, the model shows stronger expressive power and generalization ability in systemic risk classification tasks, validating the effectiveness and forward-looking nature of the proposed approach. Furthermore, this paper presents a visualization of the loss function's progression during the training process, as illustrated in Figure 2. This graph serves to demonstrate the model's learning dynamics and convergence behavior over multiple training epochs. By tracking the changes in loss values, it becomes possible to assess the model's optimization efficiency and stability. Such analysis is essential for understanding how well the model adapts to complex financial data and whether it maintains generalizability without overfitting. The loss curve also provides supporting evidence for the effectiveness of the proposed architecture in capturing underlying data patterns through iterative learning.




Figure 2. Loss function drop graph.

As shown in Figure 2, the model starts with a relatively high loss value of around 0.7 during the early training stages. With the increase in training epochs, the loss drops rapidly. A clear convergence trend is observed within the first 20 epochs. This indicates that the model can quickly capture basic feature structures from the data in the initial phase, showing high learning efficiency.

In the middle and later stages of training, the loss value becomes more stable with reduced fluctuations, indicating a steady learning process. It frequently drops below 0.05, which suggests that the model has largely converged after undergoing multiple training iterations. This level of stability implies that the model has effectively captured the underlying patterns in the data without continuing to make large adjustments. No significant overfitting is observed throughout this period, which reflects the model's ability to generalize well beyond the training set. Although a slight increase in the loss value occurs around epoch 120, this change is short-lived, and the loss quickly returns to a low level. This behavior demonstrates the model's robustness and strong convergence capability, as it can maintain performance stability even in the presence of minor fluctuations during training.

By the end of training, around epoch 200, the loss remains low and stable. The overall training process shows a clear and positive optimization trend. These results validate the effectiveness of the proposed model structure in complex financial data settings. With the support of multi-scale design and attention mechanisms, the model converges quickly and captures stable feature representations. This provides a solid foundation for subsequent risk identification tasks. Finally, this paper provides an in-depth analysis of how convolution kernels of different scales influence the model's ability to identify systemic risk. By examining various kernel configurations, the study explores how temporal resolution and feature integration affect the extraction of meaningful patterns from financial time series. The analysis highlights the importance of selecting appropriate kernel dimensions to balance local detail capture with broader contextual understanding. This component of the study offers structural insights into model design and contributes to a better understanding of how convolutional architecture choices can enhance performance in complex risk prediction tasks.

Figure 3 demonstrates the impact of the size of convolutional kernels on the performance of the models with respect to the risk detection task. When compared to compact, 1D kernels (3x1, 5x1), architectures comprising width extensions (3x3, 5x3) demonstrate improvements on all three measures of accuracy (the same as the precision) and recall. The 3x3 kernel has the highest accuracy among three methods (highest of 0.843, the best performance).

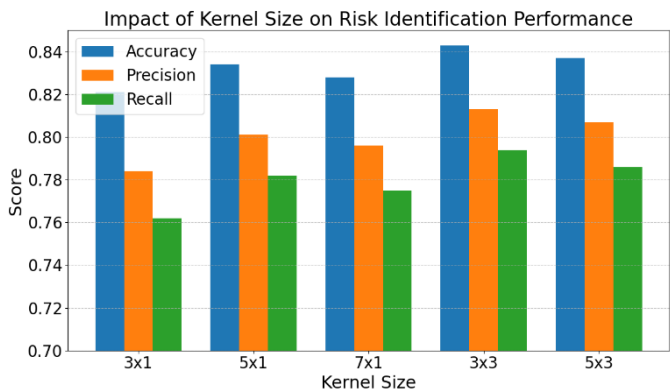


Figure 3. Analysis of the impact of different scale convolution kernels on risk identification.

Local features can be captured by single-scale convolutional kernels. However, they are not very good at capturing complex dependencies in financial time series. Larger (or piecewise) kernels, e.g., 5x3, can be applied to learn a wider range of temporal contextual features and cross-feature interactions. This allows the model to better capture early warnings of systemic risk. The multi-scale architecture can help to better represent the features and it is an important part of improving the model performance.

It is worth noting that although the 7x1 kernel improves performance to some extent, it still underperforms compared to 3x3 and 5x3. This suggests that simply extending the temporal dimension cannot replace multi-dimensional feature fusion. Therefore, choosing appropriate convolutional kernel structures—especially those with horizontal receptive fields—is crucial for improving the overall capability of risk identification models.

4. Conclusions

This study proposes a deep learning model for identifying and forecasting systemic risk in financial markets. The model integrates multi-scale convolution, attention mechanisms, and residual connections. Through multi-scale feature extraction, it captures dynamic patterns at different temporal resolutions in financial time series. The combination of channel and temporal attention enhances the model's focus on key risk factors. Residual connections ensure training stability and enable deeper feature representation. Experimental results show that the proposed method outperforms mainstream models in terms of accuracy, precision, and recall.

In comparative experiments, the model demonstrates stronger representation capabilities when dealing with complex financial time series. This confirms the effectiveness of multi-scale and attention mechanisms in systemic risk detection. In addition, the convolutional kernel sensitivity analysis reveals how different structures affect model performance. This provides structural guidance for further model optimization. The rapid convergence of the loss function also highlights the method's advantages in training efficiency and feature stability.

This research provides theoretical support and a technical pathway for applications such as financial risk warning, intelligent regulation, and investment strategy evaluation. It offers practical insights with real-world relevance. As financial data sources become more diverse and market structures evolve, traditional static risk identification methods can no longer meet the demands for real-time and precise assessment. The proposed model introduces a new solution for intelligent risk control in the field of financial technology. It also shows strong scalability and can be adapted to other high-risk forecasting tasks, such as in energy markets, portfolio optimization [24], and cross-border capital supervision [25]. Future research could further expand the model's application boundaries. One direction is to incorporate graph-structured data to build cross-market risk contagion networks. Another is to integrate multimodal data fusion to improve responsiveness to complex economic signals. In addition, domain knowledge can be introduced to constrain the model, enhancing

interpretability and policy relevance. These developments will help deepen the integration of artificial intelligence into financial stability frameworks and public policy design.

5. Use of AI

We employed AI tools solely to assist with grammar and wording. The core concepts, analysis, and writing were all the responsibility of our team.

References

1. Y. Balmaseda, M. Coronado and G. de Cadenas-Santiago, "Predicting systemic risk in financial systems using deep graph learning," Proceedings of the 2023 Intelligent Systems with Applications Conference, pp. 200240, 2023.
2. V. Kanaparthi, "Transformational application of Artificial Intelligence and Machine learning in Financial Technologies and Financial services: A bibliometric review," arXiv preprint arXiv:2401.15710, 2024.
3. Y. Cheng, "Multivariate Time Series Forecasting through Automated Feature Extraction and Transformer-Based Modeling," Journal of Computer Science and Software Applications, vol. 5, no. 5, 2025.
4. W. Cui and A. Liang, "Diffusion-Transformer Framework for Deep Mining of High-Dimensional Sparse Data," Journal of Computer Technology and Software, vol. 4, no. 4, pp. 50–66, 2025, doi: 10.5281/zenodo.15340711.
5. J. Wang, "Credit Card Fraud Detection via Hierarchical Multi-Source Data Fusion and Dropout Regularization," Transactions on Computational and Scientific Methods, vol. 5, no. 1, pp. 15–28, 2025, doi: 10.5281/zenodo.14979821.
6. B. Chen, F. Qin, Y. Shao, J. Cao, Y. Peng and R. Ge, "Fine-Grained Imbalanced Leukocyte Classification With Global-Local Attention Transformer," Journal of King Saud University - Computer and Information Sciences, vol. 35, no. 8, Article ID 101661, 2023.
7. J. Luo, W. Zhuo and B. Xu, "A deep neural network-based assistive decision method for financial risk prediction in carbon trading market," Proceedings of the 2024 Journal of Circuits, Systems and Computers Conference, vol. 33, no. 08, pp. 2450153, 2024.
8. X. Du, "Audit Fraud Detection via EfficiencyNet with Separable Convolution and Self-Attention," Transactions on Computational and Scientific Methods, vol. 5, no. 2, pp. 33–45, 2025, doi: 10.5281/zenodo.14912715.
9. O. E. Ejiofor, "A comprehensive framework for strengthening USA financial cybersecurity: integrating machine learning and AI in fraud detection systems," Proceedings of the 2023 European Journal of Computer Science and Information Technology Conference, vol. 11, no. 6, pp. 62-83, 2023.
10. Q. Bao, "Advancing Corporate Financial Forecasting: The Role of LSTM and AI in Modern Accounting," Transactions on Computational and Scientific Methods, vol. 4, no. 6, 2024, doi: https://doi.org/10.5281/zenodo.15420017
11. X. Yan, J. Du, L. Wang, Y. Liang, J. Hu and B. Wang, "The Synergistic Role of Deep Learning and Neural Architecture Search in Advancing Artificial Intelligence," Proceedings of the 2024 International Conference on Electronics and Devices, Computational Science (ICEDCS), pp. 452-456, Sep. 2024.
12. T. M. A. U. Gunathilaka, J. Zhang and Y. Li, "Fine-Grained Feature Extraction in Key Sentence Selection for Explainable Sentiment Classification Using BERT and CNN," IEEE Access, 2025.
13. J. Gong, Y. Wang, W. Xu and Y. Zhang, "A Deep Fusion Framework for Financial Fraud Detection and Early Warning Based on Large Language Models," Journal of Computer Science and Software Applications, vol. 4, no. 8, 2024, doi: https://doi.org/10.5281/zenodo.15714273.
14. Y. Sheng, "Temporal Dependency Modeling in Loan Default Prediction with Hybrid LSTM-GRU Architecture," Transactions on Computational and Scientific Methods, vol. 4, no. 8, 2024, doi: https://doi.org/10.5281/zenodo.15420061
15. J. Wei, Y. Liu, X. Huang, X. Zhang, W. Liu and X. Yan, "Self-Supervised Graph Neural Networks for Enhanced Feature Extraction in Heterogeneous Information Networks," 2024 5th International Conference on Machine Learning and Computer Application (ICMLCA), pp. 272-276, 2024.

16. Y. Lou, "Capsule Network-Based AI Model for Structured Data Mining with Adaptive Feature Representation," *Transactions on Computational and Scientific Methods*, vol. 4, no. 9, pp. 77–89, 2024, doi: 10.5281/zenodo.15420046.
17. L. Dai, W. Zhu, X. Quan, R. Meng, S. Cai and Y. Wang, "Deep Probabilistic Modeling of User Behavior for Anomaly Detection via Mixture Density Networks," *arXiv preprint arXiv:2505.08220*, 2025.
18. M. Scholkemper, X. Wu, A. Jadbabaie and M. T. Schaub, "Residual connections and normalization can provably prevent oversmoothing in GNNs," *arXiv preprint arXiv:2406.02997*, 2024.
19. P. Feng, "Hybrid BiLSTM-Transformer Model for Identifying Fraudulent Transactions in Financial Systems," *Journal of Computer Science and Software Applications*, vol. 5, no. 3, pp. 45–58, 2025, doi: 10.5281/zenodo.14985026.
20. L. Almahadeen et al., "Enhancing Threat Detection in Financial Cyber Security Through Auto Encoder-MLP Hybrid Models," *International Journal of Advanced Computer Science and Applications*, vol. 15, no. 4, pp. 924–933, 2024, doi: 10.14569/IJACSA.2024.0150495.
21. Y. Cheng et al., "A Deep Learning Framework Integrating CNN and BiLSTM for Financial Systemic Risk Analysis and Prediction," *arXiv preprint arXiv:2502.06847*, 2025.
22. Y. Wei et al., "Financial Risk Analysis Using Integrated Data and Transformer-Based Deep Learning," *Proceedings of the 2024 Journal of Computer Science and Software Applications Conference*, vol. 4, no. 7, pp. 1-8, 2024, doi: 10.5281/zenodo.15714892
23. E. Sy et al., "Fine-grained argument understanding with bert ensemble techniques: A deep dive into financial sentiment analysis," *Proceedings of the 35th Conference on Computational Linguistics and Speech Processing (ROCLING 2023)*, 2023.
24. Z. Xu et al., "Reinforcement Learning in Finance: QTRAN for Portfolio Optimization," *Journal of Computer Technology and Software*, vol. 4, no. 3, 2025, doi: 10.5281/zenodo.15164617.
25. Y. Yao, "Time-Series Nested Reinforcement Learning for Dynamic Risk Control in Nonlinear Financial Markets," *Transactions on Computational and Scientific Methods*, vol. 5, no. 1, 2025, doi: <https://doi.org/10.5281/zenodo.14677117>

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.