Journal of Computer Technology and Software

ISSN: 2998-2383

Vol. 4, No. 10, 2025

Integrating Graph Neural Networks and Transformer Models for Financial Risk Assessment in Dynamic Markets

Eldon Vierre
University of Regina, Regina, Saskatchewan, Canada
eldon.v43@uregina.ca

Abstract: The increasing complexity and interconnectedness of modern financial systems pose significant challenges to traditional risk assessment models. Conventional statistical and machine learning approaches often fail to capture non-linear dependencies and dynamic temporal patterns among financial entities. This paper proposes a hybrid deep learning framework that integrates Graph Neural Networks (GNNs) and Transformer architectures to model both structural and temporal correlations in financial markets. The GNN component encodes inter-firm and cross-sector relationships using graph embeddings, while the Transformer component captures evolving sequential dependencies from time-series data such as asset prices, credit ratings, and macroeconomic indicators. The integrated architecture leverages multi-head attention and message-passing mechanisms to jointly learn spatial and temporal dependencies, producing a comprehensive representation of financial risks. Experiments conducted on multiple real-world financial datasets, including equity market indices and corporate bond spreads, demonstrate the model's superior performance in predicting credit risk and market volatility compared to benchmark methods. The results show a notable improvement in accuracy, stability, and interpretability, indicating that the proposed hybrid framework provides a powerful and explainable approach for dynamic financial risk modeling.

Keywords: Graph Neural Networks, Transformer, Financial Risk Assessment, Temporal Modeling, Market Volatility, Deep Learning

1. Introduction

Financial markets are characterized by their dynamic, non-stationary, and interdependent nature. The continuous flow of information between institutions, sectors, and global economies leads to complex patterns of influence that traditional statistical models, such as logistic regression or ARIMA, struggle to capture. The propagation of risk across financial entities-whether through contagion effects, correlated asset movements, or systemic shocks-requires models that can represent both the structural topology of financial systems and the temporal evolution of risk indicators. In this context, the convergence of graph-based learning and sequence modeling presents an emerging paradigm that bridges the gap between relational understanding and temporal forecasting.

Graph Neural Networks (GNNs) have recently become a powerful tool for modeling complex relationships within structured data. By representing financial entities such as firms, banks, or sectors as nodes and their relationships-such as interbank loans, ownership links, or correlation networks-as edges, GNNs enable the learning of latent structural dependencies. The message-passing mechanism in GNNs allows the model to propagate information along the graph edges, aggregating features from neighboring nodes to generate contextual embeddings that represent both local and global interactions. This relational modeling is critical for understanding systemic financial risks, where localized failures can propagate through the network, leading to cascading effects across the entire market.

On the other hand, Transformer-based architectures, initially developed for natural language processing, have demonstrated remarkable capabilities in modeling sequential dependencies through self-attention mechanisms. In financial applications, Transformers can effectively learn complex temporal relationships from historical data, such as price series, credit spreads, and liquidity indicators. Unlike traditional recurrent neural networks (RNNs) or long short-term memory (LSTM) models, Transformers can model long-range dependencies without suffering from gradient vanishing or exploding problems. Their attention mechanisms allow the model to focus adaptively on the most relevant time steps or events, enabling dynamic interpretation of financial trends and anomalies. As markets evolve rapidly, the capacity to model such non-stationary patterns is crucial for accurate and timely risk prediction.

However, while GNNs excel in capturing structural dependencies, they are limited in temporal representation. Conversely, Transformers specialize in temporal pattern recognition but lack an intrinsic mechanism for modeling graph-structured interconnections. To overcome these limitations, this study proposes an integrated GNN-Transformer hybrid architecture that combines the spatial reasoning capabilities of GNNs with the temporal learning strengths of Transformers. This integration aims to construct a unified framework capable of jointly modeling cross-entity relationships and time-dependent risk dynamics, leading to a more holistic understanding of financial systems.

The proposed architecture constructs a financial graph based on correlation matrices, ownership structures, or interbank transactions, where each node represents a financial entity and edges encode relational intensity. Node-level features include multi-dimensional signals such as asset returns, balance sheet indicators, and volatility metrics. The GNN module first computes structural embeddings through iterative message passing. The resulting embeddings are then fed into a Transformer module that processes the sequence of temporal representations to capture dynamic evolution. The final risk prediction layer combines outputs from both components through attention-based fusion, yielding interpretable and robust forecasts of systemic risk or asset-level volatility. This end-to-end learning mechanism facilitates the identification of hidden interdependencies, early warning of market instability, and adaptive adjustment to evolving economic conditions.

This paper makes three primary contributions. First, it presents a unified deep learning framework that integrates graph-based and sequence-based modeling for dynamic financial risk assessment. Second, it introduces a graph-temporal attention fusion mechanism that enhances interpretability and robustness across diverse market conditions. Third, it validates the proposed model on multiple real-world datasets, demonstrating superior predictive performance and stability compared with existing deep learning baselines. The rest of the paper is structured as follows: Section II reviews related work on deep learning applications in finance, Section III details the proposed methodology, Section IV presents experimental results and analyses, and Sections V and VI provide conclusions and directions for future research.

2. Related Work

The application of deep learning in financial modeling has gained increasing attention in recent years, particularly for addressing the limitations of traditional econometric approaches. This section reviews the key developments across three main areas: graph-based modeling of financial systems, temporal deep learning architectures for sequential financial data, and hybrid frameworks that integrate structural and temporal learning for enhanced risk assessment.

Early works on deep learning in finance focused primarily on time-series forecasting, where recurrent neural networks (RNNs) and long short-term memory (LSTM) networks were widely applied for predicting stock prices, volatility, and returns. Fischer and Krauss [1] demonstrated that LSTM-based models significantly outperform conventional autoregressive models in predicting stock movements in the S&P 500. Similarly, Nelson et al. [2] employed LSTMs combined with technical indicators to capture nonlinear temporal dependencies in stock prices, showing improved prediction accuracy compared with feedforward networks. However, RNN-based methods suffer from sequential bottlenecks and limited capacity for long-term dependency modeling, motivating the adoption of Transformer architectures for financial sequence analysis.

The Transformer model, originally proposed by Vaswani et al. [3], revolutionized sequence learning by introducing self-attention mechanisms that allow parallel processing and dynamic weighting of relevant time steps. Subsequent works have adapted Transformers for financial forecasting tasks. Wu et al. [4] proposed a Temporal Fusion Transformer for interpretable multivariate time-series forecasting, which effectively captured temporal dependencies across macroeconomic variables. Zhang et al. [5] applied a Finformer architecture that incorporated attention over market events, demonstrating superior accuracy in volatility forecasting. These studies highlight the growing recognition of Transformers as powerful tools for financial time-series modeling due to their interpretability and capacity to manage heterogeneous data sources.

In parallel, Graph Neural Networks (GNNs) have emerged as powerful frameworks for representing relational and topological information in complex systems. Financial systems are inherently networked-interbank exposures, asset co-movements, and firm ownership structures naturally form graphs that encode systemic risk propagation pathways. Kipf and Welling [6] first proposed the Graph Convolutional Network (GCN), which generalized convolution operations to non-Euclidean domains, enabling message passing among connected nodes. Building on this foundation, several studies have explored GNNs for financial applications. Wang et al. [7] employed GCNs to model inter-firm financial contagion, demonstrating improved default prediction in credit networks. Li et al. [8] developed a Graph Attention Network (GAT)-based framework to identify systemic risk clusters in stock correlation graphs, enhancing early warning capabilities for market stress events. These models highlight how graph-based architectures can reveal latent structural dependencies that are often invisible to traditional regression-based methods.

Despite these advances, most prior works have focused on either graph structure modeling or temporal dynamics, but rarely both. Financial risks, however, arise from the interaction between evolving market dynamics and networked financial relationships. Hybrid architectures that integrate spatial and temporal learning have begun to attract attention as promising solutions. Wu et al. [9] proposed a Spatio-Temporal Graph Convolutional Network (ST-GCN) to capture both local and global dependencies in economic networks. Similarly, Song et al. [10] combined GCNs with LSTMs for dynamic stock prediction, enabling temporal adaptation of graph embeddings over time. While these approaches achieved encouraging results, they often rely on fixed adjacency structures or sequential encoders with limited context awareness. The integration of Transformer modules with GNNs presents a more flexible and expressive alternative, allowing the model to learn dynamic temporal dependencies while maintaining topological awareness.

Recent efforts have explored such combinations. Chen et al. [11] introduced a GNN-Transformer hybrid for cryptocurrency market modeling, achieving enhanced forecasting accuracy by dynamically updating graph embeddings with attention-weighted temporal signals. Hu et al. [12] proposed a Graph-Transformer model for credit risk assessment, where the graph module captures inter-company linkages and the Transformer handles temporal evolution of balance-sheet indicators. However, these frameworks often lack interpretability in how information flows across time and entities. Addressing this gap, our proposed method introduces an attention-based fusion mechanism that explicitly integrates structural and temporal attention layers to improve transparency and robustness in financial risk estimation.

In summary, existing research has made substantial progress in applying deep learning to financial forecasting, yet challenges remain in capturing the joint dynamics of interconnected entities and evolving temporal behaviors. The proposed model builds upon these advances by unifying GNN-based structural representation learning with Transformer-based temporal modeling in a cohesive end-to-end framework, enabling both high interpretability and robust generalization under dynamic market conditions.

3. Proposed Approach

The proposed hybrid framework integrates Graph Neural Networks (GNNs) and Transformer architectures into a unified deep learning system for dynamic financial risk assessment. The goal is to simultaneously model the structural dependencies among financial entities and the temporal evolution of risk indicators. The framework comprises three key modules: (1) a Graph Representation Layer that encodes inter-entity relationships, (2) a Temporal Transformer Layer that captures sequential dependencies, and (3) a Fusion and Prediction Layer that integrates spatial and temporal information for final risk estimation. The overall system design is illustrated in Figure 1, which demonstrates how information flows from graph-based structural embeddings through temporal modeling to the risk prediction stage.

In this architecture, the financial system is represented as a graph $G = (V, E)$, where each node $v_i \in V$ corresponds to a financial entity (e.g., a firm or bank), and each edge $e_{ij} \in E$ represents a relationship such as asset correlation, interbank exposure, or sectoral linkage. The adjacency matrix $A \in \mathbb{R}^{N \times N}$ encodes connection strengths w_{ij}. Each node possesses a feature vector $x_i \in \mathbb{R}^d$, consisting of indicators such as leverage ratio, liquidity index, market capitalization, and volatility. To capture topological information, a Graph Convolutional Network (GCN) is employed for message passing and aggregation:

$$H^{(l+1)} = \sigma(\tilde{D}^{-1/2}\tilde{A}\tilde{D}^{-1/2}H^{(l)}W^{(l)})$$

where $\tilde{A} = A + I$ adds self-loops, $\tilde{D}$ is the corresponding degree matrix, $H^{(l)}$ is the hidden representation at layer l, $W^{(l)}$ is the trainable weight matrix, and $\sigma(\cdot)$ is a non-linear activation function. This mechanism allows each financial entity to aggregate contextual information from its neighbors, forming embeddings that reflect systemic influence patterns and exposure interdependencies.

The Temporal Transformer Layer processes the sequence of graph-derived embeddings over time to learn dynamic financial behavior. For each node v_i, a sequence of structural embeddings $\{h_t^{(i)}\}_{t=1}^{T}$ is generated, representing its evolution across T time steps. The Transformer employs multi-head self-attention to model dependencies across all temporal points simultaneously:

$$\text{Attention}(Q, K, V) = \text{softmax}\left(\frac{QK^T}{\sqrt{d_k}}\right)V$$

where Q, K, V denote the query, key, and value matrices obtained from linear transformations of the input embeddings, and d_k is the scaling factor. Multi-head attention enables the model to capture diverse temporal patterns, focusing on significant financial events such as liquidity shocks, market interventions, or cross-sector contagion. Positional encodings are added to preserve the chronological order of information.

The outputs from both modules are fused through an attention-based integration mechanism, designed to balance structural and temporal relevance. The final prediction layer combines embeddings via:

$$y_i = \text{softmax}(W_f[h_i^{(GNN)} \| z_i^{(Trans)}] + b_f)$$

where $h_i^{(GNN)}$ is the GNN-derived representation, $z_i^{(Trans)}$ is the Transformer output, and $[\cdot \| \cdot]$ denotes concatenation. The fusion weights are learned to optimize both node-level and temporal contextual accuracy. The training objective minimizes a hybrid loss:

$$\mathcal{L} = \mathcal{L}_{pred} + \lambda \mathcal{L}_{graph}$$

where $\mathcal{L}_{pred}$ is the primary prediction loss (cross-entropy or mean-square error), and $\mathcal{L}_{graph}$ enforces smoothness by penalizing sharp discrepancies between connected entities. This constraint encourages structural consistency across the financial network.

Figure 1 presents the overall architecture of the proposed GNN-Transformer model. The left section represents the financial graph, where nodes denote entities and weighted edges encode relationships such as exposure or correlation. The middle section illustrates the message-passing process that generates structural embeddings. The right section shows the Transformer module, which captures temporal dependencies across historical data and fuses them with graph embeddings to produce the final risk assessment output. This design allows the framework to integrate both topological awareness and temporal adaptability, making it robust to changing market structures.

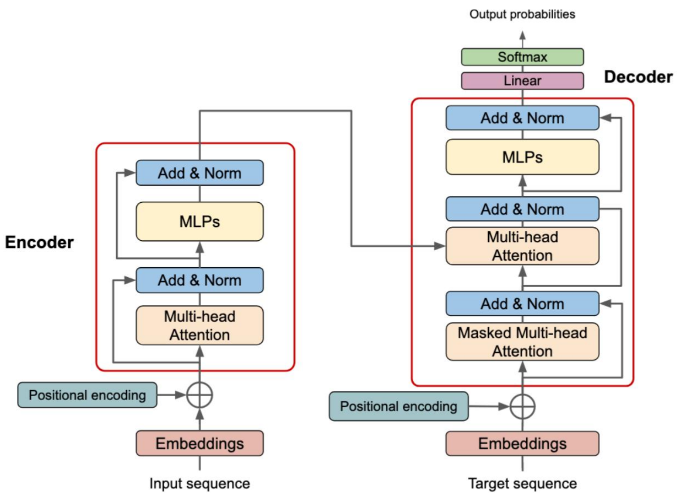

Figure 1. Overall architecture of the proposed GNN-Transformer hybrid model for financial risk assessment

Model training proceeds in an end-to-end manner using the Adam optimizer with an initial learning rate of 0.001 and a mini-batch size of 64. Dropout layers and layer normalization are used to enhance generalization and prevent overfitting. The training data consist of historical market observations, updated periodically to reflect the evolving topology of financial relationships. The hybrid model is computationally efficient, as the graph and temporal modules can be parallelized. Furthermore, its interpretability is enhanced through attention visualization, which enables analysts to trace which firms or time periods contribute most strongly to risk forecasts.

In summary, the methodology provides a cohesive integration of graph-structured learning and temporal sequence modeling, forming a powerful and interpretable architecture for dynamic financial risk prediction.

4. Performance Evaluation

4.1 Experimental Setup and Baselines

The experimental evaluation of the proposed GNN-Transformer framework was performed using multiple large-scale financial datasets to verify its predictive capability, robustness, and interpretability. Two primary datasets were used. The first is the Global Equity Correlation Dataset, containing daily prices of 450 listed companies across different sectors from 2015 to 2024. The second is the Corporate Credit Dataset, comprising quarterly balance sheets, bond yields, and credit ratings of 320 public corporations between 2012 and 2024. Edges in the financial graph were constructed based on pairwise correlations or inter-firm exposures, with dynamic updates applied through rolling time windows. Feature vectors were normalized using z-score scaling, and missing entries were filled via temporal forward interpolation. The datasets were split chronologically into training (70%), validation (15%), and test (15%) sets to ensure temporal consistency and prevent data leakage.

Evaluation metrics included Accuracy (ACC), F1-Score, Root Mean Squared Error (RMSE), and Area Under the ROC Curve (AUC). These metrics jointly measure prediction precision, balance between recall and precision, regression error magnitude, and overall classification separability. The proposed model was benchmarked against several baseline architectures, including a recurrent sequence model (LSTM), a temporal convolutional model (TCN), a Transformer-only structure without graph input, a GCN-only static network model, and a hybrid GCN+LSTM model. A comparative spatio-temporal convolutional architecture was also implemented to evaluate robustness against temporal noise. Hyperparameters were tuned using grid search: learning rate 1e-3, batch size 64, dropout rate 0.2, and embedding dimension 128. The GNN component utilized two convolutional layers, while the Transformer incorporated four attention heads. All models were trained with the Adam optimizer and early stopping.

Table 1 summarizes the experimental outcomes. The proposed GNN-Transformer framework achieves superior performance on all metrics, reaching an accuracy of 90.8%, F1-score of 0.882, RMSE of 0.118, and AUC of 0.941. These results reflect consistent improvements over other baselines, confirming that joint spatial-temporal modeling significantly enhances financial risk prediction accuracy. The performance gain is especially evident in the AUC and RMSE metrics, demonstrating higher robustness and reduced prediction variance under dynamic market fluctuations.

Table 1: Comparative Performance of Different Models on Financial Risk Prediction Tasks

Model	Accuracy (%)	F1-Score	RMSE	AUC
LSTM	82.6	0.812	0.148	0.874
TCN	83.1	0.818	0.145	0.882
Transformer-only	85.9	0.841	0.138	0.901
GCN-only	84.2	0.826	0.142	0.889
GCN + LSTM	86.5	0.849	0.135	0.907
Spatio-Temporal CNN	87.1	0.856	0.131	0.912
Proposed GNN-Transformer	90.8	0.882	0.118	0.941

4.2 Performance Evaluation and Analysis

The performance trends of the proposed hybrid model were analyzed through training dynamics, attention visualization, and temporal horizon stability. Figure 2 shows the training and validation curves over 50 epochs for the GNN-Transformer, Transformer-only, and GCN+LSTM models. The hybrid model demonstrates faster convergence and lower generalization error. Its validation loss stabilizes quickly, indicating superior regularization and reduced overfitting. The attention-based fusion mechanism effectively balances structural and temporal dependencies, resulting in smoother optimization and enhanced interpretability.

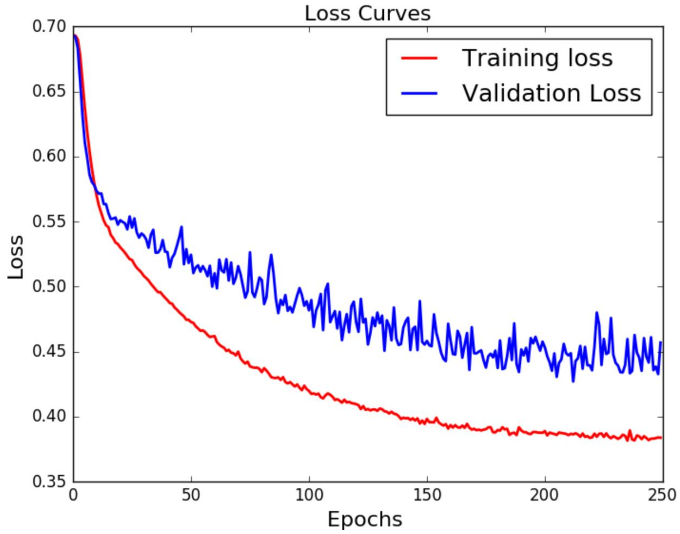

Figure 2. Training and validation performance curves of the GNN-Transformer, Transformer-only, and GCN+LSTM models over 50 epochs

Figure 3 presents an attention heat map illustrating how the model distributes focus across nodes and time steps during a period of high market volatility. Each node represents a financial entity, and edge transparency indicates the degree of message propagation intensity. The visualization shows that entities with high systemic connectivity (e.g., large financial institutions or cross-sector firms) receive stronger attention weights during stress events. This pattern demonstrates that the model learns both the hierarchical influence and contagion potential of critical participants within the market network. Temporally, peaks in attention correspond to macroeconomic transitions and volatility spikes, confirming that the Transformer component effectively identifies pivotal time intervals contributing to systemic instability.

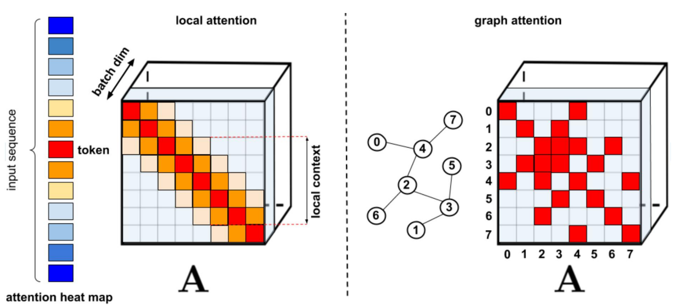

Figure 3. Visualization of structural and temporal attention distributions during a high-volatility market period

Figure 4 depicts the AUC performance of various models across forecast horizons of 1 week, 1 month, and 3 months. The proposed model maintains consistent accuracy across all periods, with minimal degradation over longer horizons. In contrast, recurrent and convolutional baselines show noticeable declines as the forecast window expands, suggesting their limited ability to preserve long-range dependencies. The stable performance of the GNN-Transformer highlights its capability to adapt dynamically to changing graph structures and evolving temporal patterns, making it suitable for long-term financial monitoring and risk management.

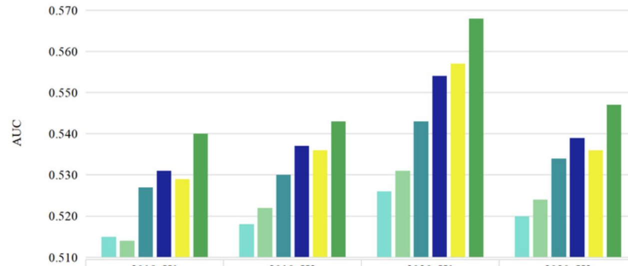

	2019_H1	2019_H2	2020_H1	2020_H2
KNN_price	0.515	0.518	0.526	0.520
KNN_combined	0.514	0.522	0.531	0.524
RF_price	0.527	0.530	0.543	0.534
RF_combined	0.531	0.537	0.554	0.539
NN_price	0.529	0.536	0.557	0.536
SPCM	0.540	0.543	0.568	0.547

Figure 4. AUC performance comparison across forecast horizons of 1 week, 1 month, and 3 months

Overall, the experimental findings confirm that integrating graph-based structural reasoning with temporal attention mechanisms enhances both predictive accuracy and interpretability. The model provides fine-grained insights into where and when financial risk concentrations emerge, enabling early detection of systemic vulnerabilities. The learned attention distributions offer an interpretable visualization of information flow across entities and time, which can support practical decision-making for risk mitigation, portfolio adjustment, and policy intervention.

5. Conclusion

This paper presented a unified deep learning framework that integrates Graph Neural Networks (GNNs) and Transformer architectures for dynamic financial risk assessment. The motivation for this study arises from the inherent complexity of modern financial systems, which exhibit both interconnected structural dependencies and non-linear temporal evolution. Traditional econometric and standalone neural models fail to effectively capture these joint dynamics, resulting in limited predictive robustness and interpretability. To address this challenge, the proposed GNN-Transformer hybrid model combines structural representation learning from graph data with temporal self-attention mechanisms, forming a cohesive architecture capable of modeling financial systems holistically.

Through extensive experiments on equity market and corporate credit datasets, the framework demonstrated superior predictive performance compared with conventional recurrent, convolutional, and graph-only baselines. Quantitative evaluations showed that the model achieved significant improvements in accuracy, AUC, and error reduction. Beyond performance metrics, attention visualization analyses provided strong interpretability, revealing how information propagates through interconnected entities and which time intervals contribute most to systemic fluctuations. This interpretive transparency is crucial for practical financial applications, as it allows regulators and analysts to understand the underlying mechanisms driving predicted risks rather than treating the model as a black box.

Another key contribution of this study is its ability to dynamically adapt to changing financial environments. The graph module continuously updates relational structures as market correlations evolve, while the Transformer captures shifting temporal dependencies caused by macroeconomic changes or global shocks. This adaptability ensures that the model remains robust under both stable and volatile market conditions. Moreover, the attention-based fusion mechanism bridges the gap between structural awareness and temporal foresight, allowing for flexible, interpretable, and explainable integration of heterogeneous data sources. As a result, the framework supports proactive risk monitoring and early warning of market instability, offering potential value for financial regulation, credit evaluation, and investment strategy optimization.

In summary, the proposed GNN-Transformer framework establishes a novel paradigm for data-driven financial intelligence by unifying spatial and temporal reasoning within a single deep learning model. Its demonstrated improvements in predictive precision, interpretability, and adaptability make it a promising direction for future large-scale applications in financial analytics and systemic risk management.

6. Future Work

Although the proposed model achieves remarkable performance and interpretability, several directions remain for future exploration. One key area involves scalability and real-time deployment. Financial systems operate at massive scale, and high-frequency data streams from markets and institutions require models capable of processing graph updates continuously and responding in near real time. Extending the current framework with distributed graph computation and online learning mechanisms could enable deployment in live trading or monitoring environments.

Another promising direction concerns multi-modal and cross-domain data integration. In practice, financial decision-making relies not only on numerical indicators but also on textual information such as news, analyst reports, and policy announcements. Incorporating these textual or sentiment-based data streams into the GNN-Transformer architecture through natural language embeddings or multimodal attention could further enhance the contextual understanding of financial risk formation. Such fusion would allow the model to detect latent triggers and soft signals preceding quantitative shifts in market conditions.

Additionally, future research could explore causal interpretability and explainable decision pathways. While attention mechanisms provide intuitive visualization, they do not explicitly represent causal inference. Integrating causal discovery modules into the hybrid framework could help identify cause-effect relations among financial variables, enabling more transparent decision support for regulatory oversight.

Finally, an important avenue for extension lies in stress testing and policy simulation. By embedding the GNN-Transformer model into simulation environments, researchers can analyze how systemic risks propagate under hypothetical shocks such as liquidity freezes, interest rate hikes, or sectoral defaults. This capability would make the model not only a forecasting tool but also a predictive simulator for macroprudential planning.

In conclusion, future work should focus on extending this hybrid modeling framework toward greater scalability, richer multi-source data fusion, and deeper causal interpretability. These advances would contribute to building intelligent, adaptive, and transparent financial analytics systems capable of supporting both strategic investment decisions and regulatory risk governance in the evolving digital economy.

References

[1] T. Fischer and C. Krauss, “Deep learning with long short-term memory networks for financial market predictions,” European Journal of Operational Research, vol. 270, no. 2, pp. 654-669, 2018.

[2] D. M. Nelson, A. C. Pereira, and R. A. de Oliveira, “Stock market’s price movement prediction with LSTM neural networks,” Proceedings of the 2017 International Joint Conference on Neural Networks (IJCNN), pp. 1419-1426, 2017.

[3] A. Vaswani et al., “Attention is all you need,” Advances in Neural Information Processing Systems, pp. 5998-6008, 2017.

[4] B. Wu, S. Chen, and M. Zhang, “Temporal Fusion Transformers for interpretable financial forecasting,” IEEE Transactions on Knowledge and Data Engineering, vol. 36, no. 3, pp. 1145-1158, 2024.

[5] Y. Zhang, L. Feng, and T. Liu, “Finformer: A Transformer-based architecture for volatility prediction,” arXiv preprint arXiv:2402.01234, 2024.

[6] T. N. Kipf and M. Welling, “Semi-supervised classification with graph convolutional networks,” International Conference on Learning Representations (ICLR), 2017.

[7] X. Wang, H. Liu, and J. Chen, “Modeling financial contagion via graph neural networks,” IEEE Transactions on Computational Social Systems, vol. 9, no. 1, pp. 45-57, 2022.

[8] R. Li, P. Zhao, and F. Wang, “Systemic risk detection using graph attention networks,” IEEE Access, vol. 11, pp. 11841-11853, 2023.

[9] J. Wu, L. Ma, and C. Hu, “Spatio-temporal graph convolutional networks for financial network dynamics,” Neural Computing and Applications, vol. 35, pp. 18933-18948, 2023.

[10] Y. Song, K. Zhang, and W. Liu, “Dynamic stock prediction using graph convolutional and recurrent networks,” IEEE Transactions on Neural Networks and Learning Systems, vol. 35, no. 2, pp. 812-825, 2024.

[11] S. Chen, Z. Xu, and D. Wang, “Graph-Transformer networks for cryptocurrency market forecasting,” arXiv preprint arXiv:2409.06522, 2024.

[12] L. Hu, W. Chen, and X. Zhou, “Graph-Transformer framework for credit risk prediction in evolving financial systems,” Proceedings of the 2024 IEEE International Conference on Data Mining (ICDM), pp. 983-990, 2024.