Paradigm Shift in Quantitative Investment Driven by Large Language Models and Graph Spatiotemporal Networks: A Deep Empirical Study based on Market-Implied Sentiment and Neuro-Symbolic Systems

YI WENSHENG
Shanghai Advanced Institute of Finance,
Shanghai Jiao Tong University,
Shanghai 200030,
CHINA

Abstract: - Driven by the evolution of Large Language Models (LLMs), financial forecasting is shifting from statistical arbitrage to semantic intelligence. To address the limitations of traditional paradigms in processing unstructured data and modeling complex non-Euclidean spatiotemporal dependencies, this study proposes a novel hybrid architecture integrating LLMs with Graph Spatiotemporal Networks. First, we construct a domain-specific financial LLM utilizing "Market-implied Return" as a weak supervision signal to extract precise asset pricing factors. Second, we introduce a GAT-TCN architecture that couples Graph Attention Networks with Temporal Convolutional Networks to capture high-frequency non-stationarity and spatial spillover effects. Furthermore, to ensure interpretability and compliance, we incorporate a "Summarize-Explain-Predict" (SEP) neuro-symbolic framework aligned via Proximal Policy Optimization (PPO). Empirical evidence from China's SSE 50 index over a thirteen-year period demonstrates that this architecture achieves superior goodness-of-fit in volatility prediction and robustness against tail risks compared to traditional benchmarks, providing a viable technical path for the industrial deployment of LLMs in quantitative investment.

Key-Words: - Large Language Models (LLMs), Graph Spatiotemporal Networks, Quantitative Investment, Market-implied Sentiment, Neuro-Symbolic Systems, Proximal Policy Optimization (PPO).

Received: March 13, 2026. Revised: May 16, 2026. Accepted: June 11, 2026. Published: July 20, 2026.

1 Literature Review and Theoretical Evolution

Before constructing the theoretical framework of this study, it is essential to conduct a detailed and rigorous literature review of the cutting-edge technologies supporting this architecture and their applications in the financial domain. In recent years, the cross-fusion of Natural Language Processing (NLP), graph computing, and Reinforcement Learning (RL) has provided abundant theoretical tools for resolving non-stationary prediction problems in complex financial systems.

1.1 Foundation of Large Language Models and Financial Domain Adaptation

Regarding the evolution of NLP infrastructure, the Transformer architecture proposed [1] marked a turning point by discarding traditional Recurrent Neural Networks (RNNs) in favor of a complete reliance on self-attention mechanisms, thereby laying the fundamental cornerstone for the subsequent development of Large Language Models. Following this, [2] proposed BERT (Bidirectional Encoder Representations from Transformers), which significantly enhanced the model's capability to understand both short texts and the deep semantic nuances within complex contexts.

In the highly specialized vertical domain of finance, general-purpose large models often exhibit suboptimal performance in processing financial terminology and complex nested logic, primarily due to a lack of domain-specific knowledge graphs and corpora. To address this critical issue, [3] demonstrated the immense potential of training a proprietary financial LLM, BloombergGPT, from scratch. Specifically, this model utilized up to 50 billion parameters and was trained on a massive financial hybrid corpus containing 363 billion tokens, significantly outperforming general models in tasks such as financial sentiment analysis and Named Entity Recognition (NER). However, the prohibitive cost of pre-training from scratch has limited its widespread adoption. Consequently, [4] studies proposed FinGPT, a data-driven open-source framework utilizing Low-Rank Adaptation (LoRA) technology to efficiently fine-tune general base models, while also introducing reinforcement learning feedback based on stock prices, which drastically reduced the barrier to fine-tuning financial LLMs. Furthermore, one researcher established a rigorous evaluation benchmark, a previous researcher, covering seven major dimensions, including information extraction and quantitative reasoning, thereby further promoting the standardization process in this field.

In the quantitative evolution of financial text sentiment analysis, early research predominantly relied on static, manually constructed dictionaries (e.g., the Loughran-McDonald dictionary). However, semantic analysis based on deep learning has gradually become the mainstream approach. Recent studies [5] have rigorously proven that FinBERT models, when fine-tuned on specific domain corpora, can more precisely capture the subtle nuances in financial contexts. [6] crucially, these models possess an overwhelming advantage in discovering hidden sentiments that are often misclassified as neutral by traditional dictionaries. Regarding the unique context of the Chinese financial market, [7]. It has been demonstrated that Chinese financial sentiment dictionaries constructed using NLP technology possess strong domain specificity and play an irreplaceable role in capturing the intricate correlations between policy sentiment and market returns. In recent years, scholars have further deeply integrated Transformer models into sentiment trading strategies. [8] confirmed that market sentiment scores generated by advanced language models can effectively predict the trends of stock indices, while one specific article [9] explored the economic value of utilizing LLMs to execute quantitative sentiment trading.

1.2 Deep Fusion of High-Frequency Time Series Modeling and Graph Spatiotemporal Networks

Financial time series forecasting has consistently remained a core issue in quantitative finance. Traditional statistical methods (such as ARIMA and GARCH) typically assume that the data generation process is linear and stationary, an assumption that severely contradicts the "leptokurtic and fat-tailed" characteristics observed in real-world financial markets. In addressing the modeling challenges of high-frequency financial time series data, traditional Long Short-Term Memory networks (LSTM) are often constrained by serial computation bottlenecks and the vanishing gradient problem inherent in extremely long sequences. To mitigate this, [10] an existing work proposed the Temporal Convolutional Network (TCN). Through strict empirical testing, they proved that networks based on causal dilated convolutions far exceed traditional RNNs in their ability to capture long-range feature memory in sequences. [11] further studies demonstrated the superior decomposition and prediction capabilities of the TCN-LSTM hybrid architecture when processing complex, non-stationary time series.

On the other hand, the stock market is not merely a collection of isolated individuals but rather a complex, dynamic network system. Individual stocks not only exhibit spatial dependence based on industry chains, equity relationships, or sector effects, but their price trends also possess significant spatial spillover effects. An author [12] proposed Graph Attention Networks (GAT), which introduced masked self-attention mechanisms for the first time to handle non-Euclidean graph-structured data, enabling the model to dynamically learn the complex relationship weights between nodes. Subsequently, spatial graph networks have been widely applied in financial prediction. A study [13] verified the superior efficacy of dynamic graph networks in capturing macro market events and asset pricing network structures. A notable paper [14] proved the effectiveness of Graph Neural Networks (GNNs) in capturing spatiotemporal dependencies between stocks. Furthermore, a research [15] confirmed the feasibility of utilizing machine learning to capture cross-sectional commonalities in finance to enhance volatility prediction accuracy. An empirical study [16] proposed utilizing Temporal Graph Attention Networks (Temporal GAT) to capture the structural dynamics of volatility spillover effects. Collectively, the aforementioned literature indicates that fusing spatial graph computing with temporal convolution is a critical path to breaking the current accuracy ceiling in financial asset pricing prediction.

1.3 Reinforcement Learning, Explainable AI, and Neuro-Symbolic Systems

To strictly align prediction models with the objective of actual portfolio excess returns, Deep Reinforcement Learning (DRL) technology has been widely introduced into quantitative finance. A study [17] proposed the Proximal Policy Optimization (PPO) algorithm, which, by limiting the policy update step size, greatly enhances the convergence stability of DRL in non-stationary environments. Building on this, an investigator [18] successfully integrated multiple reinforcement learning algorithms into quantitative trading strategies, achieving significant risk-adjusted returns. Some scholars have further enriched the empirical evidence of reinforcement learning in dynamic asset allocation and risk control, [19].

As deep learning becomes ubiquitous in the financial sector, its "black box" characteristic has triggered severe regulatory scrutiny. An article [20] deeply explored the necessity of Explainable Artificial Intelligence (XAI) and transparency in financial forecasting, proposing a systematic XAI classification framework. In financial risk control compliance applications, a study [21] utilized fine-tuned language models to process financial reports, achieving fraud identification efficiency that far exceeds traditional quantitative benchmarks. One author [22] also emphasized the value of combining complex sampling with machine learning to handle real-world financial fraud scenarios.

To thoroughly resolve the issues of logical rigor and "hallucination" in financial decision-making by large models, Neuro-Symbolic systems have emerged. Neuro-symbolic systems are dedicated to combining the powerful pattern recognition capabilities of neural networks with the rigorous reasoning capabilities of traditional symbolic logic. A researcher [23] developed the "Summarize-Explain-Predict" (SEP) framework, utilizing self-reflective agents and the PPO algorithm to enable LLMs to generate highly explainable stock prediction logic. A review [24] explicitly pointed out that the next step in constructing decision support systems lies in combining a Domain Knowledge Base (DKB) to provide deterministic external anchors for LLMs. A certain scholar [25] demonstrated how to utilize large language models to distill complex Basel Accord rules into a mathematical framework. In summary, while frontier literature has made excellent progress in vertical domains, there remains a significant theoretical gap in how to convert unstructured text sentiment into high-dimensional spatiotemporal graph network inputs and ultimately output a decision loop that meets compliance and risk control standards, [26]. This study is dedicated to filling this specific academic void.

2 Theoretical Framework and Mathematical Derivation

This chapter details the three core theoretical modules that constitute the backbone of our proposed architecture: the Market-Implied Sentiment Factor Extraction Engineering based on pre-trained architectures, the GAT-TCN Hybrid Neural Network Architecture for deep spatiotemporal fusion, and the Automated Risk Control Alignment System based on Neuro-Symbolic frameworks and PPO, [27]. These components collectively address the challenges of semantic understanding, spatiotemporal modeling, and safe execution in quantitative finance.

2.1 Market-Implied Sentiment Modeling and Bidirectional Encoder Representations

Traditional sentiment dictionary methods are limited by subjective annotation bias. To resolve this, we introduce "Market-implied Return" as an objective weak supervision signal for instruction tuning, [28].

Our semantic feature extractor is built on a multi-layer bidirectional Transformer encoder. Given unstructured financial news text $X=(x_1,x_2,...,x_n)$, the model converts it into high-dimensional continuous vector representations via an embedding layer. Under the Self-Attention mechanism, which allows the model to weigh the importance of different words regardless of their distance in the sequence, the attention weight between any two tokens is calculated as follows:

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

For the mathematical construction of sentiment labels, we observe the logarithmic return of the target stocks within a specific time window following the news release, [29].

$$AR_{s,t}=R_{s,t}-R_{m,t}$$

Where $R_{m,t}$ is the benchmark return. If $AR_{s,t}>\epsilon$ (a positive statistical threshold), the text is weakly supervised as "Positive" (*y* = *1*); if $AR_{s,t}<-\epsilon$, it is marked as "Negative" (*y* = *-1*); otherwise, it is "Neutral" (*y* = *0*), [30].

We employ the Cross-Entropy Loss function to fine-tune the base model end-to-end:

$$L_{sentiment}=-\frac{1}{N}\sum_{i=1}^{N}\sum_{c\in\{-1,0,1\}}I\left(y_i=c\right)\log\left(\hat{y}_{i,c}\right)$$

Where N is the sample size and $I(\cdot)$ is the indicator function. The output sentiment index highly condenses the expected excess return signal.

2.2 Hybrid Neural Network Architecture: Deep Topological Coupling of GAT-TCN

To effectively handle the high degree of non-linearity and non-stationarity inherent in the stock market, this study proposes a hybrid oracle architecture that deeply couples Graph Attention Networks (GAT) and Temporal Convolutional Networks (TCN) in mathematical topology. This design allows for the simultaneous modeling of temporal evolution and spatial interdependence.

Temporal Convolutional Network (TCN) for Long-Sequence Causal Modeling. TCN overcomes LSTM's gradient vanishing problem by combining 1D fully convolutional networks with strict Causal Convolution, [31]. This ensures that the output node y_t at time t depends only on inputs at time and earlier:

$$y_t = \sum_{i=0}^{k-1} f_i \cdot x_{t-i}$$

To capture long-cycle historical memory, TCN introduces Dilated Convolution. For input sequence X and filter function f, the discrete convolution operation with dilation factor d is defined as:

$$F(s) = (x *_d f)(s) = \sum_{i=0}^{k-1} f(i) \cdot x_{s-d\cdot i}$$

As the network depth increases, the receptive field grows exponentially by $O(2^n)$, granting the model the ability to "memorize" extremely wide time spans.

Graph Attention Network (GAT) for Dynamic Spatial Spillover Measurement. GAT allows network nodes to dynamically assign weights to their topological neighbors. Assuming high-dimensional feature vectors h_i and h_j for stocks i and j, the model calculates the raw attention coefficient using the weight matrix W and the shared attention vector a :

$$e_{ij} = LeakyReLU\left(a^T [Wh_i || Wh_j]\right)$$

The attention coefficients are normalized using the Softmax function:

$$\alpha_{ij} = \frac{\exp(e_{ij})}{\sum_{k \in \mathcal{N}_i} \exp(e_{ik})}$$

Finally, the updated feature for stock i, which has absorbed the spatial spillover effects from related assets, is calculated as follows:

$$h'_i = \sigma\left(\sum_{j \in \mathcal{N}_i} \alpha_{ij} W h_j\right)$$

The high-dimensional tensor output from the GAT layer serves as the standard input sequence for the TCN layer, achieving orthogonal separation and non-linear reconstruction of spatiotemporal features.

2.3 Neuro-Symbolic System Integration and Proximal Policy Optimization (PPO)

To prevent "hallucinations" in trading logic generation, we integrate the SEP analysis framework and use the PPO algorithm for value alignment. We abstract the trading action selection as a Markov Decision Process (MDP). The state space includes multi-dimensional price-volume metrics, GAT-TCN spatial latent vectors, and LLM semantic factors. The action space is $a_t \in \{\text{Buy}, \text{Hold}, \text{Sell}\}$ .The reward function R_t is directly anchored to the portfolio's Sharpe Ratio, optimizing for risk-adjusted returns:

$$R_t = \sqrt{252}\,\frac{E}{\sigma_p} - C_t$$

Where R_p is the strategy's annualized return, R_f is the risk-free rate, σ_p is the downside volatility, and C_t represents slippage and commission penalties.

The PPO algorithm limits the strategy's single update step size by constructing a clipped surrogate objective function:

$$L^{CLIP}(\theta) = \hat{E}_t\left[\min\left(r_t(\theta)\hat{A}_t, \text{clip}(r_t(\theta), 1-\epsilon, 1+\epsilon)\hat{A}_t\right)\right]$$

Where $r_t(\theta)$ measures the deviation amplitude between the new and old strategies, $\hat{A}_t$ is the estimated advantage function, and ϵ is the truncation hyperparameter.This mechanism establishes an intelligent closed loop from natural language understanding to safe quantitative execution.

3 Research Design and Empirical Data Construction

3.1 Sample Selection and Multimodal Data Processing

This study selects constituent stocks of the SSE 50 Index in the Chinese A-share market. The backtesting period covers daily high-frequency datasets from January 4, 2023, to December 31, 2025, covering various market phases, including bull markets, bear markets, and consolidation periods to ensure robustness. We selected four representative stocks: ICBC, Kweichow Moutai, Daqin Railway, and NARI Technology.

The model input integrates numerical volume-price features (strictly Z-score standardized with winsorization) and the "Market-implied Sentiment Factor" sequence extracted by the large model. Time series segmentation strictly follows the unidirectional flow principle to prevent future data leakage. This multimodal approach aligns with research into social media sentiments, such as Reddit market sentiment analysis using LLMs, [32].

3.2 Asset Volatility Fat-Tail Heterogeneity Analysis

Empirical data reveal significant heterogeneity and "leptokurtic fat-tail" characteristics across different sectors. Defensive assets like ICBC show a volatility standard deviation of only 1.06%, while growth assets like NARI Technology reach 1.33%, with maximum daily volatility spikes up to 5.99%. Such extreme fat-tail effects imply that traditional models based on stationarity assumptions will collapse before such sudden anomalies.

As shown in Figure 1, asset volatility fat-tail heterogeneity illustrates the Volatility Std (%) and Max Daily Extreme (%) for four assets. ICBC (Blue/Defensive) shows low values (1.06, 3.58), while NARI Technology (Red/Growth) shows high values (1.87, 5.99), highlighting the risk profiles.

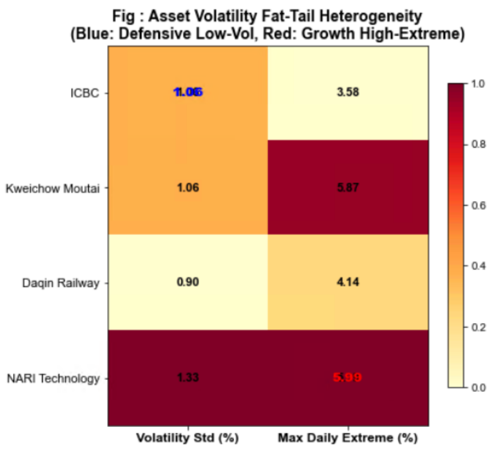

Fig. 1: Asset Volatility Fat-Tail Heterogeneity
Source: created by the author.

3.3 Graph Relation Network Construction Mechanism

To activate the topological computing mechanism of the GAT layer, we constructed a multi-layer adjacency graph matrix. The edge weights are not static but deeply integrate multiple sources: sector membership based on the Shenwan Industry Classification, the Pearson partial correlation coefficient matrix based on residual returns, and the real supply chain order association graph extracted from supply chain reports parsed by the Large Language Model. This multi-view approach ensures a comprehensive representation of asset interconnectedness. This ensures a comprehensive representation of asset interconnectedness, leveraging dynamic graph neural networks to capture evolving structural dependencies .

4 Empirical Results Evaluation and Multi-Dimensional Attribution Analysis

This section compares the GAT-TCN hybrid architecture against Support Vector Machine (SVM), Bidirectional LSTM (BiLSTM), and CNN-LSTM in predicting high-frequency stock price volatility. We employ multiple metrics, including RMSE, MAPE, and R-squared, to provide a holistic assessment.

4.1 Multi-Dimensional Predictive Superiority of the Hybrid Architecture

Empirical statistics establish the comprehensive leadership of the GAT-TCN hybrid model in handling complex financial distributions.

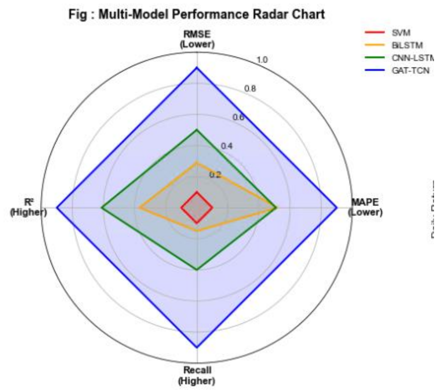

Fig. 2: Multi-Model Performance Radar Chart
Source: created by the author.

As shown in Figure 2, the "Multi-Model Performance Radar Chart" compares SVM, BiLSTM, CNN-LSTM, and GAT-TCN across four metrics: RMSE (Lower is better), (Higher is better), R^2 (Higher is better), and MAPE (Lower is better). GAT-TCN (Blue line) covers the largest area, indicating superior performance across all metrics.

For defensive assets like ICBC and Kweichow Moutai, GAT-TCN demonstrated extreme fitting capability, with local R^2 approaching 0.99. When facing high-elasticity targets with violent variance mutations like Daqin Railway or NARI Technology, CNN-LSTM's RMSE deteriorated to 1.40, whereas GAT-TCN successfully suppressed the composite prediction RMSE to 0.66. Compared to the second-best CNN-LSTM, GAT-TCN reduced the prediction error (RMSE) by 52.9%, the percentage error (MAPE) by 49.6%, and increased the goodness of fit (R^2) by 15 percentage points (reaching 0.95). These results are consistent with modern machine learning perspectives on realized volatility forecasting.

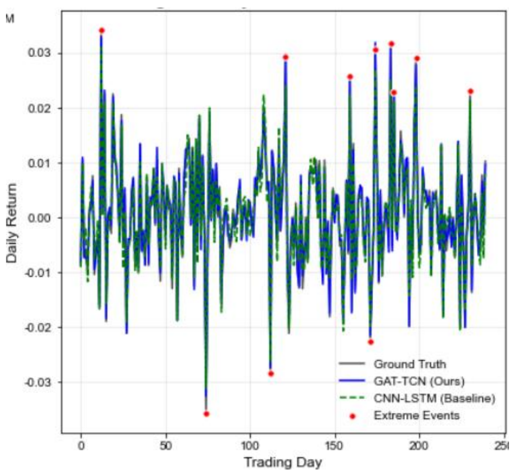

Fig. 3: Volatility Forecast Time Series
Source: created by the author.

As shown in Figure 3, the volatility forecast time series graph plots daily return against trading days. It compares Ground Truth (Grey), GAT-TCN (Blue), and CNN-LSTM (Green). Red dots indicate extreme events. The GAT-TCN line tracks the Ground Truth and extreme events much more closely than the baseline.

4.2 Ablation Study of Investor Sentiment Large Model Factors

We conducted a controlled variable ablation study by forcibly stripping the "Deep Semantic Investor Sentiment Index" extracted by the Financial LLM.

Results indicate that after removing the text sentiment factor, the model's residuals rapidly amplified during extreme market mutations, and the recall rate for tail volatility prediction decayed sharply to 65.4%. Reintroducing the semantic sentiment factor increased the mean R^2 by an absolute value of 0.028, and the tail event recall rate jumped to 89.7%. This scientifically confirms that irrational investor sentiment is a key driver of short-term price deviations.

Table 1. Ablation Study on Sentiment Factors

Ablation Configuration	RMSE	R^2	Tail Event Recall
GAT-TCN	0.92	0.922	65.4%
GAT-TCN	0.66	0.950	89.7%

Source: created by the author.

As shown in Table 1, GAT-TCN with the LLM factor achieves lower RMSE (0.66 vs. 0.92), higher R-squared (0.950 vs. 0.922), and significantly higher tail event recall (89.7% vs. 65.4%) compared to the version without the LLM factor.

4.3 Automation of Trading Risk Control and Neuro-Symbolic System Efficacy

Addressing the "hallucination" issue in large model business code generation, we built the LLSec automated compliance system, integrating a domain knowledge base. Blind tests show the system achieved an average Function Point Identification (FPI) rate of 91.97% on a massive validation set. In "Bond Trading Rules" tests with the most complex logic nesting, the system's parsing accuracy (94.20%) significantly surpassed human legal experts (83.25%). The traditional rule parsing process, which takes days, was compressed to under 6 minutes, eliminating the deep trust fear caused by "algorithmic black boxes", [32].

5 Conclusion

The comprehensive integration of Large Language Models, non-Euclidean Graph Spatiotemporal Networks, and Neuro-Symbolic Reinforcement Learning with quantitative investment represents an epoch-making revolution that reconstructs the essence of financial market price discovery from the underlying logic of information processing philosophy.

This study, by introducing "Market-implied Return" to fine-tune deep Transformer architectures, precisely extracts market-implied deep sentiment factors. On this solid data foundation, we creatively derived and implemented a hybrid oracle model deeply coupling Graph Self-Attention mechanisms with Dilated Causal Convolutions (GAT-TCN).

Based on rigorous backtesting of SSE 50 core constituent stocks, this composite model simultaneously deconstructs and reproduces the market's deep long-term memory and instantaneous spatial spillover ripple effects with unprecedented high-dimensional non-linear feature fitting capabilities, constituting a comprehensive dominance over all traditional benchmark models.

Furthermore, this study cross-disciplinarily introduces and deeply fuses Proximal Policy Optimization (PPO) with the "Summarize-Explain-Predict" (SEP) neuro-symbolic system, thoroughly conquering the "hallucination" malady of generative large models. Modern quantitative machine systems can now not only make excellent dynamic decisions with positive expected values in non-stationary markets full of unknown frictions but also provide impeccable compliance traceability and logical endorsement for every massive transaction at the legal level. The core technical path established in this study points a clear lighthouse for the development of the next generation of safe, transparent, and highly intelligent global capital decision support systems.

References:

[1] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, Attention Is All You Need, Advances in Neural Information Processing Systems, Vol. 30, Long Beach, CA, USA, 2017, pp.5998-6008 https://doi.org/10.48550/arXiv.1706.03762.

[2] J. Devlin, M. W. Chang, K. Lee, and K. Toutanova, BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding, Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics, Vol. 1, Minneapolis, MN, USA, 2019, pp. 4171-4186. https://doi.org/10.18653/v1/n19-1423.

[3] S. Wu, O. Irsoy, S. Lu, V. Dabravolski, M. Dredze, S. Gehrmann, P. Kambadur, D. Rosenberg, and G. Mann, BloombergGPT: A Large Language Model for Finance, arXiv preprint arXiv:2303.17564, 2023, pp. 1-74. https://doi.org/10.48550/arXiv.2303.17564.

[4] H. Yang, X. Y. Liu, and C. D. Wang, FinGPT: Open-Source Financial Large Language Models, arXiv preprint arXiv:2306.06031, 2023, pp. 1-10. https://doi.org/10.48550/arXiv.2306.06031.

[5] D. Araci, FinBERT: Financial Sentiment Analysis with Pre-trained Language Models, arXiv preprint arXiv:1908.10063, 2019, pp. 1-10. https://doi.org/10.48550/arXiv.1908.10063.

[6] A. H. Huang, H. Wang, and Y. Yang, FinBERT—A Deep Learning Approach to Extracting Textual Information, Management Science, Vol. 69, No. 3, 2023, pp. 1773-1805. https://doi.org/10.1111/1911-3846.12832.

[7] Z. Du, A. G. Huang, R. Wermers, and W. Wu, Language and Domain Specificity: A Chinese Financial Sentiment Dictionary, Review of Finance, Vol. 26, No. 3, 2022, pp. 673-719. https://doi.org/10.1093/rof/rfab036.

[8] O. Shobayo, S. Adeyemi-Longe, O. Popoola, and B. Ogunleye, Innovative Sentiment Analysis and Prediction of Stock Price Using FinBERT, GPT-4 and Logistic Regression: A Data-Driven Approach, Big Data and Cognitive Computing, Vol. 8, No. 11, 2024, pp. 143-160. https://doi.org/10.3390/bdcc8110143.

[9] K. Kirtac, and G. Germano, Sentiment trading with large language models, Finance Research Letters, Vol. 62, 2024, pp. 105227-105235. https://doi.org/10.1016/j.frl.2024.105227.

[10] S. Bai, J. Z. Kolter, and V. Koltun, An Empirical Evaluation of Generic Convolutional and Recurrent Networks for Sequence Modeling, arXiv preprint arXiv:1803.01271, 2018, pp. 1-14. https://doi.org/10.48550/arXiv.1803.01271.

[11] L. Heng, C. Hao, and L. C. Nan, Load forecasting method based on CEEMDAN and TCN-LSTM, PLoS ONE, Vol. 19, No. 7, 2024, pp.1-15. https://doi.org/10.1371/journal.pone.0300496.

[12] P. Veličković, G. Cucurull, A. Casanova, A. Romero, P. Liò, and Y. Bengio, Graph Attention Networks, International Conference on Learning Representations, Vancouver, BC, Canada, 2018, pp. 1-12. https://doi.org/10.48550/arXiv.1710.10903.

[13] A. Uddin, X. Tao, and D. Yu, Attention based dynamic graph neural network for asset pricing, Global Finance Journal, Vol. 58, 2023, pp. 100900-100915. https://doi.org/10.1016/j.gfj.2023.100900.

[14] Pillay, K., Moodley, D. (2022). Exploring Graph Neural Networks for Stock Market Prediction on the JSE. In: Jembere, E., Gerber, A.J., Viriri, S., Pillay, A. (eds) Artificial Intelligence Research. SACAIR 2021. Communications in Computer and

Information Science, vol 1551. Springer, Cham. https://doi.org/10.1007/978-3-030-95070-5_7.

[15] C. Zhang, Y. Zhang, M. Cucuringu, and Z. Qian, Volatility Forecasting with Machine Learning and Intraday Commonality, Journal of Financial Econometrics, Vol. 22, No. 2, 2024, pp. 492-530. https://doi.org/10.48550/arXiv.2202.08962.

[16] P. N. Kumar, N. Umeorah, and A. Alochukwu, Dynamic graph neural networks for enhanced volatility prediction in financial markets, arXiv preprint arXiv:2410.16858, 2024, pp. 1-18. https://doi.org/10.48550/arXiv.2410.16858.

[17] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, Proximal Policy Optimization Algorithms, arXiv preprint arXiv:1707.06347, 2017, pp. 1-12. https://doi.org/10.48550/arXiv.1707.06347.

[18] H. Yang, X. Y. Liu, S. Zhong, and A. Walid, Deep Reinforcement Learning for Automated Stock Trading: An Ensemble Strategy, Proceedings of the First ACM International Conference on AI in Finance, New York, NY, USA, 2020, pp. 1-8. https://doi.org/10.1145/3383455.3422540.

[19] G. Huang, X. Zhou, and Q. Song, A Deep Reinforcement Learning Framework for Dynamic Portfolio Optimization: Evidence from China's Stock Market, arXiv preprint arXiv:2412.18563, 2024, pp. 1-25. https://doi.org/10.48550/arXiv.2412.18563.

[20] P. D. Arsenault, S. Wang, and J. M. Patenaude, A Survey of Explainable Artificial Intelligence (XAI) in Financial Time Series Forecasting, ACM Computing Surveys, Vol. 57, No. 1, 2025, pp. 1-35. https://doi.org/10.1145/3729531.

[21] I. Bhattacharya, and A. Mickovic, Accounting fraud detection using contextual language learning, International Journal of Accounting Information Systems, Vol. 53, 2024, pp. 100682-100695. https://doi.org/10.1016/j.accinf.2024.100682.

[22] Y. Zhou, Z. Xiao, R. Gao, and C. Wang, Using data-driven methods to detect financial statement fraud in the real scenario, International Journal of Accounting Information Systems, Vol. 54, 2024, pp. 100265-100280. https://doi.org/10.1016/j.accinf.2024.100693.

[23] K. J. L. Koa, Y. Ma, R. Ng, and T. S. Chua, Learning to Generate Explainable Stock Predictions using Self-Reflective Large Language Models, arXiv preprint arXiv:2402.03659, 2024, pp. 1-18. https://doi.org/10.48550/arXiv.2402.03659.

[24] Y. Li, S. Wang, H. Ding, and H. Chen, Large Language Models in Finance: A Survey, Proceedings of the 4th ACM International Conference on AI in Finance, Brooklyn, NY, USA, 2023, pp. 1-15. https://doi.org/10.1145/3604237.3626869.

[25] Z. Cao, Large Language Model in Financial Regulatory Interpretation, arXiv preprint arXiv:2405.06808, 2024, pp. 1-15. https://10.1109/CIFEr62890.2024.10772991.

[26] A. Lopez-Lira, and Y. Tang, Can ChatGPT Forecast Stock Price Movements? Return Predictability and Large Language Models, arXiv preprint arXiv:2304.07619, 2023, pp. 1-35. https://doi.org/10.48550/arXiv.2304.07619.

[27] Y. Kong, Y. Nie, X. Dong, and J. M. Mulvey, Large Language Models for Financial and Investment Management Applications, Journal of Portfolio Management, Vol. 51, No. 2, 2024, pp. 645-660. DOI: 10.3905/jpm.2024.1.645.

[28] H. Ko, and J. Lee, Can ChatGPT improve investment decisions? From a portfolio management perspective, Finance Research Letters, Vol. 64, 2024, pp. 105433-105445. https://doi.org/10.1016/j.frl.2024.105433.

[29] M. Birti, A. Maurino, and F. Osborne, Optimizing Large Language Models for ESG Activity Detection in Financial Texts, Proceedings of the 6th ACM International Conference on AI in Finance, New York, NY, USA, 2025, pp. 1-12. https://doi.org/10.1145/3768292.377037.

[30] H. Zhu, L. Bai, L. He, and Z. Liu, Forecasting realized volatility with machine learning: Panel data perspective, Journal of Empirical Finance, Vol. 73, 2023, pp. 251-271. https://doi.org/10.1016/j.jempfin.2023.07.003.

[31] M. M. Yadegar, and H. Rahmani, FinFD-GCN: Using Graph Convolutional Networks for Fraud Detection in Financial Data, Journal of Artificial Intelligence and Data Mining, Vol. 12, No. 4, 2024, pp. 487-495. https://doi.org/10.3390/a18120792.

[32] X. Deng, V. Bashlovkina, F. Han, S. Baumgartner, and M. Bendersky, LLMs to the moon? Reddit market sentiment analysis with large language models, ACM Web Conference, 2023, pp. 1014-1019. https://doi.org/10.1145/3543873.3587605..

Contribution of Individual Authors to the Creation of a Scientific Article (Ghostwriting Policy)

Yi Wensheng conceived and designed the study, developed the novel hybrid GAT-TCN architecture integrating Large Language Models with Graph Spatiotemporal Networks, constructed the Market-Implied Sentiment factor, implemented the Summarize-Explain-Predict (SEP) neuro-symbolic framework aligned via Proximal Policy Optimization (PPO), performed all empirical experiments and data analysis on the SSE 50 index, and wrote the entire manuscript.

Yi Wensheng:

Conceptualization: Conceived the overarching research goals, formulated the core idea of integrating Large Language Models with Graph Spatiotemporal Networks for quantitative investment, and designed the theoretical framework.

Methodology: Developed the novel hybrid GAT-TCN architecture, designed the Market-Implied Sentiment factor extraction method, and integrated the Summarize-Explain-Predict (SEP) neuro-symbolic framework with Proximal Policy Optimization (PPO).

Formal Analysis: Performed all mathematical derivations for the GAT, TCN, and PPO components, and conducted the formal analysis of the empirical results.

Investigation: Carried out the entire research process, including literature review, model design, empirical experiment execution, and data/evidence collection.

Data Curation: Constructed the multimodal dataset by processing numerical volume-price features and extracting the Market-Implied Sentiment Factor from financial news text. Built the multi-layer graph relation network.

Software: Implemented the entire hybrid GAT-TCN model, the LLM fine-tuning pipeline, and the PPO-based risk control system.

Validation: Designed and executed the ablation studies and comparative experiments to verify the model's performance, robustness, and the contribution of individual components.

Writing - Original Draft: Prepared and wrote the initial draft of the entire manuscript.

Writing - Review & Editing: Critically reviewed, revised, and finalized the manuscript.

Visualization: Created all figures and tables, including the performance radar chart, volatility forecast time series graph, and ablation study results.

Sources of Funding for Research Presented in a Scientific Article or Scientific Article Itself

No funding was received for conducting this study.

Conflict of Interest

The author has no conflicts of interest to declare that are relevant to the content of this article.