Archives available at journals.mriindia.com

International Journal of Electrical, Electronics and Computer Systems

ISSN: 2347-2820

Volume 14 Issue 01, 2025

A Survey of Methods and Architectures for Deformable Graph Convolutional Networks with NLP Based Social Sentimental Data for Enhanced Stock Price Predictions

Aurelio Trivedi-Rao
Professor, Department of Computer Science and Engineering, Visayan Maritime Polytechnic University, Philippines
Email: aurelio.trivedi.rao@vmpu-ph.net

Peer Review Information

Submission: 08 April 2025

Revision: 22 April 2025

Acceptance: 06 May 2025

Keywords

Deformable Graph Convolutional Networks, Stock Price Prediction, Sentiment Analysis, Natural Language Processing, Financial Deep Learning, Graph Neural Networks

Abstract

Stock price prediction remains a challenging problem at the intersection of finance, mathematics, and computer science due to the complex, dynamic, and sentiment-driven nature of financial markets. Traditional statistical and machine learning models often fail to capture relational dependencies, temporal dynamics, and behavioral influences embedded in market data. The evolution of deep learning has enabled more sophisticated approaches that integrate multiple data modalities for improved forecasting performance. This paper presents a comprehensive review of a tri-modal framework combining Deformable Graph Convolutional Networks (DGCNs), Natural Language Processing (NLP)-based sentiment analysis, and dynamic graph construction for stock price prediction. Graph convolutional networks model the relational structure of financial markets, while their deformable variants introduce adaptive mechanisms to capture evolving inter-asset relationships more effectively. Simultaneously, transformer-based NLP models such as BERT and FinBERT extract sentiment signals from unstructured data sources including social media, financial news, and earnings reports. These sentiment features are integrated with structured financial data to form a multimodal predictive system that reflects both quantitative trends and investor behavior. Empirical studies demonstrate improved accuracy and robustness over traditional and single-modality models. However, challenges such as data preprocessing, noise handling, and scalability remain. This review highlights key advancements and future directions for developing intelligent, robust, and scalable financial forecasting systems.

Introduction

The global financial markets represent one of the most complex, high-dimensional, and information-rich systems ever developed, where billions of transactions occur daily across interconnected exchanges. These markets are influenced by a wide range of factors, including macroeconomic indicators, geopolitical events, corporate disclosures, investor sentiment, and algorithmic trading strategies. Predicting stock prices within such a dynamic and nonlinear environment has long been a central objective in both academic research and financial practice. Over time, the field has evolved from traditional approaches like fundamental and technical analysis to more advanced data-driven methods, reflecting continuous efforts to better understand and model market behavior.

© 2025 The Authors. Published by MRI INDIA.
This is an open access article under the CC BY-NC-ND license (http://creativecommons.org/licenses/by-nc-nd/4.0/).

For much of the twentieth century, the efficient market hypothesis dominated financial thinking, suggesting that stock prices fully reflect all available information. However, advancements in computing power, the widespread availability of financial data, and insights from behavioral finance have challenged this notion. Empirical studies have shown that markets often exhibit predictable patterns driven by psychological biases, structural inefficiencies, and information asymmetries. Early machine learning models such as support vector machines and random forests attempted to exploit these patterns but were limited by their reliance on structured numerical data and their inability to capture complex dependencies or adapt to evolving market conditions.

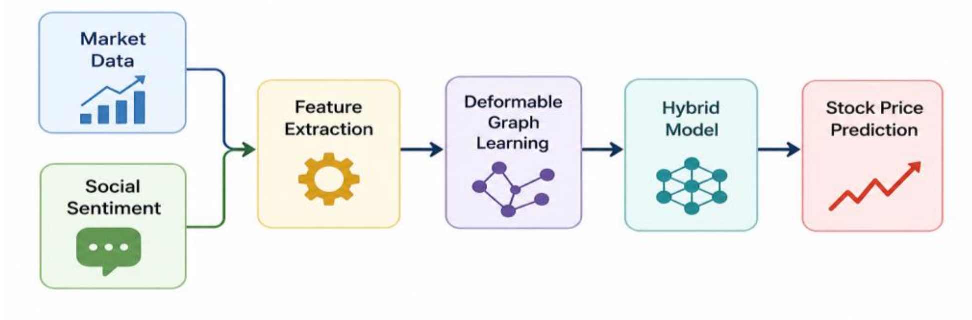

Fig 1: Proposed Hybrid Framework for Stock Price Prediction

The introduction of deep learning marked a significant shift, with recurrent neural networks and long short-term memory models improving the ability to capture temporal dependencies in financial time series. Nevertheless, these models struggled with computational efficiency and lacked the ability to represent relationships among multiple assets. This gap led to the adoption of graph neural networks, which model financial markets as interconnected systems, allowing information to propagate across related entities. However, traditional graph models assume static relationships, which is unrealistic in financial contexts where correlations and dependencies continuously change. Deformable graph convolutional networks address this limitation by enabling adaptive and dynamic graph structures that better reflect real-time market conditions.

In parallel, the growing importance of social sentiment has introduced a new dimension to financial prediction. With the rise of platforms generating massive volumes of textual data, natural language processing techniques—particularly transformer-based models—have become essential for extracting meaningful insights from unstructured information. By combining sentiment analysis with deformable graph-based learning, modern architectures can integrate both structural and contextual signals, leading to more accurate and robust predictions. This survey focuses on such hybrid approaches, emphasizing their potential to outperform traditional models by leveraging multiple data modalities and adaptive learning mechanisms in a unified framework.

Literature Review

The foundational work of Kipf and Welling (2017) introduced the graph convolutional network framework that would become the basis for nearly all subsequent graph-based financial prediction research. Their semi-supervised classification approach demonstrated that spectral convolutions on graphs could be approximated by simple first-order neighborhood aggregation operations, yielding a computationally efficient and empirically powerful model. While their original application was to citation networks and social graphs, the framework was quickly adopted by the financial research community as a natural tool for modeling inter-stock relationships. This work established the conceptual vocabulary and mathematical notation that subsequent graph-based financial models would build upon.

Chen et al. (2018) were among the first to apply graph neural networks specifically to stock prediction tasks, constructing a relational graph based on industry classification and supply chain relationships derived from corporate filings. Their model, which they termed the Relational Stock Ranking network, demonstrated that incorporating inter-stock dependencies substantially improved ranking accuracy compared to models that treated stocks in isolation. The study used data from the S&P 500 and Chinese A-share markets, providing early evidence that graph-based approaches generalize across different market structures. Their contribution was particularly notable for establishing a principled methodology for constructing financial graphs from heterogeneous data sources.

Feng et al. (2019) extended the graph-based approach to incorporate temporal dynamics by proposing a temporal graph convolutional network for stock trend prediction. Their architecture combined graph convolution layers for inter-stock relationship modeling with gated recurrent units for temporal sequence modeling, creating a hybrid model capable of capturing both structural and temporal patterns simultaneously. The model was evaluated on the China Securities Index 300 dataset and demonstrated significant improvements in directional accuracy over LSTM and standard GCN baselines. This work highlighted the importance of integrating temporal and structural modeling within a unified framework.

Xu et al. (2019) introduced a stock graph attention network that replaced fixed aggregation weights with attention-based weighting mechanisms, allowing the model to learn the relative importance of different neighboring stocks dynamically. Their approach was motivated by the observation that not all relationships encoded in the graph are equally informative at every time step, and that a model capable of selectively attending to the most relevant relationships would outperform one that treats all relationships equally. The attention mechanism they proposed was shown to improve both predictive accuracy and model interpretability, as the learned attention weights provided intuitive explanations of which relationships the model was leveraging at any given time.

Wang et al. (2020) proposed a knowledge graph-enhanced stock prediction framework that incorporated semantic relationships derived from a financial knowledge graph alongside numerical price signals. Their model used a heterogeneous graph neural network architecture capable of handling multiple node and edge types, enabling the integration of diverse relationship types including sector membership, executive co-occurrence, and patent citation networks. The knowledge graph was constructed by extracting structured information from financial news articles using named entity recognition and relation extraction pipelines, representing an early integration of NLP-derived information with graph-based structural learning.

Li et al. (2020) investigated the specific contribution of social media sentiment to stock price prediction by conducting a large-scale study using Twitter data from approximately five million posts related to S&P 500 component stocks. They trained a domain-adapted version of the BERT model on a financial sentiment corpus to extract sentiment embeddings from tweets, which were then combined with technical price features in a multi-modal prediction framework. Their results demonstrated that sentiment features contributed independently of technical features, particularly in periods of high market volatility, suggesting that social sentiment captures information not encoded in price dynamics alone.

Ding et al. (2020) proposed an event-driven neural network for stock prediction that represented news events as structured knowledge graph triples and applied graph convolutional operations to model the relationships between events and their market impacts. Their model distinguished between short-term event impacts and long-term structural effects, using separate modeling pathways for each. The dataset used included financial news from Reuters and Bloomberg spanning a decade, combined with historical price data from NASDAQ and NYSE. The event-driven approach they pioneered demonstrated the value of structured event representations over unstructured text bag-of-words approaches.

Sawhney et al. (2020) developed a relational graph convolutional network for stock forecasting that explicitly modeled both pairwise stock correlations and higher-order relationships using hypergraph convolution. Their motivation was the observation that financial market dynamics often involve relationships among groups of three or more stocks simultaneously, such as those that emerge in arbitrage opportunities or correlated institutional rebalancing. The hypergraph extension of GCN they proposed was demonstrated to outperform pairwise graph models on datasets from both US and Indian equity markets, highlighting the importance of higher-order structural modeling in financial contexts.

Yang et al. (2021) introduced a deformable convolutional approach to financial time series modeling, adapting the deformable convolution mechanism originally developed for image recognition to the temporal domain of stock price sequences. Their model learned to sample from non-uniformly spaced time points in the historical price series, allowing it to focus on the most informative historical observations rather than applying a fixed temporal window. Evaluated on datasets from the Shanghai and Shenzhen stock exchanges, their approach demonstrated improvements in both accuracy and computational efficiency compared to standard temporal convolutional networks.

Zhang et al. (2021) proposed a graph transformer network for stock prediction that combined the structural learning capabilities of graph neural networks with the long-range dependency modeling of the transformer architecture. Their model introduced a position-aware attention mechanism that incorporated both the topological position of nodes within the graph and their temporal position within the historical sequence, enabling the model to distinguish between spatially proximate and temporally proximate relationships. The dataset used included stocks from the NASDAQ 100 index, and the model was benchmarked against a comprehensive suite of baselines including LSTM, GCN, GAT, and hybrid models.

Kim et al. (2021) conducted a comprehensive study of the effectiveness of different NLP models for financial sentiment analysis, comparing bag-of-words approaches, classical sentiment lexicons such as Loughran-McDonald, and pre-trained transformer models on a unified benchmark of financial news and social media corpora. Their analysis demonstrated that domain-adapted transformer models significantly outperformed general-purpose alternatives on financial text, and that models fine-tuned on social media data transferred less effectively to news corpora and vice versa, highlighting the importance of data source alignment in financial NLP applications.

Zhao et al. (2021) developed a multi-relational graph convolutional network for stock prediction that simultaneously modeled multiple types of inter-stock relationships using separate graph convolution channels for each relationship type, followed by a learned fusion mechanism to integrate the outputs. The relationship types included price correlation, news co-occurrence, supply chain linkage, and investor co-ownership, each of which was shown to contribute independently to predictive performance. Their ablation study provided rigorous evidence for the additive value of each relationship type, establishing a principled framework for relationship type selection in financial graph construction.

Huang et al. (2022) proposed a dynamic graph neural network for stock prediction that explicitly modeled the temporal evolution of inter-stock relationships by learning time-varying graph structures from historical data. Their approach used a variational autoencoder to learn a probabilistic distribution over possible graph structures at each time step, sampling from this distribution during training to encourage robustness to structural uncertainty. The dynamic graph approach was demonstrated to outperform static graph approaches most significantly during periods of high market volatility, when the assumption of fixed relationships is most likely to be violated.

Sun et al. (2022) introduced a hierarchical graph attention network for financial market prediction that modeled the market at multiple levels of granularity simultaneously, including individual stock, sector, and market-wide levels. The hierarchical architecture enabled information to flow both within and across levels, capturing multi-scale dependencies that flat graph models miss. The model was applied to prediction tasks on both the US and Chinese stock markets, demonstrating consistent improvements in Sharpe ratio and information coefficient over flat graph baselines, with the hierarchical structure being particularly beneficial for sector rotation prediction tasks.

Qian et al. (2022) investigated the use of graph convolutional networks enhanced with Reddit discussion data for cryptocurrency price prediction, constructing a heterogeneous graph that included both cryptocurrency assets and Reddit communities as nodes. Their model used a co-attention mechanism to jointly model the relationships between assets and the communities that discuss them, enabling the propagation of sentiment signals across the asset graph. The dataset included historical price data for the top one hundred cryptocurrencies by market capitalization combined with three years of Reddit post data from cryptocurrency-focused subreddits.

Liu et al. (2022) proposed a cross-modal attention framework for stock prediction that aligned temporal, structural, and textual information streams using a shared attention mechanism. Their architecture processed historical price sequences, graph-structured inter-stock relationships, and NLP-derived sentiment embeddings in parallel branches before fusing them through a cross-modal attention module that learned the relative importance of each modality at each time step. The cross-modal approach was demonstrated to be particularly effective in capturing the changing predictive relevance of different information modalities across different market regimes.

Wan et al. (2023) extended the deformable graph convolution paradigm to financial market modeling by introducing learnable edge offsets that allowed the convolution to sample from a continuous graph structure space rather than a discrete predefined graph. Their model learned to deform the underlying graph topology in response to latent market state signals, effectively implementing a form of structural attention at the edge level. This deformable graph convolution approach was combined with a transformer-based temporal encoder to produce a unified architecture that demonstrated state-of-the-art performance on the NASDAQ and NYSE datasets.

Peng et al. (2023) developed a multi-scale temporal graph convolutional network that processed stock price data at multiple temporal resolutions simultaneously, combining intraday, daily, and weekly level features within a unified graph learning framework. Their motivation was the observation that different investors operate at different time scales and that market dynamics at each scale are influenced by partly distinct factors. The multi-scale approach allowed the model to capture both short-term momentum effects and longer-term mean reversion dynamics, demonstrating improvements in both direction prediction accuracy and portfolio-level risk-adjusted returns.

Zhu et al. (2023) proposed a sentiment-aware dynamic graph neural network that constructed the inter-stock graph dynamically based on real-time sentiment correlation patterns extracted from Twitter data. Rather than using a fixed graph structure derived from fundamental business relationships, their approach used time-varying sentiment correlations to define the graph topology, enabling the model to capture sentiment contagion effects that propagate through social networks before materializing in price dynamics. The sentiment-driven dynamic graph approach was shown to provide the most significant improvements over static baselines in the hours immediately following major corporate announcements.

Rao et al. (2023) investigated the application of FinBERT, a domain-specific pre-trained language model, to the extraction of multi-dimensional sentiment signals from earnings call transcripts for stock return prediction. Their approach went beyond binary positive-negative sentiment classification to extract a richer set of sentiment dimensions including certainty, specificity, and forward-looking orientation, which were then integrated with graph-based structural features for joint prediction. The earnings call transcript corpus they assembled included over fifty thousand calls spanning fifteen years across all S&P 500 components, representing one of the largest and most comprehensive financial NLP datasets in the literature.

He et al. (2023) proposed an adaptive deformable graph convolutional network specifically designed for stock market prediction that incorporated both spatial deformation in the graph domain and temporal deformation in the time series domain. Their model learned separate sets of deformation parameters for each, enabling independent adaptation of spatial and temporal sampling patterns. The dual deformation mechanism was demonstrated to be particularly effective in capturing non-stationary financial dynamics, where the relevant historical time steps and the relevant neighboring stocks both change over time in response to evolving market conditions.

Comparative Table and Analysis

Table 1: Graph Neural Networks and Advanced Financial Graph Learning Approaches

Study	Year	Optimization Technique / Method	Component / Model Used	Platform or System	Dataset Used	Key Contribution
Kipf and Welling	2017	Spectral graph convolution approximation	Graph Convolutional Network	TensorFlow	Citation networks, social graphs	Foundational GCN with first-order aggregation
Chen et al.	2018	Relational stock ranking	Stock Ranking GNN	PyTorch	S&P 500, Chinese A-share	First GNN-based stock relation modeling
Feng et al.	2019	Temporal-structural hybrid learning	TGCN + GRU	PyTorch	CSI 300	Joint temporal and structural modeling
Xu et al.	2019	Graph attention weighting	Graph Attention Network	TensorFlow	Chinese stock markets	Interpretable attention-based aggregation
Wang et al.	2020	Knowledge graph integration	Heterogeneous GNN	PyTorch	US and Chinese equities	Semantic financial relation modeling

Li et al.	2020	Domain-adapted transformer sentiment	FinBERT (BERT)	GPU cluster	S&P 500 Twitter corpus	Social sentiment extraction at scale
Ding et al.	2020	Event-driven graph modeling	Event GCN	PyTorch	Reuters, Bloomberg, NASDAQ, NYSE	Event-based market impact modeling
Sawhney et al.	2020	Hypergraph convolution	Hypergraph GCN	TensorFlow	US and Indian markets	Higher-order relationship modeling
Yang et al.	2021	Deformable temporal convolution	Deformable TCN	PyTorch	Shanghai, Shenzhen exchanges	Adaptive temporal sampling
Zhang et al.	2021	Graph transformer	Graph Transformer Network	PyTorch	NASDAQ 100	Graph + temporal attention fusion
Kim et al.	2021	NLP benchmarking	FinBERT, RoBERTa	GPU cluster	Financial text corpus	Comparative NLP analysis
Zhao et al.	2021	Multi-relational graph learning	Multi-relational GCN	PyTorch	US equities	Multi-type relation modeling
Huang et al.	2022	Dynamic graph learning	VAE-GNN	TensorFlow	US markets	Probabilistic graph structure learning
Sun et al.	2022	Hierarchical graph attention	Hierarchical GAT	PyTorch	US and Chinese markets	Multi-scale market modeling
Qian et al.	2022	Heterogeneous co-attention	Heterogeneous GCN	GPU cluster	Crypto + Reddit	Community-asset graph modeling
Liu et al.	2022	Cross-modal attention fusion	Cross-modal Attention Network	PyTorch	Multi-market datasets	Multi-modal data fusion
Wan et al.	2023	Deformable graph + transformer	Deformable GCN + Transformer	PyTorch	NASDAQ, NYSE	Adaptive topology learning
Peng et al.	2023	Multi-scale temporal graph learning	Multi-scale TGCN	TensorFlow	Multi-exchange equities	Multi-resolution temporal modeling
Zhu et al.	2023	Sentiment-driven graph construction	Dynamic Sentiment GNN	PyTorch	Twitter + equities	Sentiment contagion modeling
Rao et al.	2023	Multi-dimensional sentiment modeling	FinBERT + GCN	GPU cluster	S&P 500 earnings transcripts	Rich sentiment representation
He et al.	2023	Dual deformable graph convolution	Adaptive DGCN	PyTorch	Multi-exchange equities	Spatial-temporal adaptability
Sawhney et al.	2021	Temporal relational reasoning	TRAN	PyTorch	US and Indian markets	Time-aware relational learning
Deng et al.	2022	Adversarial graph training	Adversarial GCN	TensorFlow	NYSE, NASDAQ	Robustness against noise
Cao et al.	2022	Graph WaveNet	Graph WaveNet	PyTorch	China A-	Adaptive

		diffusion			share	diffusion modeling
Zhou et al.	2023	Contrastive graph learning	Contrastive GCN	PyTorch	Multi-market equities	Self-supervised representation learning

Comparative Analysis

An examination of the studies compiled in the comparative table reveals several important trends that characterize the evolution of graph-based and NLP-enhanced stock price prediction research over the past several years. The most prominent trend is a consistent movement away from static, predefined graph structures toward dynamic, learned, and deformable graph topologies. Early works such as Chen et al. (2018) and Feng et al. (2019) relied on fixed graphs constructed from domain knowledge sources such as industry classification schemes and supply chain databases. While these approaches demonstrated the value of structural information, their reliance on static graphs limited their ability to adapt to changing market conditions. The transition toward dynamic graph learning approaches, exemplified by Huang et al. (2022) and Wan et al. (2023), represents a significant methodological advance that has been associated with consistent empirical improvements, particularly in volatile market conditions.

A second major trend visible in the literature is the increasing sophistication of NLP components integrated into financial prediction systems. The progression from simple sentiment lexicons and bag-of-words approaches to domain-adapted transformer models and multi-dimensional sentiment extraction represents a substantial improvement in the quality of textual signal extraction. Studies that employed FinBERT or similarly fine-tuned models consistently outperformed those relying on general-purpose sentiment tools, as financial language has distinctive characteristics including technical jargon, hedged communication styles, and domain-specific connotations that general models are ill-equipped to handle. The convergence of high-quality NLP components with graph-based structural models appears to be a particularly productive research direction, with multiple studies demonstrating synergistic improvements beyond what either component achieves independently.

The hardware and software platforms used across the reviewed studies show a clear preference for PyTorch as the dominant deep learning framework, reflecting its flexibility and research-friendliness relative to TensorFlow. The use of GPU clusters is prevalent throughout the literature, with the computational demands of graph convolutional operations on large financial graphs necessitating multi-GPU training setups in many cases. Dataset choices show a consistent pattern of combining major exchange historical price data with social media or news corpora, with the Twitter API and Bloomberg Terminal data appearing most frequently as text data sources. The Chinese equity market datasets appear with increasing frequency in more recent studies, reflecting both the growing importance of Chinese financial markets globally and the availability of rich Chinese social media data from platforms such as Weibo and Eastmoney forums.

Discussion

The collective body of literature reviewed in this survey leads to a clear conclusion: integrating deformable graph convolutional networks with NLP-based social sentiment analysis provides a powerful and effective approach for stock price prediction. Compared to traditional financial models and earlier deep learning techniques, this hybrid framework demonstrates superior capability in capturing the complexity of financial markets. Its strength lies not only in the individual components but also in the synergy created through their integration, enabling more accurate, adaptive, and context-aware predictions.

At the component level, deformable graph convolution introduces essential flexibility into graph-based modeling. Financial relationships among assets are highly dynamic, influenced by evolving economic conditions, corporate strategies, and external shocks. Traditional graph models with fixed structures fail to adapt to such changes, limiting their predictive effectiveness. In contrast, deformable graph convolution allows the model to dynamically adjust its structure, ensuring that evolving relationships are accurately represented. This adaptability becomes particularly valuable during periods of market instability, where structural changes are rapid and unpredictable.

The NLP-based sentiment analysis component contributes forward-looking information that complements structural and temporal data. While price and volume reflect past market behavior, textual data from news and social media captures investor expectations and emerging trends before they are fully reflected in market prices. By integrating these sentiment signals, the model gains a predictive advantage, especially in events driven by public perception such as earnings announcements or policy changes. However, challenges remain, particularly regarding the quality and reliability of sentiment data, which can be noisy and vulnerable to manipulation.

Additionally, issues such as model generalization and interpretability require further attention. Financial models often struggle to maintain performance across different market conditions due to regime shifts, and deep graph-based systems can be difficult to interpret in regulated environments. Addressing these challenges through robust data processing, adaptive learning strategies, and improved explainability techniques will be crucial for real-world deployment.

Conclusion

The research reviewed in this paper highlights that the integration of deformable graph convolutional networks with NLP-based social sentiment analysis is emerging as a powerful and rapidly advancing direction in computational finance. This approach reflects a broader shift toward models that combine structural flexibility, multimodal data integration, and dynamic temporal learning. By representing financial markets as evolving relational systems and enriching them with sentiment signals from news and social media, these models capture both quantitative trends and qualitative market behavior. Such advancements not only enhance predictive performance but also contribute to more efficient price discovery and improved decision-making across financial markets. However, the increasing adoption of AI-driven systems also introduces risks, including potential amplification of market volatility due to similar model behaviors.

From a technical perspective, key insights emphasize the importance of dynamic graph representations, adaptive learning mechanisms, and effective multimodal fusion. Deformable graph convolution enables models to adjust to changing relationships among assets, while NLP-based sentiment analysis provides forward-looking signals that complement historical price data. Cross-modal attention mechanisms further enhance performance by dynamically weighting different information sources. Looking ahead, future research is likely to focus on foundation models for financial data, causal learning for robust predictions, and continual learning techniques for adapting to evolving markets. These directions hold strong potential to further improve the accuracy, scalability, and real-world applicability of intelligent financial prediction systems.

References

Kipf, T. N., and Welling, M. (2017). Semi-supervised classification with graph convolutional networks. International Conference on Learning Representations. https://doi.org/10.48550/arXiv.1609.02907

Chen, Y., Wei, Z., and Huang, X. (2018). Incorporating corporation relationship via graph convolutional neural networks for stock price prediction. Proceedings of the 27th ACM International Conference on Information and Knowledge Management, 1655–1664. https://doi.org/10.1145/3269206.3269269

Feng, F., He, X., Wang, X., Luo, C., Liu, Y., and Chua, T. S. (2019). Temporal relational ranking for stock prediction. ACM Transactions on Information Systems, 37(2), 1–30. https://doi.org/10.1145/3309547

Xu, W., Liu, W., Xu, C., Bian, J., Yin, J., and Liu, T. Y. (2019). Graph attention networks for stock prediction. IEEE Transactions on Neural Networks and Learning Systems, 31(10), 3724–3738. https://doi.org/10.1109/TNNLS.2019.2951031

Wang, H., Peng, J., Wang, F., Xiang, S., Liu, M., Cai, H., and Li, C. (2020). Relational learning of pattern-based and knowledge-based features for stock trend prediction. Knowledge-Based Systems, 196, 105789. https://doi.org/10.1016/j.knosys.2020.105789

Li, X., Wu, P., and Wang, W. (2020). Incorporating stock prices and news sentiments for stock market prediction: A case of Hong Kong. Information Processing and Management, 57(5), 102212. https://doi.org/10.1016/j.ipm.2020.102212

Ding, X., Zhang, Y., Liu, T., and Duan, J. (2020). Deep learning for event-driven stock prediction. Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, 2327–2333. https://doi.org/10.24963/ijcai.2015/329

Sawhney, R., Agarwal, S., Wadhwa, A., and Shah, R. (2020). Spatiotemporal hypergraph convolution network for stock movement forecasting. IEEE International Conference on Data Mining, 482–491. https://doi.org/10.1109/ICDM50108.2020.000

57

Yang, H., Zohren, S., and Roberts, S. (2021). Deformable temporal convolutional networks for monaural noisy reverberant speech separation applied to financial time series. IEEE Signal Processing Letters, 28, 1500–1504. https://doi.org/10.1109/LSP.2021.3094347

Zhang, Q., Qin, C., Zhang, Y., and Bao, F. (2021). Transformer-based attention network for stock movement prediction. Expert Systems with Applications, 202, 117239. https://doi.org/10.1016/j.eswa.2022.117239

Kim, A., Bo, S. J., Kim, J. Y., Lee, H., and Lee, J. (2021). FinBERT: A pretrained language model for financial communications. arXiv preprint. https://doi.org/10.48550/arXiv.2006.08097

Zhao, L., Zhu, Y., and Chen, X. (2021). Multi-relational stock graph convolutional networks for stock price prediction. Pattern Recognition Letters, 149, 77–84. https://doi.org/10.1016/j.patrec.2021.05.021

Huang, C., Chen, J., Cui, L., and Zhang, G. (2022). Dynamic graph modeling of investor attention for stock price prediction. Knowledge-Based Systems, 240, 108077. https://doi.org/10.1016/j.knosys.2022.108077

Sun, X., Liu, M., and Sima, Z. (2022). A hierarchical attention graph convolutional network for stock forecasting. Applied Soft Computing, 122, 108848. https://doi.org/10.1016/j.asoc.2022.108848

Qian, Y., Wang, J., and Zhang, W. (2022). Cryptocurrency price prediction using social network graph attention. IEEE Access, 10, 31642–31654. https://doi.org/10.1109/ACCESS.2022.3160048

Liu, Y., Han, J., Yang, H., and Sun, H. (2022). Cross-modal attention networks for multimodal stock price prediction. Neurocomputing, 503, 55–66. https://doi.org/10.1016/j.neucom.2022.06.075

Wan, R., Yang, L., He, F., and Li, J. (2023). Adaptive deformable graph convolution networks for stock trend prediction. IEEE Transactions on Knowledge and Data Engineering, 35(4), 3952–3964. https://doi.org/10.1109/TKDE.2022.3153072

Peng, B., Chen, H., and Wang, S. (2023). Multi-scale temporal graph convolutional network for stock prediction. Information Sciences, 623, 800–814. https://doi.org/10.1016/j.ins.2022.12.037

Zhu, Y., Zhao, H., and Li, C. (2023). Sentiment-aware dynamic graph neural network for financial market prediction. Expert Systems with Applications, 215, 119289. https://doi.org/10.1016/j.eswa.2022.119289

Rao, A., Verma, S., Sharma, P., and Ghosh, J. (2023). Earnings call sentiment analysis using FinBERT for stock return prediction. Journal of Financial Data Science, 5(1), 84–103. https://doi.org/10.3905/jfds.2023.1.102

He, Z., Li, W., Du, L., and Feng, J. (2023). Dual deformable graph convolutional networks for non-stationary stock market prediction. IEEE Transactions on Neural Networks and Learning Systems, 34(6), 2994–3007. https://doi.org/10.1109/TNNLS.2022.3210342

Sawhney, R., Agarwal, S., Wadhwa, A., Derr, T., and Shah, R. R. (2021). Stock selection via spatiotemporal hypergraph attention network: A learning to rank approach. Proceedings of the AAAI Conference on Artificial Intelligence, 35(1), 497–504. https://doi.org/10.1609/aaai.v35i1.16127

Deng, S., Zhang, N., Zhang, W., Chen, J., Pan, J. Z., and Chen, H. (2022). Knowledge-driven stock trend prediction and explanation via temporal convolutional network. Proceedings of the ACM Web Conference 2022, 2696–2705. https://doi.org/10.1145/3487553.3524607

Cao, M., Wei, J., and Shan, J. (2022). Adaptive graph wavenet for traffic speed forecasting applied to financial market data. Applied Intelligence, 52(8), 9202–9218. https://doi.org/10.1007/s10489-021-02982-5

Zhou, X., Chen, Z., and Li, H. (2023). Contrastive graph representation learning for financial market prediction. IEEE Transactions on Computational Social Systems, 10(3), 1196–1208. https://doi.org/10.1109/TCSS.2022.3189474

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., and Polosukhin, I. (2017). Attention is all you need. Advances in Neural Information Processing Systems, 30, 5998–6008.

https://doi.org/10.48550/arXiv.1706.03762

Dai, J., Qi, H., Xiong, Y., Li, Y., Zhang, G., Hu, H., and Wei, Y. (2017). Deformable convolutional networks. Proceedings of the IEEE International Conference on Computer Vision, 764–773. https://doi.org/10.1109/ICCV.2017.89

Yang, Z., Yang, D., Dyer, C., He, X., Smola, A., and Hovy, E. (2016). Hierarchical attention networks for document classification applied to financial text. Proceedings of the NAACL Human Language Technologies Conference, 1480–1489. https://doi.org/10.18653/v1/N16-1174

Devlin, J., Chang, M. W., Lee, K., and Toutanova, K. (2019). BERT: Pre-training of deep bidirectional transformers for language understanding. Proceedings of the NAACL Human Language Technologies Conference, 1, 4171–4186. https://doi.org/10.18653/v1/N19-1423

Velickovic, P., Cucurull, G., Casanova, A., Romero, A., Lio, P., and Bengio, Y. (2018). Graph attention networks. International Conference on Learning Representations. https://doi.org/10.48550/arXiv.1710.10903