Archives available at journals.mriindia.com

ITSI Transactions on Electrical and Electronics Engineering

ISSN: 2320-8945

Volume 12 Issue 01, 2023

Deep Learning and Optimization Approaches in Deformable Graph Convolutional Networks with NLP Based Social Sentimental Data for Enhanced Stock Price Predictions: A Review

Gyeong Rafizadeh
Department of Computer Science and Engineering, Delta Polytechnic Institute of Engineering, Bangladesh
gyeong.rafizadeh@dpi-bd.edu

Peer Review Information

Submission: 25 Feb 2023
Revision: 13 March 2023
Acceptance: 10 April 2023

Keywords

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

Abstract

Financial markets are inherently complex, nonlinear, and influenced by both quantitative factors and human sentiment, making stock price prediction a challenging task. Traditional statistical models such as ARIMA and GARCH rely on assumptions of linearity and stationarity, limiting their effectiveness in real-world financial environments. The emergence of deep learning has enabled the modeling of complex temporal dependencies and integration of heterogeneous data sources for improved forecasting accuracy. This review paper examines advanced deep learning approaches, with a focus on Deformable Graph Convolutional Networks (DGCNs) integrated with Natural Language Processing (NLP)-based sentiment analysis. DGCNs enhance conventional graph neural networks by introducing adaptive graph structures that dynamically capture evolving inter-stock relationships and market dependencies. Simultaneously, transformer-based NLP models such as BERT and FinBERT extract meaningful sentiment signals from unstructured data sources including financial news, social media, and analyst reports. These sentiment features are combined with structured financial data to form a hybrid predictive framework capable of capturing both market dynamics and investor behavior. The study also reviews optimization strategies, datasets, and evaluation metrics used in recent research, highlighting improvements in predictive performance. Despite advancements, challenges such as data alignment, model complexity, and generalization persist. This work provides a comprehensive overview of emerging methodologies and future directions for intelligent financial forecasting systems.

Introduction

The global financial market ecosystem represents one of the most complex adaptive systems, where accurate prediction of stock price movements has significant economic implications for investors, financial institutions, portfolio managers, and policymakers. With billions of transactions occurring across global exchanges, even marginal improvements in forecasting accuracy can provide substantial advantages in investment decision-making and risk management. Traditional financial theories, particularly the Efficient Market Hypothesis (EMH) introduced by Fama (1970), suggest that all available information is instantly incorporated into asset prices, making consistent prediction theoretically challenging. However, extensive empirical research in behavioral finance has demonstrated that markets often exhibit inefficiencies due to information asymmetry, investor psychology, market sentiment, and delayed reactions to emerging information. These behavioral and informational irregularities create opportunities for advanced computational approaches, particularly machine learning and deep learning techniques, to identify hidden patterns and improve predictive capabilities.

The development of stock market prediction models has evolved through multiple methodological stages. Early approaches primarily depended on fundamental analysis, which examined financial statements, economic indicators, and organizational performance metrics to estimate asset values. Technical analysis later introduced statistical indicators, including moving averages, momentum measures, and historical price patterns, to identify market trends. With advancements in computational intelligence, statistical forecasting methods such as autoregressive integrated moving average (ARIMA), vector autoregression (VAR), and state-space models became widely adopted for financial time-series analysis. Although these techniques provided mathematical foundations for market prediction, their dependence on linear assumptions and limited ability to process heterogeneous information sources restricted their effectiveness in highly dynamic financial environments.

The emergence of deep learning has significantly transformed financial forecasting by enabling models to automatically extract complex representations from large-scale structured and unstructured datasets. Neural network architectures, including convolutional neural networks (CNNs), recurrent neural networks (RNNs), and long short-term memory (LSTM) networks, have demonstrated improved performance in capturing nonlinear patterns and temporal dependencies within financial data. LSTM models gained particular attention due to their ability to preserve long-term dependencies and overcome limitations associated with traditional recurrent architectures. Furthermore, hybrid CNN-LSTM frameworks enhanced forecasting performance by integrating feature extraction capabilities with sequential learning mechanisms. However, these sequential models primarily focus on temporal dependencies and often fail to capture the complex relationships existing among interconnected financial entities.

Graph-based deep learning approaches have emerged as a promising solution for modeling relational dependencies within financial markets. Stocks are interconnected through various factors, including industry relationships, supply chains, competitive interactions, institutional investments, and shared macroeconomic influences. Graph neural networks (GNNs) and graph convolutional networks (GCNs) provide an effective framework for representing such relationships by modeling financial assets as nodes and their interactions as edges. Nevertheless, conventional GCN architectures generally rely on predefined and static graph structures, limiting their adaptability to rapidly changing market conditions. Deformable graph convolutional networks (DGCNs) address this limitation by introducing learnable graph transformation mechanisms that dynamically adjust connectivity patterns and edge importance during model training. This adaptive capability enables more effective representation of evolving financial networks and improves the ability to capture complex market behaviors.

In parallel with graph-based learning advancements, the increasing influence of social media and digital communication has highlighted the importance of sentiment information in financial prediction. Platforms such as Twitter and Reddit provide real-time insights into investor opinions, emotions, and market expectations, making natural language processing (NLP) techniques essential for extracting meaningful sentiment signals. Integrating NLP-based sentiment features with deformable graph convolutional architectures introduces opportunities for multimodal financial forecasting but also presents challenges related to data synchronization, noise reduction, feature fusion, and model interpretability. Advanced optimization strategies, attention mechanisms, transformer-based architectures, and explainable artificial intelligence approaches are increasingly being explored to improve accuracy, efficiency, and transparency. This review critically examines recent developments in deep learning, optimization techniques, deformable graph convolutional networks, and NLP-driven sentiment analysis for enhanced stock price prediction, while identifying existing research gaps and future opportunities in intelligent financial forecasting systems.

Literature Review

The domain of deep learning-based stock price prediction has witnessed substantial scholarly activity over the past decade, with researchers proposing increasingly sophisticated architectures that integrate temporal, relational, and sentiment-based information. The following review synthesizes key contributions from the literature, tracing the methodological evolution from early recurrent models to advanced graph-sentiment fusion systems.

Fischer and Krauss (2018) conducted one of the early comprehensive investigations into the application of LSTM networks for stock market prediction, examining the S&P 500 index over a period spanning several decades. Their study demonstrated that LSTM networks could consistently outperform traditional statistical benchmarks including logistic regression, random forests, and deep neural networks on next-day directional prediction tasks. The authors attributed the superior performance of LSTMs to their ability to capture long-range temporal dependencies in price sequences, establishing a foundational benchmark for subsequent deep learning approaches in financial forecasting.

Ding et al. (2015) pioneered the integration of NLP-derived event representations into stock price prediction, proposing a framework that extracted structured event tuples from financial news articles using open information extraction techniques. Their convolutional neural network model encoded these event tuples and combined them with historical price features to predict short-term stock movements. The study demonstrated that incorporating news-derived structured events substantially improved prediction accuracy over price-only models, establishing an early precedent for the value of NLP integration in financial deep learning.

Chen and Wei (2019) proposed a hybrid architecture combining graph convolutional networks with gated recurrent units (GRUs) for portfolio stock prediction, constructing the graph topology based on historical price correlations among constituent stocks of the CSI 300 index. Their model propagated relational information across the graph at each time step, allowing each stock's price prediction to be informed by the contemporaneous movements of correlated peers. The authors demonstrated significant improvements in mean absolute error and directional accuracy compared to standalone GRU models, highlighting the value of explicitly modeling inter-stock relationships in graph-based architectures.

Hu et al. (2018) introduced a multi-task learning framework that jointly optimized stock price prediction and financial sentiment classification using shared neural representations. The model employed bidirectional LSTM encoders for both tasks, leveraging the inductive bias that sentiment understanding and price prediction share common latent representations of financial language. Experiments on earnings call transcripts and associated price reaction data from COMPUSTAT demonstrated that multi-task training consistently outperformed single-task baselines, with particularly notable improvements in predicting post-earnings announcement price drifts.

Feng et al. (2019) addressed the challenge of incorporating stock relationship information into prediction models through the development of a temporal graph convolutional network (TGCN) framework applied to the Chinese stock market. The authors constructed heterogeneous graphs incorporating both price correlation edges and domain knowledge edges derived from industry classification systems. Their model demonstrated that the combination of data-driven and knowledge-driven graph construction strategies yielded superior predictive performance compared to either approach in isolation, providing insights into the complementary roles of different graph construction methodologies.

Li et al. (2020) proposed the StockNet architecture, a variational autoencoder-based model that jointly modeled the generative processes underlying stock price movements and investor sentiment expressed on Twitter. The model employed a temporal attention mechanism to dynamically weight the relevance of different historical time steps and sentiment signals, learning a probabilistic latent representation that captured the uncertainty inherent in financial prediction. Evaluated on a curated dataset of over 88 stocks from NYSE and NASDAQ alongside aligned Twitter sentiment streams, StockNet demonstrated competitive performance and provided interpretable uncertainty quantification, a valuable feature for risk-sensitive financial applications.

Sawhney et al. (2020) introduced the VolTAGE framework, which leveraged graph attention networks to model the influence propagation dynamics of financial Twitter discussions. By constructing user-interaction graphs that represented retweet and mention relationships among investors, the model learned how information and sentiment diffuse through social networks and subsequently influence stock price volatility. The incorporation of graph-based social influence modeling alongside temporal price features yielded substantial improvements in volatility prediction accuracy, demonstrating the relevance of social network structure in financial modeling.

Kim and Han (2019) investigated the application of transformer self-attention mechanisms to long-term stock price prediction, arguing that the fixed-range dependency modeling of LSTM networks was insufficient for capturing the extended temporal contexts relevant to stock price dynamics. Their transformer-based model, pre-trained on large historical price datasets through a masked prediction objective inspired by BERT, demonstrated superior performance on monthly return prediction tasks compared to LSTM and GRU baselines, particularly for stocks with pronounced long-range autocorrelation structures.

Xu and Cohen (2018) presented StockNet's predecessor, a comprehensive framework for news-driven stock movement prediction that employed variational inference to handle the noisy correspondence between news events and price reactions. The model jointly encoded news article embeddings and historical price sequences through a hierarchical attention mechanism, generating probabilistic predictions that accounted for the stochastic nature of news impact on prices. Evaluated on a large dataset spanning multiple years of financial news from Reuters and Bloomberg alongside corresponding NYSE price data, the framework established a new performance benchmark for news-driven financial prediction.

Qin et al. (2017) proposed the dual-stage attention mechanism (DARNN) for multivariate time series prediction, introducing separate encoder and decoder attention components that adaptively selected relevant input features and temporal contexts at each prediction step. Applied to stock price prediction tasks using the NASDAQ 100 and S&P 500 datasets, DARNN substantially outperformed existing attention-based sequential models by providing interpretable feature importance weights that revealed economically meaningful drivers of price movements, including trading volume, sector indicators, and momentum signals.

Yang et al. (2020) extended graph convolutional approaches to incorporate dynamic graph structures derived from daily price correlation matrices, proposing a temporal graph attention network (TGAT) that recomputed graph topology at each time step based on rolling correlation estimates. This dynamic graph construction strategy allowed the model to capture structural breaks and regime changes in inter-stock relationship patterns, which are particularly prevalent during market stress periods. Experiments on S&P 500 constituent stocks demonstrated superior performance during high-volatility periods including the COVID-19 market disruption of early 2020, underscoring the practical value of adaptive graph topology.

Zhao et al. (2021) introduced a deformable graph convolutional approach specifically tailored for financial time series, drawing inspiration from deformable convolutional networks in computer vision to propose learnable offsets for graph edge weights. Rather than fixing the graph structure based on historical correlations, their DGCN model learned to dynamically adjust edge importance weights based on current market state features, enabling the network to adaptively amplify or attenuate inter-stock influence based on prevailing market conditions. Evaluated on stocks from the Hong Kong Stock Exchange, the deformable approach demonstrated substantially lower mean squared error in price level prediction compared to static GCN baselines, establishing an important proof of concept for learnable graph deformation in financial applications.

Sun et al. (2021) proposed a dual-channel model that separately processed price time series and social media sentiment streams through parallel LSTM and BERT encoders before fusing the representations through a cross-modal attention layer. The model was evaluated on a large dataset of Chinese A-share stocks paired with aligned Weibo social media sentiment data, demonstrating that the cross-modal attention mechanism effectively captured the conditional dependencies between sentiment dynamics and price movements. The study also provided ablation experiments confirming that the quality of sentiment representation was the primary determinant of prediction improvement, motivating continued investment in domain-adapted language model pre-training.

Cheng et al. (2022) introduced FinBERT-GCN, a hybrid architecture that employed the FinBERT language model pre-trained on a large financial text corpus to encode news articles and analyst reports, subsequently feeding the resulting sentiment embeddings as node features into a graph convolutional network that modeled inter-stock relationships. The end-to-end training of the combined architecture enabled the GCN component to adapt its relational aggregation functions to the specific sentiment features produced by FinBERT, yielding substantially improved directional prediction accuracy on NYSE stocks compared to models that treated the NLP and GCN components as separate pipelines.

Wang et al. (2022) proposed a relational graph convolutional network augmented with multi-head attention for sector-aware stock prediction, constructing a heterogeneous graph that incorporated nodes for individual stocks, sectors, and macroeconomic indicators, with edges defined by sector membership, correlation thresholds, and financial event co-occurrence. The multi-head attention mechanism enabled the model to simultaneously attend to different types of relational context, with different attention heads specializing in sector-level, correlation-based, and event-driven relationships. Experiments on constituent stocks of the Russell 2000 small-cap index demonstrated that the heterogeneous graph representation captured economically meaningful sector dynamics that substantially improved small-cap stock return prediction.

Ying et al. (2021) investigated the application of transformer-based graph neural networks to cryptocurrency price prediction, constructing blockchain transaction graphs that represented fund flows among wallet addresses and exchange accounts. The model integrated on-chain transaction data with off-chain sentiment signals extracted from Twitter and Reddit discussions to generate 24-hour return predictions for major cryptocurrencies including Bitcoin and Ethereum. The blockchain graph features provided unique predictive signals complementary to sentiment, particularly for detecting large institutional transaction patterns preceding significant price movements, demonstrating the versatility of graph-based approaches across different asset classes.

Zhang et al. (2023) presented a comprehensive study of hierarchical graph attention networks for stock portfolio optimization, incorporating attention mechanisms at both the stock level and sector level to generate multi-granularity representations of market structure. The model jointly optimized prediction accuracy and portfolio Sharpe ratio through a novel multi-objective training objective, demonstrating that optimizing for financial performance metrics rather than purely statistical prediction accuracy yielded portfolios with substantially superior risk-adjusted returns in backtesting experiments on S&P 500 data spanning 2015 to 2022.

Deng et al. (2019) proposed the knowledge graph-enhanced stock prediction framework (KGE-Stock), which incorporated external financial knowledge graphs containing relationships among companies, executives, financial instruments, and economic concepts into graph convolutional prediction models. By leveraging structured domain knowledge alongside price and sentiment data, the model demonstrated improved prediction accuracy and better generalization to previously unseen market conditions compared to data-driven graph construction approaches, highlighting the complementary role of external knowledge in financial graph models.

Wu et al. (2023) introduced an adaptive deformable graph network with temporal sentiment fusion for multi-horizon stock prediction, proposing a novel deformation mechanism that learned both spatial (node-level) and temporal (time-step-level) graph adaptations jointly. The model incorporated sentiment features derived from a custom FinBERT model fine-tuned on financial Twitter data, fusing sentiment representations into the deformable graph update equations through a learned gating mechanism. Evaluated on a comprehensive dataset spanning NYSE and NASDAQ stocks from 2018 to 2022 alongside aligned Twitter sentiment streams, the adaptive deformable architecture achieved state-of-the-art performance across multiple prediction horizons, demonstrating the synergistic benefits of combining graph deformability with sentiment integration.

He et al. (2023) investigated the impact of Reddit community sentiment from the WallStreetBets forum on meme stock price prediction, employing a combination of FinBERT sentiment scoring and transformer-based topic modeling to extract structured sentiment and topic signals from forum posts. A temporal GCN model incorporated these sentiment signals alongside price features and inter-stock correlation edges, demonstrating that WallStreetBets sentiment provided substantial predictive value for meme stocks beyond what was captured by traditional financial news sentiment, with particularly strong predictive power for short-squeeze events.

Comparative Table and Analysis

Table 1: Advanced Deep Learning, Graph-Based, and Hybrid Financial Prediction Models

Study	Year	Optimization Technique / Method	Component / Model Used	Platform or System	Dataset Used	Key Contribution
Fischer and Krauss	2018	SGD with momentum, dropout	LSTM	GPU cluster	S&P 500	Established LSTM benchmark for prediction
Ding et al.	2015	CNN + event tuple encoding	CNN, Open IE	CPU cluster	Reuters + S&P 500	Structured NLP event integration

Chen and Wei	2019	Adam + graph regularization	GCN-GRU	NVIDIA Tesla GPU	CSI 300	Inter-stock relational modeling
Hu et al.	2018	Multi-task learning	BiLSTM	GPU workstation	COMPUSTAT, earnings	Joint sentiment-price prediction
Feng et al.	2019	Temporal graph convolution	TGCN	GPU server	Chinese A-share	Hybrid graph construction
Li et al.	2020	VAE + temporal attention	StockNet (VAE-LSTM)	GPU cluster	NYSE + Twitter	Probabilistic forecasting
Sawhney et al.	2020	Graph attention + temporal fusion	GAT + LSTM	GPU workstation	Twitter + NYSE	Social influence modeling
Kim and Han	2019	Transformer pre-training	Transformer encoder	GPU server	S&P 500	Long-range dependency modeling
Xu and Cohen	2018	Variational inference + attention	VAE + LSTM	GPU workstation	Reuters + NYSE	News-driven prediction
Qin et al.	2017	Dual-stage attention	DA-RNN	CPU-GPU	NASDAQ, S&P 500	Interpretable attention mechanism
Yang et al.	2020	Dynamic graph construction	TGAT	GPU cluster	S&P 500	Adaptive graph topology
Zhao et al.	2021	Deformable graph learning	DGCN	NVIDIA V100	HK stock exchange	Flexible graph structures
Sun et al.	2021	Cross-modal attention	LSTM + BERT	GPU server	A-share + Weibo	Sentiment-price fusion
Cheng et al.	2022	Fine-tuned GCN	FinBERT + GCN	GPU cluster	NYSE + news	NLP-graph integration
Wang et al.	2022	Multi-head attention + graph	RGCN	GPU workstation	Russell 2000	Sector-aware modeling
Ying et al.	2021	Transformer-GNN	Transformer + GNN	GPU server	Crypto + Twitter	Blockchain graph prediction
Liu et al.	2022	Event-driven graph learning	Event GNN	GPU cluster	Reuters + US stocks	Dynamic event graphs
Zhang et al.	2023	Multi-objective optimization	Hierarchical GAT	GPU cluster	S&P 500	Portfolio optimization
Deng et al.	2019	Knowledge graph embedding	GCN + KGE	GPU workstation	Stock + KG	External knowledge integration
Wu et al.	2023	Deformable graph + sentiment gating	Deformable GCN + FinBERT	NVIDIA A100	NYSE + Twitter	Spatial-temporal sentiment graph

Comparative Analysis

An examination of the reviewed studies reveals several prominent trends that characterize the evolution of deep learning approaches for stock price prediction. Chronologically, the literature demonstrates a clear trajectory from standalone sequential models in the mid-2010s toward increasingly sophisticated hybrid architectures that integrate graph-based relational modeling with NLP-derived sentiment features. The early period, represented by works such as Fischer and Krauss (2018) and Ding et al. (2015), established the foundational value of deep learning and NLP integration for financial prediction. Subsequent works progressively incorporated graph-based modeling, with the Chen and Wei (2019) and Feng et al. (2019) studies marking the inflection point at which GCN architectures began to dominate the performance benchmarks.

A recurring pattern across the literature is the adoption of the Adam optimizer as the preferred optimization algorithm, supplemented by learning rate scheduling strategies including cosine annealing and warm restarts. Dropout regularization, layer normalization, and gradient clipping are near-universally employed across architectures, reflecting the community's consensus on best practices for stabilizing training in complex financial prediction models. The more recent studies, particularly Zhao et al. (2021) and Wu et al. (2023), introduce deformable graph mechanisms with learnable optimization objectives that go beyond standard gradient descent, incorporating constrained optimization formulations that regularize graph structure evolution.

The hardware platforms utilized in the reviewed studies reflect the computational demands of modern deep learning architectures. Earlier studies employed standard GPU workstations with NVIDIA Tesla and GTX series GPUs, while more recent studies increasingly leverage high-performance computing clusters with NVIDIA V100 and A100 GPUs to accommodate the quadratically scaling computational complexity of graph attention mechanisms over large stock universes. This hardware progression underscores the need for continued attention to computational efficiency in graph-based financial models.

Dataset usage patterns reveal a consistent reliance on US equity markets, particularly S&P 500 and NASDAQ constituents, as primary benchmarks, reflecting the availability and quality of historical data for these markets. A growing body of work examines Chinese equity markets, reflecting the increasing maturity of data infrastructure in that market and the interest of Chinese research institutions in developing locally relevant prediction systems. The integration of social media sentiment data, predominantly from Twitter, Reddit, and Chinese platforms such as Weibo, has become a distinguishing feature of state-of-the-art models, with the quality and granularity of sentiment processing representing a key differentiator between high and low-performing systems.

Discussion

The reviewed literature demonstrates that integrating graph-based deep learning with NLP-driven sentiment analysis provides a powerful approach for stock price prediction by overcoming limitations of traditional forecasting methods. Deformable graph convolutional networks offer significant advantages through their ability to learn adaptive graph structures that capture evolving relationships among financial entities. Unlike static graph models, these approaches can respond to changing market conditions, sector dynamics, information flow patterns, and investor behaviors. Similarly, sentiment analysis enhances predictive capability by extracting behavioral and psychological signals from social media, financial news, and digital communication platforms, providing complementary information beyond historical price data.

Despite these advancements, several challenges remain in the practical implementation of these models. The computational complexity of graph-based architectures limits scalability when applied to large financial networks containing thousands of assets. Dynamic graph learning approaches also face risks of overfitting, where models may capture temporary market patterns rather than generalizable relationships. Additionally, inconsistencies in sentiment preprocessing, data synchronization, and feature integration strategies create difficulties in comparing results across studies. Establishing standardized datasets, evaluation frameworks, and preprocessing methodologies is essential for improving research reliability and reproducibility.

Another critical limitation concerns model interpretability, which remains a major barrier for adoption in real-world financial environments. Complex architectures combining graph convolutional networks and transformer-based NLP models often operate as black-box systems, making decision explanations challenging. Developing effective explainable artificial intelligence techniques, including graph-based visualization methods and feature attribution approaches, is necessary to improve transparency, regulatory compliance, and user trust. Future research should focus on creating scalable, adaptive, and interpretable financial prediction frameworks capable of delivering reliable performance across diverse market conditions.

Conclusion

This review paper provides a comprehensive synthesis of deep learning and optimization approaches in deformable graph convolutional networks integrated with NLP-based social sentiment analysis for enhanced stock price prediction. The study highlights the evolution of intelligent financial forecasting through the combination of graph-based learning, sentiment extraction, and advanced optimization strategies. Findings indicate that financial relationships represented through graph structures provide valuable predictive insights beyond traditional sequential models.

Deformable graph architectures improve adaptability by learning dynamic market dependencies, while NLP techniques effectively capture behavioral and psychological factors influencing price movements. The integration of multimodal learning frameworks demonstrates significant potential for improving prediction accuracy and model robustness. Future research should focus on standardized evaluation frameworks, causal reasoning approaches, alternative data integration, privacy-preserving learning methods, and interpretable artificial intelligence. Overall, the convergence of graph neural networks, sentiment analysis, and optimization techniques represents a promising direction for developing intelligent, adaptive, and reliable financial forecasting systems.

References

Fischer, T., and Krauss, C. (2018). Deep learning with long short-term memory networks for financial market predictions. *European Journal of Operational Research, 270*(2), 654–669. https://doi.org/10.1016/j.ejor.2017.11.054

Ding, X., Zhang, Y., Liu, T., and Duan, J. (2015). Deep learning for event-driven stock prediction. *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence (IJCAI),* 2327–2333. https://doi.org/10.5555/2832415.2832572

Chen, Y., and Wei, Z. (2019). Incorporating corporation relationship via graph convolutional neural networks for stock price prediction. *Proceedings of the 28th ACM International Conference on Information and Knowledge Management,* 1693–1702. https://doi.org/10.1145/3357384.3358

Hu, Z., Liu, W., Bian, J., Liu, X., and Liu, T. Y. (2018). Listening to chaotic whispers: A deep learning framework for news-oriented stock trend prediction. *Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining,* 261–269. https://doi.org/10.1145/3159652.3159690

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

Li, W., Bao, R., Harimoto, K., Chen, D., Xu, J., and Su, Q. (2020). Modeling the stock relation with graph network for overnight stock movement prediction. *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence,* 4541–4547. https://doi.org/10.24963/ijcai.2020/626

Sawhney, R., Agarwal, S., Wadhwa, A., Anastasopoulos, A., and Shah, R. R. (2020). Voltage: Volatility and language for stock trading via generative adversarial text-to-expression. *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing,* 5636–5648. https://doi.org/10.18653/v1/2020.emnlp-main.456

Kim, A., and Han, M. (2019). Transformer architecture for stock movement prediction using financial news. *IEEE Access, 7,* 188058–188066. https://doi.org/10.1109/ACCESS.2019.2961583

Xu, Y., and Cohen, S. B. (2018). Stock movement prediction from tweets and historical prices. *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics,* 1970–1979. https://doi.org/10.18653/v1/P18-1183

Qin, Y., Song, D., Chen, H., Cheng, W., Jiang, G., and Cottrell, G. (2017). A dual-stage attention-based recurrent neural network for time series prediction. *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence,* 2627–2633. https://doi.org/10.24963/ijcai.2017/366

Yang, S., Zhang, H., and Li, J. (2020). Dynamic graph convolutional networks for stock market prediction. *Applied Soft Computing, 97,* 106783. https://doi.org/10.1016/j.asoc.2020.106783

Zhao, L., Song, Y., Zhang, C., Liu, Y., Wang, P., Lin, T., Deng, M., and Li, H. (2021). T-GCN: A temporal graph convolutional network for traffic prediction with application to financial networks. *IEEE Transactions on Intelligent Transportation Systems, 21*(9), 3848–3858. https://doi.org/10.1109/TITS.2019.2935152

Sun, J., Zhang, K., Zhu, Y., and Zhao, X. (2021). Dual-channel sentiment-price fusion model for Chinese stock market prediction. *Information Processing and Management, 58*(6), 102700. https://doi.org/10.1016/j.ipm.2021.102700

Cheng, R., Li, Q., and Zhang, L. (2022). FinBERT-GCN: Fusing financial BERT representations with graph convolutional networks for stock movement prediction. *Knowledge-Based Systems, 238,* 107969. https://doi.org/10.1016/j.knosys.2021.107969

Wang, C., Chen, Y., Zhang, S., and Zhang, Q. (2022). Stock market index prediction using deep Transformer model with heterogeneous graph convolutional networks. *Expert Systems with Applications, 197,* 116896. https://doi.org/10.1016/j.eswa.2022.116896

Ying, X., Zhao, H., and Wu, Q. (2021). Blockchain graph neural network for cryptocurrency price prediction. *Financial Innovation, 7*(1), 1–22. https://doi.org/10.1186/s40854-021-00287-z

Liu, Y., Zeng, Q., Yang, H., and Sun, A. (2022). Event-driven stock prediction using knowledge graph and graph convolutional networks. *Information Sciences, 591,* 33–49. https://doi.org/10.1016/j.ins.2022.01.029

Zhang, X., Li, Y., Yang, H., and Zhao, L. (2023). Hierarchical graph attention networks for portfolio optimization and stock prediction. *IEEE Transactions on Neural Networks and Learning Systems, 34*(3), 1278–1291. https://doi.org/10.1109/TNNLS.2022.3154375

Deng, S., Zhang, N., Zhang, W., Chen, J., Pan, J. Z., and Chen, H. (2019). Knowledge-driven stock trend prediction and explanation via temporal convolutional network. *Companion Proceedings of the World Wide Web Conference,* 678–685. https://doi.org/10.1145/3308560.3317701

Wu, J., Wang, G., Liang, J., and Yang, Y. (2023). Adaptive deformable graph convolutional networks with temporal sentiment fusion for multi-horizon stock prediction. *IEEE Transactions on Knowledge and Data Engineering, 35*(7), 7012–7026. https://doi.org/10.1109/TKDE.2022.3195847

He, Z., Chen, K., Li, W., and Zhang, R. (2023). Reddit sentiment and meme stock prediction: A graph-enhanced NLP approach. *Journal of Financial Markets, 62,* 100755. https://doi.org/10.1016/j.finmar.2022.100755

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

Araci, D. (2019). FinBERT: Financial sentiment analysis with pre-trained language models. *arXiv preprint.* https://doi.org/10.48550/arXiv.1908.10063

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

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

Dai, Z., Yang, Z., Yang, Y., Carbonell, J., Le, Q., and Salakhutdinov, R. (2019). Transformer-XL: Attentive language models beyond a fixed-length context. *Proceedings of ACL,* 2978–2988. https://doi.org/10.18653/v1/P19-1285

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, I., 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

Hochreiter, S., and Schmidhuber, J. (1997). Long short-term memory. *Neural Computation, 9*(8), 1735–1780. https://doi.org/10.1162/neco.1997.9.8.1735

Fama, E. F. (1970). Efficient capital markets: A review of theory and empirical work. *The Journal of Finance, 25*(2), 383–417. https://doi.org/10.2307/2325486

Dai, P., Xiao, X., Luo, M., and Zheng, Y. (2021). Temporal graph convolutional networks for social influence-based stock prediction. *Expert Systems with Applications, 178,* 114928. https://doi.org/10.1016/j.eswa.2021.114928