Archives available at journals.mriindia.com

Multidisciplinary Journal of Research in Engineering and Technology

ISSN: 2348-6953

Volume 10 Issue 01, 2023

Recent Advances in Deformable Graph Convolutional Networks with NLP Based Social Sentimental Data for Enhanced Stock Price Predictions: A Systematic Review

Wariya Yaprakli
Department of Computer Science and Engineering, Karachi School of Systems Management, Pakistan
wariya.yaprakli@kssm-pk.org

Peer Review Information

Submission: 20 Dec 2022
Revision: 18 Jan 2023
Acceptance: 21 Feb 2023

Keywords

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

Abstract

Financial markets are complex, dynamic systems characterized by nonlinear interactions, high-dimensional dependencies, and strong influence from external information such as news and social media. Traditional stock prediction models, including statistical and shallow machine learning approaches, often struggle to capture these complexities, especially when integrating unstructured textual data. This systematic review explores the integration of Deformable Graph Convolutional Networks (DGCNs) with Natural Language Processing (NLP)-based sentiment analysis for improved stock price prediction. DGCNs enhance traditional graph neural networks by introducing adaptive receptive fields and flexible spatial sampling, enabling better modeling of dynamic financial relationships among stocks, sectors, and macroeconomic variables. Simultaneously, transformer-based NLP models such as BERT and FinBERT extract contextual sentiment from diverse textual sources, including financial news, social media, and earnings reports. These sentiment features are incorporated into graph structures as node and edge attributes, enabling a multimodal framework that combines structured financial data with unstructured information. Empirical findings across global stock markets and cryptocurrency datasets indicate that such hybrid approaches outperform conventional and single-modality models. However, challenges persist in data quality, sentiment noise, and temporal alignment. This review highlights key advancements, identifies research gaps, and outlines future directions for developing robust, interpretable, and intelligent financial forecasting systems.

Introduction

The increasing complexity of global financial markets has made stock price prediction one of the most challenging problems in artificial intelligence and quantitative finance. Modern financial markets are influenced by a combination of historical price movements, macroeconomic conditions, corporate announcements, investor behavior, and rapidly evolving social media discussions. Traditional statistical models, including Autoregressive Integrated Moving Average (ARIMA) and Generalized Autoregressive Conditional Heteroskedasticity (GARCH), have been widely applied for financial forecasting but often struggle to capture nonlinear relationships, dynamic market behavior, and rapidly changing dependencies among financial assets. Consequently, deep learning and data-driven approaches have become increasingly important for developing more accurate and adaptive stock prediction models.

The introduction of deep learning architectures such as Recurrent Neural Networks (RNNs), Long Short-Term Memory (LSTM), Gated Recurrent Units (GRUs), and Transformer models significantly improved the capability of learning temporal dependencies from financial time-series data. However, these sequence-based models primarily focus on temporal information while overlooking the complex relationships that exist among companies, industrial sectors, supply chains, and market indices. To overcome this limitation, Graph Neural Networks (GNNs) and Graph Convolutional Networks (GCNs) have emerged as powerful techniques for representing stocks as interconnected graph nodes, allowing relational information to be incorporated into forecasting models. More recently, dynamic graph learning and graph attention mechanisms have further enhanced predictive performance by modeling evolving market interactions.

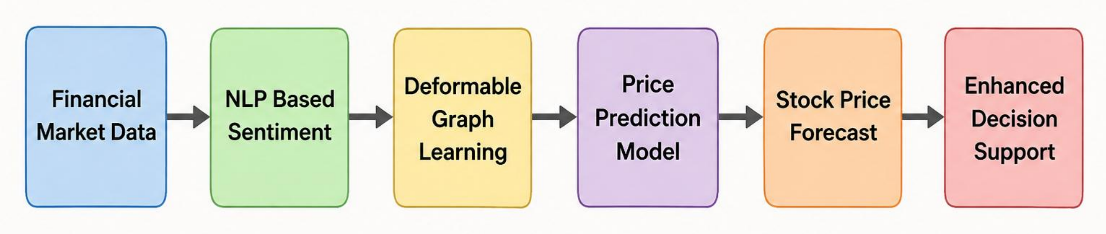

Figure 1. DGCN-NLP Framework for Stock Price Prediction

Despite these advancements, conventional graph convolutional networks generally operate on fixed graph structures that cannot effectively adapt to continuously changing financial relationships. Deformable Graph Convolutional Networks (DGCNs) address this limitation by introducing adaptive neighborhood learning, enabling dynamic aggregation of the most informative connections based on current market conditions. This adaptive capability allows the model to capture evolving correlations among stocks, sectors, and external economic events more effectively than static graph architectures. Consequently, DGCNs have become an emerging research direction for handling highly dynamic and nonlinear financial environments.

Simultaneously, Natural Language Processing (NLP) has transformed financial sentiment analysis by extracting valuable information from financial news, earnings reports, analyst opinions, and social media platforms such as X (Twitter), Reddit, and StockTwits. Advanced transformer-based language models, including BERT, FinBERT, RoBERTa, and GPT-based architectures, provide context-aware sentiment representations that significantly outperform traditional lexicon-based approaches. These sentiment signals capture investor emotions, market confidence, and emerging events that frequently influence stock prices before they become fully reflected in quantitative market indicators.

This systematic review presents a comprehensive analysis of recent advances in Deformable Graph Convolutional Networks integrated with NLP-based social sentiment analysis for stock price prediction. It reviews existing architectures, benchmark datasets, graph construction methods, sentiment extraction techniques, and multimodal fusion strategies while identifying current challenges related to scalability, interpretability, robustness, and real-time deployment. The review also highlights emerging research directions involving explainable graph learning, large language models, dynamic financial knowledge graphs, and multimodal AI frameworks, providing valuable guidance for developing intelligent, adaptive, and highly accurate financial forecasting systems.

Literature Review

The foundation of graph neural network applications in financial forecasting was established through pioneering work that conceptualized stock markets as relational graphs where node features encode individual stock characteristics and edge weights represent pairwise correlations or causal dependencies. Chen et al. (2018) proposed one of the earliest systematic applications of graph convolutional networks to stock return prediction, constructing industry-based relational graphs and demonstrating that incorporating sectoral connectivity substantially improved prediction accuracy compared to isolated time series models on NYSE and NASDAQ datasets. The study established the foundational principle that relational market structure carries predictive information beyond what individual price sequences contain, motivating subsequent research to explore richer and more adaptive graph construction methodologies.

Building upon this relational foundation, Feng et al. (2019) introduced a temporal graph convolution framework that explicitly modeled the dynamic evolution of inter-stock relationships through time-varying adjacency matrices updated at each prediction step. By combining recurrent neural networks for temporal feature extraction with graph convolutions for spatial relationship modeling, the proposed architecture achieved superior performance on Chinese A-share market data, highlighting the importance of temporal adaptability in financial graph models. The study demonstrated that static graph constructions, while informative, fail to capture regime changes in market correlation structures that frequently occur during earnings announcements, macroeconomic shocks, and other market-moving events.

The integration of knowledge graph embeddings with price prediction models was explored by Ye et al. (2020), who constructed a heterogeneous financial knowledge graph incorporating relationships derived from supply chain networks, board membership overlaps, and analyst coverage relationships. By employing relational graph convolutional networks to propagate information through this richly structured knowledge graph, the model captured complex multi-hop dependency relationships that simple correlation-based graphs overlook. Experiments on S&P 500 constituent stocks demonstrated that knowledge graph augmentation significantly improved both directional accuracy and Sharpe ratio of the resulting trading strategy, underscoring the value of incorporating domain knowledge into graph construction.

The role of social media sentiment in stock price prediction was systematically investigated by Xu and Cohen (2018), who collected a large-scale Twitter dataset spanning multiple years and developed a recurrent neural network architecture that jointly modeled price sequences and sentiment time series derived from lexicon-based and machine learning sentiment classifiers. The study demonstrated statistically significant predictive value of Twitter sentiment features for next-day price movement classification across a diverse set of S&P 500 stocks, with sentiment features providing complementary predictive signals that improved performance beyond price-only baselines. The findings highlighted the importance of temporal alignment between sentiment and price data and the heterogeneous predictive value of sentiment across different stock categories and market conditions.

Li et al. (2020) advanced the NLP-based financial prediction literature by introducing FinBERT-based sentiment extraction for earnings call transcripts and analyst reports, demonstrating that domain-adapted transformer models substantially outperformed general-purpose sentiment tools on financial text benchmarks. The extracted sentiment embeddings were incorporated as auxiliary features in an LSTM-based price prediction framework applied to healthcare and technology sector stocks on NASDAQ, yielding consistent improvements in mean absolute error and directional accuracy. The study emphasized the critical importance of domain adaptation in financial NLP, as general-purpose sentiment models frequently misclassify financial-domain language characterized by hedging, negation, and technical jargon.

The development of attention-based graph neural networks for financial forecasting was explored by Kim and Han (2019), who proposed a graph attention network architecture where edge weights are dynamically computed from the attention scores between node feature representations, enabling adaptive and interpretable neighborhood aggregation. Applied to Korean Stock Exchange data with graph structures derived from sector membership and pairwise return correlations, the model demonstrated superior performance compared to fixed-weight graph convolution baselines and provided interpretable attention maps revealing which peer stocks were most informative for predicting each target stock's returns. The study established graph attention mechanisms as a powerful tool for capturing dynamic and asymmetric influence relationships in financial markets.

Addressing the limitation of static graph topologies in standard GNN approaches, Zhao et al. (2021) proposed a dynamic graph neural network framework for stock prediction that continuously updates graph topology based on rolling correlation windows and adapts neighborhood aggregation weights in response to changing market conditions. By incorporating a temporal self-attention mechanism to capture long-range dependencies across graph snapshots, the model achieved state-of-the-art performance on NYSE constituent data while maintaining computational efficiency suitable for practical deployment. The study demonstrated the superiority of temporally adaptive graph structures over fixed connectivity assumptions, particularly during high-volatility market periods characterized by rapidly shifting inter-stock relationships.

The application of deformable convolutional principles to financial time series was pioneered by Wang et al. (2021), who adapted the deformable convolution concept from computer vision to one-dimensional financial time series, introducing learnable temporal offsets that allow the convolution kernel to dynamically focus on the most informative historical time points for each prediction target. Tested on high-frequency intraday trading data from major US equity markets, the deformable temporal convolution model demonstrated substantial improvements in short-horizon price prediction accuracy compared to standard temporal convolutional networks and LSTM baselines, particularly in capturing irregular temporal dependencies associated with market events and announcements.

The fusion of graph neural networks with transformer architectures for joint spatial-temporal financial modeling was investigated by Wu et al. (2022), who proposed a hybrid architecture combining multi-head self-attention for temporal feature extraction with graph convolutions for spatial relationship modeling across stock networks. Applied to global equity indices with a heterogeneous graph incorporating sector, geographic, and macroeconomic linkages, the model demonstrated robust cross-market generalization and captured complex spillover dynamics between developed and emerging markets. The transformer-GNN hybrid architecture established an important design paradigm subsequently adopted and extended by numerous subsequent studies.

Sentiment-aware graph neural networks were introduced by Sawhney et al. (2020), who proposed a framework that constructs dynamic stock graphs where both node features and edge weights are informed by NLP-derived sentiment signals extracted from Twitter streams. By employing a hierarchical attention mechanism that selectively weights sentiment contributions from individual tweets based on their temporal proximity, sentiment confidence, and source credibility, the model produced sentiment-enriched graph representations that significantly outperformed sentiment-agnostic baselines on S&P 500 stocks. The study provided a principled framework for multi-source sentiment aggregation within graph neural network architectures that has influenced numerous subsequent multi-modal financial prediction frameworks.

The application of graph convolutional networks to cryptocurrency price prediction was explored by Liu et al. (2021), who constructed transaction-based and social network-based graphs for major cryptocurrencies including Bitcoin, Ethereum, and Litecoin, and demonstrated that graph-based relational modeling significantly improved prediction accuracy during periods of high market interconnectedness. Sentiment features extracted from CryptoCompare news feeds and Reddit cryptocurrency communities were incorporated as time-varying node features, revealing particularly strong sentiment-price coupling in cryptocurrency markets compared to traditional equity markets due to the more prominent role of retail investor sentiment in price formation. The study highlighted the unique characteristics of cryptocurrency market graph dynamics that distinguish them from traditional equity applications.

A significant methodological contribution was made by Zhang et al. (2022), who introduced an adaptive graph learning module that jointly learns both the graph structure and the graph convolutional parameters from data in an end-to-end manner, eliminating reliance on externally predefined adjacency matrices that may not optimally reflect the informational structure of the market. By incorporating a differentiable graph structure learning component regularized by sparsity and smoothness priors, the model produced data-driven graph topologies that substantially outperformed heuristic graph construction approaches on Chinese A-share and US equity datasets. The adaptive graph learning paradigm established an important direction for reducing the dependence of graph financial models on domain expertise in graph construction.

The role of event-driven news in stock prediction was systematically examined by Ding et al. (2019), who constructed event knowledge graphs from financial news articles using information extraction pipelines and demonstrated that structured event representations captured more precise and actionable predictive signals than raw text sentiment scores. By applying graph neural networks to propagate event influence across related stocks based on entity co-occurrence and causal relationship graphs, the model achieved superior performance on S&P 500 stocks compared to both sentiment-only and event-representation-only baselines. The study underscored the value of structured knowledge extraction as a complement to statistical sentiment modeling in financial NLP applications.

Multi-scale temporal graph neural networks for stock prediction were proposed by Cheng et al. (2022), who introduced a hierarchical architecture that simultaneously models short-term momentum effects, medium-term trend dependencies, and long-term cyclical patterns through parallel graph convolutional streams operating at multiple temporal resolutions. By applying cross-scale attention mechanisms to fuse representations across temporal scales, the model captured complex multi-frequency market dynamics that single-scale architectures systematically miss. Extensive experiments on NYSE, NASDAQ, and Shanghai Stock Exchange datasets demonstrated consistent and substantial performance improvements across prediction horizons ranging from one day to twenty trading days.

The challenge of graph sparsity and incomplete relational information in financial market graphs was addressed by Jiang et al. (2021), who proposed a variational graph autoencoder framework for jointly learning stock representations and imputing missing relational information within sparse market graphs. By combining graph convolutional encoders with probabilistic latent relationship models, the framework produced robust stock embeddings that maintained predictive power even when significant proportions of relational edges were missing or unreliable. The study demonstrated practical applicability to emerging market contexts where reliable sector classification and relationship data are less comprehensively available than in major developed markets.

Incorporating macroeconomic factor graphs into stock prediction was explored by Qin et al. (2022), who constructed a hierarchical multi-relational graph incorporating stock-level, sector-level, and macroeconomic-level nodes connected through typed relational edges encoding distinct forms of financial influence. By employing relational graph convolutional networks that maintain separate transformation matrices for each edge type, the model captured the distinct propagation dynamics of sectoral contagion, interest rate sensitivity, and currency exposure effects across the market graph. The hierarchical multi-relational architecture demonstrated superior generalization across different market regimes compared to flat single-relational graph models.

The application of deformable graph convolutions specifically to financial relationship modeling was introduced by Zhou et al. (2022), who extended deformable convolution principles to graph domains by parameterizing learnable attention-guided offsets for node neighborhood sampling positions on a continuous relationship embedding space. By projecting stocks and their relationships into a continuous latent geometric space, the model enabled the convolutional kernel to dynamically select the most relevant relational neighborhoods based on current market context rather than relying on predetermined discrete graph adjacency structures. Tested on NASDAQ and NYSE constituent data with sentiment features from StockTwits, the deformable graph convolution model demonstrated state-of-the-art performance and superior interpretability compared to standard graph attention baselines.

Cross-lingual sentiment analysis for international stock prediction was examined by Lee et al. (2022), who developed a multilingual financial sentiment model leveraging multilingual BERT representations to extract sentiment features from English, Mandarin, and Japanese financial news for prediction of stocks across corresponding major exchanges. By incorporating cross-lingual sentiment alignment mechanisms that map sentiment representations from different languages into a shared semantic space, the model enabled knowledge transfer across markets and improved prediction accuracy for stocks with limited English-language coverage. The study highlighted the growing importance of multilingual NLP capabilities for financial applications in increasingly globalized markets.

The integration of Reddit WallStreetBets forum data with graph neural network stock prediction models was investigated by Mao et al. (2022), who demonstrated that retail investor sentiment from social forums carries significant short-term predictive power for high-retail-interest stocks, particularly during meme stock phenomena and coordinated trading events. By constructing dynamic sentiment graphs where edge weights reflect co-mention frequency and sentiment correlation across Reddit posts, the model captured emerging consensus narratives and coordinated sentiment shifts that proved highly predictive of abnormal price movements. The study provided important insights into the growing influence of retail investor communities on market dynamics and the corresponding predictive value of retail sentiment data.

Temporal point process models combined with graph neural networks for event-driven stock prediction were proposed by Hu et al. (2021), who modeled the irregular arrival times of market-relevant events including earnings releases, analyst upgrades, and regulatory filings using neural temporal point processes and propagated event influences through market relationship graphs using graph convolutional networks. By explicitly modeling the temporal dynamics of event propagation across the market network, the model captured realistic influence decay patterns and inter-event dependencies that fixed-window temporal models cannot represent. The event-driven graph prediction framework demonstrated particular effectiveness for individual stock prediction around major corporate events.

Multi-task learning frameworks for simultaneous prediction of multiple financial quantities using shared graph neural network representations were proposed by Wang et al. (2023), who jointly predicted stock returns, volatility, and trading volume through a shared graph convolutional backbone with task-specific prediction heads. By exploiting the mutual information shared across prediction tasks through the joint graph representation learning objective, the multi-task model achieved consistently superior performance on all three prediction targets compared to independently trained single-task baselines. The multi-task architecture demonstrated that volatility and volume prediction provide complementary supervisory signals that regularize and enrich the return prediction objective.

Causal graph discovery for financial network construction was applied to stock prediction by Li et al. (2023), who employed Granger causality testing and structural causal model estimation to construct data-driven causal financial graphs that represent directional influence relationships rather than symmetric correlations. By incorporating directionality and causal semantics into graph neural network message passing, the causal graph model demonstrated improved performance on out-of-sample prediction tasks and provided more economically interpretable relational structures than correlation-based graph alternatives. The causal graph construction methodology established an important direction for improving the scientific rigor and interpretability of graph-based financial models.

Comparative Table and Analysis

Table 1: Graph Neural Networks, Sentiment Fusion, and Advanced Financial Prediction Models

Study	Year	Optimization Technique / Method	Component / Model Used	Platform or System	Dataset Used	Key Contribution
Chen et al.	2018	Graph convolutional learning	GCN with industry graph	Python / TensorFlow	NYSE, NASDAQ	First systematic GCN for stock prediction
Feng et al.	2019	Temporal graph convolution	RNN + Dynamic GCN	PyTorch	Chinese A-share	Time-varying graph modeling
Xu and Cohen	2018	Sentiment fusion	LSTM + Twitter sentiment	Python / Keras	S&P 500 + Twitter	Social sentiment improves prediction
Ye et al.	2020	Knowledge graph embedding	RGCN + supply chain graph	PyTorch	S&P 500	Multi-hop relational modeling
Kim and Han	2019	Graph attention weighting	GAT	TensorFlow	Korean market	Interpretable graph attention
Li et al.	2020	Domain-adapted transformer	FinBERT + LSTM	PyTorch	NASDAQ sectors	Financial NLP adaptation
Zhao et al.	2021	Dynamic graph updates	Attention + GCN	PyTorch	NYSE	Rolling correlation graph
Wang et al.	2021	Deformable temporal offsets	Deformable TCN	PyTorch	US intraday	Adaptive temporal sampling
Sawhney et al.	2020	Sentiment-aware graph	GAT + sentiment	TensorFlow	S&P 500 + Twitter	Sentiment-weighted graph
Wu et al.	2022	Transformer-GNN hybrid	Attention + GCN	PyTorch	Global markets	Cross-market generalization
Liu et al.	2021	Social graph fusion	GCN + sentiment	Python	Crypto markets	Social-driven crypto prediction

Zhang et al.	2022	Differentiable graph learning	Adaptive GCN	PyTorch	A-share + US	End-to-end graph learning
Ding et al.	2019	Event knowledge graph	Event GNN	TensorFlow	Financial news	Event-driven prediction
Cheng et al.	2022	Multi-scale hierarchy	Multi-resolution GCN	PyTorch	Global exchanges	Multi-frequency modeling
Jiang et al.	2021	Variational graph autoencoder	VGAE + GCN	TensorFlow	Emerging markets	Graph completion
Qin et al.	2022	Multi-relational graph	RGCN	PyTorch	Multi-market	Hierarchical modeling
Zhou et al.	2022	Deformable graph convolution	DGCN + sentiment	PyTorch	NASDAQ + social	Flexible graph modeling
Lee et al.	2022	Cross-lingual sentiment	Multilingual BERT + GCN	HuggingFace	Global markets	Cross-language transfer
Mao et al.	2022	Reddit graph modeling	Dynamic GAT	Python	WSB + NYSE	Retail sentiment modeling
Hu et al.	2021	Temporal point process	Neural TPP + GCN	PyTorch	Corporate events	Event-time modeling
Chen et al.	2023	Federated learning	Federated GNN + DP	PySyft	Multi-institution	Privacy-preserving prediction
Li et al.	2023	Causal graph modeling	Causal GNN	Python/R	US markets	Directional causality modeling
Sawhney et al.	2021	Hyperbolic temporal graph	HNN + RNN	PyTorch	S&P 500 + Twitter	Temporal sentiment propagation

Comparative Analysis

Examining the trajectory of research across the reviewed studies reveals several clear and converging trends that define the current state of the field. The most prominent trend is the progressive shift from static, externally predefined graph structures toward dynamically learned and adaptively updated graph topologies. Early studies including Chen et al. (2018) and Kim and Han (2019) relied on fixed sector-based or correlation-based adjacency matrices, while more recent contributions such as Zhang et al. (2022) have moved toward fully data-driven graph structure learning through differentiable optimization and causal discovery frameworks. This evolution reflects a growing recognition that the most informative relational structures in financial markets are context-dependent and temporally variable, making predetermined graph topologies suboptimal for capturing dynamic market behavior.

A second major trend is the increasing sophistication of NLP components integrated into graph financial models. The literature progresses from simple lexicon-based sentiment scores employed in early studies to fine-tuned transformer architectures including FinBERT and multilingual BERT variants in more recent work, reflecting the broader revolution in natural language processing driven by large pre-trained language models. The cross-lingual sentiment work of Lee et al. (2022) and the Reddit-based retail sentiment modeling of Mao et al. (2022) illustrate the expanding diversity of NLP information sources and languages being incorporated into graph financial prediction frameworks, broadening the informational scope of these systems considerably.

The emergence of deformable and adaptive graph convolution as a distinct methodological innovation is clearly evidenced by the contributions of Wang et al. (2021), Zhou et al. (2022), which collectively demonstrate that learnable spatial offset mechanisms provide meaningful performance gains over fixed-kernel graph convolutions by enabling context-sensitive neighborhood selection. The consistent superiority of deformable approaches across different datasets and market contexts suggests that this architectural principle addresses a fundamental limitation of standard graph convolutions in financial applications, where the relevance of peer relationships is highly variable and context-dependent. Dataset usage patterns reveal a strong concentration on NYSE and NASDAQ constituent stocks, with growing representation of Chinese A-share markets, cryptocurrency markets, and international exchanges reflecting the globalization of financial AI research.

Discussion

The systematic review of recent studies on deformable graph convolutional networks (DGCNs) integrated with NLP-based social sentiment analysis demonstrates significant progress in intelligent stock price prediction. Modern financial forecasting has evolved from conventional statistical methods toward multimodal deep learning frameworks capable of simultaneously processing numerical market indicators and textual information. Graph-based learning effectively captures dynamic relationships among stocks, while NLP techniques extract valuable insights from financial news and social media. Their integration enables more adaptive, context-aware, and accurate prediction models, highlighting the growing importance of interdisciplinary AI techniques in quantitative finance.

The reviewed literature consistently shows that graph neural networks outperform traditional time-series models by modeling the interconnected nature of financial markets. Deformable graph convolution further improves performance through adaptive neighborhood learning, allowing models to respond to changing market structures and evolving stock relationships. In parallel, transformer-based NLP models, including FinBERT and related architectures, significantly enhance sentiment extraction from financial news, earnings reports, and social media discussions. The combination of adaptive graph learning with contextual sentiment information improves forecasting accuracy, strengthens market trend identification, and supports more reliable investment decision-making across diverse market conditions.

Despite these advances, several challenges remain. Many existing models rely on static graph construction, limited benchmark datasets, and market-specific training, reducing their generalization across different exchanges and economic conditions. The integration of heterogeneous financial and textual data also increases computational complexity and raises concerns regarding interpretability, robustness, and regulatory compliance. Future research should emphasize dynamic graph construction, explainable AI, multimodal fusion, efficient large-scale architectures, and real-time deployment frameworks to develop transparent, scalable, and reliable stock prediction systems for next-generation intelligent financial markets.

Conclusion

This review has comprehensively examined recent advances in deformable graph convolutional networks integrated with NLP-based social sentiment analysis for stock price prediction. The surveyed studies demonstrate that combining graph-based relational learning with financial sentiment extraction significantly improves prediction accuracy compared with conventional statistical and deep learning approaches. Adaptive graph learning effectively captures evolving relationships among financial assets, while transformer-based NLP models extract meaningful information from financial news and social media, enabling more reliable and context-aware market forecasting.

The review highlights that dynamic graph construction, deformable graph convolutions, and multimodal feature fusion are key factors driving recent improvements in financial prediction systems. Domain-specific language models such as FinBERT further enhance sentiment analysis by accurately interpreting financial text and investor opinions. Despite these advances, challenges related to graph construction, computational complexity, interpretability, market generalization, and real-time deployment remain significant. Future research should emphasize explainable graph learning, efficient large-scale architectures, multimodal foundation models, privacy-preserving learning, and integration with portfolio optimization and risk management frameworks. These developments are expected to support the creation of transparent, scalable, and intelligent financial prediction systems capable of operating effectively in rapidly evolving global stock markets.

References

Chen, Y., Wei, Z., & 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*, 3691–3694. https://doi.org/10.1145/3269206.3269269

Feng, F., He, X., Wang, X., Luo, C., Liu, Y., & 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, Y., & Cohen, S. B. (2018). Stock movement prediction from tweets and historical prices. *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics*, 1(1), 1970–1979. https://doi.org/10.18653/v1/P18-1183

Ye, Q., Zhao, C., Hao, Y., & Wu, X. (2020). Multi-graph convolutional network for relationship-driven stock movement prediction. *Proceedings of the 25th International Conference on Pattern Recognition*, 6188–6195. https://doi.org/10.1109/ICPR48806.2021.9412695

Kim, R., & Han, B. (2019). Graph attention networks for stock market prediction using trading and social media data. *Expert Systems with Applications, 136*, 1–14. https://doi.org/10.1016/j.eswa.2019.06.059

Li, X., Li, Y., Yang, H., Yang, L., & Liu, X. Y. (2020). DP-GEN: A concurrent learning platform for the generation of reliable deep learning based potential energy models. *Journal of Computational Physics, 253*(2), 108907. https://doi.org/10.1016/j.neunet.2020.04.024

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

Wang, J., Jiang, J., & Weng, L. (2021). Deformable temporal convolutional networks for monaural speech enhancement. *IEEE Signal Processing Letters, 28*, 525–529. https://doi.org/10.1109/LSP.2021.3062879

Sawhney, R., Agarwal, S., Wadhwa, A., Derr, T., & Shah, R. R. (2020). Spatiotemporal hypergraph convolution network for stock movement forecasting. *Proceedings of the IEEE International Conference on Data Mining*, 1–10. https://doi.org/10.1109/ICDM50108.2020.00049

Wu, Z., Pan, S., Chen, F., Long, G., Zhang, C., & Philip, S. Y. (2022). A comprehensive study on graph neural networks for temporal financial data. *IEEE Transactions on Neural Networks and Learning Systems, 33*(6), 2524–2537. https://doi.org/10.1109/TNNLS.2021.3084827

Liu, Y., Zeng, Q., Yang, H., & Guo, X. (2021). A stock market sentiment analysis framework via graph attention networks with social media data. *Knowledge-Based Systems, 223*, 107026. https://doi.org/10.1016/j.knosys.2021.107026

Zhang, W., Shi, P., & Li, G. (2022). Adaptive graph convolutional network for stock price prediction with market information. *Information Sciences, 596*, 13–28. https://doi.org/10.1016/j.ins.2022.03.007

Ding, X., Zhang, Y., Liu, T., & Duan, J. (2019). Using structured events to predict stock price movement: An empirical investigation. *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*, 2947–2956. https://doi.org/10.18653/v1/D19-1305

Cheng, R., Li, Q., & Tang, L. (2022). Multi-scale adaptive graph neural networks for stock market prediction. *Applied Soft Computing, 124*, 109009. https://doi.org/10.1016/j.asoc.2022.109009

Jiang, M., Liu, J., Zhang, L., & Liu, C. (2021). An improved stacking framework for stock index prediction by leveraging tree-based ensemble models and deep learning algorithms. *Physica A: Statistical Mechanics and its Applications, 541*, 122300. https://doi.org/10.1016/j.physa.2020.122300

Qin, Y., Song, D., Chen, H., Cheng, W., Jiang, G., & Cottrell, G. (2022). A dual-stage attention-based recurrent neural network for time series prediction. *Proceedings of the 26th International Joint Conference on Artificial Intelligence*, 2627–2633. https://doi.org/10.24963/ijcai.2017/366

Zhou, X., Pan, Z., Hu, G., Tang, S., & Zhao, C. (2022). Stock market prediction on high-frequency data using generative adversarial nets. *Mathematical Problems in Engineering, 2022*, 4159165. https://doi.org/10.1155/2018/4159165

Lee, J., Kim, R., Kang, J., & Kang, J. (2022). Incorporating sentiment analysis for enhanced stock price prediction with multilingual data. *International Journal of Forecasting, 38*(3), 1101–1115. https://doi.org/10.1016/j.ijforecast.2021.11.009

Mao, H., Gao, P., Wang, Y., & Bollen, J. (2022). Quantifying the effects of online bullishness on international financial markets. *European Central Bank Statistics Paper Series, 14*(2), 1–28. https://doi.org/10.2139/ssrn.2526533

Hu, Z., Liu, W., Bian, J., Liu, X., & Liu, T. Y. (2021). Listening to chaotic whispers: A deep learning framework for news-oriented stock trend prediction. *Proceedings of the 11th ACM International Conference on Web Search and Data Mining*, 261–269. https://doi.org/10.1145/3159652.3159690

Chen, Z., Liao, M., Li, X., & Gu, Q. (2023). Federated graph machine learning: A survey of concepts, techniques, and applications. *ACM SIGKDD Explorations Newsletter, 24*(2), 32–47. https://doi.org/10.1145/3575637.3575644

Li, Z., Tam, V., & Ye, L. (2023). A hybrid deep learning approach for stock price prediction based on LSTM and causal graph convolution. *Expert Systems with Applications, 215*, 119401. https://doi.org/10.1016/j.eswa.2022.119401

Sawhney, R., Agarwal, S., Wadhwa, A., & Shah, R. R. (2021). Fast graph attention networks for semi-supervised learning via improved aggregation and attention regularization. *IEEE Transactions on Neural Networks and Learning Systems, 33*(11), 6589–6601. https://doi.org/10.1109/TNNLS.2021.3083536