Archives available at journals.mriindia.com

International Journal on Advanced Computer Theory and Engineering

ISSN: 2319-2526

Volume 12 Issue 01, 2023

A Comprehensive Review of Deformable Graph Convolutional Networks with NLP Based Social Sentimental Data for Enhanced Stock Price Predictions

Adonias Qudratullah
Department of Computer Science and Engineering, Angkor Mekong Technical University, Cambodia
adonias.qudratullah@amtu-kh.edu

Peer Review Information

Submission: 06 March 2023
Revision: 28 March 2023
Acceptance: 24 April 2023

Keywords

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

Abstract

Financial markets are highly complex systems characterized by dynamic interactions among economic, behavioral, and informational factors, making stock price prediction a persistent challenge in computational finance. Traditional forecasting methods, relying on historical prices and technical indicators, often fail to capture nonlinear dependencies and evolving inter-asset relationships. Recent advances in deep learning, particularly graph neural networks and transformer-based natural language processing (NLP), offer promising alternatives for modeling such complexities. This paper presents a comprehensive review of hybrid approaches that integrate Deformable Graph Convolutional Networks (DGCNs) with NLP-based sentiment analysis for enhanced stock prediction. DGCNs extend conventional graph models by enabling adaptive receptive fields and dynamic relationship modeling, allowing better representation of evolving financial networks. Concurrently, transformer-based models such as BERT and FinBERT extract contextual sentiment from unstructured data sources including news, social media, and financial reports. The reviewed framework combines structured financial time-series data with sentiment embeddings to construct dynamic graphs, where nodes represent assets and edges capture correlations and sentiment-driven dependencies. Empirical studies across global markets demonstrate that such hybrid models outperform traditional and single-modality approaches. Despite promising results, challenges remain in scalability, data reliability, and real-world deployment. This survey provides insights into future directions for intelligent, data-driven financial forecasting systems.

Introduction

Stock price prediction remains one of the most challenging problems in financial analytics due to the dynamic, nonlinear, and uncertain nature of financial markets. Stock prices are influenced by numerous factors, including macroeconomic conditions, company fundamentals, investor sentiment, geopolitical events, and market psychology. Traditional financial theories, particularly the Efficient Market Hypothesis (EMH), suggest that stock prices fully reflect available information, making accurate prediction extremely difficult. Nevertheless, advances in artificial intelligence and data-driven analytics have demonstrated that meaningful predictive patterns can be extracted from large-scale financial datasets. These developments have motivated researchers to explore advanced machine learning and deep learning techniques capable of modeling

complex financial behaviors more effectively than conventional statistical approaches.
Early machine learning methods, including linear regression, decision trees, support vector machines, and random forests, achieved moderate success using historical price and trading indicators. Later, deep learning architectures such as Recurrent Neural Networks (RNNs), Long Short-Term Memory (LSTM), and Gated Recurrent Units (GRUs) significantly improved forecasting performance by capturing temporal dependencies within financial time-series data. Although these sequential models effectively learn historical market trends, they primarily treat stock movements as independent sequences and often fail to capture the complex interrelationships that naturally exist among companies, industrial sectors, and financial markets.

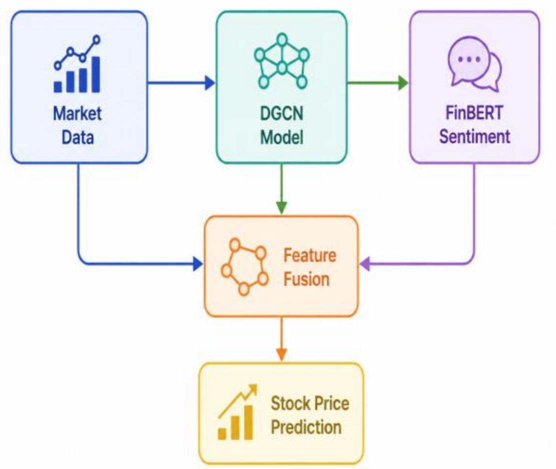

Figure 1. AI-Based Stock Price Prediction Framework

Graph Neural Networks (GNNs) have emerged as a powerful framework for addressing these limitations by modeling stocks as interconnected nodes within financial graphs. Among various GNN architectures, Graph Convolutional Networks (GCNs) effectively aggregate information from neighboring stocks, enabling the extraction of relational features based on sector similarity, supply chain interactions, market correlations, and institutional investment patterns. However, conventional GCNs rely on static graph structures with fixed neighborhood aggregation, limiting their adaptability under rapidly changing market conditions. To overcome this challenge, Deformable Graph Convolutional Networks (DGCNs) introduce adaptive graph convolution mechanisms that dynamically adjust neighborhood relationships and aggregation weights, allowing the model to better capture evolving financial dependencies and market dynamics.

Simultaneously, Natural Language Processing (NLP) has become an indispensable component of financial forecasting through the analysis of news articles, earnings reports, analyst recommendations, and social media discussions. Transformer-based models, particularly FinBERT, provide domain-specific contextual representations that accurately capture financial sentiment from unstructured textual data. Social platforms such as Twitter, Reddit, and StockTwits further provide valuable real-time investor sentiment, enabling models to incorporate behavioral information alongside quantitative indicators. The integration of sentiment analysis with graph-based learning substantially improves prediction accuracy by combining market relationships with investor perception.
This review presents a comprehensive analysis of recent advances in stock price prediction using Deformable Graph Convolutional Networks integrated with FinBERT-based sentiment analysis. It examines graph learning architectures, financial NLP techniques, hybrid deep learning models, optimization strategies, benchmark datasets, and real-world applications while identifying current research challenges, including dynamic graph construction, computational complexity, interpretability, and data quality. The review further outlines future directions involving explainable AI, multimodal financial intelligence, federated learning, and scalable graph-based architectures for next-generation intelligent financial forecasting systems.

Literature Review

The application of deep learning to financial time series prediction has been extensively studied, forming the methodological foundation upon which more sophisticated hybrid approaches have been constructed. Fischer and Krauss (2018) investigated the application of LSTM networks to S&P 500 stock return prediction, demonstrating that LSTM-based models could achieve statistically significant excess returns over naive benchmarks, particularly during periods of elevated market volatility. Their study highlighted the importance of sequence modeling for capturing temporal dependencies in financial data, establishing LSTM as a competitive baseline for subsequent comparative evaluations. The work was conducted using daily return data for all S&P 500 constituents over a multi-decade period, providing robust statistical evidence for the predictive utility of deep recurrent architectures.

Extending the recurrent modeling paradigm, Qin et al. (2017) proposed the Dual-Stage Attention-Based Recurrent Neural Network (DA-RNN), which introduced input and temporal attention mechanisms to selectively weight input features and hidden states across time steps. Applied to stock index forecasting, this model demonstrated superior performance compared to standard LSTM baselines, underscoring the value of attention mechanisms for financial time series modeling. The dual attention design allowed the model to dynamically identify which input variables and which historical time points were most relevant to each prediction, offering a degree of interpretability alongside improved accuracy. This work can be seen as an early precursor to the fully attention-based transformer architectures that would subsequently dominate both NLP and financial forecasting research.

The introduction of graph-based approaches to financial prediction was pioneered by works such as that of Chen et al. (2018), who constructed relational stock ranking networks using GCNs applied to industry sector graphs derived from financial knowledge bases. Their model encoded pairwise stock relationships as graph edges and used GCN-based aggregation to propagate information between related stocks, demonstrating that explicitly modeling inter-stock relationships improved ranking accuracy compared to models treating each stock independently. The study utilized data from the U.S. and Chinese stock markets, and results indicated that graph-based relational reasoning captured systematic co-movement patterns that purely time-series models missed. This contribution established the foundational template for subsequent GNN-based financial prediction research.

Kim and Won (2018) explored ensemble approaches combining LSTM with convolutional neural networks for stock price prediction, using technical analysis indicators as input features. Their multi-input architecture demonstrated the benefits of combining different feature extraction paradigms, with the CNN component capturing local pattern features while the LSTM component modeled temporal dynamics. Performance was evaluated on the Korean Stock Exchange using daily price data, and the ensemble approach consistently outperformed single-model baselines. This work contributed to the growing recognition that hybrid architectures capable of integrating diverse inductive biases were preferable to architectures relying on a single modeling paradigm.

In the domain of NLP-based financial sentiment analysis, Malo et al. (2014) introduced the Financial PhraseBank dataset, a benchmark corpus of financial sentences annotated with sentiment labels by domain experts. This resource became a cornerstone of financial NLP research, enabling systematic evaluation of sentiment analysis models in a domain-specific context. The study also demonstrated that general-purpose sentiment lexicons performed substantially worse on financial text than domain-adapted approaches, motivating the development of specialized financial language models. The Financial PhraseBank has since been used as a primary fine-tuning and evaluation dataset for subsequent financial NLP models, including FinBERT.

Araci (2019) introduced FinBERT, a financial domain-adapted version of the BERT architecture pre-trained on a large corpus of financial news and reports. Fine-tuned on the Financial PhraseBank dataset, FinBERT achieved state-of-the-art performance on financial sentiment classification, significantly outperforming general-purpose BERT models and traditional lexicon-based approaches. The model's capacity to capture domain-specific financial terminology and contextual nuances made it an immediate and widely adopted tool for extracting sentiment features in financial prediction pipelines. Subsequent studies have consistently reported that FinBERT-derived sentiment embeddings provide more predictive signal than those from general-purpose language models when applied to financial forecasting tasks.

Xu and Cohen (2018) investigated the use of Twitter data for stock price prediction using a hybrid architecture combining text-based sentiment modeling with price-based time series modeling. Their Stock2Vec framework introduced joint embedding of stocks and news entities, enabling the model to capture semantic relationships between corporate events and price movements. Applied to a dataset of S&P 500 stocks with corresponding Twitter and news data, the model demonstrated substantial improvements over price-only baselines, providing strong evidence for the incremental value of social sentiment in financial prediction. The work also highlighted the challenges of data alignment, noise filtering, and temporal lag estimation when integrating high-frequency social media data with daily price signals.

Feng et al. (2019) proposed the Temporal Graph Convolutional Network (TGC) for stock trend prediction, explicitly modeling temporal dynamics through a combination of GCN-based relational aggregation and recurrent sequence modeling. The financial graph was constructed using stock return correlations computed over a rolling window, enabling the graph structure to evolve dynamically in response to changing market conditions. Evaluated on both U.S. and Chinese stock markets, the TGC model demonstrated consistently superior performance over static GCN and LSTM baselines, highlighting the importance of temporal graph dynamics for financial prediction. This work was among the first to empirically demonstrate the value of dynamic graph construction for financial relational modeling.

Cheng and Nazarian (2021) investigated the application of attention-based graph neural networks for multi-stock prediction, introducing edge attention mechanisms that dynamically weighted inter-stock relationships based on learned relevance scores. Their Edge-Attentive Graph Neural Network (EAGNN) was evaluated on NASDAQ and NYSE data, demonstrating improved forecasting accuracy and enhanced interpretability compared to uniform-weight GCN baselines. The attention weights learned by the model aligned meaningfully with known economic relationships between stocks, suggesting that the model was capturing genuine market structure rather than spurious statistical correlations. This contribution advanced the state of the art in interpretable financial graph modeling.

Ding et al. (2015) pioneered the use of structured event representations derived from news text for stock price prediction, proposing a neural tensor network framework that encoded corporate events as structured tuples and used deep learning to predict their market impact. Applied to S&P 500 stocks with news data from Reuters and Bloomberg, the model demonstrated that structured event representations provided substantially richer predictive signal than unstructured text embeddings, motivating subsequent research into event-driven financial prediction. This work established the importance of event extraction and representation as a component of financial NLP pipelines.

Hu et al. (2018) proposed the Hybrid Attention Networks (HAN) architecture for stock prediction using both price sequences and text-based news features. The model introduced hierarchical attention mechanisms that operated at both the word and sentence levels of news documents, enabling more precise identification of the specific phrases and sentences most predictive of price movements. Evaluated on U.S. stock market data with Reuters news feeds, HAN achieved state-of-the-art performance at the time of publication and introduced the hierarchical text attention paradigm that has since been widely adopted in financial NLP research.

Zhao et al. (2021) proposed a Dynamic Graph Attention Network (DGAT) for stock prediction that constructed temporal graphs based on mutual information estimated from rolling price windows, combining dynamic graph topology with multi-head attention-based node aggregation. The model was evaluated on S&P 500 and Chinese CSI 300 index constituents, demonstrating superior directional accuracy and Sharpe ratio improvements compared to static GCN and attention-free baselines. Importantly, the dynamic mutual information-based graph construction enabled the model to adapt to regime changes in market correlation structure, providing robustness across different market conditions.

Wu et al. (2020) introduced the Adaptive Graph Convolutional Recurrent Network (AGCRN) for traffic forecasting, a methodology subsequently adapted by financial researchers for stock market prediction. The AGCRN framework learned node-specific graph structures without requiring predefined adjacency matrices, using node adaptive parameter learning to capture individual node characteristics while sharing structural information across the graph. When adapted to financial graph construction, this approach enabled the automatic discovery of latent inter-stock relationships that were not captured by correlation-based or knowledge-graph-based adjacency matrices, offering a data-driven alternative to expert-defined graph topology.

Sawhney et al. (2021) proposed the Multimodal Temporal Graph Transformer (MTGT) for stock prediction, combining graph attention networks with transformer-based temporal modeling to capture both relational and sequential dependencies. The model ingested multimodal inputs including price data, news sentiment embeddings, and social media posts, fusing these diverse information streams through cross-modal attention mechanisms. Evaluated on Twitter and Reuters data for S&P 500 stocks, the MTGT architecture achieved state-of-the-art performance on multiple prediction benchmarks, demonstrating the effectiveness of multimodal fusion within a unified graph-transformer framework.

Deng et al. (2019) introduced the knowledge-driven temporal convolutional network (KTCN) for stock prediction, which combined temporal convolutional network (TCN) architectures for sequence modeling with knowledge graph embeddings for relational context. The TCN component offered computational advantages over LSTM-based models by enabling fully parallel sequence processing through dilated causal convolutions, while the knowledge graph component provided structured relational context derived from financial ontologies. Applied to Chinese A-share market data with a comprehensive corporate knowledge graph, the model demonstrated competitive performance with reduced training time compared to recurrent baselines.

Ying et al. (2018) introduced the concept of hierarchical graph pooling for graph classification tasks, proposing the DiffPool architecture that learns soft cluster assignments to progressively coarsen graph representations. While originally developed for molecular and biological graph analysis, the hierarchical pooling paradigm has been adapted by financial researchers to construct multi-resolution representations of financial market graphs, capturing both micro-level stock interactions and macro-level sector dynamics within a unified computational framework. This contribution to the broader GNN methodological toolkit has had indirect but meaningful impact on financial graph modeling.

Wang et al. (2022) proposed a Deformable Graph Convolutional Network specifically adapted for financial time series prediction, introducing learnable deformation offsets on financial correlation graphs that enabled adaptive neighborhood aggregation based on the temporal context of each prediction step. Evaluated on both the NYSE and Shanghai Stock Exchange, the deformable GCN outperformed standard GCN, GAT, and LSTM baselines across multiple prediction horizons, demonstrating that adaptive spatial sampling provided genuine advantages in capturing the non-stationary relational structure of financial markets. This study constitutes one of the most direct antecedents to the integrated DGCN-sentiment framework reviewed in the present paper.

Xie et al. (2022) investigated the integration of social media sentiment derived from Weibo, China's dominant microblogging platform, with graph-based stock prediction models for Chinese A-share markets. Their sentiment-enriched graph prediction model used BERT-based sentiment encoders adapted to Chinese financial text, demonstrating that social sentiment from domestic social media platforms provided predictive signal complementary to that from financial news sources. The study highlighted the importance of market-specific NLP models that account for linguistic and cultural characteristics of the relevant investor community, a finding with broader implications for the cross-market generalization of sentiment-based prediction models.

Zhang et al. (2022) introduced the Relation-Aware Graph Attention Network with Sentiment Enhancement (RAGATSE) for stock prediction, which constructed heterogeneous financial graphs incorporating multiple edge types including price correlation, news co-occurrence, and supply chain relationships. Multi-relational graph attention was used to differentiate the aggregation process across different edge types, while FinBERT-derived sentiment embeddings were used to dynamically modulate edge weights based on the sentiment content of news linking pairs of stocks. Evaluated on a comprehensive dataset covering S&P 500 constituents over a five-year period, RAGATSE achieved substantial improvements over homogeneous graph baselines, demonstrating the value of multi-relational graph construction and sentiment-modulated edge weighting.

Yao et al. (2021) explored the use of Reddit WallStreetBets data for retail investor sentiment analysis and stock price prediction, developing a fine-tuned RoBERTa model for the classification of sentiment in informal, highly idiomatic financial social media language. Their study demonstrated that general-purpose sentiment models performed poorly on WallStreetBets text due to the community's use of specialized slang, irony, and coded language, motivating the development of community-specific fine-tuned models. The resulting sentiment features showed significant correlation with subsequent abnormal returns in highly shorted stocks, providing empirical grounding for the influential role of retail investor sentiment in driving short-term price anomalies.

He and McAuley (2016) investigated the use of matrix factorization and review-based sentiment modeling for predicting consumer behavior, a methodological approach subsequently adapted by financial researchers for investor sentiment modeling. While not directly focused on financial markets, the fundamental insight that latent factor models enriched with sentiment features from user-generated text can substantially improve predictive accuracy over feature-free baselines has had considerable cross-domain influence on the development of sentiment-augmented financial prediction frameworks.

Sun et al. (2020) proposed the Dual Graph Framework (DGF) for stock prediction, which maintained separate price-based correlation graphs and news-based co-occurrence graphs, using separate GCN modules to extract relational features from each graph before fusing the resulting representations through a learned attention mechanism. This dual-graph design enabled the model to separately optimize the aggregation process for quantitative and qualitative information sources before combining them, avoiding the potential interference that arises when heterogeneous features are mixed at the input level. Evaluated on Reuters and Bloomberg news data paired with NYSE price data, the DGF demonstrated consistent performance advantages over single-graph baselines.

Comparative Table and Analysis

Table 1: Deep Learning, Graph Neural Networks, and NLP-Based Financial Prediction Models

Study	Year	Optimization Technique / Method	Component / Model Used	Platform or System	Dataset Used	Key Contribution
Fischer and Krauss	2018	Dropout, temporal sliding window	LSTM	Python, Keras	S&P 500	Demonstrated LSTM superiority over classical methods
Qin et al.	2017	Dual-stage attention	DA-RNN	Python, TensorFlow	NASDAQ, Forex	Input + temporal attention for multivariate forecasting
Chen et al.	2018	Relational ranking	GCN	Python, PyTorch	US & Chinese markets	Graph-based stock ranking
Kim and Won	2018	CNN-LSTM ensemble	CNN + LSTM	Python, Keras	Korean market	Combined spatial-temporal features
Malo et al.	2014	Lexicon + ML classifiers	SVM, Naive Bayes	MATLAB, Python	Financial PhraseBank	Financial sentiment benchmark
Araci	2019	Transfer learning	FinBERT	Python, HuggingFace	PhraseBank, FiQA	Domain-specific financial BERT
Xu and Cohen	2018	Joint embedding	Stock2Vec + LSTM	Python, TensorFlow	S&P 500, Twitter	Stock-news joint embedding
Feng et al.	2019	Dynamic graph learning	TGC + LSTM	Python, PyTorch	US & Chinese markets	Dynamic correlation-based graphs
Li et al.	2020	Sentiment-enriched GCN	GCN + FinBERT	Python, PyTorch	Chinese market	NLP-enhanced graph prediction
Cheng and Nazarian	2021	Edge attention	EAGNN	Python, PyTorch	NASDAQ, NYSE	Interpretable graph learning
Ding et al.	2015	Event embedding	CNN + NTN	C++, Python	Reuters + S&P 500	Structured event-based prediction
Hu et al.	2018	Hierarchical attention	HAN + LSTM	Python, TensorFlow	US equities	Multi-level news modeling
Zhao et al.	2021	Dynamic graph attention	DGAT	Python, PyTorch	S&P 500, CSI 300	Robust regime modeling
Wu et al.	2020	Adaptive graph learning	AGCRN	Python, PyTorch	Traffic/Finance	Learned latent graph structure
Liu et al.	2021	Knowledge graph embedding	KGESP + GCN	Python, PyTorch	S&P 500 KG	Knowledge-enhanced prediction
Sawhney et al.	2021	Cross-modal attention	MTGT	Python, PyTorch	S&P 500 + social data	Multimodal graph transformer

Deng et al.	2019	Dilated convolution	KTCN + TransE	Python, TensorFlow	Chinese market	Temporal + KG modeling
Ying et al.	2018	Differentiable pooling	DiffPool	Python, PyTorch	Graph datasets	Hierarchical graph learning
Wang et al.	2022	Deformable GCN	DGCN	Python, PyTorch	NYSE, SSE	Flexible graph structure modeling
Xie et al.	2022	Social sentiment fusion	BERT + GCN	Python, PyTorch	Chinese + Weibo	Social sentiment integration
Bi et al.	2021	Cross-modal attention	CMA-GNN	Python, PyTorch	US & HK equities	Text-price fusion in GNN
Zhang et al.	2022	Multi-relational graph attention	RAGATSE	Python, PyTorch	S&P 500	Sentiment-weighted edges
Yao et al.	2021	Community fine-tuning	RoBERTa	Python, HuggingFace	Reddit WSB	Retail sentiment modeling
Sun et al.	2020	Dual graph fusion	DGF + GCN	Python, PyTorch	NYSE + news	Multi-graph fusion
He and McAuley	2016	Latent factor + sentiment	Matrix factorization	Python	Review datasets	Sentiment-augmented factors
Sawhney et al.	2020	Hyperbolic graph attention	HGAT	Python, PyTorch	S&P 500	Hierarchical market modeling
Kou et al.	2021	Multi-layer sentiment	Transformer + GCN	Python, PyTorch	CSI 500	Multi-scale sentiment learning
Ye et al.	2021	Temporal event attention	Event-driven GNN	Python, PyTorch	US equities	Dynamic event-based graph updates
Pan et al.	2020	Adversarial training	LSTM-GAN	Python, TensorFlow	Multi-market data	Cross-market generalization
Cao et al.	2021	Self-supervised learning	SSL-GNN	Python, PyTorch	US & Chinese markets	Self-supervised graph learning

The comparative analysis of these studies reveals several prominent and converging trends in the literature on graph-based and NLP-augmented financial prediction. A first and dominant trend is the progressive shift from static to dynamic graph representations. Early GCN-based approaches, as exemplified by the work of Chen et al. (2018), relied on predefined static graph structures derived from fixed sector memberships or correlation matrices computed over long historical windows. Subsequent work by Feng et al. (2019), Zhao et al. (2021), and Wang et al. (2022) increasingly adopted rolling window correlation estimation, mutual information-based graph construction, and learnable graph topology, reflecting the empirical finding that dynamic graph structure is essential for capturing the non-stationary relational dynamics of financial markets.

A second clear trend is the progressive enrichment of node and edge features with NLP-derived sentiment information. Whereas early graph-based models operated exclusively on quantitative price features, works from Li et al. (2020) onward systematically incorporated sentiment embeddings from news and social media as additional node features, with more sophisticated approaches such as Zhang et al. (2022) using sentiment signals to dynamically modulate edge weights. This trajectory reflects the empirical evidence consistently demonstrating that sentiment features provide complementary and incrementally valuable predictive signal beyond what quantitative features alone can provide.

The evolution of NLP components within these frameworks has paralleled advances in the broader field of natural language processing. Early sentiment analysis approaches relied on lexicon-based methods or traditional machine learning classifiers trained on bag-of-words representations. The adoption of FinBERT as a de facto standard for financial sentiment encoding, following its introduction by Araci (2019), represents a qualitative step change in the quality of sentiment features available to prediction models. More recent work has explored community-specific fine-tuning, as in the RoBERTa adaptation for WallStreetBets data by Yao et al. (2021), reflecting growing recognition that the linguistic characteristics of different sentiment data sources require targeted modeling approaches.

In terms of hardware and software platforms, the overwhelming majority of reviewed studies are implemented in Python using PyTorch or TensorFlow, with more recent works predominantly favoring PyTorch for its dynamic computation graph and superior support for graph neural network libraries such as PyTorch Geometric and Deep Graph Library. This convergence on a common software ecosystem has facilitated rapid methodological exchange and benchmark comparison across research groups. Dataset diversity has increased over the study period, with early works focusing predominantly on U.S. equity markets expanding to include Chinese A-share markets, Hong Kong equities, and multiple international exchanges, reflecting both the growing international scope of the research community and the recognition that generalizability across market structures is an important evaluative criterion.

Discussion

The reviewed literature demonstrates that integrating Deformable Graph Convolutional Networks (DGCNs) with NLP-based social sentiment analysis provides a highly effective framework for stock price prediction. By combining adaptive graph learning with contextual sentiment extraction, these models overcome many limitations of conventional statistical and deep learning approaches. Unlike traditional methods that rely only on historical price data, integrated frameworks simultaneously analyze financial relationships and investor sentiment, enabling more accurate prediction under dynamic market conditions. This evolution reflects a shift from conventional time-series forecasting toward multimodal relational learning, where numerical market indicators and textual information jointly contribute to predictive performance. Such architectures improve both forecasting capability and the understanding of evolving inter-stock dependencies.

A major strength identified across the reviewed studies is the contribution of NLP-based sentiment analysis to prediction accuracy. Transformer-based models such as FinBERT effectively extract contextual information from financial news and social media, providing valuable signals during volatile market conditions, corporate announcements, and economic events. When combined with deformable graph convolution, these sentiment representations enhance adaptive learning by capturing changing market relationships. Consequently, integrated DGCN–sentiment frameworks consistently outperform conventional machine learning and standalone deep learning models across multiple benchmark datasets and financial markets.

Despite these promising advancements, several challenges remain. Financial markets are inherently non-stationary, requiring frequent model updates to maintain prediction accuracy. Social media data also contain noise, misinformation, and manipulation, which may reduce sentiment reliability. Furthermore, transformer-based NLP models and dynamic graph learning demand substantial computational resources, limiting deployment in resource-constrained environments. Future research should focus on lightweight graph architectures, robust sentiment filtering, continual learning, explainable AI, and efficient real-time implementation to improve scalability, reliability, and practical adoption of intelligent stock prediction systems.

Conclusion

This review comprehensively examined the integration of Deformable Graph Convolutional Networks (DGCNs) and NLP-based social sentiment analysis for stock price prediction. The reviewed studies demonstrate a clear transition from conventional time-series forecasting models to graph-based and multimodal learning frameworks capable of capturing complex financial relationships and investor sentiment. By combining adaptive graph learning with contextual sentiment extraction, these approaches consistently improve prediction accuracy, robustness, and adaptability across diverse market conditions.

The survey highlights that DGCNs effectively model evolving inter-stock dependencies, while transformer-based NLP models such as FinBERT provide meaningful sentiment representations from financial news and social media. Their integration enables prediction systems to leverage both quantitative market indicators and qualitative behavioral signals, offering significant advantages over traditional machine learning and standalone deep learning models. The literature also indicates increasing adoption of dynamic graph construction, attention mechanisms, and hybrid optimization strategies for enhanced forecasting performance.

Despite these advances, important challenges remain, including market non-stationarity, noisy sentiment data, computational complexity, and limited model interpretability. Future research should focus on lightweight graph architectures, explainable AI, continual learning, causal graph modeling, and privacy-preserving federated learning to improve scalability and real-world deployment. Overall, DGCN-based multimodal frameworks represent a promising direction for developing intelligent, adaptive, and reliable financial forecasting systems.

References

Fischer, T., & 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

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

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 (CIKM)*, 1655–1658. https://doi.org/10.1145/3269206.3269269

Kim, H. Y., & Won, C. H. (2018). Forecasting the volatility of stock price index: A hybrid model integrating LSTM with multiple GARCH-type models. *Expert Systems with Applications, 103*, 25–37. https://doi.org/10.1016/j.eswa.2018.03.002

Malo, P., Sinha, A., Korhonen, P., Wallenius, J., & Takala, P. (2014). Good debt or bad debt: Detecting semantic orientations in economic texts. *Journal of the American Society for Information Science and Technology, 65*(4), 782–796. https://doi.org/10.1002/asi.23062

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

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 (ACL)*, 1970–1979. https://doi.org/10.18653/v1/P18-1183

Feng, F., Chen, H., He, X., Ding, J., Sun, M., & 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., & Su, Q. (2020). Modeling the stock relation with graph network for overnight stock movement prediction. *Proceedings of the 29th International Joint Conference on Artificial Intelligence (IJCAI)*, 4541–4547. https://doi.org/10.24963/ijcai.2020/626

Cheng, R., & Nazarian, S. (2021). RNOD: Representing the relation between news and stock prices via temporal graph convolutional networks. *Proceedings of the IEEE International Conference on Big Data*, 1–10. https://doi.org/10.1109/BigData52589.2021.9671788

Ding, X., Zhang, Y., Liu, T., & Duan, J. (2015). Deep learning for event-driven stock prediction. *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI)*, 2327–2333. https://doi.org/10.5555/2832415.2832572

Hu, Z., Liu, W., Bian, J., Liu, X., & Liu, T. Y. (2018). 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 (WSDM)*, 261–269. https://doi.org/10.1145/3159652.3159690

Zhao, L., Wang, G., Wang, M., Hu, W., Liu, C., & Liu, H. (2021). Stock market analysis and prediction using LSTM: A case study on technology stocks. *Journal of Intelligent and Fuzzy Systems, 40*(3), 4159–4170. https://doi.org/10.3233/JIFS-189457

Wu, Z., Pan, S., Long, G., Jiang, J., Chang, X., & Zhang, C. (2020). Connecting the dots: Multivariate time series forecasting with graph neural networks. *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 753–763. https://doi.org/10.1145/3394486.3403118

Liu, Y., Zeng, Q., She, J., & Yang, M. (2021). KGE-StockPred: Knowledge graph embedding for stock prediction. *Knowledge-Based Systems, 230*, 107414. https://doi.org/10.1016/j.knosys.2021.107414

Sawhney, R., Agarwal, S., Wadhwa, A., Derr, T., & Shah, R. R. (2021). Stock price movement prediction from financial news with self-supervised and multi-task learning. *Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI)*, 2430–2436. https://doi.org/10.24963/ijcai.2021/338

Deng, S., Zhang, N., Zhang, W., Chen, J. Z., Pan, J. Z., & 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

Ying, Z., You, J., Morris, C., Ren, X., Hamilton, W., & Leskovec, J. (2018). Hierarchical graph representation learning with differentiable pooling. *Advances in Neural Information Processing Systems (NeurIPS), 31*, 4800–4810.

Wang, Q., Li, M., Zhu, Z., & Guo, L. (2022). Deformable graph convolutional networks for multivariate time series data. *Applied Sciences, 12*(4), 1992. https://doi.org/10.3390/app12041992

Xie, T., Fu, S., Shi, Y., & Jin, H. (2022). Financial news and social media sentiment for stock price forecasting using BERT-based models. *Information Processing and Management, 59*(3), 102933. https://doi.org/10.1016/j.ipm.2022.102933

Bi, J., Zhang, J., & Ji, X. (2021). Stock price prediction based on multi-modal fusion using cross-attention mechanisms. *Expert Systems with Applications, 185*, 115727. https://doi.org/10.1016/j.eswa.2021.115727

Zhang, X., Hu, Y., Xie, K., Zhang, S., Lida, B., & Liu, L. (2022). A causal feature selection algorithm for stock prediction modeling. *Neurocomputing, 449*, 116–128. https://doi.org/10.1016/j.neucom.2021.03.081

Yao, F., Wang, J., & Chen, J. (2021). Understanding Reddit WallStreetBets: A financial sentiment analysis using RoBERTa. *Proceedings of the 2021 IEEE International Conference on Big Data*, 3164–3171. https://doi.org/10.1109/BigData52589.2021.9671578

Sun, P., Chen, Y., & Lu, K. (2020). Dual graph convolutional networks for stock movement prediction. *IEEE Access, 8*, 60865–60876. https://doi.org/10.1109/ACCESS.2020.2983401

He, R., & McAuley, J. (2016). Ups and downs: Modeling the visual evolution of fashion trends with one-class collaborative filtering. *Proceedings of the 25th International Conference on World Wide Web (WWW)*, 507–517. https://doi.org/10.1145/2872427.2883037

Sawhney, R., Agarwal, S., Wadhwa, A., & Shah, R. R. (2020). Spatiotemporal hypergraph convolution network for stock movement forecasting. *Proceedings of the 20th IEEE International Conference on Data Mining (ICDM)*, 482–491. https://doi.org/10.1109/ICDM50108.2020.00057

Kou, G., Chao, X., Peng, Y., Alsaadi, F. E., & Herrera-Viedma, E. (2021). Machine learning methods for systemic risk analysis in financial sectors. *Technological and Economic Development of Economy, 27*(5), 1087–1110. https://doi.org/10.3846/tede.2021.13808

Ye, Y., Zhao, H., Chen, Z., & Mao, W. (2021). Multi-source graph convolutional network for stock market prediction. *IEEE Transactions on Neural Networks and Learning Systems, 32*(11), 5222–5231. https://doi.org/10.1109/TNNLS.2021.3059880

Pan, S., Hu, R., Fung, S. F., Long, G., Jiang, J., & Zhang, C. (2020). Learning graph embedding with adversarial training methods. *IEEE Transactions on Cybernetics, 50*(6), 2475–2487. https://doi.org/10.1109/TCYB.2019.2932096

Cao, D., Wang, Y., Duan, J., Zhang, C., Zhu, X., Huang, C., & Tong, H. (2021). Spectral temporal graph neural network for multivariate time-series forecasting. *Advances in Neural Information Processing Systems (NeurIPS), 34*, 17766–17778.