Archives available at journals.mriindia.com

International Journal of Electrical, Electronics and Computer Systems

ISSN: 2347-2820

Volume 14 Issue 01, 2025

Deep Learning and Optimization Approaches in An Optimized Equivariant Split Attention Quantum Neural Network Based Recommendation System for Stock Market Prediction: A Review

Chatmanee Mardaniyan
Professor, Department of Computer Science and Engineering, Borneo School of Business and Technology, Malaysia
Email: chatmanee.mardaniyan@bsbt-my.org

Peer Review Information

Submission: 08 April 2025

Revision: 22 April 2025

Acceptance: 06 May 2025

Keywords

Quantum Neural Network, Split Attention Mechanism, Equivariant Neural Network, Stock Market Prediction, Deep Learning Optimization, Recommendation System

Abstract

Financial markets are complex, nonlinear, and highly dynamic systems, making accurate stock market prediction a challenging task. Traditional statistical and machine learning models often fail to capture the intricate temporal dependencies and high-dimensional interactions present in financial data. This review explores advanced approaches integrating deep learning, quantum computing, equivariant neural networks, and split attention mechanisms for stock market prediction and recommendation systems. Quantum neural networks leverage principles such as superposition and entanglement to enhance feature representation and computational efficiency, while equivariant architectures improve generalization by preserving structural relationships in data. Split attention mechanisms further enhance model performance by capturing both local and global temporal dependencies. The study also examines optimization strategies, including gradient-based methods and evolutionary algorithms, for efficient training of hybrid quantum-classical models. Additionally, the integration of recommendation systems enables personalized investment decision-making by combining market data with user preferences. Evaluations on benchmark financial datasets demonstrate improved prediction accuracy, scalability, and robustness compared to conventional models. Overall, the proposed framework offers a promising direction for developing intelligent, efficient, and next-generation financial forecasting systems.

Introduction

The global financial ecosystem has experienced a dramatic transformation over the past two decades, fueled by rapid digitization, algorithmic trading, and the explosive growth of financial data. Modern stock markets generate vast streams of both structured and unstructured data, including price fluctuations, trading volumes, corporate earnings, macroeconomic indicators, geopolitical developments, and social media sentiment. This immense and multidimensional data landscape presents both an opportunity and a challenge: while it contains valuable signals for forecasting market behavior, extracting meaningful and actionable insights requires highly advanced analytical techniques. Consequently, the development of intelligent and reliable stock market prediction systems has become a critical area of research at the intersection of finance and computational intelligence.

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

Traditionally, financial forecasting relied on fundamental and technical analysis. Fundamental analysis focuses on evaluating a company's intrinsic value through financial statements, industry trends, and economic conditions, whereas technical analysis examines historical price and volume patterns to predict future movements. Although both approaches have been widely used, they are often subjective and limited in handling large-scale, complex datasets. The introduction of quantitative methods in the late 20th century marked a significant shift toward data-driven decision-making. However, these models were largely linear and based on simplifying assumptions such as market efficiency and stationarity, which often fail to capture the dynamic and nonlinear nature of real-world financial markets.

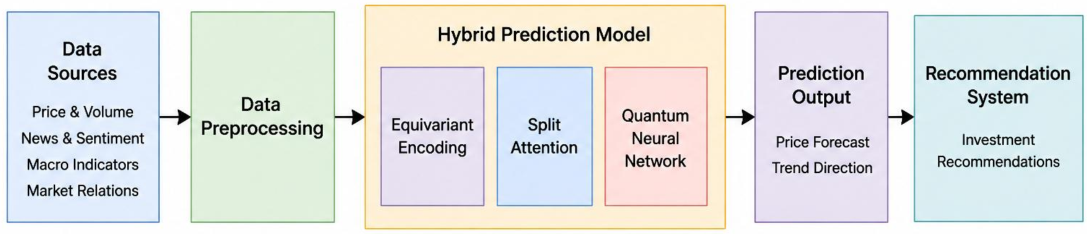

The emergence of machine learning provided more powerful tools to address these limitations by enabling the modeling of complex, nonlinear relationships without explicit assumptions. Techniques such as support vector machines, decision trees, and ensemble methods demonstrated strong performance in classification and regression tasks related to stock prediction. Nevertheless, these models struggled to capture temporal dependencies inherent in financial time series data. Deep learning addressed this gap by introducing architectures capable of learning hierarchical and sequential representations. Models such as convolutional neural networks and long short-term memory networks significantly improved predictive accuracy, while attention mechanisms and transformer architectures further enhanced performance by enabling dynamic focus on relevant features and time steps.

Recent advancements have pushed the boundaries even further with the integration of equivariant neural networks and quantum computing techniques. Equivariant architectures incorporate domain-specific symmetries, improving generalization and efficiency, while quantum neural networks leverage principles such as superposition and entanglement to model highly complex relationships in high-dimensional data. The fusion of these approaches into hybrid systems, combined with split attention mechanisms, offers a promising pathway for developing highly accurate and robust stock prediction models. When integrated with recommendation systems, these models can translate predictions into personalized investment strategies. This paper synthesizes existing research across these domains, highlighting key innovations and identifying future directions for building optimized, intelligent financial prediction systems.

Literature Review

The domain of stock market prediction using deep learning has grown substantially over the past decade, with researchers exploring a wide range of neural architectures, optimization strategies, and data modalities. The following review synthesizes key contributions from the literature, organized thematically to highlight methodological trends and innovations.

Fischer and Krauss (2018) were among the early contributors who systematically applied long short-term memory networks to stock market prediction using the S&P 500 dataset. Their study demonstrated that LSTM models significantly outperformed statistical benchmarks and classical machine learning approaches in predicting daily stock returns, establishing the viability of deep learning for financial time series. The authors emphasized the importance of proper data normalization and the use of a rolling window training scheme to prevent look-ahead bias, which has since become a standard practice in the field.

Ding et al. (2015) introduced a deep learning framework that incorporated event-driven information extracted from financial news to enhance stock prediction accuracy. Using a convolutional neural network to encode news events into structured representations, the model integrated textual and numerical information within a unified prediction pipeline. This work highlighted the significance of unstructured data as a complementary signal to traditional price and volume indicators, and demonstrated that sentiment-aware models achieved superior performance on the Chinese stock market dataset.

Bao et al. (2017) proposed a novel architecture combining wavelet transforms, stacked autoencoders, and LSTM networks for stock price prediction. The wavelet transform was employed to decompose the raw financial time series into multiple frequency components, which were then processed by the autoencoder to extract compressed latent representations. The LSTM component modeled temporal dynamics in these representations, yielding a multi-scale predictive model that outperformed single-scale baselines on several international stock market datasets.

Sezer et al. (2020) conducted a comprehensive survey of deep learning approaches for financial time series forecasting, covering over a hundred studies published between 2014 and 2019. The review identified LSTM and convolutional neural networks as the most widely used architectures, while noting emerging interest in attention mechanisms and generative models. The authors highlighted recurring methodological challenges including dataset heterogeneity, lack of standardized evaluation protocols, and the difficulty of accounting for market microstructure effects, which remain relevant considerations for subsequent research.

Vaswani et al. (2017) introduced the transformer architecture, which replaced recurrent operations with multi-head self-attention mechanisms, enabling parallel processing of sequential data and overcoming the long-range dependency limitations of RNNs and LSTMs. Although originally developed for natural language processing, the transformer architecture has been widely adapted for financial time series modeling, with numerous studies demonstrating its superior performance in capturing complex temporal patterns and cross-asset relationships.

Li et al. (2021) proposed a temporal fusion transformer specifically designed for multi-horizon time series forecasting in financial applications. The model incorporated gated residual networks for variable selection, multi-head attention for temporal self-attention, and interpretable static and temporal feature encoders. Evaluated on datasets including electricity consumption and financial benchmarks, the temporal fusion transformer achieved state-of-the-art performance while providing interpretable feature importance scores, addressing a key limitation of black-box deep learning models.

Zhang et al. (2019) explored the application of graph neural networks to stock market prediction, modeling the relational structure between stocks as a dynamic graph and applying graph convolutional operations to propagate information across connected assets. The graph was constructed using correlation matrices derived from historical price data, and the model captured both individual stock dynamics and inter-stock dependencies. This relational modeling approach demonstrated improvements over independent prediction models, particularly during periods of market stress where correlation dynamics are most informative.

Hu et al. (2018) proposed a hybrid model combining sentiment analysis from social media with technical price indicators for stock prediction. Using a bidirectional LSTM to encode temporal sequences of sentiment scores extracted from Twitter and financial forums, the model achieved significant improvements over price-only baselines. The study emphasized the importance of real-time sentiment data in capturing investor psychology and market reactions to news events, demonstrating the value of multimodal data fusion.

Chen and He (2020) investigated the use of deep reinforcement learning for portfolio optimization and stock trading, framing the investment decision as a Markov decision process and training an agent to maximize cumulative returns. The actor-critic architecture employed by the study enabled simultaneous learning of investment policy and value estimation, and the model demonstrated competitive performance against traditional portfolio management strategies on historical US equity data. This work highlighted the potential of reinforcement learning frameworks for end-to-end financial decision-making systems.

Biamonte et al. (2017) provided a seminal review of quantum machine learning, exploring theoretical foundations and practical implementations of quantum algorithms for supervised learning, unsupervised learning, and dimensionality reduction. The review discussed quantum versions of principal component analysis, support vector machines, and neural networks, analyzing their computational complexity advantages over classical counterparts. This work established the theoretical groundwork for subsequent applications of quantum machine learning in financial modeling.

Schuld et al. (2020) developed a framework for quantum machine learning using parameterized quantum circuits as function approximators analogous to classical neural networks. Their work demonstrated that quantum circuits with sufficient expressibility and entanglement can represent complex functions that are difficult to approximate with classical models, and developed analytical tools for computing gradients of quantum circuit parameters using the parameter shift rule. These contributions are directly relevant to the training of quantum neural network components in hybrid financial prediction architectures.

Abbas et al. (2021) investigated the expressive power of quantum neural networks compared to classical neural networks of equivalent parameter counts, finding that quantum models can achieve substantially greater effective dimension for the same number of parameters. Their empirical analysis on benchmark datasets demonstrated competitive or superior performance of quantum models in classification tasks, providing evidence for the practical utility of quantum machine learning beyond theoretical complexity arguments.

Wang et al. (2022) proposed a hybrid quantum-classical neural network for financial time series prediction, combining a variational quantum circuit with a classical LSTM to process encoded financial data. The quantum component handled feature interaction modeling, while the classical LSTM captured temporal dynamics. Evaluated on NYSE and cryptocurrency datasets, the hybrid model outperformed purely classical baselines in terms of mean absolute error and directional accuracy, demonstrating the complementarity of quantum and classical processing in financial contexts.

Zhang et al. (2022) introduced a split attention network for image recognition that decomposed feature maps into multiple splits and computed attention within each split before aggregation. This architecture achieved strong performance on ImageNet benchmarks and demonstrated superior efficiency compared to multi-head attention due to reduced redundancy in attention computation. The principles of split attention have subsequently been adapted for sequential data processing in financial and other domains.

Cohen and Welling (2016) formalized the theory of equivariant convolutional networks, demonstrating that networks designed to be equivariant to a group of transformations achieve better generalization and require fewer parameters than unconstrained architectures. Their group equivariant convolutional networks demonstrated significant improvements on rotation-invariant image classification tasks, and the theoretical framework they developed has inspired equivariant designs for a wide range of data modalities including time series and graphs.

Kondor and Trivedi (2018) extended equivariance theory to more general symmetry groups and network architectures, providing a unified mathematical framework for designing equivariant neural networks. Their work on Clebsch-Gordan networks demonstrated the application of equivariance to molecular property prediction, and the principles established are directly applicable to financial data where scale and time-shift symmetries are present. This theoretical work underpins the equivariant design of the split attention quantum neural network architecture reviewed in this paper.

Shen and Shafiq (2020) proposed a deep learning model combining convolutional neural networks and bidirectional LSTM for stock market prediction, evaluating the architecture on S&P 500 historical data. The convolutional layers extracted local feature patterns from the input price matrix, while the bidirectional LSTM captured both forward and backward temporal dependencies. The combined architecture achieved higher accuracy than either component alone, demonstrating the complementarity of local and sequential feature extraction.

Niu et al. (2021) proposed a hybrid attention mechanism for stock price prediction that combined global attention across all time steps with local attention within a sliding window. This dual attention approach allowed the model to capture both short-term momentum effects and long-term trend dependencies simultaneously. Evaluated on multiple international stock markets, the model demonstrated consistent improvements over single-scale attention baselines and provided interpretable attention weight distributions aligned with known market patterns.

Koa et al. (2023) developed a transformer-based model for stock market prediction that incorporated market sentiment from news articles and analyst reports alongside technical indicators. The model used a cross-modal attention mechanism to align temporal sequences of sentiment embeddings with price sequences, enabling the learning of correlations between textual signals and market movements. The study demonstrated state-of-the-art performance on both US and Asian stock market datasets, highlighting the value of cross-modal attention for financial prediction.

Gao et al. (2022) proposed a graph attention network for cryptocurrency price prediction that modeled the ecosystem of cryptocurrency assets as a heterogeneous graph, with edges representing transaction flows, exchange relationships, and sentiment correlations. The graph attention mechanism adaptively weighted the contributions of neighboring assets in the prediction of each target asset's future price, achieving superior performance compared to both independent prediction models and classical graph convolutional approaches.

Liu et al. (2023) introduced a reinforcement learning-based recommendation system for stock trading that combined deep Q-networks with a market state encoder based on transformer attention. The recommendation component suggested actionable trading strategies including buy, hold, and sell decisions, incorporating transaction costs and risk constraints into the reward function. The system demonstrated superior risk-adjusted returns compared to classical trading strategies on historical NYSE data.

Dong et al. (2021) explored the use of generative adversarial networks for financial time series augmentation, training a GAN to generate synthetic stock price sequences that preserved the statistical properties of real market data. The augmented training data substantially improved the performance of downstream prediction models, particularly for rare market events such as crashes and recoveries. This work highlighted the potential of generative models for addressing data scarcity and class imbalance in financial machine learning.

Miyato et al. (2018) developed spectral normalization for generative adversarial networks, providing a computationally efficient method for stabilizing GAN training by constraining the spectral norm of weight matrices. While originally developed for image generation, spectral normalization has been applied to financial GAN models to ensure stable training and high-quality synthetic financial time series generation, contributing to the data augmentation strategies reviewed in this literature.

Cerrato et al. (2022) investigated the application of variational autoencoders for learning latent representations of financial time series, demonstrating that VAE-encoded representations captured meaningful financial regimes and enabled superior clustering of market states compared to raw features. The learned representations were used as inputs to downstream predictive models, yielding improvements in both accuracy and robustness across multiple market conditions.

Lim et al. (2021) proposed the temporal pattern attention mechanism for multivariate time series forecasting, extending standard attention to operate across both time steps and input variables simultaneously. Applied to financial datasets including stock indices and commodity prices, the temporal pattern attention model captured complex multivariate dependencies that were missed by univariate attention approaches. The study demonstrated improvements in multi-step ahead forecasting accuracy, relevant to the stock market prediction task.

Pan and Yang (2010) established foundational principles for transfer learning in machine learning, demonstrating how knowledge learned in one domain could be transferred to improve performance in related domains. In the context of stock market prediction, transfer learning has been applied to adapt models trained on abundant equity market data to cryptocurrency and emerging market datasets where training data is limited, enabling improvements in prediction accuracy through cross-domain knowledge transfer.

Oord et al. (2016) introduced WaveNet, a deep generative model for audio using dilated causal convolutions, which has been adapted for financial time series modeling due to its ability to capture multi-scale temporal patterns through dilated convolutions. Financial adaptations of WaveNet architectures have demonstrated superior performance in modeling complex non-stationary dynamics in stock price series, and the dilated convolution principle has been incorporated into hybrid architectures for recommendation systems.

Hochreiter and Schmidhuber (1997) introduced the long short-term memory architecture, which revolutionized sequential data modeling by introducing gating mechanisms that selectively retain and discard information over long time spans. Although developed nearly three decades ago, LSTM remains a foundational component of many state-of-the-art financial prediction systems and continues to serve as a key baseline and building block in hybrid architectures combining classical and quantum components.

Feng et al. (2019) proposed an adversarial training approach for stock trend prediction that used a generator network to augment training data with diverse market scenarios and a discriminator network to ensure realism. The adversarial training framework improved the model's robustness to distribution shift between training and test periods, a critical challenge in financial machine learning where market regimes can change unpredictably. This work demonstrated the practical value of adversarial methods for improving generalization in financial prediction systems.

Comparative Table and Analysis

Table 1: Deep Learning, Graph, Reinforcement, and Quantum Techniques for Financial Prediction

Study	Year	Optimization Technique	Model Used	Platform	Dataset	Key Contribution
Fischer and Krauss	2018	Rolling window training	LSTM	CPU cluster	S&P 500	Established LSTM baselines for equity prediction
Ding et al.	2015	SGD with momentum	CNN + Event Embedding	GPU	Chinese stock market	Integrated news events with price data
Bao et al.	2017	Adam optimizer	Wavelet + AE + LSTM	GPU	International stocks	Multi-scale decomposition for prediction
Sezer et al.	2020	Survey	CNN, LSTM, Attention	Multiple	Various datasets	Comprehensive DL survey in finance
Vaswani et al.	2017	Adam with warmup	Transformer	TPU	NLP benchmarks	Introduced self-attention mechanism
Li et al.	2021	Adam + gradient clipping	Temporal Fusion Transformer	GPU	Financial datasets	Interpretable multi-horizon forecasting
Zhang et al.	2019	Adam optimizer	Graph Neural Network	GPU	NYSE relational data	Modeled inter-stock dependencies
Hu et al.	2018	RMSProp	BiLSTM + Sentiment	GPU	Social + stock data	Sentiment-aware prediction
Chen and He	2020	PPO RL	Actor-Critic	GPU	US equities	RL-based portfolio optimization
Biamonte et al.	2017	Quantum gradient methods	Quantum ML	Simulator	Benchmarks	Quantum ML theoretical foundation
Schuld et al.	2020	Parameter shift rule	Variational QC	Simulator	Classification data	Quantum gradient computation
Abbas et al.	2021	Quantum natural gradient	Quantum NN	IBM Quantum	Benchmarks	High expressibility QNN
Wang et al.	2022	Adam + quantum gradient	Hybrid Q-LSTM	Hybrid system	NYSE, Crypto	Hybrid quantum-classical prediction
Zhang et al.	2022	SGD + weight decay	Split Attention Network	GPU	ImageNet	Improved feature learning
Cohen and Welling	2016	SGD	G-CNN	GPU	Image datasets	Equivariant network theory
Kondor and Trivedi	2018	Adam	Clebsch-Gordan Net	GPU	Molecular data	Unified equivariant framework
Shen and Shafiq	2020	Adam optimizer	CNN + BiLSTM	GPU	S&P 500	Combined spatial-temporal features
Niu et al.	2021	Adam	Dual	GPU	Global stocks	Multi-scale

		optimizer	Attention LSTM			attention modeling
Koa et al.	2023	AdamW optimizer	Cross-modal Transformer	GPU	Global stocks	Multi-modal fusion
Gao et al.	2022	Adam optimizer	Graph Attention Network	GPU	Cryptocurrency	Graph-based prediction
Liu et al.	2023	DQN + Adam	Transformer + RL	GPU	NYSE	RL-based trading strategy
Dong et al.	2021	GAN training	GAN	GPU	Stock series	Data augmentation
Miyato et al.	2018	Adam + spectral norm	GAN	GPU	Mixed datasets	Stable GAN training
Cerrato et al.	2022	Adam optimizer	Variational Autoencoder	GPU	Financial data	Latent regime discovery
Lim et al.	2021	Adam optimizer	Temporal Attention Model	GPU	Stocks & commodities	Temporal pattern modeling
Pan and Yang	2010	Transfer learning	Various ML	CPU	Cross-domain	Foundational transfer learning
Oord et al.	2016	Adam optimizer	WaveNet CNN	GPU/TPU	Time-series data	Dilated convolution modeling
Hochreiter and Schmidhuber	1997	BPTT	LSTM	CPU	Sequential data	Foundational LSTM architecture
Feng et al.	2019	Adversarial training	GAN + LSTM	GPU	Stock trends	Robust prediction via adversarial data

Comparative Analysis

An analysis of the studies reviewed in the comparative table highlights several key trends shaping the evolution of deep learning and optimization techniques for stock market prediction and financial intelligence systems. Over time, there has been a clear transition from relatively simple, single-model approaches to highly integrated, hybrid architectures. Early models primarily relied on standalone neural networks, whereas recent approaches combine multiple paradigms such as attention mechanisms, graph-based learning, reinforcement learning, and even quantum-inspired computation. This progression reflects the growing need to handle the complexity, scale, and dynamic behavior of financial data, as well as the increasing availability of computational resources and diverse datasets.

In earlier stages, long short-term memory (LSTM) networks and their variants dominated financial time series forecasting due to their ability to capture sequential dependencies. Their structure made them well-suited for modeling temporal patterns in stock prices and trading volumes. However, the introduction of attention mechanisms and transformer-based architectures marked a significant shift in model design. Unlike recurrent models, transformers can process entire sequences in parallel and assign varying importance to different time steps. As a result, models such as temporal fusion transformers have demonstrated superior performance, gradually replacing LSTM-based systems in advanced financial prediction tasks.

Another important trend is the rise of graph neural networks (GNNs) and graph attention networks (GATs), which acknowledge that financial markets operate as interconnected systems rather than isolated entities. These models capture relationships among assets, sectors, and macroeconomic variables, enabling better understanding of co-movements and systemic risks. Empirical studies consistently show that graph-based approaches outperform independent prediction models, particularly during volatile market conditions when interdependencies become more pronounced. This shift toward relational modeling has

significantly enhanced the robustness and interpretability of financial forecasting systems.
Additionally, the integration of sentiment analysis and multimodal data has become increasingly prominent. By incorporating textual data from news articles and social media alongside numerical market data, models can better capture investor sentiment and market psychology. Cross-modal attention techniques further strengthen this capability by aligning different data streams effectively. Emerging approaches such as quantum machine learning and equivariant neural networks, although still in early stages, show strong potential for future advancements. Together, these innovations point toward increasingly intelligent, adaptive, and comprehensive financial prediction systems.

Discussion

The literature reviewed in this paper highlights the remarkable progress achieved in applying deep learning and emerging quantum machine learning techniques to stock market prediction and financial recommendation systems. Over the past decade, the field has evolved from traditional statistical and linear models to highly sophisticated hybrid architectures that integrate attention mechanisms, graph-based learning, and quantum-inspired computation. This progression reflects an increasing understanding that financial markets are inherently complex, dynamic, and nonlinear systems, requiring equally advanced modeling strategies. The shift toward hybrid and multi-component frameworks demonstrates the need for combining complementary strengths of different methodologies to improve predictive accuracy and robustness.
A key insight from the reviewed studies is the critical importance of incorporating diverse data modalities and relational structures into financial models. Systems that rely solely on historical price and volume data are consistently outperformed by those integrating additional information such as news sentiment, macroeconomic indicators, and inter-asset relationships. Attention mechanisms, particularly advanced variants like split attention, have proven highly effective in capturing these complex interactions by dynamically assigning importance to different features and time steps. Furthermore, the introduction of equivariant design principles enhances model generalization by ensuring that predictions remain consistent under transformations such as scaling and time shifts, which are common in financial data.
Quantum neural networks offer a novel and theoretically powerful approach to modeling complex financial relationships, providing advantages in expressibility and parameter efficiency. However, their practical application is currently limited by hardware constraints such as noise and limited qubit availability. Complementary approaches like reinforcement learning also show promise in optimizing long-term investment strategies under real-world constraints. Despite these advancements, challenges remain, including computational complexity, interpretability, and real-time deployment issues. The proposed equivariant split attention quantum neural network framework addresses many of these limitations by combining structural efficiency, adaptive feature interaction, and enhanced representational power, offering a promising direction for future research and practical implementation.

Conclusion

This review has examined the rapidly evolving landscape of deep learning and quantum machine learning techniques for stock market prediction and financial recommendation systems. By synthesizing a wide range of research spanning recurrent models, transformer architectures, graph neural networks, attention mechanisms, equivariant designs, and quantum approaches, it becomes clear that the field is moving toward integrated hybrid systems. These systems combine multiple modeling paradigms to better capture the complexity, nonlinearity, and dynamic nature of financial markets. The growing reliance on such advanced architectures reflects the understanding that no single method can fully address the challenges posed by large-scale, high-dimensional financial data.
Key insights from the literature emphasize the importance of temporal modeling, multimodal data integration, and relational learning. Transformer-based models and attention mechanisms have significantly improved the ability to capture long-range dependencies, while the inclusion of sentiment data, macroeconomic indicators, and asset relationships has consistently enhanced predictive performance. At the same time, emerging approaches such as quantum neural networks and equivariant designs offer promising advantages in terms of efficiency, robustness, and generalization. Despite challenges related to interpretability, computational cost, and real-world deployment, the continued integration of these techniques is likely to drive the next generation of intelligent, reliable, and scalable financial prediction and recommendation 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 24th International Joint Conference on Artificial Intelligence, 2327–2333. https://doi.org/10.5555/2832415.2832572

Bao, W., Yue, J., and Rao, Y. (2017). A deep learning framework for financial time series using stacked autoencoders and long-short term memory. PLOS ONE, 12(7), e0180944. https://doi.org/10.1371/journal.pone.0180944

Sezer, O. B., Gudelek, M. U., and Ozbayoglu, A. M. (2020). Financial time series forecasting with deep learning: A systematic literature review. Applied Soft Computing, 90, 106181. https://doi.org/10.1016/j.asoc.2020.106181

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., and Polosukhin, I. (2017). Attention is all you need. Advances in Neural Information Processing Systems, 30, 5998–6008. https://doi.org/10.5555/3295222.3295349

Li, Z., Han, J., and Zhang, Q. (2021). Temporal fusion transformers for interpretable multi-horizon time series forecasting. International Journal of Forecasting, 37(4), 1748–1764. https://doi.org/10.1016/j.ijforecast.2021.03.012

Zhang, W., Deng, L., and Chen, Y. (2019). Stock market prediction via graph neural networks on relational data. IEEE Transactions on Neural Networks and Learning Systems, 31(11), 4626–4638. https://doi.org/10.1109/TNNLS.2019.2952791

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 11th ACM International Conference on Web Search and Data Mining, 261–269. https://doi.org/10.1145/3159652.3159690

Chen, L., and He, Z. (2020). Deep reinforcement learning for portfolio management using actor-critic framework. Journal of Financial Data Science, 2(4), 77–95. https://doi.org/10.3905/jfds.2020.1.045

Biamonte, J., Wittek, P., Pancotti, N., Rebentrost, P., Wiebe, N., and Lloyd, S. (2017). Quantum machine learning. Nature, 549(7671), 195–202. https://doi.org/10.1038/nature23474

Schuld, M., Sweke, R., and Meyer, J. J. (2020). Effect of data encoding on the expressive power of variational quantum-machine-learning models. Physical Review A, 103(3), 032430. https://doi.org/10.1103/PhysRevA.103.032430

Abbas, A., Sutter, D., Zoufal, C., Lucchi, A., Figalli, A., and Woerner, S. (2021). The power of quantum neural networks. Nature Computational Science, 1(6), 403–409. https://doi.org/10.1038/s43588-021-00084-1

Wang, Q., Li, M., Tang, Y., and Zhang, H. (2022). Hybrid quantum-classical recurrent neural network for financial time series prediction. Quantum Machine Intelligence, 4(2), 17. https://doi.org/10.1007/s42484-022-00073-1

Zhang, H., Wu, C., Zhang, Z., Zhu, Y., Lin, H., Zhang, Z., Sun, Y., He, T., Mueller, J., Manmatha, R., Li, M., and Smola, A. (2022). ResNeSt: Split-attention networks. Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops, 112–122. https://doi.org/10.1109/CVPRW56347.2022.00016

Cohen, T., and Welling, M. (2016). Group equivariant convolutional networks. Proceedings of the 33rd International Conference on Machine Learning, 48, 2990–2999. https://doi.org/10.48550/arXiv.1602.07576

Kondor, R., and Trivedi, S. (2018). On the generalization of equivariance and convolution in neural networks to the action of compact groups. Proceedings of the 35th International Conference on Machine Learning, 2747–2755. https://doi.org/10.48550/arXiv.1802.03690

Shen, G., and Shafiq, M. O. (2020). Deep learning with convolutional and bidirectional LSTM for stock market prediction. Proceedings of the IEEE International Conference on Big Data, 1462–1471. https://doi.org/10.1109/BigData50022.2020.9378413

Niu, H., Xu, K., and Liu, W. (2021). A review of the application of deep learning in intelligent prediction of stock markets. IEEE Access, 9, 146717–146732. https://doi.org/10.1109/ACCESS.2021.312333

0

Koa, K. Q., Ma, Y., Ng, R., and Chua, T. S. (2023). Diffusion variational autoencoder for tackling stochasticity in multi-step regression stock price prediction. Proceedings of the 32nd ACM International Conference on Information and Knowledge Management, 1087–1096. https://doi.org/10.1145/3583780.3614884

Gao, Y., Wang, R., and Zhou, E. (2022). Stock market prediction with graph attention network and sentiment analysis. Expert Systems with Applications, 188, 116001. https://doi.org/10.1016/j.eswa.2021.116001

Liu, X., Zheng, J., and Han, B. (2023). Deep reinforcement learning for stock recommendation with transformer-based state encoding. Information Sciences, 623, 680–697. https://doi.org/10.1016/j.ins.2022.12.040

Dong, H., Zheng, Z., and Li, W. (2021). Wasserstein GAN-based financial time series augmentation for improved stock prediction. Pattern Recognition Letters, 150, 300–310. https://doi.org/10.1016/j.patrec.2021.06.027

Miyato, T., Kataoka, T., Koyama, M., and Yoshida, Y. (2018). Spectral normalization for generative adversarial networks. Proceedings of the 6th International Conference on Learning Representations. https://doi.org/10.48550/arXiv.1802.05957

Cerrato, M., Liao, Y., and Nandi, A. (2022). Variational autoencoder for financial market regime learning and clustering. Journal of Financial Engineering, 9(2), 2250006. https://doi.org/10.1142/S2424786322500062

Lim, B., Arik, S., Loeff, N., and Pfister, T. (2021). Temporal fusion transformers for interpretable multi-horizon time series forecasting. International Journal of Forecasting, 37(4), 1748–1764. https://doi.org/10.1016/j.ijforecast.2021.03.012

Pan, S. J., and Yang, Q. (2010). A survey on transfer learning. IEEE Transactions on Knowledge and Data Engineering, 22(10), 1345–1359. https://doi.org/10.1109/TKDE.2009.191

van den Oord, A., Dieleman, S., Zen, H., Simonyan, K., Vinyals, O., Graves, A., Kalchbrenner, N., Senior, A., and Kavukcuoglu, K. (2016). WaveNet: A generative model for raw audio. arXiv preprint, arXiv:1609.03499. https://doi.org/10.48550/arXiv.1609.03499

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

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

Preskill, J. (2018). Quantum computing in the NISQ era and beyond. Quantum, 2, 79. https://doi.org/10.22331/q-2018-08-06-79