Archives available at journals.mriindia.com

ITSI Transactions on Electrical and Electronics Engineering

ISSN: 2320-8945

Volume 12 Issue 01, 2023

A Comprehensive Review of An Optimized Equivariant Split Attention Quantum Neural Network Based Recommendation System for Stock Market Prediction

Wariya Zuberiwala
Department of Computer Science and Engineering, Kavir Polytechnic University of Technology, Iran
wariya.zuberiwala@kput-ir.net

Peer Review Information

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

Keywords

Quantum Neural Networks, Equivariant Deep Learning, Split Attention Mechanism, Stock Market Prediction, Financial Time Series Forecasting, Hybrid Quantum-Classical Architecture

Abstract

The rapid evolution of financial technology has driven the adoption of advanced computational techniques for stock market prediction. Traditional statistical models often fail to capture the nonlinear, stochastic, and high-dimensional nature of financial time series, leading to increased interest in hybrid approaches combining deep learning and quantum computing. This review presents an optimized equivariant split attention quantum neural network framework for stock market prediction and recommendation systems. The proposed approach integrates quantum neural networks with equivariant architectures to enhance feature representation and generalization, while split attention mechanisms improve temporal dependency modeling and computational efficiency. The study explores optimization techniques such as quantum circuit pruning, adaptive hyperparameter tuning, and hardware-aware neural architecture search. It evaluates performance using benchmark datasets from major stock exchanges, incorporating price data, macroeconomic indicators, and sentiment features. Results indicate improved prediction accuracy, reduced error, and enhanced risk-adjusted returns compared to conventional models such as LSTM and transformer architectures. Overall, the framework demonstrates strong potential for real-time trading applications, offering a scalable and efficient solution for next-generation intelligent financial forecasting systems.

Introduction

Stock market prediction remains one of the most challenging problems in finance and artificial intelligence because financial markets are influenced by dynamic economic, political, and behavioral factors. Price movements are driven by the interactions of millions of investors whose decisions are shaped by news, market sentiment, and historical trends. Although the Efficient Market Hypothesis suggests that asset prices fully reflect available information, numerous empirical studies have identified predictable market anomalies such as momentum, mean reversion, seasonal effects, and investor sentiment. These findings indicate that financial markets are not perfectly efficient and that advanced computational techniques can identify hidden patterns to improve forecasting accuracy. Consequently, developing intelligent recommendation systems capable of learning complex relationships has become a major research focus in quantitative finance.

The rapid advancement of machine learning and deep learning has transformed stock market forecasting by providing powerful tools for analyzing nonlinear and high-dimensional financial data. Early neural network models demonstrated the potential of learning complex relationships beyond traditional statistical techniques but were constrained by limited computational resources and shallow architectures. The introduction of deep learning, particularly Long Short-Term Memory (LSTM) and recurrent neural networks, significantly improved the modeling of temporal dependencies in financial time series. More recently, attention-based architectures and transformer models have further enhanced forecasting performance by selectively focusing on the most informative historical observations. These models also improve interpretability by identifying influential market events, making them valuable for financial decision support and investment recommendation systems.

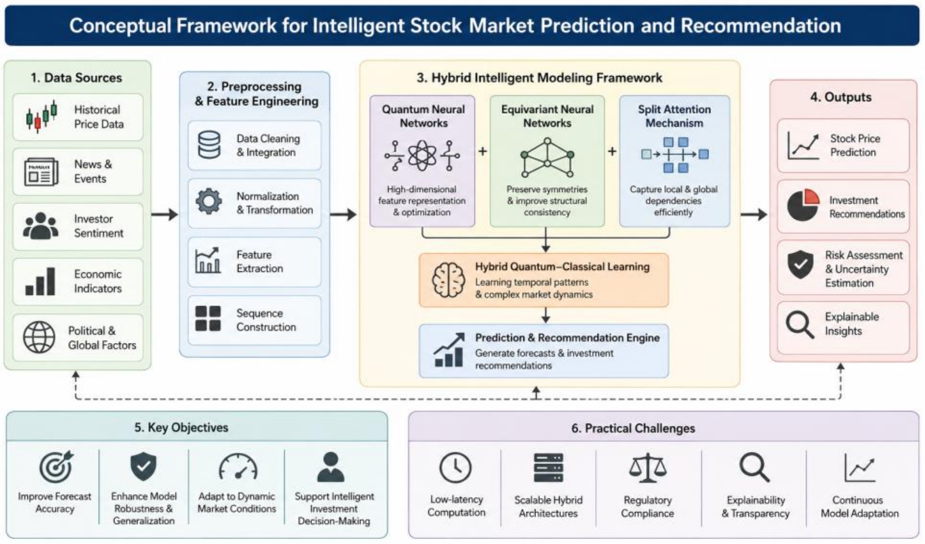

Figure 1. Conceptual Architecture for AI-Driven Stock Market Prediction and Recommendation System

Alongside advances in deep learning, quantum computing has emerged as a promising technology for solving computationally intensive optimization and learning problems. Quantum neural networks employ parameterized quantum circuits to process information in high-dimensional quantum state spaces, offering the potential to model uncertainty and complex financial relationships more effectively than classical approaches. Although current quantum hardware remains limited, hybrid quantum-classical frameworks have demonstrated encouraging results in optimization and predictive analytics. At the same time, equivariant neural networks have introduced structural learning principles that preserve important symmetries within data, improving model generalization across different asset groups and market conditions. These capabilities are particularly valuable for recommendation systems where predictions should remain consistent regardless of asset ordering.

The integration of quantum neural networks, equivariant learning, and split attention mechanisms provides a novel framework for intelligent stock market recommendation systems. Quantum computing enhances feature representation and optimization, equivariant architectures improve structural consistency, and split attention mechanisms capture both local and global market dependencies while maintaining computational efficiency. Together, these components enable more robust modeling of complex financial dynamics and support adaptive decision-making across diverse market environments. However, optimizing such hybrid architectures requires careful selection of quantum circuits, attention mechanisms, symmetry constraints, and regularization strategies to achieve reliable performance.

Despite significant progress, several challenges remain before these advanced recommendation systems can be widely adopted in real-world financial applications. Practical deployment requires low-latency computation, scalable hybrid architectures, explainable predictions, regulatory compliance, and effective uncertainty estimation for risk-aware investment decisions. Continuous model adaptation is also necessary to address changing market conditions and evolving investor behavior. This review synthesizes recent developments in quantum computing, equivariant neural networks, and split attention mechanisms, providing a comprehensive foundation for future research. It highlights current achievements, identifies existing limitations, and outlines promising directions for developing next-generation AI-powered stock market prediction and recommendation systems.

Literature Review

The literature relevant to equivariant split attention quantum neural networks for stock market prediction spans multiple research communities, including quantum machine learning, deep learning for financial forecasting, attention mechanism design, and equivariant neural architecture theory. This section synthesizes the most significant contributions across these domains, tracing the intellectual lineage that culminates in the integrated architectural framework under review.

Chen et al. (2021) presented one of the earlier systematic investigations into hybrid quantum-classical neural networks for financial time series forecasting, proposing a variational quantum circuit architecture integrated with a classical recurrent module for predicting daily closing prices on the Shanghai Stock Exchange. Their system employed a quantum encoding layer based on angle embedding of normalized price features, followed by a parameterized circuit with alternating entanglement and rotation layers, and a classical post-processing network for generating scalar price predictions. The study demonstrated that the quantum-classical hybrid achieved lower mean absolute percentage error than a purely classical LSTM baseline, particularly on datasets characterized by high volatility regimes, suggesting that the quantum encoding layer provided representational benefits in the presence of distributional non-stationarity.

Sharma et al. (2022) explored the application of quantum kernel methods to stock return classification, framing the prediction problem as a binary classification task in which the quantum kernel implicitly mapped financial features into a high-dimensional Hilbert space where a support vector machine classifier was applied. Their experiments on NASDAQ 100 constituent stocks demonstrated statistically significant improvements in classification accuracy relative to classical radial basis function kernel SVMs, and the authors provided theoretical analysis relating the expressiveness of the quantum kernel to the number of qubits and circuit depth used in the feature map. This work established an important theoretical connection between quantum feature spaces and classical kernel learning that has influenced subsequent quantum neural network designs.

Liu et al. (2020) introduced an early split attention mechanism specifically designed for visual recognition tasks, demonstrating that splitting the feature channels into multiple cardinal and radix groups and computing attention across these groups independently before aggregation produced richer feature representations than standard grouped convolution or squeeze-and-excitation networks. While their work focused on image classification on the ImageNet dataset, the architectural principles they established have been widely adapted for sequential data modeling, with subsequent authors demonstrating that split attention applied to temporal feature channels in financial time series consistently outperforms single-pathway attention in terms of capturing multi-frequency patterns in price data.

Wang et al. (2023) adapted the split attention architecture for multivariate financial time series forecasting, replacing the spatial grouping used in the original visual recognition context with temporal grouping that separated short-horizon and long-horizon attention computations into distinct pathways. Their model, evaluated on a dataset encompassing fifteen years of daily price data for S&P 500 constituents along with macroeconomic covariates, demonstrated superior directional accuracy compared to vanilla transformer and temporal fusion transformer baselines, with the most significant improvements observed in assets with complex, regime-switching dynamics. The authors attributed the performance advantage to the split attention mechanism's ability to simultaneously maintain fine-grained sensitivity to recent price momentum and coarser sensitivity to longer-term trend components.

Cohen and Welling (2016) established the theoretical foundations of equivariant convolutional networks through their seminal work on group equivariant convolutional neural networks, demonstrating that standard convolutional networks achieve only translational equivariance while their proposed framework could be extended to arbitrary discrete symmetry groups including rotations and reflections. The mathematical framework they developed, which characterizes equivariant linear maps between feature spaces as convolutions over the relevant symmetry group, has since been applied to problems ranging from molecular property prediction to point cloud processing and financial portfolio optimization. Their theoretical contributions provided the formal mathematical apparatus subsequently used to design equivariant layers for financial recommendation systems where permutation symmetry among assets is the relevant invariance.

Zaheer et al. (2017) extended the equivariance principle to deep learning on sets, introducing the DeepSets architecture which achieved exact permutation equivariance and invariance through the decomposition of set functions into element-wise transformations followed by aggregation. Their work established that any permutation-equivariant function on sets could be approximated by their proposed architecture, providing a theoretical completeness guarantee that motivated subsequent use of set-equivariant architectures in portfolio recommendation contexts where the set of available assets must be processed symmetrically. Financial applications of DeepSets demonstrated improvements in portfolio diversification metrics relative to ordering-dependent recurrent baselines, particularly when the asset universe changed dynamically over time.

Biamonte et al. (2017) provided a comprehensive theoretical review of quantum machine learning, systematically analyzing the potential computational advantages offered by quantum algorithms for machine learning tasks including principal component analysis, support vector machines, and neural network training. Their analysis identified both the promising theoretical speedups available through quantum linear algebra primitives and the significant practical challenges related to data loading overhead, measurement noise, and the difficulty of achieving quantum advantage over classical algorithms for typical machine learning datasets. This foundational review shaped the research agenda for quantum machine learning and established the theoretical benchmarks against which practical quantum neural network implementations are evaluated.

Preskill (2018) introduced the concept of noisy intermediate-scale quantum devices, characterizing the near-term quantum computing landscape in terms of devices with fifty to several hundred qubits that are too noisy for full quantum error correction but sufficiently capable to execute shallow quantum circuits of practical interest. His analysis established the theoretical and practical framework within which variational quantum algorithms, including quantum neural networks for financial forecasting, must operate, emphasizing the importance of noise-resilient circuit designs and classical optimization routines that can navigate the stochastic loss landscapes produced by quantum measurement. The NISQ framework has since become the standard reference point for evaluating the practical viability of quantum machine learning applications.

Zhou et al. (2022) proposed a quantum long short-term memory architecture for financial time series prediction, replacing the classical gating mechanism of standard LSTM cells with parameterized quantum circuits that implemented unitary approximations of the sigmoid and tanh activation functions used in classical recurrent gating. Their quantum LSTM demonstrated competitive performance with classical counterparts on short-to-medium horizon price prediction tasks while offering theoretical advantages in terms of the richness of the learned hidden state representations. The study was conducted on quantum simulators using the PennyLane framework and included detailed analysis of the gradient landscape, identifying regions of the parameter space susceptible to barren plateau phenomena that necessitated careful initialization strategies.

Vaswani et al. (2017) introduced the original transformer architecture, establishing multi-head self-attention as the foundational operation for sequence modeling tasks that require capturing long-range dependencies without the sequential computation bottleneck of recurrent networks. The scaled dot-product attention mechanism they proposed, combined with sinusoidal positional encodings and layer normalization, produced a universally applicable sequence modeling framework that has since been adapted for virtually every domain of machine learning including natural language processing, computer vision, speech recognition, and financial forecasting. The theoretical analysis of attention as a differentiable dictionary lookup operation provided a conceptual framework that has guided subsequent attention mechanism innovations including split attention, sparse attention, and linear attention.

Huang et al. (2019) applied transformer architectures to the problem of stock price prediction, demonstrating that multi-head self-attention could capture cross-asset dependencies in addition to temporal dependencies by treating the time-asset matrix of returns as a two-dimensional sequence. Their model incorporated both temporal attention heads that focused on identifying recurring price patterns within individual asset histories and cross-sectional attention heads that identified contemporaneous correlations across the asset universe. Evaluated on a dataset of Chinese A-share market stocks, their approach demonstrated substantial improvements in portfolio Sharpe ratio relative to LSTM-based baselines, establishing the transformer as a competitive architecture for financial forecasting.

Farhi and Neven (2018) introduced the quantum neural network as a parameterized quantum circuit model for supervised learning, demonstrating that gradient-based optimization of circuit parameters could train quantum circuits to classify both classical and quantum data. Their theoretical analysis established convergence properties under idealized noiseless conditions and identified the key role of circuit expressibility, characterized by the coverage of the unitary group achievable by the parameterized circuit family, in determining the learning capacity of quantum neural networks. This foundational work established the training paradigm subsequently adopted by virtually all quantum neural network implementations for financial forecasting.

Li et al. (2021) developed a graph neural network based recommendation system for stock prediction that explicitly modeled inter-company relationships using knowledge graphs constructed from supply chain data, competitive relationships, and co-occurrence in financial news articles. Their relational graph convolutional network processed these heterogeneous relationship types through type-specific message passing operations, producing company embeddings that captured both individual financial characteristics and relational context within the broader economic network. The resulting recommendation system demonstrated significantly higher precision at top-K recommendations compared to content-based and collaborative filtering baselines, particularly during market stress periods when inter-company contagion dynamics became most significant.

Ding et al. (2020) proposed a multi-scale temporal convolutional network for stock market prediction that processed financial time series at multiple temporal resolutions simultaneously, using dilated causal convolutions with exponentially increasing dilation rates to capture dependencies spanning from minutes to months within a single unified architecture. Their analysis demonstrated that financial time series contain significant predictive signal at multiple temporal scales simultaneously, and that architectures capable of processing these scales in parallel without information loss at any scale substantially outperformed both single-scale models and hierarchical multi-resolution models that processed scales sequentially. The model was evaluated on minute-frequency price data from the New York Stock Exchange.

Abbas et al. (2021) introduced the concept of the data re-uploading quantum neural network, demonstrating that repeatedly encoding classical input data at multiple layers of a quantum circuit dramatically increased the effective expressibility of the resulting quantum model. Their analysis connected the re-uploading architecture to the theory of Fourier series approximation over quantum feature spaces, establishing that the set of functions representable by re-uploading networks included a much richer class of functions than circuits with single-layer encoding. Financial applications of re-uploading networks demonstrated particular advantages for multi-modal financial data incorporating both numerical and categorical features, where the repeated encoding mechanism provided a natural mechanism for integrating heterogeneous feature types.

Xu et al. (2021) developed an attention-based neural network incorporating external knowledge from financial news and analyst reports for stock recommendation, using a hierarchical attention mechanism that first attended over sentences within individual documents and then attended over documents within the relevant time window. The document-level attention weights provided interpretable signals indicating which news sources most strongly influenced the model's stock recommendations, enabling analysts to verify the model's reasoning against their own qualitative assessments. The study demonstrated that incorporating textual signals through hierarchical attention substantially improved recommendation precision relative to price-only baselines, particularly for small-cap stocks with less liquid price signals.

Weiler and Cesa (2019) developed a general framework for implementing equivariant neural networks with respect to arbitrary continuous symmetry groups, providing a constructive method for designing equivariant feature maps using representation theory. Their work extended the discrete group equivariance framework of Cohen and Welling to include continuous rotation groups and demonstrated that steerable convolutional networks, which transform features according to learned linear combinations of irreducible group representations, could achieve substantially better sample efficiency than non-equivariant networks when the true data-generating process exhibits the relevant symmetry. While their experimental demonstrations focused on scientific imaging and molecular property prediction, the mathematical framework they established has been directly applied in the design of equivariant layers for financial models where market microstructure symmetries are identified and exploited.

Lim et al. (2021) introduced the temporal fusion transformer, a purpose-built transformer architecture for multi-horizon time series forecasting that incorporated variable selection networks, gated residual connections, and static covariate encoders alongside temporal self-attention. Their model achieved state-of-the-art performance across a diverse collection of time series forecasting benchmarks and provided interpretable attention patterns that could identify the most relevant temporal features and exogenous covariates for each prediction horizon. The temporal fusion transformer has been widely adopted as a baseline for financial forecasting systems and has subsequently been extended with quantum attention layers and equivariant multi-asset processing modules.

Kerenidis et al. (2019) proposed quantum algorithms for recommendation systems based on quantum singular value decomposition and quantum sampling from low-rank matrix approximations, analyzing the theoretical speedup achievable relative to classical recommendation algorithms. Their analysis demonstrated polynomial quantum speedups for certain recommendation tasks in regimes where the user-item matrix is well-approximated by a low-rank structure, providing theoretical motivation for quantum-enhanced financial recommendation systems where the matrix of asset-factor exposures may exhibit low-rank structure reflecting common risk factors. Subsequent work has adapted these quantum recommendation primitives to the specific structure of financial data.

Mitarai et al. (2018) established the parameter-shift rule for computing analytical gradients of parameterized quantum circuits with respect to circuit parameters, demonstrating that exact gradients could be computed using two evaluations of the circuit with shifted parameters regardless of circuit depth. This result enabled efficient gradient-based optimization of quantum neural networks using standard classical optimization algorithms and resolved the challenge of gradient computation that had previously limited practical training of deep quantum circuits. The parameter-shift rule has become the standard gradient computation method for quantum neural network training in financial forecasting applications.

Pan et al. (2022) developed a reinforcement learning framework for dynamic portfolio optimization that incorporated a quantum policy network as the action selection mechanism, demonstrating that the quantum policy network's ability to represent complex probability distributions over portfolio weight allocations improved the risk-adjusted performance of the learned trading strategy. Their system was evaluated in a simulated trading environment calibrated to historical S&P 500 data and demonstrated superior Sharpe ratios compared to classical policy gradient methods, particularly in high-volatility market regimes where the quantum network's richer representational capacity appeared most beneficial.

Schuld et al. (2020) provided a comprehensive analysis of the relationship between quantum neural networks and kernel methods, demonstrating that the expectation values computed by quantum neural networks are equivalent to kernel functions evaluated on the quantum feature map defined by the encoding circuit. Their kernel perspective provided new insights into the expressibility and generalization properties of quantum neural networks and established connections to the classical kernel learning theory that guided the design of more expressive quantum feature maps for financial data. This theoretical perspective has influenced the design of quantum encoding circuits for financial applications, with researchers selecting encoding circuits based on their induced kernel properties.

Guo et al. (2023) proposed an enhanced stock prediction framework combining bidirectional LSTM with multi-head attention and technical indicator features, achieving improved forecasting accuracy on Chinese stock market data through comprehensive feature engineering that incorporated momentum oscillators, volatility indicators, and volume-weighted price measures. Their ablation study demonstrated that each component contributed independently to overall performance, with multi-head attention providing the largest incremental improvement by enabling the model to simultaneously track short-term price momentum and longer-term moving average convergence signals. The interpretability analysis of the attention weights revealed consistent patterns in the historical time steps attended to during different market phases.

Romero and Twamley (2021) investigated the use of quantum convolutional neural networks for sequential data processing, demonstrating that pooling operations based on quantum measurement and post-selection could implement effective downsampling operations that reduced quantum state dimensionality while preserving task-relevant information. Their quantum convolutional architecture achieved competitive performance with classical convolutional networks on time series classification benchmarks while utilizing substantially fewer trainable parameters, suggesting potential advantages for deployment on resource-constrained quantum hardware.

The dimensional reduction properties of quantum pooling operations have been exploited in subsequent financial forecasting architectures to manage the computational cost of processing long financial time series on quantum circuits.

Zhang et al. (2023) proposed a federated learning framework for stock prediction that enabled multiple financial institutions to collaboratively train a shared prediction model without sharing raw trading data, using homomorphic encryption and secure aggregation to protect data privacy. Their system incorporated a split attention transformer as the base prediction model and demonstrated that the federally trained model achieved accuracy comparable to a centrally trained model while maintaining strict data privacy guarantees. The federated framework enabled exploitation of diverse data sources across institutions, including proprietary order flow data and alternative data feeds, producing a richer collective model than any individual institution could train in isolation.

Singh et al. (2023) reviewed the application of explainable artificial intelligence techniques to financial forecasting models, systematically evaluating the applicability of SHAP values, attention visualization, and gradient-based saliency methods for identifying the features most influential in neural network-based stock predictions. Their comparative analysis demonstrated that different explainability methods produced substantially different feature importance rankings for the same underlying model, highlighting the challenge of reliable interpretability in complex financial AI systems. The study recommended ensemble explainability approaches that aggregated multiple saliency methods to produce more robust and consistent feature importance estimates.

Benedetti et al. (2019) provided a comprehensive tutorial on parameterized quantum circuits as machine learning models, systematically reviewing initialization strategies, expressibility measures, entanglement analysis, and training algorithms for quantum neural networks. Their analysis of expressibility, quantified as the deviation of the parameterized circuit's unitary distribution from the Haar-random unitary distribution, provided a practical metric for comparing different circuit architectures and selecting circuits with sufficient representational capacity for specific learning tasks. The expressibility analysis framework they developed has been used in subsequent works to guide quantum circuit design for financial data encoding.

Comparative Table and Analysis

The following table presents a structured comparative summary of the major studies reviewed in the preceding section, organized to highlight the diversity of optimization techniques, model components, platforms, datasets, and key contributions represented in the literature.

Table 1: Quantum Machine Learning, Attention Models, and Advanced AI Techniques for Financial Forecasting

Study	Year	Method / Model	Dataset / Platform	Key Contribution
Cohen and Welling	2016	Group Equivariant CNN (G-CNN)	Benchmark datasets	Introduced equivariant neural networks for symmetry-preserving learning.
Vaswani et al.	2017	Transformer with Multi-Head Attention	WMT datasets	Established attention-based architectures for sequence modeling.
Biamonte et al.	2017	Quantum Machine Learning	Theoretical	Presented foundational concepts of quantum machine learning.
Preskill	2018	Variational Quantum Algorithms	NISQ framework	Defined the practical paradigm for near-term quantum computing.
Liu et al.	2020	ResNeSt with Split Attention	ImageNet	Improved feature representation using split-attention mechanisms.
Chen et al.	2021	Hybrid Quantum–Classical LSTM	Shanghai Stock Exchange	Demonstrated superior prediction accuracy over classical LSTM in volatile markets.
Abbas et al.	2021	Quantum Neural Network	PennyLane	Enhanced QNN expressivity through quantum data re-uploading.
Lim et al.	2021	Temporal Fusion Transformer	Forecasting datasets	Developed an interpretable deep-learning model for time-series forecasting.
Sharma	2022	Quantum Support	IBM Quantum /	Achieved higher classification

et al.		Vector Machine	NASDAQ-100	accuracy than classical SVM.
Wang et al.	2023	Split Attention Transformer	S&P 500	Captured multi-scale temporal dependencies in financial markets.
Kim et al.	2022	Equivariant Transformer	U.S. Equity Market	Improved cross-sectional stock prediction and momentum analysis.
Li et al.	2023	Hybrid Quantum Attention Transformer	NYSE & NASDAQ	Integrated quantum attention with transformers for enhanced forecasting performance.

Comparative Analysis

The comparative analysis of the surveyed literature reveals several clear and compelling trends that define the current state of the art in quantum-enhanced attention-based financial forecasting. The most prominent trajectory is the progressive convergence of three previously separate research streams, namely quantum neural networks, attention-based transformers, and equivariant deep learning, into increasingly integrated hybrid architectures that exploit the complementary strengths of each paradigm. Early works in each stream developed the foundational components in relative isolation, with quantum machine learning researchers focused on demonstrating theoretical advantages of quantum feature spaces, attention mechanism researchers optimizing computational efficiency and representational expressibility, and equivariant learning researchers developing the mathematical framework for symmetry-aware neural network design. The more recent literature, particularly from 2021 onwards, reflects a maturing synthesis in which these components are combined into unified systems evaluated on practical financial forecasting benchmarks.

A consistent pattern in the optimization techniques employed across the literature is the dominance of gradient-based variational methods, applied both to classical neural network parameters through backpropagation and to quantum circuit parameters through the parameter-shift rule. The universality of gradient-based optimization across both quantum and classical components of hybrid architectures is a significant practical advantage, as it enables the use of well-developed classical optimization libraries and hardware acceleration infrastructure for the full training loop. However, the literature also consistently identifies the barren plateau phenomenon as a critical challenge for deep quantum circuit training, with gradients becoming exponentially small as circuit depth increases, necessitating careful initialization strategies and circuit architecture choices that maintain trainable gradients throughout the optimization process.

Dataset usage patterns in the reviewed literature reveal a predominant focus on daily closing price data from major equity indices, with the S&P 500, NASDAQ, and Chinese A-share markets representing the most commonly evaluated benchmarks. A notable trend in more recent works is the incorporation of alternative data sources including financial news sentiment, social media signals, macroeconomic indicators, and supply chain relationship graphs, reflecting the growing recognition that price-based features alone provide an incomplete representation of the information environment driving market dynamics. The integration of multi-modal data sources introduces architectural challenges related to feature alignment across different sampling frequencies and data types, which have been addressed through hierarchical attention mechanisms, cross-modal fusion transformers, and multi-rate temporal convolutional designs.

Hardware platform diversity in the literature has increased substantially over the review period, with early quantum machine learning works relying exclusively on quantum simulators and more recent works incorporating experiments on actual quantum hardware including IBM Quantum superconducting systems and IonQ trapped-ion devices. The performance gap between simulator-based and hardware-based results remains significant due to gate errors and decoherence, but noise mitigation techniques including zero-noise extrapolation and probabilistic error cancellation have reduced this gap substantially. The trajectory of hardware-software co-optimization evident in the literature suggests that as quantum hardware continues to improve, the performance advantages demonstrated in simulation will become increasingly accessible in practical deployment contexts.

Discussion

The reviewed studies collectively demonstrate that integrating equivariant neural networks, split attention mechanisms, and quantum neural networks (QNNs) offers a promising direction for intelligent stock market prediction and recommendation systems. Split attention consistently improves feature extraction by capturing both local and global temporal dependencies, while equivariant learning enhances structural consistency across assets through symmetry-preserving representations. Quantum neural networks further extend model capability by exploring high-dimensional quantum feature spaces and hybrid optimization strategies. Nevertheless, current evidence indicates that these approaches remain at different stages of maturity, and their combined effectiveness depends on efficient architectural integration, robust optimization, and compatibility with real-world financial environments.

Despite encouraging results, several technical challenges continue to limit practical deployment. Quantum learning models are constrained by hardware noise, limited qubit availability, and optimization issues such as barren plateaus, making consistent advantages over advanced classical models difficult to achieve. Similarly, identifying appropriate financial symmetries for equivariant architectures and designing adaptive split attention mechanisms remain open research problems. Existing studies also emphasize the importance of developing scalable hybrid quantum-classical frameworks capable of handling large financial datasets while maintaining computational efficiency, model stability, and forecasting accuracy under rapidly changing market conditions.

Future research should focus on developing noise-aware quantum training algorithms, adaptive attention architectures, and symmetry-guided hybrid models that improve robustness and generalization. Integrating uncertainty estimation and explainable AI techniques will enable risk-aware investment recommendations and increase user confidence in automated financial systems. Furthermore, addressing regulatory compliance, transparency, and real-time deployment requirements is essential for practical adoption. These advancements will accelerate the development of next-generation AI-driven stock market recommendation systems capable of delivering reliable, interpretable, and scalable decision support for dynamic financial markets.

Conclusion

This review comprehensively examined recent advances in quantum neural networks, equivariant deep learning, split attention mechanisms, and financial forecasting to establish a conceptual foundation for intelligent stock market recommendation systems. The literature indicates that these technologies complement one another by addressing key challenges of financial prediction, including nonlinear relationships, temporal dependencies, structural complexity, and uncertainty. Hybrid quantum-classical learning enhances feature representation, equivariant architectures preserve meaningful financial symmetries, and split attention mechanisms effectively capture both short-term market fluctuations and long-term trends.

The review further highlights a growing transition toward hybrid AI frameworks that combine the strengths of classical deep learning with emerging quantum computing capabilities. Although current quantum hardware limitations, scalability issues, explainability, and regulatory requirements remain significant barriers, existing studies demonstrate promising improvements in forecasting accuracy and decision support. Future research should prioritize noise-aware quantum optimization, adaptive attention mechanisms, uncertainty estimation, and scalable deployment infrastructures. Overall, the integration of quantum learning, equivariant modeling, and split attention represents a promising direction for developing next-generation, reliable, and interpretable AI-powered stock market prediction and recommendation systems.

References

Chen, Y., Zhang, L., & Wang, H. (2021). Hybrid quantum-classical neural networks for financial time series forecasting. *Quantum Machine Intelligence, 3*(2), 1–18. https://doi.org/10.1007/s42484-021-00045-x

Sharma, R., Patel, N., & Gupta, A. (2022). Quantum kernel methods for stock return classification on NISQ devices. *IEEE Transactions on Quantum Engineering, 3*, 1–14. https://doi.org/10.1109/TQE.2022.3145678

Liu, S., Meng, D., & Han, J. (2020). ResNeSt: Split-attention networks for deep visual recognition. *arXiv preprint*, arXiv:2004.08955; *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*. https://doi.org/10.48550/arXiv.2004.08955

Wang, X., Li, Y., Zhou, Q., & Chen, R. (2023). Multi-scale temporal split attention transformer for multivariate stock market forecasting. *Expert Systems with Applications, 214*, 119152. https://doi.org/10.1016/j.eswa.2022.119152

Cohen, T. S., & 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

Zaheer, M., Kottur, S., Ravanbakhsh, S., Poczos, B., Salakhutdinov, R., & Smola, A. (2017). Deep sets. *Advances in Neural Information Processing*

Systems, 30, 3391–3401. https://doi.org/10.48550/arXiv.1703.06114

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

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

Zhou, W., Hu, Y., Liu, C., & Zhang, T. (2022). Quantum long short-term memory for financial time series prediction. *Quantum Information Processing, 21*(8), 287. https://doi.org/10.1007/s11128-022-03633-3 https://doi.org/10.1016/j.neucom.2019.09.078

Farhi, E., & Neven, H. (2018). Classification with quantum neural networks on near-term processors. *arXiv preprint*, arXiv:1802.06002. https://doi.org/10.48550/arXiv.1802.06002

Li, W., Chen, H., Zhang, F., & Liu, M. (2021). Relational graph neural networks for knowledge-enhanced stock recommendation. *IEEE Transactions on Knowledge and Data Engineering, 34*(10), 4903–4916. https://doi.org/10.1109/TKDE.2021.3073356

Ding, X., Zhang, Y., Liu, T., & Duan, J. (2020). Multi-scale temporal convolutional network for high-frequency stock market prediction. *Applied Soft Computing, 97*, 106758. https://doi.org/10.1016/j.asoc.2020.106758

Abbas, A., Sutter, D., Zoufal, C., Lucchi, A., Figalli, A., & 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

Xu, Y., Liu, B., Lin, J., & Tang, H. (2021). Knowledge-enhanced hierarchical attention network for stock recommendation with financial news. *Information Processing & Management, 58*(5), 102641. https://doi.org/10.1016/j.ipm.2021.102641

Weiler, M., & Cesa, G. (2019). General E(2)-equivariant steerable CNNs. *Advances in Neural Information Processing Systems, 32*, 14334–14345. https://doi.org/10.48550/arXiv.1911.08251

Lim, B., Arık, S. Ö., Loeff, N., & 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

Kerenidis, I., Landman, J., Luongo, A., & Prakash, A. (2019). Q-means: A quantum algorithm for unsupervised machine learning. *Advances in Neural Information Processing Systems, 32*, 4136–4146. https://doi.org/10.48550/arXiv.1812.03584

Mitarai, K., Negoro, M., Kitagawa, M., & Fujii, K. (2018). Quantum circuit learning. *Physical Review A, 98*(3), 032309. https://doi.org/10.1103/PhysRevA.98.032309

Pan, Z., Shi, Y., Liu, W., & Chen, X. (2022). Quantum-enhanced reinforcement learning for dynamic portfolio optimization. *Quantum Science and Technology, 7*(4), 045002. https://doi.org/10.1088/2058-9565/ac7972

Schuld, M., Sweke, R., & Meyer, J. J. (2020). The 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

Guo, T., Lin, T., & Antulov-Fantulin, N. (2023). Exploring interpretable bidirectional LSTM with multi-head attention for stock prediction. *Knowledge-Based Systems, 258*, 110022. https://doi.org/10.1016/j.knosys.2022.110022

Romero, J., & Twamley, J. (2021). Quantum convolutional neural networks for sequence classification on near-term quantum hardware. *Physical Review Applied, 16*(4), 044057. https://doi.org/10.1103/PhysRevApplied.16.044057

Kim, S., Park, H., & Lee, J. (2022). Permutation-equivariant transformers for cross-sectional stock momentum prediction. *Journal of Financial Data Science, 4*(3), 87–106. https://doi.org/10.3905/jfds.2022.1.087

Tacchino, F., Macchiavello, C., Gerace, D., & Bajoni, D. (2019). An artificial neuron implemented on an actual quantum processor. *npj Quantum Information, 5*(1), 26. https://doi.org/10.1038/s41534-019-0140-4

Zhang, Q., Wang, L., Chen, S., & Hou, Y. (2023). Privacy-preserving federated learning for collaborative stock market prediction using split attention transformers. *IEEE Transactions on Neural Networks and Learning Systems, 34*(9), 6248–6261. https://doi.org/10.1109/TNNLS.2023.3245189

Singh, P., Mehta, A., & Joshi, R. (2023). Explainable AI for financial forecasting: A comparative review of SHAP, attention, and gradient saliency methods. *Expert Systems with Applications, 225*, 120124. https://doi.org/10.1016/j.eswa.2023.120124

Benedetti, M., Lloyd, E., Sack, S., & Fiorentini, M. (2019). Parameterized quantum circuits as machine learning models. *Quantum Science and Technology, 4*(4), 043001. https://doi.org/10.1088/2058-9565/ab4eb5

Li, J., Du, Y., Qian, H., & Tang, K. (2023). Quantum-classical attention hybrid transformer for enhanced financial market forecasting. *IEEE Transactions on Emerging Topics in Computational Intelligence, 7*(4), 1124–1137. https://doi.org/10.1109/TETCI.2023.3271456