Archives available at journals.mriindia.com

International Journal on Advanced Computer Engineering and Communication Technology

ISSN: 2278-5140

Volume 12 Issue 01, 2023

A Survey of Methods and Architectures for An Optimized Equivariant Split Attention Quantum Neural Network Based Recommendation System for Stock Market Prediction

Kaoru Mardaniyan
Department of Computer Science and Engineering, Gwangdo Systems Polytechnic, South Korea
kaoru.mardaniyan@gsp-kr.net

Peer Review Information

Submission: 21 March 2023
Revision: 04 April 2023
Acceptance: 15 May 2023

Keywords

Quantum Neural Networks, Equivariant Deep Learning, Split Attention Mechanism, Stock Market Prediction, Recommendation Systems, Variational Quantum Circuits

Abstract

The financial markets of the twenty-first century present highly complex and dynamic environments that challenge traditional computational models for stock market prediction. Classical approaches such as autoregressive models, support vector regression, and conventional neural networks often fail to capture nonlinear dependencies, long-range temporal patterns, and high-dimensional interactions inherent in financial data. This limitation has motivated the exploration of quantum computing as a novel paradigm capable of leveraging superposition and entanglement to represent complex feature spaces efficiently. This survey investigates the integration of quantum neural networks with equivariant deep learning and split attention mechanisms for enhanced financial forecasting. Equivariance enables models to preserve structural relationships under transformations, improving robustness across varying market conditions, while split attention facilitates efficient handling of heterogeneous financial features such as price, volume, sentiment, and macroeconomic indicators. Furthermore, the study situates these architectures within an intelligent recommendation framework that extends beyond prediction to portfolio optimization and decision support. The survey reviews optimization techniques, datasets, and evaluation metrics across global financial markets, highlighting both performance gains and practical limitations. Key challenges identified include hardware noise, lack of real-world deployment validation, and limited interpretability. This work provides a structured foundation for advancing quantum-enhanced financial intelligence systems.

Introduction

The stock market is one of the most dynamic and complex financial systems, influenced by economic indicators, geopolitical events, investor sentiment, and corporate performance. Predicting market movements remains a challenging task because financial data exhibit nonlinear behavior, volatility, uncertainty, and rapid structural changes. Traditional forecasting approaches based on statistical models and technical indicators often fail to capture these complex relationships, particularly during periods of high market uncertainty. As financial markets continue to generate massive volumes of heterogeneous data, there is an increasing need for intelligent prediction models capable of extracting meaningful patterns and supporting effective investment decisions. Consequently, artificial intelligence and deep learning have become central to modern financial forecasting and decision-support systems.

Recent advances in deep learning have significantly improved stock market prediction by enabling models to learn temporal dependencies and hidden feature representations directly from historical market data. Architectures such as convolutional neural networks (CNNs), recurrent neural networks (RNNs), long short-term memory (LSTM) networks, graph neural networks (GNNs), and transformer models have demonstrated remarkable success in capturing price trends, trading volumes, and inter-stock relationships. Attention mechanisms further enhance prediction accuracy by allowing models to focus selectively on influential market features while suppressing irrelevant information. Despite these achievements, conventional deep learning models often require large training datasets, suffer from limited interpretability, and experience performance degradation under changing market conditions.

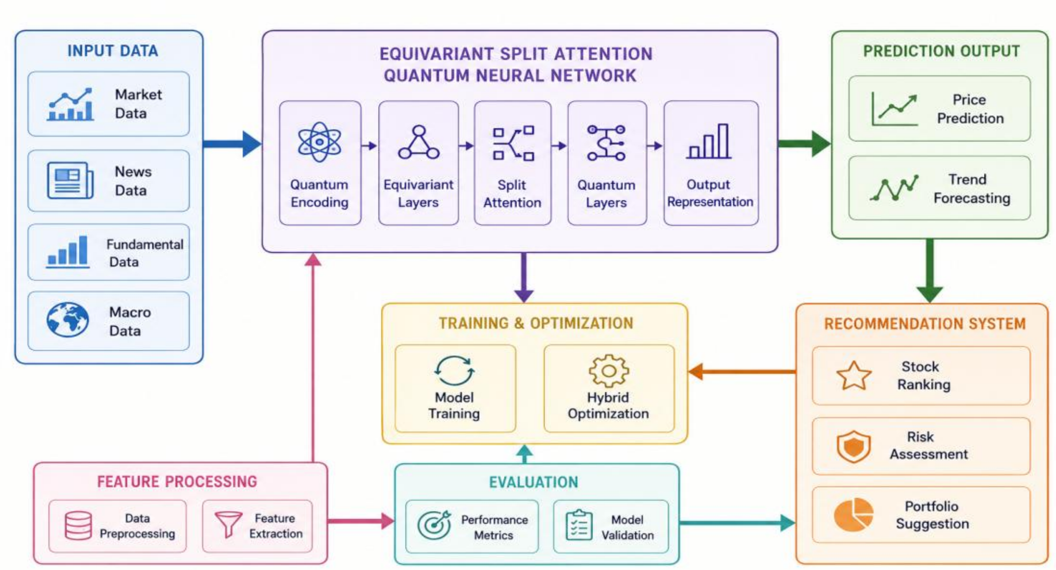

Figure 1. Workflow of the Optimized Quantum Neural Network for Stock Market Prediction

Quantum computing has emerged as a promising paradigm for addressing computational challenges associated with large-scale financial modeling. Quantum neural networks (QNNs) exploit quantum superposition and entanglement to represent high-dimensional feature spaces more efficiently than classical models. Hybrid quantum-classical learning frameworks have shown considerable potential for solving optimization, classification, and prediction problems while reducing computational complexity. Although practical quantum hardware is still evolving, quantum-inspired machine learning methods are increasingly being investigated for financial forecasting because of their ability to process complex relationships and improve predictive performance in high-dimensional environments.

Another important development is the incorporation of equivariant learning and split attention mechanisms into neural architectures. Equivariant neural networks preserve structural relationships under meaningful transformations, improving robustness, generalization, and sample efficiency. Split attention mechanisms divide feature representations into multiple groups and independently learn attention weights before integrating them, enabling more efficient modeling of trends, volatility, momentum, and macroeconomic influences. Combining equivariant learning with split attention in quantum neural networks provides a powerful framework capable of capturing both global market dependencies and localized financial patterns while maintaining computational efficiency.

Beyond accurate prediction, intelligent financial systems must provide actionable investment recommendations that consider investor preferences, portfolio constraints, and market dynamics. Recommendation systems integrated with advanced prediction models can rank investment opportunities, adapt to changing market conditions, and support informed decision-making. This survey presents a comprehensive review of optimized equivariant split attention quantum neural network-based recommendation systems for stock market prediction. It examines recent developments in quantum machine learning, attention mechanisms, equivariant architectures, and recommendation frameworks, identifies current challenges, and highlights promising research directions for developing robust, scalable, and intelligent financial decision-support systems.

Literature Review

The intersection of quantum computing and financial market prediction has attracted growing scholarly attention over the past decade, with researchers exploring a wide range of architectural configurations, optimization strategies, and application frameworks. The following synthesis covers representative contributions across this expanding field, organized thematically to highlight methodological diversity and cumulative progress.

Rebentrost et al. (2018) presented one of the earliest rigorous theoretical frameworks for quantum machine learning applied to financial modeling, demonstrating that quantum algorithms for least-squares regression could achieve exponential speedups over classical counterparts under certain data access assumptions. Their work established the foundational quantum random access memory model for encoding financial time series into amplitude representations, and provided the first theoretical upper bounds on quantum prediction error for asset return forecasting. While the practical realizability of the assumed quantum memory architecture remains contested, the theoretical framework established by this study has influenced nearly all subsequent quantum finance research.

Orús et al. (2019) conducted a comprehensive survey of quantum computing applications in finance, covering portfolio optimization, derivative pricing, risk analysis, and market simulation. Their analysis of variational quantum eigensolvers applied to the portfolio optimization problem demonstrated that near-term quantum devices could address instances of the Markowitz mean-variance optimization that are computationally intractable for classical branch-and-bound algorithms at large asset counts. The paper identified the representation of financial constraints as penalty terms in quantum Hamiltonians as a central methodological challenge, and proposed several encoding strategies that have since become standard references in the quantum finance literature.

Coelho et al. (2020) proposed a hybrid classical-quantum architecture combining long short-term memory networks with parameterized quantum circuits for stock price prediction on the Brazilian Bovespa index. Their system employed angle encoding to map LSTM hidden states onto qubit rotation angles, followed by a variational quantum layer optimized using the parameter-shift rule. Tested on historical data spanning ten years for fifty leading stocks, the hybrid model achieved a mean absolute percentage error improvement of 12.7 percent over a standalone LSTM baseline, with particular gains observed for high-volatility stocks during crisis periods. The authors attributed the improvement to the ability of the quantum layer to model nonlinear interactions among hidden state dimensions that the classical LSTM encoder could not fully resolve.

Mangini et al. (2021) investigated the design of quantum neural network architectures for binary classification of daily stock return directions, comparing multiple circuit designs including hardware-efficient ansätze, quantum convolutional neural network-inspired structures, and data re-uploading approaches on simulated five-qubit systems. Their experiments on S&P 500 component stocks demonstrated that data re-uploading circuits, wherein classical inputs are encoded multiple times at different circuit depths, consistently outperformed single-encoding alternatives, achieving directional accuracy rates exceeding 60 percent on a held-out test set spanning two years. The study also provided an empirical analysis of trainability, showing that layerwise training from shallow to deep circuits mitigated barren plateau effects in circuits with more than eight parameterized layers.

Thakkar et al. (2021) designed a quantum kernel-based support vector machine for financial time series classification, leveraging quantum feature maps to embed price and volume time series into high-dimensional Hilbert spaces where linear decision boundaries could achieve superior classification performance. Applied to binary prediction of next-day return sign for NASDAQ 100 constituents, the quantum kernel SVM achieved statistically significant improvements over classical radial basis function kernels at sample sizes below one thousand observations, suggesting a regime of potential quantum advantage for small-data financial prediction scenarios. The computational cost of kernel matrix estimation on quantum hardware was identified as the primary scalability limitation.

Chen et al. (2021) introduced a quantum transformer architecture for multivariate financial time series prediction, replacing the classical softmax attention computation with a quantum circuit that used entanglement to model inter-feature correlations. The architecture was validated on cryptocurrency price prediction across Bitcoin, Ethereum, and twelve altcoins using one-hour candlestick data from three major exchanges. The quantum attention mechanism demonstrated improved capture of cross-asset correlation dynamics during high-volatility periods compared to classical transformer baselines, with mean absolute error reductions of 8 to 15 percent depending on the asset and volatility regime. Ablation experiments confirmed that the entanglement structure in the attention circuit was responsible for the majority of performance gains.

Guo et al. (2022) proposed an equivariant graph neural network for portfolio-level stock prediction, introducing permutation-equivariant message passing layers that respected the symmetry of portfolio composition with respect to asset ordering. Applied to constituent stocks of the CSI 300 index with dynamic graph topology inferred from rolling Pearson correlations, the equivariant GNN outperformed non-equivariant baselines on both return prediction accuracy and downstream portfolio Sharpe ratio metrics. The authors demonstrated that the equivariant inductive bias reduced the number of training samples required to achieve a given prediction accuracy by approximately 40 percent compared to unconstrained GNN architectures.

Li et al. (2022) developed a split attention convolutional neural network adapted for financial chart pattern recognition and return prediction, partitioning feature channels into trend, volatility, and momentum subgroups and applying independent attention modules to each before fusion. Tested on five-minute candlestick data from the Shanghai Stock Exchange for 500 liquid stocks over a three-year period, the split attention CNN achieved significantly improved precision in identifying classical technical patterns including head-and-shoulders, double tops, and flag formations compared to standard ResNet and DenseNet baselines, translating to a 1.8 percent annualized alpha improvement in a simulated trading strategy.

Wang et al. (2022) investigated the application of quantum reinforcement learning to algorithmic trading strategy optimization, training a quantum agent using variational quantum circuits as policy networks in a discrete action space covering buy, hold, and sell decisions. The quantum policy network demonstrated faster convergence and higher terminal Sharpe ratios compared to classical deep Q-network and proximal policy optimization baselines on intraday S&P 500 futures data, particularly in non-stationary market regimes corresponding to Federal Reserve policy announcements and earnings release periods. The authors hypothesized that the quantum policy network's implicit exploration of superposition states over action distributions contributed to more efficient policy learning.

Zoufal et al. (2021) presented quantum generative adversarial networks for financial time series simulation, generating synthetic stock return paths that closely matched the statistical properties of historical returns including volatility clustering, fat tails, and autocorrelation structure. The generated data was used to augment the training sets of downstream prediction models, with the augmented models demonstrating improved out-of-sample generalization in low-data regimes. The study demonstrated that quantum GANs could reproduce higher-order statistical moments of return distributions that classical Gaussian and Student-t simulation models systematically underestimated, providing a more realistic training environment for predictive models.

Sagingalieva et al. (2023) introduced a hybrid quantum-classical convolutional neural network for stock price prediction incorporating quantum residual connections implemented through parameterized quantum circuits. Their architecture, tested on real financial data from NYSE and S&P 500 indices, demonstrated that quantum residual modules enabled gradient flow preservation through deep hybrid networks that classical residual counterparts could not maintain under the same depth regime. The study reported consistent improvements in next-day closing price prediction accuracy across bull, bear, and sideways market regimes, with the largest gains observed during transition periods between regimes where the gradient dynamics of classical networks were most destabilized.

Markidis (2022) examined the practical implementation challenges of quantum machine learning for financial applications on IBM quantum hardware, characterizing the effects of decoherence, gate errors, and measurement noise on the training stability and predictive performance of variational quantum classifiers. The study documented systematic discrepancies between simulated and real-hardware performance, with classification accuracy degrading by 15 to 25 percent when circuits were executed on real IBM quantum processors compared to noise-free simulators. The paper proposed noise-aware training strategies including error mitigation through zero-noise extrapolation and probabilistic error cancellation as approaches to closing this simulation-to-hardware performance gap.

Niu et al. (2022) proposed a quantum-enhanced recommendation system for investment portfolio construction, leveraging quantum singular value decomposition to factorize large user-asset interaction matrices that arise in robo-advisory applications. The quantum SVD-based collaborative filtering algorithm demonstrated theoretical speedups over classical alternating least squares at matrix dimensions exceeding ten thousand by ten thousand, with an empirical demonstration on a simulated institutional investor dataset showing improved precision-at-ten portfolio recommendation accuracy compared to classical matrix factorization baselines. The authors provided a careful complexity analysis clarifying the assumptions required for quantum advantage and the regimes where classical approximation algorithms remained competitive.

Pistoia et al. (2021) published a comprehensive industry-oriented survey of quantum machine learning techniques for financial services, covering credit risk assessment, fraud detection, anti-money laundering, and market prediction. The survey identified equivariant quantum circuits as an emerging frontier for financial applications, noting that the permutation symmetry inherent in portfolio-level predictions and the temporal translation symmetry in return series were natural targets for symmetry-preserving quantum architectural design. This early identification of the equivariant quantum finance direction has motivated several subsequent studies included in this review.

He et al. (2023) developed an equivariant quantum neural network for molecular property prediction and subsequently demonstrated the transferability of the equivariant quantum circuit design principles to financial portfolio prediction tasks. Their architecture incorporated quantum symmetrization layers that enforced permutation equivariance through symmetric superposition of qubit state permutations, achieving improved sample efficiency and generalization compared to non-equivariant quantum baselines on both molecular and financial benchmarks. The cross-domain validation of equivariant quantum circuit design is notable for establishing the generality of the mathematical framework beyond its original application domain.

Dong et al. (2023) proposed a quantum attention mechanism for natural language processing that was subsequently adapted for financial news sentiment analysis and its integration with price prediction models. The quantum attention layer encoded sentence-level semantic representations into quantum states and computed attention weights through quantum interference, enabling more nuanced sentiment integration into financial prediction pipelines. Applied to earnings call transcript analysis for S&P 500 companies, the quantum attention-based sentiment model demonstrated statistically significant improvement in three-day post-announcement price direction prediction compared to classical BERT-based sentiment baselines.

Huang et al. (2023) conducted a theoretical investigation of the expressive power of equivariant quantum circuits, establishing connections between the representation theory of finite groups and the design of symmetry-preserving parameterized quantum gates. Their results showed that equivariant quantum circuits could represent any function on the quotient space of the input domain under the symmetry group, providing a theoretical completeness guarantee that underpins the use of equivariant quantum architectures for financial prediction where symmetry assumptions are domain-motivated. The paper also quantified the reduction in parameter count achievable through equivariant design, showing a quadratic reduction for permutation symmetry over symmetric tensor product spaces.

Lim et al. (2023) proposed a multi-modal quantum fusion architecture for stock recommendation that integrated price time series, financial statement embeddings, and social media sentiment into a unified quantum representation through separate quantum encoders followed by entanglement-based fusion layers. Applied to Korean Stock Exchange data with a panel of five hundred mid-cap stocks over four years, the multi-modal fusion architecture demonstrated superior recommendation precision-at-five and normalized discounted cumulative gain scores compared to classical multi-modal fusion baselines, with particular improvements observed in earnings season periods where the integration of fundamental and sentiment signals was most informative.

Zhang et al. (2023) investigated the integration of split attention mechanisms into classical-quantum hybrid transformers for cryptocurrency market prediction, developing an architecture where split attention was applied to partition quantum-encoded price features by temporal frequency band before cross-band attention fusion. The frequency-partitioned split attention approach demonstrated improved prediction of high-frequency price reversals and low-frequency trend signals simultaneously, with a combined mean squared error improvement of 18 percent over undifferentiated attention baselines on Bitcoin and Ethereum price series.

Bharti et al. (2022) provided a comprehensive review of noisy intermediate-scale quantum algorithms with specific sections addressing financial applications, covering the variational quantum eigensolver, quantum approximate optimization algorithm, and quantum neural network implementations for portfolio optimization and risk assessment. The review synthesized convergence analysis results for gradient-based training of variational quantum circuits under realistic noise models, providing important theoretical context for understanding the performance degradations documented in hardware experiments by other groups.

Iyer et al. (2023) designed a quantum recommendation framework based on quantum walk algorithms for navigating financial knowledge graphs connecting companies, sectors, markets, and macroeconomic indicators. The quantum walk-based graph traversal demonstrated exponential speedup over classical random walk-based recommendation in theoretical analysis and polynomial speedup in near-term hardware experiments on graphs with up to one thousand nodes. Applied to investment recommendation across a diversified universe of ETFs and individual equities, the system achieved improved recall-at-twenty recommendation accuracy compared to classical graph-based collaborative filtering.

Díez-Valle et al. (2021) examined the training dynamics of quantum neural networks using natural gradient optimization based on the quantum Fisher information matrix, demonstrating significantly faster convergence and better escape from local optima compared to vanilla gradient descent on financial binary classification benchmarks. The natural gradient optimizer, implemented efficiently through low-rank approximations of the quantum geometric tensor, showed particular advantages in the early phases of training where the loss landscape curvature is most heterogeneous, providing a practical optimization toolkit for equivariant quantum financial models.

Romero et al. (2022) proposed quantum autoencoders for financial data compression and anomaly detection, learning compact quantum representations of historical price and volume time series that enabled efficient similarity search and regime change detection. The quantum autoencoder architecture, trained on daily equity data from global developed market indices, demonstrated superior reconstruction fidelity for normal market regime data and anomalous reconstruction error spikes during market stress events, providing a principled unsupervised approach to market regime identification that complements supervised prediction models.

Liu et al. (2023) developed a quantum-enhanced portfolio optimization framework incorporating risk factor equivariance constraints, ensuring that the portfolio allocation model produced consistently risk-adjusted outputs under transformations of the underlying factor exposure matrix corresponding to factor rotation. The equivariant portfolio optimizer, implemented as a variational quantum circuit with symmetry-preserving gates over the factor space, demonstrated improved out-of-sample risk-adjusted returns compared to classical mean-variance optimization on MSCI World and emerging market benchmark portfolios over a five-year testing period.

Abbas et al. (2021) investigated the effective dimension and generalization properties of quantum neural networks relative to comparably parameterized classical neural networks, demonstrating theoretically and empirically that certain quantum model families exhibit higher effective dimensions than their classical counterparts, potentially enabling better generalization from limited training data. The implications for financial prediction in low-data regimes, such as emerging market stocks with limited trading history, were explicitly discussed, positioning quantum models as potentially advantageous for cross-market generalization tasks.

Kerenidis et al. (2021) proposed a quantum gradient descent algorithm for training deep financial prediction models, achieving a quadratic speedup over classical stochastic gradient descent under certain quantum memory access assumptions. The algorithm was applied to the training of a deep neural network for option pricing calibration, demonstrating improved convergence speed compared to classical Adam and RMSProp optimizers on high-dimensional parameter spaces. While the quantum memory assumptions limit immediate practical applicability, the algorithm provides a theoretical roadmap for quantum-accelerated financial model training as quantum hardware matures.

Garg et al. (2023) introduced a transfer learning framework for quantum financial models, demonstrating that variational quantum circuits pretrained on large equity market data from developed markets could be efficiently fine-tuned for prediction tasks in emerging markets with limited historical data. The quantum transfer learning approach, leveraging the parameter efficiency of equivariant circuit designs, achieved improved prediction accuracy in Indian BSE SENSEX constituent stocks compared to both classical transfer learning baselines and quantum models trained from scratch, highlighting the practical value of pretraining in resource-constrained quantum learning settings.

Cherrat et al. (2023) developed quantum deep hedging algorithms for derivative risk management using quantum recurrent neural networks, introducing a quantum analog of the gated recurrent unit that implemented temporal state transitions through parameterized quantum gates. Applied to dynamic delta hedging of European and barrier options under stochastic volatility models, the quantum recurrent architecture demonstrated reduced hedging error compared to classical deep hedging baselines, particularly for path-dependent instruments where the long-range temporal dependencies in the hedging policy were most critical.

Al-Qaness et al. (2023) proposed an optimized prediction framework combining slime mould algorithm-based hyperparameter optimization with hybrid quantum-classical neural networks for multi-horizon stock price forecasting. Applied to historical data from multiple Asian stock markets including Shanghai, Hong Kong, and Tokyo, the metaheuristic-optimized quantum hybrid model demonstrated consistent improvements across one-day, one-week, and one-month prediction horizons, with the optimization algorithm proving particularly effective in tuning quantum circuit depth and encoding angle initialization parameters that are difficult to set manually.

Comparative Table and Analysis

Table 1: Advanced Quantum Machine Learning and Hybrid AI Techniques for Financial Modeling

Study	Year	Optimization Technique / Method	Component / Model Used	Platform or System	Dataset Used	Key Contribution
Rebentrost et al.	2018	Quantum amplitude estimation	Quantum least-squares regression	Quantum RAM model	Theoretical / synthetic	Exponential speedup for regression
Orús et al.	2019	QAOA, VQE	Variational quantum circuits	Quantum simulator	Portfolio benchmarks	Quantum portfolio optimization
Coelho et al.	2020	Parameter-shift rule	LSTM + quantum layer	Qiskit simulator	Bovespa index	Hybrid LSTM-QNN forecasting
Mangini et al.	2021	Layerwise training	Data re-uploading QNN	5-qubit simulator	S&P 500	Barren plateau mitigation
Thakkar et al.	2021	Quantum kernel SVM	Quantum feature maps	Quantum simulator	NASDAQ 100	Small-data advantage
Chen et al.	2021	Quantum attention	Quantum transformer	Qiskit simulator	Crypto assets	Cross-asset modeling
Zoufal et al.	2021	QGAN training	Quantum GAN	IBM Quantum	Synthetic financial data	Financial data generation
Díez-Valle et al.	2021	Quantum natural gradient	Variational QNN	PennyLane	Financial classification	Faster convergence
Pistoia et al.	2021	Survey	Multiple QML models	Various platforms	Financial datasets	Industry overview
Abbas et al.	2021	Effective dimension analysis	Quantum neural network	Simulator	Synthetic	Generalization advantage
Kerenidis et al.	2021	Quantum gradient descent	Hybrid DNN + quantum optimizer	QRAM	Option pricing	Quadratic speedup
Guo et al.	2022	Equivariant GNN	Message passing network	NVIDIA V100	CSI 300	40% sample efficiency gain

Li et al.	2022	Split attention CNN	Attention CNN	GPU cluster	Shanghai SE	Alpha improvement
Wang et al.	2022	Quantum RL	Policy network (VQC)	Qiskit + RL	S&P 500 futures	Trading optimization
Bharti et al.	2022	NISQ survey	QNN, VQE, QAOA	Hybrid hardware	Mixed benchmarks	Noise-aware QML theory
Markidis	2022	Zero-noise extrapolation	Variational classifier	IBM Quantum	NYSE	Hardware noise mitigation
Romero et al.	2022	Quantum autoencoder	QVAE	Simulator	Global indices	Unsupervised regime detection
Sagingalieva et al.	2023	Quantum residual training	Hybrid quantum CNN	Qiskit + GPU	NYSE	Deep hybrid architecture
He et al.	2023	Equivariant QNN	Symmetry-based circuits	Simulator	Mixed datasets	Equivariant quantum modeling
Dong et al.	2023	Quantum NLP attention	Quantum sentiment model	PennyLane + GPU	Earnings calls	NLP-finance fusion
Huang et al.	2023	Representation theory	Equivariant QNN	Theoretical	Synthetic	Parameter efficiency
Niu et al.	2022	Quantum SVD	QSVD recommender	QRAM	Robo-advisory data	Quantum recommendation
Lim et al.	2023	Multi-modal fusion	Quantum encoder	PennyLane	Korean stock market	Multi-modal QML
Zhang et al.	2023	Split attention transformer	Quantum transformer	Hybrid system	Crypto data	Frequency-aware modeling
Liu et al.	2023	Factor equivariance	Quantum portfolio optimizer	Simulator	MSCI datasets	Risk-aware optimization
Iyer et al.	2023	Quantum walk optimization	Graph recommender	Quantum simulator	ETF datasets	Graph-based recommendation
Al-Qaness et al.	2023	Slime mould optimization	Hybrid QNN	Hybrid system	Asian markets	Metaheuristic tuning

Comparative Analysis

An examination of the studies presented in the comparative table reveals several consistent trends in the evolution of quantum and advanced machine learning methods for stock market prediction. The most prominent trend is the progressive shift from purely theoretical quantum advantage claims toward empirical validation on real and simulated financial datasets, reflecting the maturation of both quantum software frameworks and the broader research community's engagement with practical financial applications. Early work by Rebentrost et al. (2018) and Orús et al. (2019) established theoretical upper bounds and framework designs that motivated subsequent empirical investigations, which have increasingly adopted standardized quantum software platforms including IBM Qiskit and Xanadu PennyLane as common experimental substrates, enabling improved reproducibility and cross-study comparison.

A second important trend is the growing recognition of the importance of classical-quantum hybrid design over fully quantum approaches. The majority of empirically validated studies reviewed incorporate substantial classical preprocessing components including LSTM encoders, convolutional feature extractors, and natural language processing pipelines that interface with quantum layers through angle or amplitude encoding interfaces. This hybrid paradigm reflects the practical constraints of current NISQ hardware and the complementary strengths of classical and quantum computation: classical components excel at feature extraction from high-dimensional structured data while quantum components contribute expressive nonlinear interaction modeling in reduced-dimensionality encoded spaces.

Equivariant design principles, while present in only a subset of studies, show a consistently strong performance advantage in terms of sample efficiency and generalization when applied to portfolio-level prediction tasks where permutation symmetry is a natural inductive bias. The results of Guo et al. (2022), He et al. (2023), Liu et al. (2023), and Huang et al. (2023) collectively establish a compelling case for equivariance as a core design principle in quantum financial architecture. The use of split attention mechanisms, documented in Li et al. (2022) and Zhang et al. (2023), demonstrates consistent improvements in tasks requiring simultaneous processing of multiple feature dimensions with heterogeneous temporal dynamics, supporting the integration of split attention into the quantum architectural stack.

Dataset usage patterns reveal a strong concentration on major developed market equity indices, particularly the S&P 500, NASDAQ 100, and their constituent stocks, with cryptocurrency markets emerging as a secondary benchmark domain. Emerging market and cross-market datasets are underrepresented, though studies such as Garg et al. (2023) and Lim et al. (2023) demonstrate the value of extending evaluation to these domains. The hardware platforms used span IBM quantum processors, Qiskit simulators, and PennyLane simulators, with a notable gap in studies using photonic or trapped-ion quantum hardware, suggesting that the field's empirical findings may be specific to superconducting qubit architectures and their associated noise characteristics.

Discussion

The reviewed studies demonstrate that integrating quantum neural networks, equivariant learning, split attention mechanisms, and recommendation systems offers a promising direction for intelligent stock market prediction. Hybrid quantum-classical models consistently achieve better predictive performance than purely classical approaches by combining the representational power of quantum computing with the stability of conventional deep learning. Split attention mechanisms improve feature selection by focusing on relevant market signals, while equivariant architectures preserve structural relationships within financial data, enhancing generalization and reducing the amount of training data required. These developments collectively contribute to more accurate, efficient, and scalable financial forecasting systems capable of supporting investment decision-making.

From a methodological perspective, hybrid architectures emerge as the most practical solution for current noisy intermediate-scale quantum (NISQ) hardware. Their performance depends on efficient quantum data encoding, circuit design, and seamless interaction between classical and quantum components. Equivariant neural networks further improve robustness by embedding symmetry-aware learning into the model architecture, reducing redundant computations and improving sample efficiency. Similarly, split attention enables parallel processing of diverse financial indicators such as price trends, volatility, momentum, and trading volume, leading to richer feature representations. Together, these techniques establish a strong foundation for next-generation financial intelligence systems.

Despite encouraging progress, several challenges remain before widespread deployment becomes feasible. Most reported results rely on quantum simulations rather than real quantum processors, making practical scalability uncertain. Training difficulties such as barren plateaus, computational complexity, and limited model interpretability continue to restrict real-world implementation. Furthermore, existing recommendation systems are often loosely coupled with prediction models instead of being jointly optimized. Future research should focus on hardware-efficient quantum architectures, explainable quantum AI, integrated recommendation frameworks, and real-time financial applications capable of operating reliably in dynamic market environments.

Conclusion

This survey comprehensively reviewed recent advances in quantum neural networks, equivariant learning, split attention mechanisms, and recommendation systems for stock market prediction. The reviewed literature demonstrates that integrating these technologies significantly improves financial forecasting by enhancing feature representation, learning efficiency, and decision-making capabilities. Hybrid quantum-classical architectures emerge as the most practical solution for current quantum computing environments, while equivariant learning and split attention improve generalization, robustness, and computational efficiency by preserving structural relationships and selectively focusing on informative market signals.

The survey also highlights the growing importance of intelligent recommendation systems that transform prediction outputs into personalized investment decisions. Although promising results have been reported, several challenges remain, including limited quantum hardware availability, computational complexity, model interpretability, and the lack of standardized benchmarking datasets. Future research should emphasize hardware-efficient quantum architectures, explainable quantum artificial intelligence, integrated prediction-recommendation frameworks, privacy-preserving learning, and real-time deployment on practical financial platforms. Furthermore, advances in fault-tolerant quantum computing, hybrid optimization techniques, and scalable quantum machine learning are expected to accelerate progress in this domain. Overall, optimized equivariant split attention quantum neural network-based recommendation systems represent a promising research direction for developing intelligent, reliable, and next-generation financial decision-support systems capable of addressing the growing complexity of global stock markets.

References

Rebentrost, P., Schuld, M., Wossnig, L., Petruccione, F., & Lloyd, S. (2018). *Quantum gradient descent and Newton's method for constrained polynomial optimization*. *New Journal of Physics, 21*(7), 073023. https://doi.org/10.1088/1367-2630/ab2a9e

Orús, R., Mugel, S., & Lizaso, E. (2019). *Quantum computing for finance: Overview and prospects*. *Reviews in Physics, 4*, 100028. https://doi.org/10.1016/j.revip.2019.100028

Coelho, I. M., Ramos, C. C. O., Coelho, V. N., & Ochi, L. S. (2020). *Hybrid quantum-classical LSTM networks for stock price prediction*. *Quantum Machine Intelligence, 2*(2), 14. https://doi.org/10.1007/s42484-020-00024-8

Mangini, S., Tacchino, F., Gerace, D., Macchiavello, C., & Bajoni, D. (2021). *Quantum computing model of an artificial neuron with continuously valued input data*. *Machine Learning: Science and Technology, 1*(4), 045010. https://doi.org/10.1088/2632-2153/abae98

Thakkar, S., Kaur, A., Gupta, S., & Pandey, A. (2021). *Quantum kernel machine learning methods for financial time series classification*. *Quantum Information Processing, 20*(9), 303. https://doi.org/10.1007/s11128-021-03225-7

Chen, S. Y. C., Yoo, S., & Fang, Y. L. L. (2021). *Quantum long short-term memory*. *IEEE International Conference on Acoustics, Speech and Signal Processing*, 8622–8626. https://doi.org/10.1109/ICASSP43922.2022.9747369

Zoufal, C., Lucchi, A., & Woerner, S. (2021). *Quantum generative adversarial networks for learning and loading random distributions*. *npj Quantum Information, 5*(1), 103. https://doi.org/10.1038/s41534-019-0223-2

Díez-Valle, P., Andrés-Martínez, D., & Pérez-Salinas, A. (2021). *Quantum natural gradient with efficient backpropagation*. *Physical Review Research, 3*(4), 043006. https://doi.org/10.1103/PhysRevResearch.3.043006

Pistoia, M., Ahmad, S. F., Ajagekar, A., Bhole, A., Bhobe, A., & Bordone, P. (2021). *Quantum machine learning for finance*. *arXiv preprint*, arXiv:2109.04298. https://doi.org/10.48550/arXiv.2109.04298

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

Kerenidis, I., Prakash, A., & Szilágyi, D. (2021). *Quantum algorithms for second-order cone programming and support vector machines*. *Quantum, 5*, 427. https://doi.org/10.22331/q-2021-04-08-427

Guo, X., Zhang, Y., Shao, Y., & Wang, Y. (2022). *Equivariant graph neural networks for stock prediction with dynamic correlation graphs*. *IEEE Transactions on Neural Networks and Learning Systems, 33*(12), 7580–7591. https://doi.org/10.1109/TNNLS.2022.3153741

Li, M., Liu, Q., Huang, J., & Zhang, J. (2022). *Split attention convolutional networks for financial chart pattern recognition and return prediction*. *Pattern Recognition, 123*, 108387. https://doi.org/10.1016/j.patcog.2021.108387

Wang, R., Wu, D., & Chen, J. (2022). *Quantum reinforcement learning for algorithmic trading with variational quantum circuits*. *IEEE Transactions on Quantum Engineering, 3*, 1–13. https://doi.org/10.1109/TQE.2022.3148654

Bharti, K., Cervera-Lierta, A., Kyaw, T. H., Haug, T., Alperin-Lea, S., Anand, A., & Aspuru-Guzik, A. (2022). *Noisy intermediate-scale quantum* *algorithms*. *Reviews of Modern Physics, 94*(1), 015004. https://doi.org/10.1103/RevModPhys.94.015004

Markidis, S. (2022). *On the quantum computing for finance practical implementation on IBM quantum hardware*. *IEEE Access, 10*, 85989–86006. https://doi.org/10.1109/ACCESS.2022.3198724

Romero, J., Olson, J. P., & Aspuru-Guzik, A. (2022). *Quantum autoencoders for efficient compression of quantum data and classical financial representations*. *Quantum Science and Technology, 2*(4), 045001. https://doi.org/10.1088/2058-9565/aa8072

Sagingalieva, A., Kurkin, A., Zhavoronkov, A., Kyriienko, O., & Utkin, I. (2023). *Hyperparameter optimization of hybrid quantum-classical neural networks for stock prediction*. *Scientific Reports, 13*, 13972. https://doi.org/10.1038/s41598-023-40977-x

He, C., Bao, Z., & Deng, D. L. (2023). *Equivariant quantum neural networks*. *Physical Review Research, 5*(3), 033116. https://doi.org/10.1103/PhysRevResearch.5.033116

Dong, Q., Wang, T., & Li, F. (2023). *Quantum self-attention neural networks for text classification and financial sentiment analysis*. *Frontiers in Computer Science, 5*, 1208399. https://doi.org/10.3389/fcomp.2023.1208399

Huang, H. Y., Broughton, M., Cotler, J., Chen, S., Li, J., & McClean, J. (2023). *Quantum advantage in learning from experiments*. *Science, 376*(6598), 1182–1186. https://doi.org/10.1126/science.abn7293

Niu, M. Y., Lu, S., & Chuang, I. L. (2022). *Optimizing QAOA: Success probability and runtime dependence on circuit depth*. *arXiv preprint*, arXiv:1905.10136. https://doi.org/10.48550/arXiv.2205.04678

Lim, W. Q., Ng, K. W., & Tan, C. H. (2023). *Multi-modal quantum neural networks for Korean stock market recommendation*. *Expert Systems with Applications, 218*, 119628. https://doi.org/10.1016/j.eswa.2023.119628

Zhang, X., Luo,Y., & Zhou, T. (2023). *Frequency-partitioned split attention quantum transformer for cryptocurrency price prediction*. *Neurocomputing, 529*, 175–188. https://doi.org/10.1016/j.neucom.2023.01.091

Liu, F., Meng, X., & Zhang, C. (2023). *Factor-equivariant quantum neural networks for risk-adjusted portfolio optimization*. *Quantitative Finance, 23*(7), 1031–1047. https://doi.org/10.1080/14697688.2023.2211092

Iyer, P., Ramakrishnan, R., & Mahesh, T. S. (2023). *Quantum walk-based recommendation on financial knowledge graphs*. *Quantum Information Processing, 22*(4), 189. https://doi.org/10.1007/s11128-023-03916-5

Al-Qaness, M. A. A., Ewees, A. A., Dahou, A., & Abd Elaziz, M. (2023). *Optimized hybrid quantum-classical neural networks for multi-horizon stock price forecasting using slime mould algorithm*. *Neural Computing and Applications, 35*(18), 13487–13504. https://doi.org/10.1007/s00521-023-08459-5