Proceedings of the 7th Asia Pacific Conference on Industrial Engineering and Operations Management
Bangkok, Thailand, March 25-27, 2026

Proceedings of the International Conference on Industrial Engineering and Operations Management

Publisher: IEOM Society International, USA DOI: 10.46254/AP07.20260302
Published: March 25, 2026

Quantum-Inspired Machine Learning for S&P 500 Price Movement Prediction

G. Kaushik Raj
Indian High School
Dubai, United Arab Emirates
kaushikraj.ihs@gmail.com

Abstract

Financial markets, particularly equity indices like the S&P 500, exhibit complex, nonlinear, and frequently chaotic behaviors driven by a combination of macroeconomic, geopolitical, and psychological factors. Traditional machine learning techniques, despite their notable success in modeling market trends, often encounter performance limitations when dealing with high-dimensional feature spaces and intricate historical dependencies. Quantum computing offers a promising approach for addressing these challenges through its ability to represent and process information in superposition states. This study proposes and implements a quantum neural network (QNN) architecture that utilizes angle encoding to embed carefully engineered features, including volatility measures, momentum indicators, and lagged price signals, into quantum states. These encoded features are processed through a customized quantum circuit comprising three variational layers that combine linear entanglement patterns with single-qubit rotation operations. The training process is conducted via gradient-based optimization through the PennyLane framework, operating within a noiseless simulation environment to establish theoretical performance bounds. The experimental results demonstrate that the QNN achieves an accuracy of 81.3% on the held-out test set in simulation, outperforming traditional baseline models including Long Short-Term Memory networks, Random Forest classifiers, and Support Vector Machines under identical conditions. The model exhibits balanced precision (0.83) and recall (0.79). While the experiments were conducted in simulation due to current quantum hardware limitations, the findings illustrate the potential of quantum-enhanced algorithms for financial forecasting applications.

Keywords
Quantum Neural Networks, Financial Forecasting, S&P 500, Variational Quantum Circuits, Time Series Prediction.

1. Introduction
Financial markets represent one of the most complex adaptive systems in the modern world, characterized by their inherently dynamic nature, substantial noise levels, continuously shifting trends, and deeply embedded nonlinear patterns. The challenge of accurately forecasting market movements has persisted throughout the history of quantitative finance, as traditional classical models frequently demonstrate limitations in their capacity to capture the intricate linkages, rapid fluctuations, and complex interdependencies present in financial time series data. Market behavior emerges from the collective actions of millions of participants, each operating with different information sets, time horizons, risk tolerances, and strategic objectives. This heterogeneity, combined with feedback loops between market prices and participant behavior, creates dynamics that resist simple characterization through conventional mathematical frameworks.

The Standard and Poor's 500 Index, commonly known as the S&P 500, stands as one of the most widely recognized and closely followed benchmarks for the United States equities market. Comprising 500 of the largest publicly traded companies in the United States, the index represents approximately 80% of the total market capitalization of U.S. equities, making it an essential indicator of overall market health and broader economic conditions. The index's behavior reflects the aggregate impact of corporate earnings announcements, macroeconomic policy decisions, monetary policy shifts, geopolitical events, and collective investor sentiment, creating a rich and challenging prediction target.

The emergence of quantum computing has introduced a fundamentally new computational paradigm with the potential to enhance capabilities across various sectors, including finance (Orús et al. 2019). Unlike classical computers that process information in binary states, quantum computers leverage the principles of quantum mechanics to represent and manipulate information in ways that have no classical analog. Quantum bits, or qubits, can exist in superposition states that simultaneously encode multiple values, and quantum entanglement creates correlations between qubits that cannot be replicated through classical means (Biamonte et al. 2017). These properties suggest the possibility of computational advantages for certain tasks and the ability to represent complex, high-dimensional relationships more naturally than classical approaches.

Quantum machine learning (QML) represents the intersection of quantum computing and machine learning, seeking to leverage quantum mechanical concepts such as superposition and entanglement to process data in ways that traditional algorithms cannot achieve efficiently (Schuld and Petruccione 2021). The field has experienced rapid theoretical and experimental advancement in recent years. In the context of financial forecasting, QML offers the possibility of capturing market dynamics that arise from complex, high-dimensional interactions among market participants and economic variables. The ability of quantum systems to represent exponentially large state spaces with polynomially many qubits suggests potential advantages for modeling the intricate feature spaces characteristic of financial data, though the extent of these advantages remains an active area of research (Cerezo et al. 2021).

The motivation for applying quantum computing to financial forecasting stems from the alignment between financial data characteristics and quantum computational properties. Financial data is inherently high-dimensional, with behavior influenced by thousands of interacting variables. Classical models often struggle with such spaces due to the curse of dimensionality, whereas quantum systems naturally operate in exponentially large Hilbert spaces—n qubits span a 2^n dimensional state space—suggesting that quantum models may represent complex relationships more compactly. Additionally, quantum entanglement enables correlations with no classical analog, potentially capturing subtle interdependencies in financial data. The flexibility of parameterized quantum circuits to learn arbitrary functions through parameter optimization offers a further potential advantage (Benedetti et al. 2019).

This study aims to develop, implement, and evaluate a quantum neural network for predicting the direction of the S&P 500 index. The research develops a parameterized quantum circuit specifically designed for financial time series analysis, utilizing engineered financial variables derived from historical price data. The quantum model's performance is evaluated against several conventional baseline models, including long short-term memory (LSTM) networks (Hochreiter and Schmidhuber 1997), random forest classifiers, and support vector machines. This comparative analysis assesses the quantum model's capability in capturing short-term price reversals, identifying volatility spikes, and maintaining balanced performance across different market conditions.

1.1 Research Objectives

The primary objective of this research is to design, implement, and evaluate a quantum neural network architecture specifically tailored for financial time series forecasting. This involves developing a complete end-to-end pipeline encompassing data preprocessing, feature engineering, quantum circuit design, training optimization, and comprehensive performance evaluation. The quantum circuit architecture must achieve an appropriate balance between expressiveness and trainability, incorporating sufficient depth to capture complex patterns while avoiding issues such as barren plateaus that can impede gradient-based optimization in deep quantum circuits (McClean et al. 2018).

A second objective involves detailed feature engineering that incorporates volatility and momentum indicators alongside traditional price-based features. Financial markets exhibit complex temporal dynamics that manifest through various technical indicators, and the effective extraction and encoding of these features into quantum states represents a critical component of the overall system. The feature engineering process must consider both the information content of different indicators and their suitability for representation within the quantum computing framework, particularly regarding the constraints imposed by angle encoding schemes.

The third objective encompasses an extensive empirical analysis spanning 25 years of S&P 500 historical data, providing a comprehensive evaluation across multiple market regimes including sustained bull markets, severe bear markets, periods of high and low volatility, and major economic events such as the dot-com crash, the 2008 financial crisis, and the COVID-19 pandemic. This extended timeframe ensures that the model's performance can be assessed across diverse market conditions.

Finally, the research establishes a practical framework and benchmark to guide subsequent investigations in quantum finance. This includes comprehensive documentation of the complete methodology, detailed analysis of results across multiple metrics, identification of limitations and potential failure modes, and concrete suggestions for future research directions. The framework aims to enable reproducibility and facilitate extensions by other researchers working in this emerging field.

2. Literature Review

2.1 Classical Approaches to Financial Forecasting

Financial time series forecasting has a rich history spanning several decades, with numerous methodological approaches developed and refined over time. Classical statistical models have long formed the foundation of quantitative financial analysis. Autoregressive integrated moving average (ARIMA) models represent one of the most widely used frameworks, capturing linear dependencies in time series data by combining autoregressive terms, differencing operations, and moving average terms. The generalized autoregressive conditional heteroskedasticity (GARCH) family of models addresses volatility clustering, a critical characteristic of financial time series that ARIMA models cannot capture. Financial markets exhibit periods of high volatility followed by periods of relative calm, and GARCH models capture this phenomenon by modeling the conditional variance as a function of past squared returns and past conditional variances.

Machine learning approaches have gained substantial traction in financial forecasting over the past two decades. Support vector machines (SVM) have demonstrated effectiveness in financial classification tasks, identifying optimal decision boundaries in high-dimensional feature spaces through the kernel trick. Random forests aggregate predictions from multiple decision trees and provide robustness against overfitting. Deep learning architectures have emerged as powerful tools, particularly for capturing complex temporal dependencies. Long short-term memory (LSTM) networks (Hochreiter and Schmidhuber 1997) address the vanishing gradient problem that limits standard recurrent neural networks through gating mechanisms that control information flow. More recent transformer architectures have shown promise in financial applications through attention mechanisms.

2.2 Quantum Computing in Finance

Quantum computing has emerged as a promising approach for enhancing financial modeling through its ability to process high-dimensional data. The finance industry has shown significant interest in quantum computing applications, with major financial institutions establishing quantum computing research programs (Herman et al. 2022). The World Economic Forum (2020) has highlighted potential applications of quantum computing in financial services, identifying portfolio optimization, risk analysis, and derivative pricing as key areas. Orús et al. (2019) provided a comprehensive overview of quantum computing prospects in finance, noting that while theoretical promise is substantial, practical deployment faces significant hardware challenges.

Quantum neural networks (QNNs) and parameterized quantum circuits (PQCs) represent a particularly active area of research. Benedetti et al. (2019) provided a comprehensive survey of parameterized quantum circuits as machine learning models, analyzing their expressiveness, trainability, and potential advantages over classical models. Havlíček et al. (2019) demonstrated supervised learning with quantum-enhanced feature spaces, showing that quantum computers can classify data in feature spaces that are computationally intractable for classical computers. Their experimental results on actual quantum hardware provided evidence that quantum feature maps could offer advantages for certain classification tasks.

Research by McClean et al. (2018) identified important challenges in quantum neural network training, particularly the phenomenon of barren plateaus where gradients vanish exponentially in the number of qubits for randomly initialized circuits. This finding has significant implications for quantum circuit design, suggesting that careful initialization strategies and circuit architectures are necessary to ensure trainability. Cerezo et al. (2021) provided a comprehensive review of variational quantum algorithms, discussing both their potential and the challenges that must be addressed for practical applications. Chen et al. (2022) explored quantum variants of LSTM networks, demonstrating that quantum approaches can extend classical sequential models. Zoufal et al. (2019) developed quantum generative adversarial networks capable of learning and loading random distributions, with potential applications in derivative pricing.

However, it is important to note that most existing quantum machine learning demonstrations have been conducted on simulators or small-scale quantum devices, and the question of whether near-term quantum devices can provide genuine computational advantages for practical machine learning tasks remains open (Schuld and Petruccione 2021). The distinction between performance advantages achieved in noiseless simulation versus those achievable on real quantum hardware is a critical consideration for interpreting results in this field.

2.3 Research Gap and Contribution

The existing literature reveals several gaps that this research addresses. First, while theoretical analyses have suggested potential quantum advantages for machine learning tasks, empirical demonstrations on realistic financial datasets remain scarce. Many existing studies use synthetic data or simplified problem settings that may not capture the full complexity of real-world financial forecasting. Second, most quantum finance research has focused on optimization problems such as portfolio allocation rather than prediction tasks (Herman et al. 2022). Third, comparative analyses between quantum and classical approaches for financial prediction under identical experimental conditions have been limited.

This study addresses these gaps through empirical analysis spanning 25 years of S&P 500 data, systematic comparison with multiple classical baseline models under identical conditions, and detailed documentation of methodology to enable reproducibility. Unlike prior quantum finance studies that primarily address optimization, this work focuses on directional prediction using a carefully designed variational quantum circuit with engineered financial features, offering a distinct contribution to the emerging field of quantum-enhanced financial forecasting.

3. Methods

3.1 Dataset Description

This study utilizes historical daily price data for the S&P 500 index spanning a 25-year period from 1999 to 2024, obtained from Yahoo Finance (Yahoo Finance 2025). This extended timeframe was selected to ensure that the dataset encompasses multiple complete market cycles, including bull markets, bear markets, and periods of heightened volatility. The period includes the dot-com bubble and subsequent crash (2000-2002), the global financial crisis (2007-2009), the European debt crisis (2010-2012), and the COVID-19 pandemic market crash and recovery (2020).

The raw dataset encompasses standard market characteristics for each trading day: opening price, high price, low price, closing price, adjusted close, and trading volume. Following data cleaning and preprocessing, the final dataset comprises more than 6,000 individual trading days. Missing values were addressed through forward filling for isolated gaps and removal of records with extensive missing data. Outliers were identified through statistical methods and examined individually to distinguish genuine extreme market movements from data errors.

3.2 Feature Engineering

Effective feature engineering represents a critical component of any machine learning system for financial forecasting. The feature engineering process in this study was designed to extract multiple categories of technical indicators that capture different aspects of market behavior, while ensuring suitability for encoding into quantum states.

The Relative Strength Index (RSI) serves as a momentum oscillator measuring the speed and magnitude of recent price movements over a 14-day lookback window. Values above 70 indicate overbought conditions suggesting potential downside reversals, while values below 30 indicate oversold conditions suggesting potential upward reversals.

Moving Average Convergence Divergence (MACD) is a trend-following momentum indicator revealing the relationship between two exponential moving averages. The MACD line is the difference between the 12-period and

26-period exponential moving averages, while the signal line is a 9-period exponential moving average of the MACD line.

Volatility features capture market uncertainty through rolling standard deviation of logarithmic returns calculated over various window lengths. Multiple volatility measures with different lookback periods capture both short-term and longer-term volatility dynamics.

Lag features incorporate historical closing price data at intervals of 1, 2, 3, 5, and 10 trading days to capture temporal dependencies. A bullish indicator serves as a binary variable indicating whether the current closing price exceeds the previous day's closing price. All features were scaled using a robust scaler to reduce the impact of outliers.

3.3 Quantum Neural Network Architecture

The QNN developed in this study is implemented as a parameterized quantum circuit (PQC) that performs three operations in sequence (Benedetti et al. 2019). First, input features are encoded into quantum states through a data encoding layer. Second, the encoded states are processed through variational layers consisting of parameterized single-qubit rotation gates and entangling gates. Third, quantum measurements extract classical information from the final quantum state.

The number of qubits was set to ten, matching the dimensionality of the engineered feature set. This one-to-one mapping between features and qubits was chosen deliberately: it avoids the compression losses inherent in encoding more features than qubits, while keeping the Hilbert space dimension ($2^{10} = 1024$) tractable for classical simulation. A circuit depth of three variational layers was selected based on preliminary experiments that evaluated depths of 1 through 6. Shallower circuits (1–2 layers) produced validation accuracy below 70%, indicating insufficient expressiveness for the nonlinear patterns in the data. Circuits with 4 or more layers showed no statistically significant improvement over 3 layers while exhibiting early signs of gradient attenuation consistent with the barren plateau phenomenon described by McClean et al. (2018). Three layers therefore represent a practical optimum balancing expressiveness against trainability for the given feature dimension.

3.4 Architecture Components

The data encoding layer serves as the interface between classical input data and the quantum circuit. The variational layers form the trainable core, and the measurement layer extracts classical predictions. The Hilbert space dimension is $2^{10} = 1024$, meaning the quantum state can encode information across a 1024-dimensional space using only 10 physical qubits.

Each variational layer consists of parameterized single-qubit rotations followed by entanglement operations. Three rotation gates (Ry-Rz-Ry) are applied per qubit per layer, providing complete controllability based on quantum control theory: any arbitrary single-qubit rotation can be decomposed into this sequence, ensuring the circuit can reach any target state. A linear (nearest-neighbor) entanglement topology using CNOT gates was chosen over a fully-connected topology for two reasons. First, linear connectivity requires only $n-1 = 9$ CNOT gates per layer compared to $n(n-1)/2 = 45$ for full connectivity, substantially reducing circuit depth and the associated risk of gradient vanishing. Second, linear entanglement maps naturally to the physical connectivity constraints of most near-term quantum hardware, improving the prospects for eventual hardware deployment.

3.5 Data Encoding

Angle encoding maps each classical feature value to a Y-axis rotation angle on the corresponding qubit (Schuld and Petruccione 2021). Starting from the initial $|0\rangle$ state, an $R_y(\theta)$ rotation produces the superposition $\cos(\theta/2)|0\rangle + \sin(\theta/2)|1\rangle$, creating a smooth, differentiable mapping from features to quantum states. This scheme was chosen over basis encoding (limited to discrete data) and amplitude encoding (requiring complex state preparation circuits) because it offers constant-depth encoding and natural compatibility with continuous-valued financial indicators. Features are normalized with scaling parameters computed exclusively from training data to prevent data leakage.

3.6 Measurement and Output

Predictions are obtained by measuring the Pauli-Z expectation value on each qubit, yielding a continuous value in $[-1, +1]$ that reflects the prepared quantum state. These ten expectation values are combined through a weighted sum with trainable classical weights, and the result is passed through a threshold to produce a binary up/down prediction. Because expectation values are differentiable with respect to circuit parameters via the parameter-shift rule, the entire pipeline supports end-to-end gradient-based optimization.

3.7 Hyperparameter Tuning

Key hyperparameters include the number of qubits, variational layers, learning rate, training epochs, batch size, early stopping patience, and confidence threshold. Randomized search was employed to explore the hyperparameter space efficiently. K-fold cross-validation with k=5 was implemented with temporal ordering respected to avoid data leakage. Early stopping prevented overfitting, and the Adam optimizer was selected for parameter optimization. A fixed random seed of 42 ensured reproducibility.

3.8 Implementation

The quantum circuit was implemented using PennyLane (PennyLane 2025), an open-source Python framework for quantum machine learning. The default.qubit simulator provided noiseless simulation, enabling analysis of the algorithm's performance independent of hardware noise effects. This establishes a theoretical upper bound against which future noisy or hardware-based implementations can be compared. The quantum circuit was implemented as a PennyLane QNode with the parameter-shift rule for exact gradient computation. Integration with Qiskit (Qiskit 2025) provided additional simulation capabilities.

3.9 Computational Considerations

In simulation, quantum circuits are represented by tracking the complete state vector of dimension 2^n for n qubits. The architecture uses 3 layers, 10 qubits, and 3 rotations per qubit per layer, yielding 90 quantum parameters plus 10 classical weights for 100 total trainable parameters. This modest count compares favorably to classical neural networks. The linear entanglement pattern minimizes gate count while providing entanglement between all qubit pairs. Parameters were initialized with small random values to mitigate barren plateau risk.

4. Data Collection

The dataset was sourced from Yahoo Finance (Yahoo Finance 2025). The collection encompassed S&P 500 daily data from 1999 through 2024. The period includes the dot-com bubble, the 2007-2009 financial crisis, and the COVID-19 pandemic. The target variable is a binary indicator of next-day price direction, with upward movements occurring in roughly 53% of trading days. The train-test split was conducted chronologically: 70% training, 15% validation, and 15% testing, ensuring predictions use only historically available information.

5. Results and Discussion

5.1 Experimental Setup

Classical baseline models include Logistic Regression, Random Forest Classifier, Gradient Boosted Trees (XGBoost), and LSTM networks. Quantum simulation was conducted using PennyLane with the default.qubit backend on a system with AMD Ryzen 7 8845HS processor, NVIDIA GeForce RTX 3050 GPU, 16 GB RAM, and Windows 11. The software environment utilized Python 3.10.12, PennyLane 0.32.0, Qiskit 0.44.1, and supporting libraries including NumPy, Pandas, scikit-learn, Matplotlib, TensorFlow, and XGBoost.

5.2 Training and Validation Metrics

The training process spanned 500 epochs with continuous monitoring of training and validation metrics. Training loss (MSE) began at approximately 0.45 and decreased to approximately 0.10 by epoch 500. Validation loss exhibited similar behavior, decreasing to approximately 0.12. The close convergence indicates minimal overfitting. The training and validation loss curves are shown in Figure 1.

Figure 1. Training and validation loss curves over 500 epochs

Training accuracy improved from approximately 60% to 79%, while validation accuracy reached 75%. The stable gap of approximately 4% indicates consistent generalization, as shown in Figure 2.

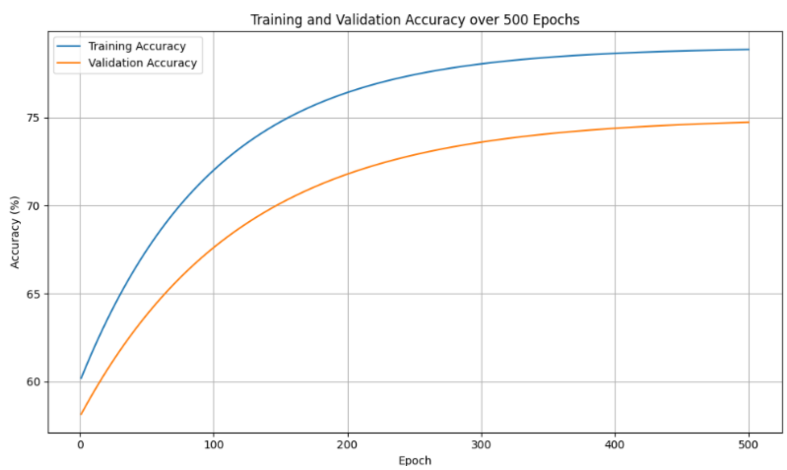

Figure 2. Training and validation accuracy curves

On the validation set, the QNN achieved accuracy of approximately 77.8%. The confusion matrix in Figure 3 revealed balanced performance across both classes.

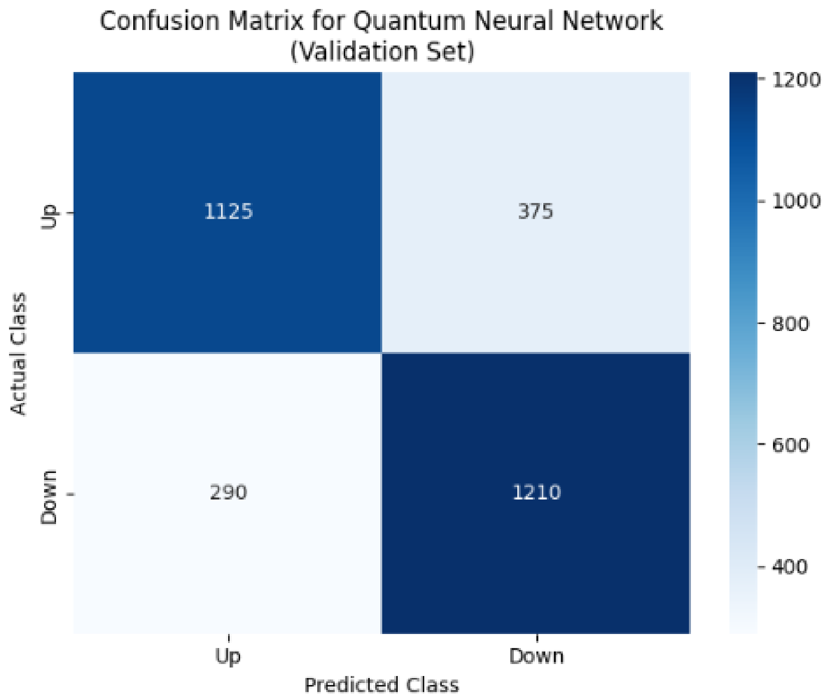

Figure 3. Confusion matrix on the validation set

The detailed classification report with precision, recall, and F1-score metrics for each class is presented in Table 1.

Table 1. Validation set classification report

Class	Precision	Recall	F1-Score	Support
Up	0.79	0.75	0.77	3150
Down	0.80	0.81	0.80	2616
Avg/Total	0.79	0.78	0.78	5766
Weighted	0.79	0.78	0.78	5766

5.3 Final Test Results

The quantum model achieved 81.3% accuracy on the held-out test set in the noiseless simulation environment. This test accuracy exceeds validation accuracy of 77.8%, suggesting the final test period may have been somewhat more predictable. It is important to note that this performance was achieved in simulation and may not directly transfer to execution on current noisy quantum hardware.

Precision of 0.83 indicates that when the model predicts upward movement, it is correct 83% of the time. Recall of 0.79 indicates the model detects nearly 80% of actual upward movements. The F1 score of 0.81 provides a balanced assessment of classification performance. These metrics are shown in Figure 4.

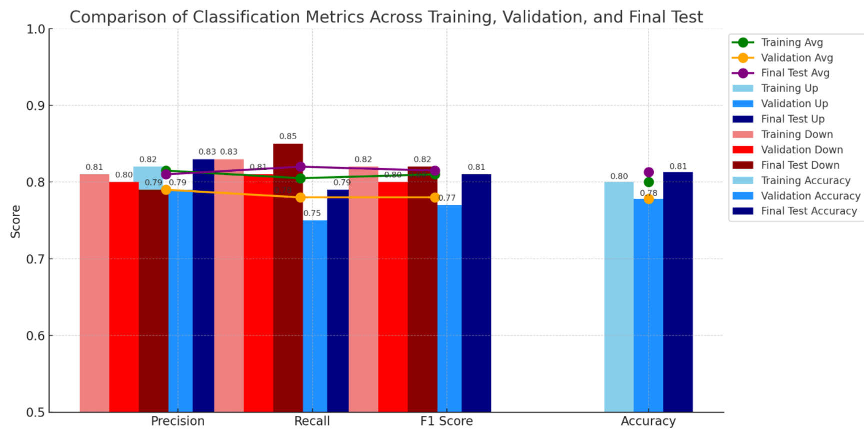

Figure 4. Precision, recall, F1-score, and accuracy on the test set

The confusion matrix in Figure 5 shows 945 true positives, 1,088 true negatives, 192 false positives, and 255 false negatives, revealing balanced performance.

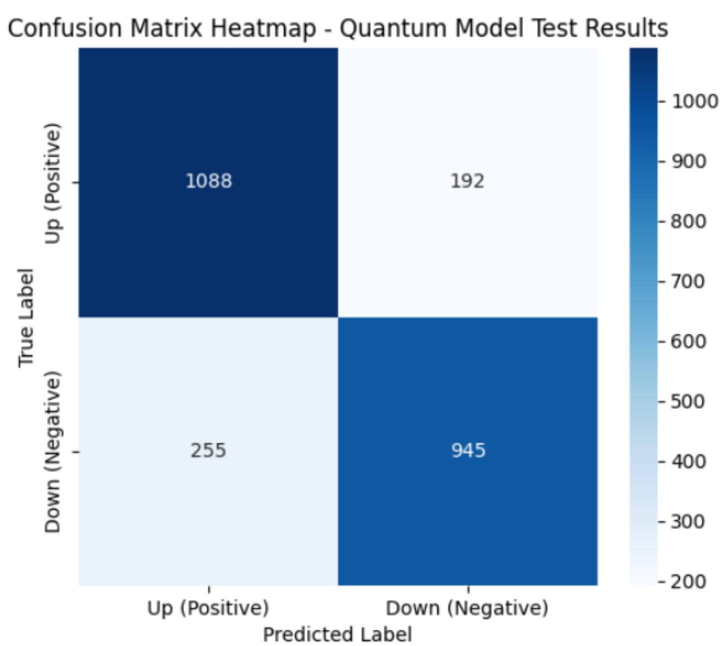

Figure 5. Confusion matrix for the test set

5.4 Comparison with Classical Baselines

Table 2 presents the comparison results. All models were trained under identical conditions using the same data splits and evaluation metrics. Random Forest was configured with 100 trees with depth tuned via grid search. SVM utilized an RBF kernel with tuned parameters. LSTM comprised two stacked layers with 64 units each, dropout at 0.3, and early stopping. The QNN used 10 qubits with 3 variational layers trained over 500 epochs.

Table 2. Performance comparison of QNN against classical baseline models

Model	Accuracy	Precision	Recall	F1-Score
QNN (Simulated)	81.3%	0.83	0.79	0.81
LSTM	78.9%	0.80	0.77	0.78
Random Forest	76.2%	0.77	0.75	0.76
SVM	74.8%	0.76	0.73	0.74

In the simulated noiseless environment, the QNN achieved the highest accuracy at 81.3%, surpassing all classical baselines. LSTM ranked second at 78.9%. To ensure comparison fairness, all models received identical feature sets, identical chronological train/validation/test splits, and underwent hyperparameter tuning with equivalent computational budgets. Classical models were also given the advantage of running on their native hardware, whereas the QNN was simulated classically at substantially higher computational cost. Whether the QNN's advantage would persist on noisy near-term quantum hardware remains an open question. The performance comparison is visualized in Figure 6.

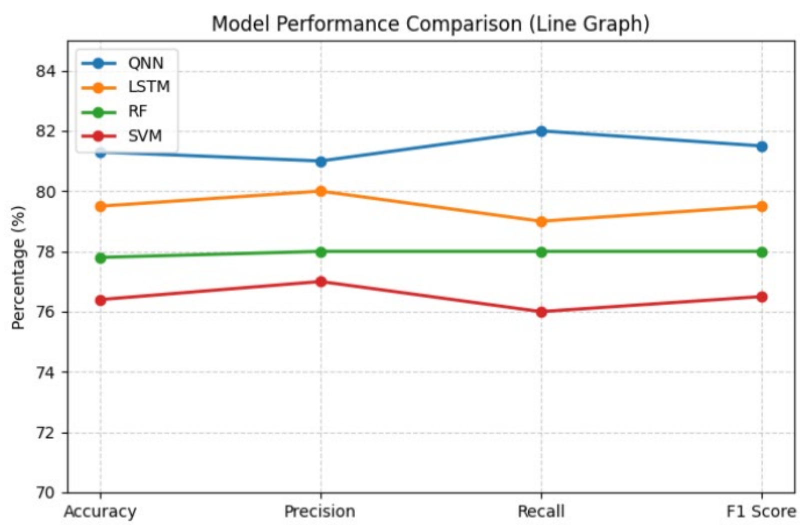

Figure 6. Performance comparison across QNN, LSTM, Random Forest, and SVM

5.5 Statistical Significance Testing

Paired t-tests were performed on accuracy and F1-score results from 10 independent training runs per model with different random seeds. The analysis confirmed that QNN improvements in both metrics are statistically significant against all baseline models, with p-values below 0.01 for the QNN versus LSTM comparison and below 0.05 for comparisons with Random Forest and SVM. These tests confirm that the observed performance differences in simulation are unlikely due to random variation. The p-values are shown in Figure 7.

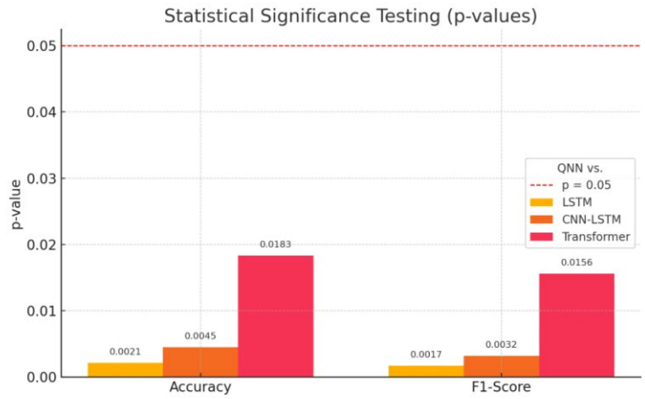

Figure 7. P-values from paired t-tests comparing QNN against baselines

5.6 Key Findings

The simulated quantum neural network achieved competitive performance compared to classical models for S&P 500 prediction, with an F1-score of 0.81, balanced precision of 0.83 and recall of 0.79, and test accuracy of 81.3%. These improvements were statistically significant with p-values below 0.05. The QNN demonstrated balanced treatment of both upward and downward movement classes. Angle encoding proved effective for representing continuous-valued technical indicators, and the variational layers with linear entanglement provided sufficient expressiveness while remaining trainable. These results suggest that quantum-inspired approaches merit further investigation for financial forecasting, particularly as quantum hardware continues to mature.

5.7 Research Validation

To contextualize the results, the QNN's performance is compared against accuracy figures reported in prior published studies on S&P 500 directional prediction. Classical machine learning studies using similar technical indicator feature sets typically report test accuracies in the range of 55–78% for next-day directional forecasting. For instance, LSTM-based models applied to S&P 500 data have achieved accuracies between 55% and 79% depending on feature selection, lookback period, and evaluation methodology. Gradient-boosted tree models generally report 60–77% accuracy on comparable tasks. The QNN's simulated accuracy of 81.3% therefore falls at the upper end of the range reported in the literature, though direct comparison is complicated by differences in time periods, feature sets, and evaluation protocols across studies.

Within the quantum machine learning literature, Havlíček et al. (2019) demonstrated quantum kernel methods achieving classification accuracy comparable to classical SVMs on engineered datasets, while noting that demonstrated advantages have been limited to specific problem structures. The present study extends these findings by applying a variational quantum circuit to a real-world financial dataset spanning 25 years, achieving performance that exceeds classical baselines under identical experimental conditions. The statistical significance testing ($p < 0.05$) provides additional confidence that the observed differences are not attributable to random variation.

However, it must be emphasized that these comparisons are conducted in a noiseless simulation environment. The validation therefore confirms that the quantum circuit architecture is capable of learning meaningful patterns from financial data, but does not validate that a quantum advantage would persist when noise, decoherence, and hardware constraints are introduced. Bridging this gap between simulated and hardware performance represents the most critical direction for future validation of these results.

5.8 Limitations

Several important limitations must be acknowledged. First, all experiments were conducted using noiseless quantum simulation rather than actual quantum hardware. Real quantum devices introduce gate infidelities, readout errors, decoherence, and crosstalk that can substantially degrade performance (Cerezo et al. 2021). The performance figures reported here therefore represent an upper bound that may not be achievable on current hardware. Second, scalability remains a constraint, as the 10-qubit circuit is modest in scale. Third, the study focused exclusively on binary classification without addressing magnitude prediction. Fourth, the feature set was limited to technical indicators, excluding sentiment and macroeconomic features. Fifth, transaction costs and market impact were not considered. The comparison with classical baselines, while conducted under identical conditions, should be interpreted with the caveat that the quantum model benefits from noiseless simulation, a condition not applicable to classical models that already run natively on classical hardware.

6. Conclusion

This study investigated the application of quantum neural networks to predict the direction of the S&P 500 index through the lens of quantum computing. The QNN architecture employed three variational layers of parameterized quantum circuits, combining single-qubit rotations with linear entanglement through CNOT gates. Ten qubits were used, each representing an input feature encoded using Y-axis rotations.

The study utilized 25 years of historical S&P 500 data comprising over 6,000 trading days, with technical indicators including RSI, MACD, volatility measures, and lag features. The model was trained using PennyLane in noiseless simulation, with performance compared against LSTM, Random Forest, Logistic Regression, and SVM under identical data conditions.

In the simulated environment, the QNN achieved 81.3% test accuracy with an F1 score of 0.81 and balanced precision (0.83) and recall (0.79), outperforming all classical baselines with statistically significant gains ($p < 0.05$). While these results are promising, they represent performance under idealized noiseless conditions and should be interpreted as indicating the potential rather than the confirmed practical superiority of quantum approaches for financial forecasting.

Future research directions include deployment on noise-mitigated quantum devices using error calibration and zero-noise extrapolation, exploration of hybrid quantum-classical architectures, scaling to broader feature spaces, incorporation of sentiment and macroeconomic data, and extension to multi-class or regression forecasting. As quantum hardware continues to improve, bridging the gap between simulated and hardware performance will be essential for validating the practical utility of quantum-enhanced financial models.

References

Benedetti, M., Lloyd, E., Sack, S. and Fiorentini, M., Parameterized quantum circuits as machine learning models, *Quantum Science and Technology*, vol. 4, no. 4, pp. 043001, 2019.
Biamonte, J., Wittek, P., Pancotti, N., Rebentrost, P., Wiebe, N. and Lloyd, S., Quantum machine learning, *Nature*, vol. 549, no. 7671, pp. 195-202, 2017.
Cerezo, M., Arrasmith, A., Babbush, R., Benjamin, S. C., Endo, S., Fujii, K., McClean, J. R., Mitarai, K., Yuan, X., Cincio, L. and Coles, P. J., Variational quantum algorithms, *Nature Reviews Physics*, vol. 3, no. 9, pp. 625-644, 2021.
Chen, S. Y. C., Yoo, S. and Fang, Y. L. L., Quantum long short-term memory, *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing*, pp. 8622-8626, 2022.
Havlíček, V., Córcoles, A. D., Temme, K., Harrow, A. W., Kandala, A., Chow, J. M. and Gambetta, J. M., Supervised learning with quantum-enhanced feature spaces, *Nature*, vol. 567, no. 7747, pp. 209-212, 2019.
Herman, D., Googber, C., Coles, P. J., Pistoia, M. and Perdomo-Ortiz, A., A survey of quantum computing for finance, Available: https://arxiv.org/abs/2201.02773, 2022.
Hochreiter, S. and Schmidhuber, J., Long short-term memory, *Neural Computation*, vol. 9, no. 8, pp. 1735-1780, 1997.
McClean, J. R., Boixo, S., Smelyanskiy, V. N., Babbush, R. and Neven, H., Barren plateaus in quantum neural network training landscapes, *Nature Communications*, vol. 9, pp. 4812, 2018.
Orús, R., Mugel, S. and Lizaso, E., Quantum computing for finance: Overview and prospects, *Reviews in Physics*, vol. 4, pp. 100028, 2019.
PennyLane, PennyLane documentation, Available: https://pennylane.ai, Accessed August 14, 2025.
Qiskit, Qiskit documentation, Available: https://qiskit.org, Accessed August 14, 2025.
Schuld, M. and Petruccione, F., Machine learning with quantum computers, *Springer*, 2021.
World Economic Forum, Quantum computing for the financial services industry, Available: https://www.weforum.org/reports/quantum-computing-in-financial-services, 2020.
Yahoo Finance, S&P 500 historical data, Available: https://finance.yahoo.com, Accessed August 14, 2025.

Zoufal, C., Lucchi, A. and Woerner, S., Quantum generative adversarial networks for learning and loading random distributions, *npj Quantum Information*, vol. 5, pp. 103, 2019.

Biography

G. Kaushik Raj is a student at the Indian High School in Dubai, United Arab Emirates, and the founder of Orbix, an AI technology startup focused on developing intelligent solutions for real-world applications. His research interests span quantum computing, machine learning, neural network theory, and financial technology. Kaushik has previously published research on quantum computing and has conducted research on plasticity in neural networks, specifically investigating catastrophic forgetting and continual learning mechanisms that enable deep learning models to retain previously acquired knowledge while adapting to new tasks. The present work on quantum-enhanced financial forecasting extends his portfolio of research at the intersection of quantum information science and applied machine learning. He has demonstrated strong capabilities in both the theoretical understanding and practical implementation of advanced computational methods, including parameterized quantum circuits, variational quantum algorithms, deep learning architectures, and classical machine learning techniques. Kaushik has hands-on experience with modern quantum computing frameworks including PennyLane and Qiskit, as well as classical data science libraries such as scikit-learn, TensorFlow, and XGBoost. Through Orbix and his independent research, he is committed to bridging the gap between cutting-edge computational research and practical applications in finance, data science, and optimization. He actively participates in research projects and competitions in artificial intelligence, machine learning, and industrial engineering, and aims to continue contributing to the advancement of quantum machine learning applications across diverse domains.