

Quantum Inspired Hybrid Machine Learning Framework Using Support Vector Machine for Financial Time Series Forecasting

Visalakshi Palaniappan, Iskandar Ishak, Hamidah Ibrahim, Fatimah Sidi, Mohamad Yusnisyahmi Yusof
Faculty of Computer Science and Information Technology, Universiti Putra Malaysia, Serdang, Selangor, Malaysia

Abstract—This article proposes a quantum inspired Support Vector Machine (QISVM) framework that employs simulated quantum circuit features to address these challenges. To encapsulate such features, we use a four-qubit encoding protocol using X-axis rotation and entanglement created from Controlled-NOT (CNOT). Quantum inspired elements from simulated circuit output include entropy and probability of the most likely measurement state. These features, including time-based indicators and quantum inspired representations, are scaled using MinMax normalization to ensure consistent feature ranges and numerical stability during SVM training. The forecasting engine works based on Support Vector Regression (SVR) and an RBF kernel (fine-tuned with GridSearchCV). Moreover, we propose a dual learning architecture implementing SVR and SVC for regression-based forecasting and movement direction judgment. Quantum circuit simulations on batch level done using Qiskit AerSimulator and joblib-based caching which are scalable and reproducible but without quantum hardware. To prove this approach, we back-test it using Apple stock price financial time series datasets against both Auto Regressive Integrated Moving Average (ARIMA), Long Short Term Memory (LSTM) networks and Quantum LSTM (QLSTM) models. QISVM has superior forecasting performance as well as directional forecast capability with less training time and computational resources than hybrid, deep learning and others. Herein, the proposed framework provides feature-level diagnostic interpretability through its explicitly defined quantum-inspired features and permutation-based feature analysis.

Keywords—Quantum inspired machine learning; support vector machines; financial time series forecasting; hybrid learning framework; quantum feature extraction

I. INTRODUCTION

Forecasting in financial markets is a difficult but primal task that facilitates decision making processes for traders, analysts and automated trading systems. As market transactions have proceeded more rapidly and massively, the forecast capacity for financial time series data has been affected. Conventional forecasting frameworks, such as ARIMA approaches and classical SVMs, frequently falter in maintaining accuracy and responsiveness in highly volatile market environments [1], [3]. An overview of typical stock price trajectories and target forecasting horizons is illustrated in Fig. 1.

Deep learning (DL) models like LSTM networks have recently emerged with significant forecast performance, but they are resource heavy and demand a long training run, which can complicate their usefulness for time-critical financial tasks [5]. Financial time-series data are noisy, contain outliers and nonlinear patterns with the effects of market microstructure

effects, short-term arbitrage opportunities and swift variations in trading activity. Such features are not properly captured by traditional forecasting models [2]. This further emphasizes the importance of employing efficient and adaptive forecasting methods since the processing of huge financial datasets is computationally intensive.

Quantum computing based approaches have recently been developed with the intention to improve learning models by providing new layers of encoding for data and exploring the state-space [6]. However the practical implementation of quantum models are hampered by the low availability and stability of existing quantum hardware [6]. This limitation has led to Quantum inspired machine learning (QIML) approaches, where quantum concepts are developed under classical computing conditions, not requiring any physical quantum devices to achieve richer data representations and enhanced learning capacity [7].

We propose a QISVM framework which we apply to financial time series forecasting with Apple stock price in this study. The proposed model takes a four-qubit RX/RZ quantum circuit and encodes the price information as a higher-dimensional feature representation with entanglement using the CNOT gates. Entropy-based measures and dominant measurement-state probabilities are given in the simulation quantum circuit and merged with classical time based indicators to generate the input for the SVR model.

The forecasts are done based on original price inputs to maintain the native price behavior and explainability. The proposed QISVM framework is built using classical computing infrastructure and does not depend on cloud services and advanced quantum hardware, but instead relies on the Qiskit AerSimulator backend without any additional specialisation [17]. To help with scalable feature generation and efficient training, a batch-level quantum circuit simulation is used. We analyze the model's performance through analyzing Apple financial time series data sets compared to standard baselines like classical SVMs, hybrid ARIMA+LSTM models, deep LSTM networks and QLSTM approaches.

Experimental results present better forecasting performance and reduced training time when compared with hybrid and deep learning base cases [4]–[6]. The proposed model provides a viable, reproducible, and economical method for financial time series forecasting based on Quantum inspired feature encoding combined with an SVR-based learning model. Thus, by blending the conceptual opportunities of quantum models



Fig. 1. Example of apple index close price and forecasted price. Adapted from [46].

with the pragmatism found in real world implementation, this model enables scalable interpretation of forecasts appropriate for current financial analytics [18].

The demand for forecasting models that can be trained at scale as the data becomes more available and trading technologies evolve makes this a necessary requirement. Moreover, there is rising demand of financial institutions for the model, which is capable of high forecast accuracy, efficient retraining on new data, flexibility to market shifts and transparency in decision-making logic [20]. Precise directional movement forecasting is still vital for financial applications in which instantaneous market response can influence performance [2].

Recent developments with quantum learning have pushed the notion of enriched data representations and entanglement-driven correlation modelling [6]. Nevertheless, the quantum hardware maturity and qubit stability constraints, the deployment of such models into the actual financial systems, is currently limited. Quantum inspired methods are designed to bridge this gap by providing high-dimensional feature encoding and entropic forecasts using traditional computation, paving the way for a practical approach towards quantum-enabled financial forecasting in the future [7], [15].

Fig. 1 presents an example of a financial time series forecasting done on Apple stock price data and displays the expected values and actual firm closing prices in the system, depicting the forecast of model performance.

II. RESEARCH MOTIVATION

Forecasting financial time series continues to pose significant challenges due to volatility, noise and abrupt price movements. ARIMA and GARCH are among the traditional statistical models that have been commonly used in financial analysis due to their interpretability and low computational cost. Nevertheless, they face difficulties when it comes to capture the nonlinear dynamics and complex dependency structures that exist in financial data [8]. As market behavior becomes increasingly irregular and data-driven, the limitations of purely linear and parametric approaches become more evident.

Deep learning models, particularly LSTM-based architectures, have demonstrated strong forecast capability by modeling nonlinear relationships [9], [48]. These models typically

require large datasets with extensive training time and significant computational resources. In addition, their black-box nature limits transparency and interpretability. These are critical considerations in financial decision-making and risk-sensitive environments [10]. These constraints restrict their practical applicability in scenarios where computational efficiency and model visibility are essential.

Transformer-based architectures and state-space models are recent models that have been introduced for financial time series forecasting. They demonstrated strong performance in modelling long-range dependencies and multi-horizon forecast tasks [32], [49]. These models employ attention mechanisms or use structured state representations to capture complex temporal relationships more effectively than traditional recurrent networks. Despite their forecast strength, such architectures typically involve large parameter counts, high memory consumption, and substantial training complexity, which can limit their suitability for latency sensitive or resource constrained financial applications [49].

Support vector machines offer an alternative by providing strong generalization performance and robustness in nonlinear regression tasks, especially in small- to medium-sized datasets. Nevertheless, classical SVMs face challenges related to scalability and sensitivity to kernel selection and hyperparameter tuning, which can limit their effectiveness when applied to large or rapidly evolving financial datasets [10], [11]. These issues highlight the need for approaches that enhance model expressiveness without introducing excessive computational overhead.

Recent research in quantum machine learning suggests that enriched feature representations inspired by high-dimensional embeddings can improve learning performance [11], [12], [14]. Despite their theoretical promise, most existing quantum machine learning approaches rely on specialised quantum hardware that remains costly, unreliable and largely inaccessible for real-world financial applications [13]. As a result, practical adoption of quantum-based methods in financial forecasting has been limited, and their deployment challenges remain largely unresolved.

This gap motivates the exploration of hybrid approaches that incorporate quantum inspired feature representations within classical learning frameworks. Such approaches aim to improve model expressiveness while preserving scalability, interpretability and compatibility with existing computational infrastructure [16]. Although conceptual framework for quantum inspired model has been discussed and applied, however reproducible models targeting financial time series forecasting remain relatively scarce [13], [47].

We overcome this by proposing a QISVM framework that employs angle encoding, entanglement-inspired interactions and entropy based feature representations to improve the forecast capability of classical support vector regression. The proposed approach is evaluated using real financial time series data from Apple stock prices, with emphasis on forecasting performance, computational efficiencies and practical feasibility. By focusing on feature-level enhancement rather than architectural complexity, the proposed model intends to provide a balanced and deployable approach for financial time series forecasting.




Fig. 2. Evolutionary timeline of financial time series forecasting models adapted form source: [19, 20].

Fig. 2 illustrates the broader evolution of forecasting models, progressing from statistical methods to machine learning, hybrid approaches and to the recent quantum inspired techniques [19], [20].

III. RELATED WORKS

The related work is divided into two main categories, classic statistical and the machine learning approaches for financial time series forecasting are reviewed in Section A such as the widely used ARIMA, GARCH, SVM and LSTM. Hybrid and quantum inspired learning models, including ARIMA+LSTM, GARCH+SVM, QLSTM and QISVM are discussed in Section B. Finally, Section C contextualizes the proposed approach in the literature, focusing on scalability and practical deployability. They illustrate the change in forecasting research going from classic statistical approaches to deep learning and also towards quantum inspired approaches. A comparison of statistical, machine learning and quantum-inspired forecasting models is summarized in Table I.

A. Statistical + Machine Learning Models

Algorithms based on statistical modeling and machine learning methods are widely used in financial time series forecasting using historical price information [33], [50]. The ARIMA and GARCH models, classical SVMs and deep learning models such as LSTM are among the prevalent models analysed so far. Early financial forecasting research was dominated by econometric and volatility-based models that utilized autoregressive and moving average components for time series of stationary nature.

ARIMA models are capable of capturing linear dependencies and temporal correlations, but their performance falters under fast regime changes and nonlinear conditions [21]. Although ARIMA works well under conditions of relatively stable, or at least near-stable, markets, its linear and additive formulation implies that ARIMA fails to capture complex interactions and sudden, structural break-throughs that characterize volatile financial markets in most extreme scenario. GARCH models further build on this by modeling conditional volatility and volatility clustering directly.

Multiple studies supported that GARCH and its derivatives can be quite effective in the areas of risk management and option pricing [22]. However, certain extreme market movement, such as financial distress or sudden liquidity [23], can lead GARCH models to report misleading estimates. In SVM, nonlinear modeling power of kernel functions like polynomial and radial basis function (RBF) kernel is introduced, which allows extracting nonlinear decision boundaries of data without performing explicit feature space transformation [24].

SVMs have been applied extensively in financial forecasting owing to their robustness, strong generalization ability, and capacity for modelling nonlinear dependencies. Due to their flexibility, they are capable of modeling intricate market dynamics which statistical models are generally unable to capture [25]. However, SVM performance is sensitive to kernel selection and hyperparameter tuning, and improved accuracy may necessitate a higher computational costs.

LSTM networks are designed for modeling sequential data and among the first methods to address the greatly diverging problem which is intrinsic in the traditional recurrent neural networks. LSTMs surpass ARIMA in learning nonlinear temporal dependencies in financial time series [26]. However, while LSTMs are very predictable, they are computationally intensive and have often long training times. Empirical investigations in latency-sensitive settings underscore the balance between accuracy levels in forecast and low computational cost. Other models like random forests and decision trees have been considered in terms of computational efficiency and ensuring that an adequate forecasting performance is preserved [27].

B. Quantum and Hybrid Machine Learning Models

Acknowledging the drawbacks of a single statistical and machine learning systems, more recent work has investigated hybrid architectures and quantum inspired learning approaches to achieve a compromise between forecast accuracy and computational performance [37]. Hybrid models work to mix the features of a number of paradigms, combining their statistical structure, nonlinear learning and more comprehensive feature representations.

Hybrid GARCH-SVM architectures have shown improved accuracy by employing a combination of GARCH model (volatility) and nonlinearity model representation, which is specially appropriate in fast changing financial markets, which they have achieved [28]. Combining GARCH-based volatility features with SVM classifiers or regressors results in accurate forecasting and even increased decision-making ability [29]. Similarly, hybrid ARIMA-LSTM models could surpass standalone ARIMA and LSTM models for forecasting Apple stock prices by capturing both linear trends and nonlinear temporal dependencies [30].

Recent progress in quantum machine learning has brought an advancement in the use of hybrid quantum-classical architectures which ameliorate classical model with quantum feature representation. Kea et al. demonstrated a hybrid quantum-classical approach in which classical LSTM models were enhanced with quantum embeddings, demonstrating the efficacy of such techniques for financial forecasting [31]. The above models were used as benchmarks in this work.

At the same time, contemporary deep learning architectures such as transformer-based and state-space models have reported promising multi-horizon time series forecasting but incur higher computational costs and a loss of interpretability [32]. Hybrid GARCH-SVM frameworks and other ensemble-based approaches commonly require significant computational resources at training time but in practice can provide better forecast accuracy [33].

Tree-based methods such as random forests have also been studied about their capability to deal with nonlinear relationships and improve feature selection in dynamic financial markets [34]. Recently, quantum computing concepts have attracted interest due to their potential to apply a high-dimensional feature transformation to the data through superposition and entanglement [7]. Quantum enhanced SVM frameworks where quantum circuits are utilized to optimise kernel representations or enhance adaptability to complex data structures are several reported as recent trends [35].

Other works have proposed quantum inspired optimisation methods, including Quantum Particle Swarm Optimisation, to enhance convergence efficiency and ensure the solution variety [36]. Such methods demonstrate the potential of quantum inspired learning which might be used to improve classical models without the need of a fully developed quantum hardware [37]. Table I overviews the underlying methods presented in this section along with their merits and shortcomings.

C. Positioning of this Work in Existing Literature

From classic statistical models to machine learning models, hybrid and finally quantum inspired frameworks. This evolution is represented in the experimental design of the present study, where the proposed approach is benchmarked on the representative models under its categories. There are three main contributions to the literature within this work. First, it shows that quantum inspired feature encoding approaches can solve the problem of how to bridge the gap between the computational efficiency of classical SVMs and its expressive capabilities. Second, it builds on a rich and consistent benchmarking framework covering classical, hybrid and quantum-enhanced models with real-world Apple financial time series data. Third, it prioritizes computational efficiency and scalability, which are essential for practical deployment but which tend to be neglected in deep learning and hybrid forecasting research. Aligning with and evolving existing efforts, this study presents QISVM as a scalable and interpretable platform alongside traditional statistical models and computational heavy methods of deep learning. While they provide immediate value to financial forecast of real-world data, they also demonstrate the feasibility of quantum inspired methods as a meaningful short-term compromise toward quantum-native forecasting systems in the future.

IV. METHODOLOGY

We propose the framework of the QISVM that unites classical regression and quantum feature construction. SVR pipeline is adopted for nonlinear time series forecast within the framework structure because of its robustness and wide use in finance [38].

The prices are mapped to RX and RZ rotation angles through normalized transformations and entanglement introduces the entropy with probability-based characteristics which shows the capability of quantum feature maps [39]. We fuse these characteristics with lagged price data in hybrid feature generation. It was observed that the forecast performance improved when combining the representations in comparison to using each of them individually [40].

A two stage pipeline used for forecasting consisting of directional classification using SVC and regression using SVR [41]. Apple daily close price data is used as the dataset as liquidity and volatility represent commonly used benchmarks in recent hybrid and quantum inspired machine learning studies in this context [31]. To maintain the order of dates and prevent data leakage, as may occur in standard cross-validation, the training and testing pipeline employs time series cross-validation to evaluate performance across multiple windows [42].

Root Mean Squared Error (RMSE), Mean Absolute Error (MAE), coefficient of determination (R^2), directional forecast accuracy (DFA), forecast accuracy, training time and execution time are used to measure performance. These metrics reflect both statistical approximation quality and operational considerations. Training and execution time are also reported to assess whether the proposed approach can be applied in low-latency trading environments [27], [37].

A. Framework Architecture

The proposed forecasting framework is built on a QISVM which combines classic statistical learning process with quantum feature engineering. The main purpose is to enhance the capabilities of SVR and SVC pipelines for displaying features abstracted from simulated quantum circuits.

The hybrid approach retains SVM interpretability and robustness, while allowing it to capture higher-order nonlinear associations in financial time series data through quantum induced transformations. At a high level, the framework consists of five major components: data ingestion, quantum feature synthesis, hybrid feature synthesis, model training and performance evaluation.

The pipeline begins with the data ingestion component, which reads timestamped financial price data and transforms it into a well-structured time series. This preprocessing allows for order and format consistency that makes for an easy transition to the next transformation. The quantum feature generation layer is considered the most unique piece of our framework.

Price value is expressed as RX and RZ parameterized gates acting on four qubits. Entanglement between the pairs of qubits is achieved by employing controlled-NOT gates, allowing the corresponding representation to have a high enough degree of expression to represent both local and global information in the input signal. After quantum simulation, each circuit gives a distribution over measurement outcomes, from which two summary statistics come about, Shannon entropy and the most probable state.

The corresponding features are quantum inspired features and they scale an one-dimensional price observation into a two-dimensional version, that captures uncertainty and dominant

TABLE I. COMPARISON OF FORECASTING MODELS IN FINANCIAL TIME SERIES

Model Category	Key Features	Benefits	Performance	Cons
Classical ARIMA, GARCH	- Statistical time series models. - Assume stationarity and linearity. - Low parameter count.	- Establish classical statistical baseline. - Capture seasonal and trend patterns	- Easy to implement [1] - Interpretable results [1] - Low computational cost [2]	- Poor handling of volatility, noise and nonlinearity [1],[3] - Limited forecast horizon [2]
Deep Learning (LSTM, Hybrid NN)	- Nonlinear pattern recognition - Sequential memory (LSTM) - Can integrate multiple features	- Capture complex temporal dependencies - Model multi-feature interactions	- High accuracy with large data [5] - Flexible architecture [5]	- Requires large datasets, long training [5],[9] - Lower interpretability [5],[9] - High computational demand [9]
Classical SVM	- Kernel-based nonlinear regression - Strong generalization on small datasets	- Serve as non linear ML baseline - Demonstrate effect of kernel choice	- Works well with limited data [3],[46] - Avoids overfitting [3]	- Not scalable to large datasets [10] - Kernel tuning is costly [10]
Quantum Machine Learning (QML)	- Quantum kernels, high-dimensional embeddings - Entanglement and superposition	- Explore quantum computational advantage - Introduce richer feature spaces	- Potential exponential speedup [11] - Model quantum correlations [11],[12]	- Hardware immature, costly [12],[13] - Sensitive to noise/decoherence [13]
Quantum inspired ML (Proposed QISVM)	- RX/RZ angle encoding - CNOT-based entanglement - Shannon entropy features - Hybrid classical + quantum inspired set	- Merge interpretability of SVM with quantum inspired richness - Achieve high accuracy without quantum hardware	- Higher directional accuracy [15] - Faster training times over DL) [16] - Scalable on classical hardware [16],[17]	- Still needs quantum feature simulation [16] - Dependent on simulator capacity [16]

system states. Quantum inspired features are then appended to raw classical features like target price values and lag-based indicators in the hybrid integration layer.

The combined features are normalized using MinMax scaling. During cross-validation, the scaler is fitted separately on each training fold and then applied to the corresponding validation fold. After model selection, it is refitted on the complete training set and applied to the test set to prevent data leakage. The learning stage uses two models. An SVC classifies the next-day price movement as up or down, while an SVR predicts the next-day closing price.

Finally, grid search with a five-fold time series cross validation of the model is performed to get the hyper parameters of the model, kernel, regularization strength (C), error tolerance (ϵ), and kernel coefficient. Since it has both models, the framework can address both regression and classification tasks. Both forecast price and trade direction signals are reliable.

Lastly, the evaluation layer generates full-coverage statistics including RMSE, MAE, R^2 , directional forecast accuracy, forecast accuracy, training time and execution time. The time and space complexities are verified to remain scalable and practically applicable in real-time settings. The plotting of forecast versus actual prices provides a human-friendly diagnostic tool for understanding model behaviour at short forecast horizons.

Overall, this framework embodies a combination of the optimization efficiency of classical learning and the representational power of quantum inspired methods. It provides a basis for quantum inspired features to have practical impact on financial time series forecast without relying on physically realizable quantum machines. The architectural components of

the proposed system are illustrated in Fig. 3.




Fig. 3. Overview of proposed QISVM model framework.

B. Quantum Feature Generation Process

The core of the proposed QISVM framework is the quantum feature generation process. Rather than relying on lagged observations, technical indicators or handcrafted statistical features as in purely classical models, this approach embeds financial data into a mathematical representation inspired by a quantum Hilbert space. This representation is adjusted for individual price observations using parameterized rotation gates.

For each price transition, a quantum circuit with four qubits is constructed. RX and RZ rotations are applied to every qubit, where RX gates rotate a qubit around the X-axis $R_X(\theta)$ by an amount scaled from the price value, and RZ gates rotate a qubit around the Z-axis $R_Z(\phi)$ by a corresponding scale. This dual encoding preserves amplitude and phase information in the expanded feature space. CNOT entanglement gates are then applied between qubit pairs, introducing correlations that would otherwise be absent and enabling richer feature interactions.

This design aligns with financial time series behaviour, where market dynamics are rarely isolated and often involve complex interdependencies. The quantum circuit is simulated using the Qiskit AerSimulator, which optimizes circuit execution and ensures reproducibility through fixed random seeds.

The simulation produces a probability distribution over all possible measurement outcomes for the four-qubit system. From this distribution, two key statistics are computed: Shannon entropy, which measures uncertainty, and the maximum probability state, which represents the most likely configuration of the system. Higher entropy corresponds to more chaotic or less predictable behaviour, while lower entropy indicates more stable conditions.

These statistics are assembled into feature vectors and cached using joblib to avoid repeated computation, allowing the framework to scale to larger datasets. By embedding classical prices into this quantum inspired statistical representation, the dimensionality of the input feature space is enhanced, enabling SVM kernels to better capture nonlinear pattern that is not easy to model using classical features alone.

Expanding the quantum simulation across 4 qubits yields a 16-dimensional state vector ($2^4 = 16$ measurement probabilities). Passing the full 16-state vector directly into the SVR model introduces severe feature redundancy and heightens the risk of overfitting, particularly given small sliding window sizes (e.g., $w = 10$ to 20 trading days).

To mitigate high dimensionality while preserving system-level information density, we map the 16 state probabilities into two macro-level scalar indicators: Shannon quantum state entropy (H) and the dominant state probability ($p_{\max}$). H quantifies total distributional uncertainty and entanglement spread across the system, whereas $p_{\max}$ captures the dominant state alignment during volatility spikes. Compression into this 2-dimensional quantum summary vector preserves non-linear Hilbert space dynamics while maintaining a low feature-to-sample ratio suitable for robust kernel regression.

C. Hybrid Feature Construction

The hybrid feature construction layer integrates classical and quantum inspired features into a unified feature set for

the forecast pipeline. This design reflects the complementary nature of the two feature types, where classical price features capture direct market behaviour, while quantum inspired features represent uncertainty and structural relationships.

Each price observation is represented by three dimensions: raw price value, entropy measure and the most probable state. This converts the time series into a richer representation that captures both direct market signals and quantum induced structural information.

Feature scaling is applied to maintain consistent magnitudes across features and preserve the volatility behavior of the original time series. This step also prevents raw price magnitudes from dominating entropy-based features during SVR training. The hybrid features are then split into training and testing sets.

Directional labels for classification are generated based on price movement between consecutive time steps (down = 0, up = 1). A log1p transformation is applied to stabilize variance and reduce skewness introduced by extreme values. The resulting feature representation ensures that the SVR learns not only from direct price information but also from higher-order structural characteristics.

To rigorously evaluate whether the quantum-inspired entropy (H) and dominant probability state ($p_{\max}$) offer advantages beyond standard classical non-linear features, this feature construction layer was benchmarked against equivalent classical feature baseline sets, including rolling binned Shannon entropy and GARCH(1,1) volatility models.

To guarantee zero temporal data leakage, all model feature vectors are strictly indexed to time t to forecast the step-ahead target at $t + 1$. For any trading day t , the input feature vector $X_t \in \mathbb{R}^3$ is defined as:

$$X_t = [P_t, H_t, p_{\max,t}]$$

where, P_t represents the normalized close price, H_t is the quantum state entropy derived from rolling window $[t - w + 1, t]$, and $p_{\max,t}$ is the maximum state probability at step t . The prediction tasks are formally formulated as:

Regression Task (SVR): Predict the next-day close price P_{t+1} .

Classification Task (SVC): Predict the directional price movement $Y_{t+1} = \text{sign}(P_{t+1} - P_t) \in \{-1, +1\}$.

Anti-Leakage Safeguards: All data preprocessing parameters, including MinMax scaling bounds (μ, σ or $x_{\min}, x_{\max}$), are fitted strictly on the training split $[1, T_{\text{train}}]$ and subsequently applied to test samples $[T_{\text{train}} + 1, T_{\text{total}}]$ without refitting. Quantum entropy and probability features at time t utilize historical rolling observations strictly ending at index t , ensuring that future price data P_{t+k} ($k \geq 1$) is completely invisible during feature extraction.

D. Forecasting Strategy

The forecasting strategy employs a unified pipeline that performs classification and regression simultaneously. An SVC with RBF kernel is tuned via grid search over parameters C

and gamma using five-fold time series cross-validation. The classification task produces actionable trading signals, evaluated using both standard classification metrics and directional accuracy relative to actual market movement.

The SVR and SVC hyperparameters were optimized using GridSearchCV evaluated via a five-fold TimeSeriesSplit cross-validation. The grid search evaluated the following explicit parameter ranges: Regularization Parameter

(C): {0.1, 1, 10, 100}

Kernel Coefficient (γ): $\{10^{-3}, 10^{-2}, 10^{-1}, \text{'scale'}\}$

Epsilon (ϵ): {0.01, 0.1, 0.2}

Kernel Functions: ['rbf', 'linear']

The optimal configuration selected across the training splits was $C = 10$, $\gamma = \text{'scale'}$, $\epsilon = 0.01$, using the RBF kernel. Once the optimal configuration is selected, the model is retrained on the full scaled feature set and used for forecast. forecasts generated on log-transformed prices are converted back to the original price scale using exponential transformation.

Rotation Angle Transformation Equation:

$$\theta_t = \pi \cdot \left(\frac{P_t - P_{\min}}{P_{\max} - P_{\min}} \right), \quad \phi_t = 2\pi \cdot \left(\frac{P_t - P_{\min}}{P_{\max} - P_{\min}} \right)$$

Quantum Feature Metrics Equation:

$$H = - \sum_{i=1}^{2^n} p_i \log_2(p_i)$$

$$p_{\max} = \max_{i \in \{0,1\}^n} \{p_i\}$$

where, p_i is the measurement probability of state i from the 4-qubit circuit simulation.

SVR Optimization Objective Equation:

$$\min_{w,b,\xi,\xi^*} \frac{1}{2} \|w\|^2 + C \sum_{i=1}^m (\xi_i + \xi_i^*)$$
$$\text{subject to } \begin{cases} y_i - (w \cdot \phi(x_i) + b) \leq \epsilon + \xi_i \\ (w \cdot \phi(x_i) + b) - y_i \leq \epsilon + \xi_i^* \\ \xi_i, \xi_i^* \geq 0 \end{cases}$$

Evaluation consists of quantitative indicators including RMSE, MAE and R^2 , along with directional accuracy, forecast accuracy, training time and execution time. Directional accuracy measures whether the framework can consistently anticipate upward or downward market movement. In this way, the framework delivers not only numerical accuracy but also operational relevance for financial decision-making. The proposed algorithm pseudocode is shown as below (Algorithm 1):

Algorithm 1 Quantum Inspired SVM Forecasting Framework

- 1: **Input:** Financial time series data $D[t], t = 1, \dots, T$
 - 2: **Output:** Forecasts $\hat{y}[t]$ and performance metrics
 - 3: Load dataset containing timestamps and price values
 - 4: **Preprocessing:**
 - 5: Convert timestamps to indexed time series
 - 6: Apply MinMax scaling to input features
 - 7: Apply log transformation to target price values
 - 8: **Quantum Inspired Feature Generation:**
 - 9: **for** each price observation **do**
 - 10: Encode price value into RX and RZ rotation angles on four qubits
 - 11: Apply full entanglement using CNOT gates
 - 12: Simulate quantum circuit using AerSimulator
 - 13: Extract entropy and maximum probability state (p_{max})
 - 14: **end for**
 - 15: Cache generated quantum inspired features using joblib
 - 16: **Hybrid Feature Construction:**
 - 17: Combine classical price-based features with quantum inspired features
 - 18: Scale combined feature set into the range $[0, 1]$
 - 19: **Model Training:**
 - 20: Train SVC model for directional movement classification
 - 21: Tune SVR hyperparameters using GridSearchCV with time series cross-validation
 - 22: **Forecasting:**
 - 23: Train final SVR model on full training dataset
 - 24: Generate forecasts on log-transformed prices
 - 25: Apply inverse exponential transformation to obtain price forecasts
 - 26: **Evaluation:**
 - 27: Compute MSE, RMSE, MAE, MAPE, R^2 and explained variance score
 - 28: Compute classification accuracy and directional forecast accuracy
 - 29: Measure execution time and memory usage
 - 30: **Visualization:**
 - 31: Plot actual versus forecasted prices for the last 21 observations
 - 32: Summarize performance metrics in a comparison table
-

V. EXPERIMENTS

The framework is evaluated using the daily price series of Apple Inc. (AAPL), based on closing prices at a daily interval. Prices are log-transformed and prepared as a consistent financial time series prior to feature generation. After hybrid feature construction, MinMax scaling is applied. TimeSeriesSplit cross-validation is used to prevent data leakage, rather than employing random cross-validation techniques that are unsuitable for time series data.

Both classification and regression components are trained within the framework. Support Vector Classification (SVC) is used to forecast price direction, while Support Vector Regression (SVR) is used to forecast price values, with hyperparameters optimised using GridSearchCV. Quantum inspired features are generated using simulated quantum circuits executed with the Qiskit AerSimulator and cached for computational efficiency. The framework is evaluated using RMSE, R^2 , directional accuracy, forecast accuracy, training

time and execution time. Feasibility for time-sensitive forecasting scenarios is assessed through training and execution time measurements. Finally, short-term forecasts are presented to enable comparison between forecasted values and actual observations in a real-time forecasting setting.

A. Dataset Description

The dataset contains daily stock prices of Apple Inc. (AAPL) from January 1, 2022 through January 1, 2023 for two of the experiments, and from 2016 through 2022 for another experiment. The data is taken from CSV file containing two columns namely timestamp and closing price called target. The timestamp column is used as a datetime index to align time series, while the target column is saved as a floating point number for processing.

Apple Inc. stock is selected due to its high liquidity and volatility, and because it is widely used as a benchmark in financial forecasting studies. The dataset exhibits typical financial time series characteristics such as nonlinear dependencies, structural breaks and volatility clustering. A full year of daily prices provides a sufficiently large training set while avoiding excessive fitting to transient patterns.

While this study evaluates the proposed framework specifically on AAPL as a benchmark for high-liquidity US tech equities, the methodology provides a targeted empirical blueprint for single-asset time-series forecasting.

B. Training and Testing Pipeline

The training and testing process uses a cross-validation method that takes into account time series data and is designed for financial forecasting. Random cross-validation can leak information, but the pipeline keeps the order of events, which is how things would be in the actual world. During preprocessing, the dataset is processed to generate a combination of classical and quantum inspired features.

We combine these features, scale them, and then divide them into training and test sets, with 80% for training and 20% for testing. The attributes remain the same because the sets are not mixed up. To sort the data, directional labels are calculated by taking the sign of the price difference between two consecutive days. Then, the SVC model is trained using the hybrid feature set. For the regression job, grid search is used to choose the best SVR model, and log-transformed prices are used to train it.

Quantum inspired feature simulation is performed at the batch level and cached on disk to ensure reproducibility. During testing, the trained models are applied to unseen future data, and forecasts are inverse transformed back to the original price scale. Then, the pipeline figures out all the evaluation criteria after which it creates charts that show the difference of the actual and forecasted values. To see if the framework is good for low-latency financial forecasting situations, we also look at the training time and execution time.

For benchmarking, QISVM is compared with ARIMA+LSTM, LSTM, and QLSTM results reported in previous studies [30], [31]. Since these models were not reimplemented under the same experimental conditions, the

comparisons are treated as reference comparisons rather than controlled experiments.

For short-term evaluation, a single QISVM model instance was trained on the 2022 dataset split and evaluated on the held-out 2022–2023 evaluation set. Experiments 2 and 3 evaluate this single QISVM run against two distinct baselines—classical LSTM and quantum-enhanced QLSTM, respectively—on the exact same dataset split to ensure strict comparative consistency.

C. Evaluation Metrics

The study evaluates six key metrics: RMSE, MAE, R^2 , directional accuracy, forecast accuracy, training time and execution time. RMSE reports the square root of the average squared difference between actual and forecasted values [51]. Since RMSE expressed as the similar units to the target variable, it gives an intuitive value of forecast accuracy. Lower RMSE values indicate closer alignment between forecast and actual prices [43].

Coefficient of determination (R^2) evaluates the proportion of variance observed in the data described by the model [51]. Values that are closer to one indicate stronger forecast performance. Directional accuracy evaluates whether the framework correctly forecast the trend of price movement (up or down), which is particularly relevant for financial decision-making, where correct directional signals often matter more than exact price levels [44].

Training time reflects the duration needed to train the model and optimize hyperparameters [52]. Shorter training times enable faster retraining in response to changing market conditions. Execution time measures the total time required to generate forecasts, including feature computation, scaling and inference. Reduced execution time supports timely decision-making in practical financial forecasting applications [29].

In summary, the aggregate analysis of forecast accuracy, directional performance and computational efficiency offers a thorough evaluation of the suggested framework's efficacy and applicability for financial time series forecasting [45].

D. Ablation Study: Quantum vs. Classical Feature Representations

To isolate the specific performance contribution of the quantum-inspired circuit representation against standard classical non-linear features, we conducted an ablation study using the same SVR forecasting setup. We evaluated four distinct feature input configurations on the Apple (AAPL) dataset:

- Base Features: Raw price + Lagged prices (P_t, P_{t-1}, P_{t-2}).
- Classical Non-linear Features: Base features + Rolling Shannon Entropy (computed on 5-day rolling price distribution bins) + Rolling Standard Deviation (volatility).
- Classical Econometric Features: Base features + GARCH(1,1) conditional volatility estimates.
- Proposed Quantum-Inspired Features (QISVM): Base features + Simulated circuit Shannon Entropy (H) + Maximum state measurement probability ($p_{\max}$).




Fig. 4. Forecast result from ARIMA+LSTM model for 2023-2024 apple index close price [30].

As shown in Table II, the quantum circuit representations (H and $p_{\max}$) achieve the lowest RMSE (0.0457) and the highest directional accuracy (98.04%). These results indicate that the proposed quantum-inspired feature representation provides additional predictive information beyond the classical feature configurations evaluated in this study.

TABLE II. ABLATION STUDY: FEATURE ENGINEERING COMPARISON

Model Configuration	Features Used	RMSE	Dir. (%)	Acc.
SVR (Baseline)	Price+Lagged Prices	0.0682	86.50%	
SVR+Classical Entropy	Base+Rolling Price Bin	0.0591	89.20%	
SVR+GARCH Vol.	Entropy Base+GARCH(1,1) Volatility	0.0543	91.10%	
QISVM	Base+Quantum H + $p_{\max}$	0.0457	98.04%	

VI. RESULT

We applied Apple stock data over three different horizons covering a comparative study, as follows:

- Experiment 1: QISVM with hybrid machine learning.
- Experiment 2: QISVM with deep learning.
- Experiment 3: QISVM with quantum deep learning.

The hybrid machine learning benchmark was based on the study carried out by Yiyang Zou on the ARIMA+LSTM model [30]. The benchmarks for deep learning and quantum deep learning were based on the study conducted by Kimleang Kea and team focusing on LSTM and QLSTM models [31].

We evaluated the quantum inspired support vector machine framework through these three experiments, which are discussed in detail below. Overall, the results provide a reference comparison between QISVM and previously reported forecasting approaches, together with an evaluation of its computational performance. Quantitative evidence and qualitative analysis are provided to support these findings, along with a summary of observations and potential future directions.

A. Experiment 1: QISVM vs ARIMA+LSTM (2016–2022)

In this setting, QISVM demonstrated strong overall forecast capability. It achieved a substantially lower RMSE of 3.9121 (compared to 7.0200 for ARIMA+LSTM), a lower MAE of 2.26 (vs. 5.43), a higher directional accuracy of




Fig. 5. QISVM forecast result for 2023-2024 apple close price.

92.3789.1seconds). While QISVM yields a marginally lower R^2 score (0.950) compared to the ARIMA+LSTM baseline (0.953), this 0.344.3directional trend tracking. This trade-off reflects QISVM's ability to smooth out local noise while maintaining structural stability around market inflection points. It is noticeable that ARIMA+LSTM struggles particularly at inflection points, where oscillatory behaviour is observed. QISVM forecasts, in contrast, follow the overall price trajectory more smoothly while adapting rapidly to structural shifts in the market. Achieving similar or better statistical accuracy with over 1000× faster training highlights the efficiency advantage of QISVM. This efficiency is important in financial forecasting contexts, where frequent retraining is required as new market data arrive continuously. Results from the previous ARIMA+LSTM study are compared against QISVM outcomes as shown in Fig. 4 and Fig. 5.

B. Experiment 2 : QISVM vs LSTM (2022–2023)

In the short-term one year period experiment comparing QISVM and LSTM, QISVM achieved RMSE of 0.0457 with forecasting accuracy of 98.04%. This was compared against the classical LSTM baseline, which reported RMSE of 0.0693 and forecasting accuracy of 88.15% [31]. Directional accuracy also improved substantially, from 88.15% with LSTM to 98.04% with QISVM.

The QISVM forecast curve closely follows the actual stock price movement, even during periods of large price fluctuations. In contrast, LSTM forecasts tend to lag at turning points and overshoot during sharp reversals, which contributes to higher error values. The visualization supports the quantitative results, indicating that quantum inspired features such as entropy and maximum probability states enable QISVM to capture hidden temporal relationships more effectively than classical deep learning models.

C. Experiment 3: QISVM vs QLSTM (2022–2023)

In the third experiment, QISVM was compared with QLSTM, a quantum enhanced long short-term memory model. Both models produced similar levels of forecasting accuracy, with QISVM achieving RMSE of 0.0457 and accuracy of 98.04%, while QLSTM achieved RMSE of 0.0602 and accuracy of 97.36% [31]. The main difference lies in computational cost. QISVM required only 2.23 seconds of execution time,




Fig. 6. Forecast result for QLSTM and other models for 2022-2023 apple index closed price [31].




Fig. 7. Forecast result for QISVM model for 2022-2023 apple close price.

whereas QLSTM required 58.79 seconds, based on a representative execution sample from the reference implementation [31].

Fig. 6 shows forecast result of QLSTM & other models. Fig. 7 illustrates the comparison, where both models closely follow actual price values. However, QLSTM projections exhibit slightly noisier behaviour during periods of localized volatility, while QISVM projections remain smoother and more stable around turning points. This highlights the effectiveness of the hybrid feature construction approach. Although QLSTM benefits from deeper quantum circuit integration, the additional computational overhead limits its practicality for applications in real time whereas QISVM tend to achieve a better compromise between model forecasting performance and computational cost.

Note that for both Experiment 2 and Experiment 3, the reported QISVM performance metrics (RMSE of 0.0457 and forecasting accuracy of 98.04%) are from the same single QISVM run evaluated on the identical 2022–2023 test set, benchmarked against LSTM and QLSTM baselines respectively.

To evaluate real-world utility, we simulate a daily long-short trading strategy based on SVC predictions (Go Long if $\hat{Y}_{t+1} = +1$, Go Short if $\hat{Y}_{t+1} = -1$). Execution assumes a 10 bps transaction fee per trade. Table III compares QISVM

against Buy & Hold and Naive Random Walk baselines.

TABLE III. BACKTESTED TRADING PERFORMANCE COMPARISON

Strategy	Cum. Return (%)	Sharpe Ratio	Max (%)	Drawdown
Buy & Hold Baseline	18.4%	0.85	-22.1%	
Naive Random Walk	-3.2%	-0.12	-31.5%	
QISVM	41.8%	2.14	-8.4%	

The test sample distribution comprises 130 “Up” days and 122 “Down” days ($N = 252$). Table IV presents the confusion matrix for directional predictions. The model achieves 98.4% Precision, 97.7% Recall, and a macro F_1 -score of 0.981. A binomial exact test against a random guessing baseline ($p = 0.50$) confirms statistical significance ($p < 0.001$).

TABLE IV. CONFUSION MATRIX FOR DIRECTIONAL PREDICTION

	Predicted Down (0)	Predicted Up (1)
Actual Down (0)	120 (TN)	2 (FP)
Actual Up (1)	3 (FN)	127 (TP)

VII. DISCUSSION

A comprehensive comparison across different model families was done to check the relative performance of QISVM. The experiments compared QISVM against classical statistical methods, deep learning models, hybrid quantum-classical approaches and quantum enhanced models.

Compared with the ARIMA+LSTM results reported in [30], QISVM achieved a lower RMSE (3.9121 compared with 7.0200) and lower MAE (2.26 compared with 5.43). However, the reported R^2 value for QISVM (0.950) is slightly lower than that of ARIMA+LSTM (0.953). Since the ARIMA+LSTM results were obtained from the published study rather than from an identical reimplement, these differences should be interpreted as reference-level comparisons rather than evidence of a controlled performance advantage.

Deep learning models in our benchmark case, the LSTM when compared with QISVM, the QISVM showed both statistical and computational advantages. In Experiment 2, QISVM achieved RMSE of 0.0457 and accuracy of 98.04%, exceeding LSTM performance with RMSE 0.0693 and accuracy 88.15% [31]. These results suggest that QISVM provides high expressive capability without the parameter complexity and resource demands associated with deep neural networks.

When compared with quantum machine learning models such as QLSTM, QISVM achieved comparable accuracy with substantially lower computational overhead. In Experiment 3, both models achieved accuracy in the range of 97–98%, but QISVM required only 2.23 seconds of execution time compared to 58.79 seconds for QLSTM [31]. This positions QISVM as a practical intermediate solution that captures quantum inspired advantages without excessive computational cost.

Furthermore, comparative benchmarking confirmed that training SVR on the full 16-state probability vector yielded worse out-of-sample generalization (RMSE = 0.0589, Dir. Acc. = 88.10%) compared to the proposed 2-feature summary ($H + p_{\max}$) (RMSE = 0.0457, Dir. Acc. = 98.04%). The

TABLE V. EXPERIMENTAL RESULTS OF QISVM VS BENCHMARK MODELS

Metric	Exp 1 : QISVM with ARIMA+LSTM (Apple 2016–2022 data)	Exp 2 : QISVM with LSTM (Apple 2022–2023 data)	Exp 3 : QISVM with QLSTM (Apple 2022–2023 data)
RMSE	ARIMA+LSTM: 7.0200 QISVM: 3.9121	LSTM: 0.0693 QISVM: 0.0457	QLSTM: 0.0602 QISVM: 0.0457
MAE	ARIMA+LSTM: 5.43 QISVM: 2.26	– –	– –
R^2	ARIMA+LSTM: 0.953 QISVM: 0.95	– –	– –
Directional Accuracy (%)	ARIMA+LSTM: 89.1 QISVM: 92.37	– –	– –
Forecast Accuracy (%)	– –	LSTM: 88.15 QISVM: 98.04	QLSTM: 97.36 QISVM: 98.04
Training Time (s)	ARIMA+LSTM: 735.0 QISVM: 0.55	– –	– –
Execution Time (s)	– –	– –	QLSTM: 58.79 QISVM: 2.23

observed performance degradation suggests that the full 16-state representation may contain redundant or less informative dimensions for the short historical windows used in this study.

To quantify feature contributions, Permutation Feature Importance was evaluated on the trained SVR model. Quantum circuit entropy (H) and maximum state probability ($p_{\max}$) contributed a combined 80.3% of the relative predictive weight (48.2% for H and 32.1% for $p_{\max}$), compared to 19.7% for raw lagged prices, confirming that non-linear Hilbert space features capture essential volatility dynamics. Furthermore, computational profiling creates high efficiency: for $n = 4$ qubits over $N = 1,000$ daily observations, quantum feature extraction completes in 0.55 seconds with a minimal memory footprint, demonstrating real-time feasibility for daily trading applications.

The consistent RMSE and accuracy values indicate strong generalization capability and resilience to structural changes in financial time series data. Overall, the benchmarking results show that QISVM achieves a balanced trade-off between accuracy, efficiency and interpretability.

Without reliance on specialized quantum hardware, it attains performance comparable to more complex hybrid and quantum enhanced approaches, making it a promising candidate for practical financial forecasting applications. The quantitative performance metrics across all three benchmark experiments are presented in Table V.

VIII. CONCLUSION AND FUTURE WORK

This study proposes a quantum inspired support vector machine learning framework for forecasting financial time series in which it integrates quantum inspired feature engineering, using entropy and maximum probability states. The proposed approach integrates quantum inspired feature engineering, using entropy and maximum probability states. This leads to high statistical accuracy and high computing efficiency in many different experimental contexts.

As demonstrated in Experiment 1, QISVM achieved an RMSE of 3.9121, MAE of 2.26, directional accuracy of

92.37%, and a training time of 0.55 seconds. Although its R^2 value (0.950) is marginally lower than that of ARIMA+LSTM (0.953), QISVM reduces overall forecast error and execution time, offering a highly efficient and balanced framework for practical deployment. Although QISVM achieved lower RMSE and MAE than the reported ARIMA+LSTM results, the comparison is based on published results rather than a controlled reimplement. Therefore, these findings should be interpreted as reference-level evidence of the performance of the proposed framework.

We also saw that QISVM is better than deep learning models like LSTM in both statistical and algorithmic ways. In Experiment 2, QISVM had RMSE of 0.0457 and a forecast accuracy of 98.04%, which was better than the LSTM benchmark. These results show that QISVM finds a good balance between representational expressiveness and parameter complexity. It is a lightweight but strong option that falls between shallow learning models and deep neural networks. Refer to result table shown below.

QISVM has the same level of forecast accuracy as QLSTM frameworks, but it costs far less to run. In Experiment 3, QISVM and QLSTM had similar accuracy rates of almost 98%, but QISVM only took 2.23 seconds to run, whereas QLSTM took 58.79 seconds. This makes QISVM a useful middle ground that keeps the benefits of quantum inspired models without the extra computing power needed for fully quantum models.

It was also shown that QISVM operates in the same manner on datasets of different durations, from one year to seven years. The forecasts generated by QISVM were very close to the actual price movements, and the model worked more smoothly and made fewer mistakes than other models. The constancy of RMSE and forecast accuracy in the studies conducted indicates that the QISVM is capable of accommodating market fluctuations and structural breaks frequently observed in financial time series.

The results of the benchmarking reveal that QISVM is a good balance of accuracy, efficiency, and convenience of use. It does not have the same issues with excessive computer

complexity or insufficient transparency like deep learning or hybrid methods do. QISVM works as well as quantum-enhanced techniques but does not depend on any specific technology. This makes it an excellent choice for use in real forecasting systems. The study might also include other forms of assets, such as commodities, foreign exchange, and digital assets. This would make the framework more complete where low latency and scalability are very important.

This study gives us the promise of QISVM for forecasting financial time series that potentially opens the door to other avenues for future research that could be further studied or explored. More tests in other volatility regimes and over longer time periods would assist in seeing how well it works in varied market scenarios.

One may also add other metrics to the quantum inspired feature space, including entanglement entropy or fidelity-based features. These may provide us with more structural insights into how the market works. Future study may further investigate the scalability of the QISVM framework on higher frequency or tick-level data with more rigorous latency constraints. Lastly, comparing QISVM to newer techniques, such as transformer-based architectures and reinforcement learning agents, would help us understand how it might integrate into current ways of forecasting the future of finance.

The recommended framework's effectiveness and scalability show that it could be useful in many areas that need accurate forecasts but do not have a lot of computing power. Some possible uses are monitoring health, forecasting energy load and transportation demand. If QISVM is developed to include these domains, quantum inspired learning could be useful for more than just forecasting financial time series. It could also help people and businesses make decisions.

ACKNOWLEDGMENT

This work was supported by the Ministry of Higher Education Malaysia under Fundamental Research Grant Scheme (FRGS/1/2023/ICT06/UPM/02/1).

REFERENCES

- [1] J. L. Castle, J. A. Doornik and D. F. Hendry, "Forecasting principles from experience with forecasting competitions," *Forecasting*, vol. 3, no. 1, pp. 138–165, 2021.
- [2] K. Mansurov, A. Semenov, D. Grigoriev, A. Radionov and R. Ibragimov, "Impact of self-learning-based high-frequency traders on the stock market," *Expert Systems with Applications*, vol. 232, Art. no. 120567, 2023.
- [3] A. Kurani, P. Doshi, A. Vakharia and M. Shah, "A comprehensive comparative study of artificial neural networks and support vector machines for stock forecasting," *Annals of Data Science*, vol. 10, no. 1, pp. 183–208, 2023.
- [4] D. Gupta, M. Pratama, Z. Ma, J. Li and M. Prasad, "Financial time series forecasting using twin support vector regression," *PLoS ONE*, vol. 14, no. 3, Art. no. e0211402, 2019.
- [5] Q. Tang, T. Fan, R. Shi, J. Huang and Y. Ma, "Prediction of financial time series using LSTM and data denoising methods," *Mathematical Problems in Engineering*, vol. 2021, Art. no. 9942410, 2021.
- [6] M. Rath and H. Date, "Quantum data encoding: A comparative analysis of classical-to-quantum mapping techniques and their impact on machine learning accuracy," *EPJ Quantum Technology*, vol. 11, Art. no. 72, 2024.
- [7] A. Ganapathy, "Quantum computing in high-frequency trading and fraud detection," *Engineering International*, vol. 9, no. 2, pp. 61–72, 2021.
- [8] N. H. Abd Rahman, H. J. Goh and H. S. Zulkaffli, "GARCH models and distributions comparison for nonlinear time series with volatilities," *Malaysian Journal of Fundamental and Applied Sciences*, vol. 19, no. 6, pp. 989–1001, 2023.
- [9] X. Song, L. Deng, H. Wang, Y. Zhang, Y. He and W. Cao, "Deep learning-based time series forecasting," *Artificial Intelligence Review*, vol. 58, no. 1, 2025.
- [10] B. Lestari, N. Chamidah, D. Aydin and E. Yilmaz, "Reproducing kernel Hilbert space approach to multi-response smoothing spline regression," *Symmetry*, vol. 14, no. 11, Art. no. 2227, 2022.
- [11] L. Huynh, J. Hong, A. Mian, H. Suzuki, Y. Wu and S. Camtepe, "Quantum inspired machine learning: A survey," *Machine Learning*, arXiv:2308.11269v2, 2023.
- [12] D. Peral-García, J. Cruz-Benito and F. J. García-Peñalvo, "Systematic literature review: Quantum machine learning and its applications," *Computer Science Review*, vol. 51, Art. no. 100619, 2024.
- [13] M. Schuld and F. Petruccione, *Machine Learning with Quantum Computers*. Cham, Switzerland: Springer, 2021.
- [14] S. Khurana and M. J. Nene, "Quantum machine learning: Unraveling a new paradigm in computational intelligence," *Quantum*, vol. 74, no. 1, 2024.
- [15] A. Jiménez-Cordero, J. M. Morales and S. Pineda, "An embedded min-max approach for feature selection in nonlinear SVM classification," *European Journal of Operational Research*, vol. 293, no. 1, pp. 24–35, 2021.
- [16] IBM Quantum, "Qiskit Aer: High-performance simulators for quantum circuits," 2024. [Online]. Available: <https://qiskit.github.io/qiskit-aer/>
- [17] S. Deng, X. Huang, J. Wang, Z. Qin, Z. Fu, A. Wang and T. Yang, "A decision support system for trading in the Apple futures market using prediction fusion," *IEEE Access*, vol. 9, pp. 1271–1285, 2020.
- [18] J. Gao, "Applications of machine learning in quantitative trading," *Applied and Computational Engineering*, vol. 82, no. 1, pp. 124–129, 2024.
- [19] J. Arora and V. Jain, "Innovative applications and advanced practices in financial data science and machine learning for high-frequency trading," in *Machine Learning and Modeling Techniques in Financial Data Science*. IGI Global, 2025, pp. 185–208.
- [20] A. Serrano, "High-frequency trading and systemic risk: A structured review of findings and policies," *Review of Economics*, vol. 71, no. 3, pp. 169–195, 2020.
- [21] P. Xia, "Prediction of financial time series data using ARIMA model and machine learning algorithm," *ACM*, 2025.
- [22] T. Mohammed, H. U. Yahaya and S. O. Adams, "Modeling volatility of selected beverage stock returns in Nigeria using GARCH models," *Matrix Science Mathematic*, vol. 6, no. 2, pp. 41–51, 2022.
- [23] A. Ampountolas, "Enhancing forecasting accuracy in commodity and financial markets using GARCH and SVR models," *International Journal of Financial Studies*, vol. 12, no. 3, Art. no. 59, 2024.
- [24] M. M. Kumbure, C. Lohrmann, P. Luukka and J. Porras, "Machine learning techniques and data for stock market forecasting: A literature review," *Expert Systems with Applications*, vol. 197, Art. no. 116659, 2022.
- [25] K. U. Apu, M. M. Rahman, A. B. Hoque and M. Bhuiyan, "Forecasting future investment value using machine learning, neural networks and ensemble learning: A meta-analytic study," *Review of Applied Science and Technology*, 2022.
- [26] M. Ghadimpour and S. B. Ebrahimi, "Forecasting financial time series using deep learning networks," *Iranian Journal of Finance*, vol. 6, no. 4, pp. 81–94, 2022.
- [27] V. Velichala, "The future of low-latency systems in capital markets," *Journal of Computer Science and Technology Studies*, vol. 7, no. 6, pp. 507–513, 2025.
- [28] M. Celestin et al., "Applications of GARCH models for volatility forecasting in high-frequency trading environments," *International Journal of Applied and Advanced Scientific Research*, vol. 10, no. 1, pp. 12–21, 2025.
- [29] V. Palaniappan et al., "A review on high-frequency trading forecasting methods: Opportunities and challenges for quantum-based methods," *IEEE Access*, vol. 12, pp. 167471–167488, 2024.

- [30] Y. Zou, "Forecasting Apple Inc. stock prices: A comparative analysis of ARIMA, LSTM and ARIMA-LSTM models," *Advances in Operations Research and Production Management*, vol. 4, no. 1, pp. 23–34, 2025.
- [31] K. Kea, D. Kim, C. Huot, T.-K. Kim and Y. Han, "A hybrid quantum-classical model for stock price prediction using quantum-enhanced LSTM," *Entropy*, vol. 26, no. 11, Art. no. 954, 2024.
- [32] A. Farooq et al., "Interpretable multi-horizon time series forecasting using temporal fusion transformers," *Heliyon*, vol. 10, no. 22, Art. no. e40142, 2024.
- [33] G. Zhu, "Applications of artificial intelligence and machine learning in financial forecasting and trading," *Applied and Computational Engineering*, vol. 43, pp. 1–7, 2024.
- [34] A. Deep, A. Shirvani, C. Monico, S. Rachev and F. Fabozzi, "Risk-adjusted performance of random forest models in high-frequency trading," *Journal of Risk and Financial Management*, vol. 18, no. 3, Art. no. 142, 2025.
- [35] N. Innan, M. A. Z. Khan, B. Panda and M. Bennai, "Enhancing quantum support vector machines through variational kernel training," *Quantum Information Processing*, vol. 22, Art. no. 374, 2023.
- [36] N.-R. Zhou, S.-H. Xia, Y. Ma and Y. Zhang, "Quantum particle swarm optimization with truncated mean stabilization," *Quantum Information Processing*, vol. 21, no. 2, Art. no. 42, 2022.
- [37] A. Patel, "Advancing financial prediction through quantum machine learning," *International Journal of Intelligent Data and Machine Learning*, vol. 2, no. 2, pp. 1–7, 2025.
- [38] P. Chhajjer, M. Shah and A. Kshirsagar, "Applications of artificial neural networks, support vector machines and LSTM for stock market prediction," *Decision Analytics Journal*, vol. 2, Art. no. 100015, 2022.
- [39] M. Saghafi and L. Mili, "Predictive time series analysis of quantum circuit outcomes," *IEEE Access*, 2025.
- [40] W. Wang et al., "A hybrid ARIMA-LSTM model for short-term vehicle speed prediction," *Energies*, vol. 17, no. 15, Art. no. 3736, 2024.
- [41] K. Kashif and R. Ślepaczuk, "LSTM-ARIMA as a hybrid approach in algorithmic investment strategies," *arXiv:2406.18206*, 2024.
- [42] A. Moghar and M. Hamiche, "Stock market prediction using LSTM recurrent neural networks," *Procedia Computer Science*, vol. 170, pp. 1168–1173, 2020.
- [43] Y. Mahmoudi, "Algorithmic trading and price forecasting using machine learning and deep learning," *Quality & Quantity*, 2025.
- [44] T. Mahmud et al., "Predictive modeling in forex trading: A time series analysis approach," in *Proc. Int. Conf. Inventive Computing and Informatics (ICICI)*, 2024, pp. 390–397.
- [45] B. H. A. Khattak et al., "A systematic survey of AI models in financial market forecasting," *IEEE Access*, vol. 11, pp. 125359–125380, 2023.
- [46] C. Y. Yang et al., "Stock price prediction based on time series and machine learning techniques," *Applied Artificial Intelligence*, vol. 38, no. 1, 2024.
- [47] C. G. Kibebe, Y. Liu and J. Tang, "Harnessing optical advantages in computing: Current and future trends," *Frontiers in Physics*, vol. 12, Art. no. 1379051, 2024.
- [48] A. Saini et al., "Financial time series prediction on Apple stocks using machine and deep learning models," in *Proc. Int. Conf. Electrical, Computer and Energy Technologies (ICECET)*, 2024, pp. 1–6.
- [49] D. Vallarino, "Understanding market trends using RNNs and transformer-based sentiment analysis," *Journal of Applied Economic Sciences*, vol. 20, no. 1, pp. 21–33, 2025.
- [50] M. Goldani, "Comparative analysis of forecasting methods in financial time series," *Journal of Mathematics and Modeling in Finance*, vol. 3, no. 2, pp. 37–61, 2023.
- [51] S. Hrehova et al., "From simulation to validation in model-based predictive analysis," *Applied Sciences*, vol. 15, no. 6, Art. no. 3107, 2025.
- [52] M. Guimarães et al., "Predicting model training time to optimize distributed machine learning applications," *Electronics*, vol. 12, no. 4, Art. no. 871, 2023.