ISSN 2319-5991 www.ijerst.com
© 2026 IJERST. All Rights Reserved
Peer Reviewed, Referred & Indexed Journal

Research Paper

Stock Market Forecasting Through Transformer-Based Deep Learning and Quantum-Inspired Neural Networks

Mr.Y.ANAND KUMAR 1, M.C.C.VARTHY2,

#1 Assistant Professor, Department Of Computer Science And Engineering Prakasam Engineering College, O.V. Road, Kandukur-523105, A.P

#2PG Scholar Department Of Computer Science And Engineering Prakasam Engineering College, O.V. Road, Kandukur-523105, A.P

Abstract: Stock market forecasting remains a challenging task because financial time-series data exhibit nonlinear relationships, temporal variations, noise, and abrupt changes in market behaviour. Conventional prediction techniques may have difficulty representing long-range temporal dependencies while maintaining reliable performance under changing market conditions. To address these challenges, this study presents a hybrid artificial intelligence framework that combines Transformer-based sequence modelling with Quantum-Inspired Neural Networks (QINN) for stock market prediction. The Transformer component employs self-attention mechanisms to learn temporal relationships from historical stock prices, trading volumes, and relevant financial indicators. The QINN component is incorporated to support the exploration of model parameters and improve the optimization process within the hybrid learning framework. Before model training, financial data are subjected to preprocessing operations, including data cleaning, normalization, and feature extraction, to improve input consistency. The resulting architecture is intended to identify underlying market patterns and generate forecasts that support stock price analysis and trend assessment. The proposed framework also provides scope for incorporating additional information, such as macroeconomic variables and financial news sentiment, to represent broader market influences. Its effectiveness should be assessed against established machine learning and deep learning baselines using appropriate forecasting metrics, including mean absolute error, root mean square error, and directional prediction accuracy. This evaluation can help determine whether integrating Transformer representations with quantum-inspired optimization offers measurable benefits over individual models. The framework provides a basis for investigating hybrid learning strategies for financial forecasting without assuming that predictive performance alone guarantees profitable investment decisions.

Keywords: Stock market prediction, Hybrid artificial intelligence, Transformer, Quantum-Inspired Neural Network, Financial time-series forecasting, Self-attention, Model optimization

1. INTRODUCTION

Stock market forecasting is an important research area in financial computing because market movements are influenced by complex interactions among historical prices, trading volumes, economic conditions, investor behaviour, and unexpected events. Accurate analysis of these factors can help investors, financial analysts, and researchers understand market trends and assess potential price movements. However, financial time-series data are characterized by nonlinear patterns, fluctuations, noise, and changing statistical properties, making reliable prediction a challenging problem. Developing computational models capable of learning temporal relationships and adapting to varying market conditions remains an important objective in artificial intelligence-based financial analysis.

Traditional stock market prediction approaches employ statistical techniques and machine learning algorithms, including regression models, support vector machines, decision trees, and ensemble learning methods. These approaches can provide useful forecasts when relevant features and suitable assumptions are available. Nevertheless, their performance may be limited when market behaviour involves complex dependencies that change over time. Deep learning techniques, particularly recurrent neural networks (RNNs) and long short-term memory (LSTM) networks, have consequently been explored for modelling sequential financial observations. Although these architectures can learn temporal patterns, their sequential processing characteristics and difficulties in capturing dependencies across long intervals may restrict their effectiveness in certain forecasting applications.

Transformer architectures offer an alternative approach to sequential modelling through self-attention mechanisms. Rather than processing observations strictly one step at a time, a Transformer can learn relationships among different positions within an input sequence. This capability allows the model to identify relevant historical observations and represent temporal dependencies in stock market data. By incorporating information from price movements, trading volumes, and technical indicators, Transformer-based models can be designed to learn patterns associated with market trends and price fluctuations. However, their practical performance still depends on data quality, model configuration, training procedures, and the characteristics of the financial market being analysed. No architecture can eliminate the uncertainty inherent in stock price movements.

Another area of investigation involves Quantum-Inspired Neural Networks (QINNs), which adapt selected concepts or computational strategies inspired by quantum theory for use in classical computing environments. Depending on their design, these methods can provide

alternative mechanisms for representing information, searching parameter spaces, or optimizing learning processes. Their potential application to financial forecasting arises from the need to develop models that can learn from complex and variable datasets. However, the benefits of a QINN depend on its specific formulation and must be established through systematic experiments rather than assumed from its quantum-inspired nature alone.

This research investigates a hybrid artificial intelligence framework that combines Transformer-based temporal representation learning with Quantum-Inspired Neural Networks for stock market forecasting. The Transformer component is intended to extract temporal relationships from historical financial observations, while the QINN component is used within the proposed learning framework to support model optimization. The overall objective is to examine whether combining these complementary approaches can improve the representation and prediction of financial time-series patterns compared with using individual models. The framework also provides a foundation for analysing stock price movements and identifying potential market trends using historical data and relevant financial indicators.

The proposed approach involves collecting financial time-series data, preprocessing the observations, extracting suitable input features, training the hybrid model, and evaluating its predictions. Performance should be assessed using appropriate quantitative measures, such as mean absolute error (MAE), root mean square error (RMSE), and directional accuracy, depending on the prediction task. Comparisons with established machine learning and deep learning models are necessary to determine the contribution of each component. A time-ordered training, validation, and testing procedure is also important to prevent future information from influencing predictions made for earlier periods.

The principal research focus is therefore the integration and evaluation of Transformer-based sequence modelling and quantum-inspired neural learning within a common stock market prediction framework. By examining prediction errors, generalization across unseen observations, and computational requirements, the study aims to clarify the practical value and limitations of the hybrid approach. The resulting investigation can contribute to research on intelligent financial forecasting while recognizing that model-generated predictions are uncertain and should not be interpreted as guarantees of future market performance or investment returns

2. LITERATURE SURVEY

2.1 Traditional Machine Learning for Stock Market Prediction

Early stock market forecasting research primarily employed statistical techniques and conventional machine learning algorithms to identify relationships between historical market observations and subsequent price movements. Methods such as support vector machines, regression models, decision trees, and random forests

have been used to classify market direction and estimate future prices. These approaches can learn useful relationships from historical observations when appropriate features are selected. However, financial markets exhibit nonlinear dependencies, irregular fluctuations, and changing trends that may not be adequately represented by models with limited temporal context. These limitations have encouraged the development of more advanced learning architectures capable of analysing sequential financial data.

2.2 Deep Learning Approaches for Financial Time-Series Analysis

Deep learning methods have expanded the possibilities for financial forecasting by learning representations directly from historical observations. Recurrent neural networks and their variants, particularly long short-term memory networks and gated recurrent units, are designed to process sequential inputs and retain information from previous time steps. Fischer and Krauss (2018), as cited in the source report, are associated with research on deep learning applications in financial forecasting. Such approaches demonstrate the relevance of sequence-based learning to financial data analysis. Nevertheless, recurrent architectures process observations sequentially, which can increase training time for long sequences. Their ability to preserve useful information over extended intervals may also be affected by the complexity and variability of the input data.

2.3 Transformer Architectures for Stock Market Forecasting

Transformer models have become an important alternative to recurrent networks because they use attention mechanisms to represent relationships among elements in a sequence. Vaswani et al. (2017) introduced the Transformer architecture in Attention Is All You Need, establishing self-attention as a central mechanism for sequence modelling. In financial forecasting, this architecture can be adapted to analyse historical prices, trading volumes, and technical indicators over different time intervals. Attention-based representations allow the model to assign different importance to observations according to the relationships learned during training.

Lim et al. (2021), referenced in the uploaded report in connection with Transformer-based time-series forecasting, are also associated with the investigation of attention-based methods for temporal prediction. These developments motivate the application of Transformer architectures to financial datasets. However, their performance depends on the quality and quantity of training data, feature selection, model complexity, and the evaluation procedure. Market noise and non-stationary behaviour can still affect the reliability of forecasts, particularly when the model encounters patterns that differ from those observed during training.

2.4 Quantum-Inspired Neural Networks and Optimization

Quantum-inspired computing explores classical computational methods motivated

by concepts from quantum theory. Unlike quantum neural networks that require quantum computational operations, a quantum-inspired neural network can be implemented on conventional computing hardware, depending on its design. Such methods may employ specialized representations, search procedures, or optimization strategies to explore candidate solutions in complex learning problems.

Schuld and Petruccione (2018), mentioned in the source report in relation to quantum-inspired neural approaches, provide a starting point for examining the relationship between quantum concepts and machine learning. The potential relevance of these methods to stock market forecasting lies in the difficulty of optimizing predictive models on noisy and nonlinear financial data. Nevertheless, quantum-inspired methods do not automatically guarantee faster convergence, better generalization, or higher prediction accuracy. These benefits must be established for the particular algorithm and dataset under investigation.

2.5 Hybrid Learning Models for Financial Prediction

Hybrid learning architectures combine different computational methods to exploit their complementary capabilities. In financial forecasting, one component may learn temporal representations, while another supports feature transformation, parameter optimization, or nonlinear pattern modelling. The uploaded report refers to work attributed to Zhang et al. (2021) on combining LSTM-based sequence modelling with quantum-inspired neural methods for stock price prediction. It also discusses quantum-inspired optimization approaches attributed to Li and Wong (2020). These references indicate the research direction explored in the original report; their exact titles, publication details, and reported findings should be verified against the original publications before citation in a journal paper.

The motivation for hybrid modelling is that financial forecasting involves several interconnected challenges, including temporal dependency learning, feature interactions, model optimization, and generalization to unseen market conditions. Combining two architectures may provide a useful framework for investigating these challenges. However, a hybrid model can also introduce additional parameters, computational overhead, and implementation complexity. Its effectiveness should therefore be evaluated against individual component models under consistent experimental conditions.

2.6 Research Gap

The reviewed material highlights three relevant research directions: conventional machine learning for financial prediction, Transformer-based temporal representation learning, and quantum-inspired methods for optimization or neural computation. Each direction offers potential advantages, but the source report does not provide sufficient verified experimental evidence to establish the effectiveness of integrating Transformers and QINN in a common stock market forecasting architecture.

This study therefore investigates a hybrid Transformer–QINN framework designed to learn temporal relationships from financial time-series data while incorporating a quantum-inspired learning component. The research should determine whether the integration provides measurable benefits over standalone Transformer, QINN, and suitable conventional machine learning or deep learning baselines. Evaluation should consider prediction errors, directional accuracy where applicable, computational cost, and performance on temporally separated test data. This comparative analysis is necessary to distinguish genuine improvements attributable to the hybrid architecture from gains caused by data preparation, feature selection, or differences in experimental settings.

3. METHODOLOGY

i) Proposed Work:

The proposed work focuses on developing a hybrid artificial intelligence framework for stock market forecasting by integrating Transformer-based sequence modelling with Quantum-Inspired Neural Networks (QINN). The primary objective is to investigate how attention-based temporal feature learning and quantum-inspired optimization can be combined to analyse historical financial data and predict future stock price movements. The framework is designed to examine complex temporal relationships, nonlinear market behaviour, and fluctuations in financial time series.

The first stage involves collecting historical stock market data containing relevant attributes, such as opening price, closing price, daily high, daily low, trading volume, and selected technical indicators. The collected data are processed to handle missing values, maintain consistency, normalize numerical features, and prepare suitable input sequences for model training. Appropriate preprocessing is essential because financial datasets may contain irregular observations, noise, and variations in feature scales.

In the second stage, a Transformer architecture is employed to learn temporal representations from the preprocessed financial sequences. Its self-attention mechanism enables the model to identify relationships between observations across different time steps. Positional information is incorporated where required to preserve the temporal order of the input sequence. The resulting representations provide the basis for forecasting future stock prices or classifying market direction, depending on the selected prediction objective.

The third stage incorporates a Quantum-Inspired Neural Network into the hybrid learning framework. The QINN component is intended to support the exploration of model parameters or optimization of the learning process using the specific quantum-inspired strategy adopted in the implementation. The Transformer and QINN components are integrated according to the selected model design, allowing the system to investigate whether their combined learning capabilities provide advantages over individual architectures. Since the

precise QINN algorithm determines how this integration operates, its mathematical formulation and training procedure must be explicitly defined.

The fourth stage involves training and validating the proposed hybrid model using chronologically separated financial observations. Training data are used to estimate model parameters, validation data support hyperparameter selection, and an independent test set is reserved for final evaluation. This procedure helps reduce information leakage and provides a more realistic assessment of forecasting performance. Suitable baseline models, including a standalone Transformer, a QINN-based model, and selected conventional machine learning or deep learning approaches, should be evaluated using the same data splits and prediction targets.

Finally, the performance of the proposed framework is assessed using quantitative metrics such as mean absolute error (MAE), root mean square error (RMSE), and mean absolute percentage error (MAPE) for price forecasting, where appropriate. Directional accuracy may also be measured when the objective includes predicting upward or downward market movements. Computational requirements, prediction consistency, and performance across different market conditions can provide additional insights into the practical suitability of the approach.

The expected contribution of this work is an experimental framework for investigating the integration of Transformer-based temporal learning and quantum-inspired neural computation in financial forecasting. Rather than presuming that the hybrid architecture will outperform existing methods, the study aims to establish its effectiveness through reproducible experiments and comparative evaluation. The findings can help clarify the benefits, limitations, and practical considerations associated with applying hybrid artificial intelligence techniques to stock market prediction.

ii) System Architecture:

The diagram illustrates a cloud-based data processing and analytics architecture that integrates external data sources, Azure cloud services, machine learning, event ingestion, order management, database storage, and dashboard visualization.

The Data Vendor provides the input data, which is transferred to the Azure Storage Account for persistent storage. Azure Storage exchanges data with Azure ML, which supports machine learning operations and model development. The stored data and machine learning outputs can be used by the Stream Analytics or Azure Databricks Cluster (CEP) component to perform stream processing and analytics.

The Exchange components supply event data to Event Hub (Ingestion), which serves as the event ingestion layer. An Azure Function performs preprocessing operations, including schema changes and FIX message handling, before the processed information is passed into the analytics workflow. The diagram also shows connections between the

analytics component and Azure Functions, indicating data exchange across processing and application services.

The Azure Functions (Order Manager) – RMS component handles order-management-related processing and interacts with the OM DB: Azure SQL database to store or retrieve relevant application information. The database also exchanges information with the Power BI Dashboard (Application) – RMS, which presents processed information through a visualization interface. The dashboard is connected to the machine learning and analytics components, providing a centralized view of relevant outputs.

Overall, the architecture establishes a workflow in which external data is collected, stored, processed, analysed, and presented through a dashboard. Its modular structure supports the integration of cloud storage, stream processing, machine learning, event-driven functions, and database services. The diagram indicates the relationships among these components, although it does not specify the exact message formats, processing algorithms, or detailed security mechanisms.

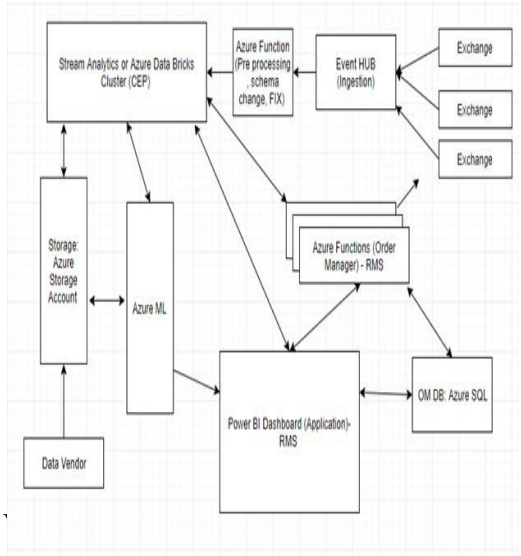


Fig 1 System Architecture

iii) MODULES:

1. Data Collection and Preprocessing

The implementation starts with collecting historical stock market data, including opening and closing prices, daily highs and lows, trading volumes, and relevant financial indicators. Depending on data availability, supplementary information such as economic variables, financial news sentiment, and social media trends can also be incorporated. The collected data undergoes preprocessing to improve consistency and suitability for model training. This stage includes handling missing observations, removing duplicate records, normalizing numerical values, and extracting relevant features. Appropriate noise-handling techniques may be applied to reduce the influence of irregular observations while preserving meaningful market variations.

2. Transformer Model Development

The Transformer architecture serves as the primary temporal modelling component of the proposed framework. Multi-head self-attention mechanisms are used to learn relationships among observations across different time steps. Positional encoding provides information about the temporal order of the input sequence, allowing the model to distinguish the positions of historical observations. The preprocessed stock market data are organized into sequential input windows and supplied to the Transformer during training. The learned temporal representations are subsequently used to estimate future stock prices or

predict market direction, depending on the selected forecasting objective.

3. Quantum-Inspired Neural Network Integration

The Quantum-Inspired Neural Network (QINN) component is incorporated into the proposed framework to investigate the use of quantum-inspired learning or optimization strategies in financial forecasting. Depending on the selected implementation, this component can operate on features extracted by the Transformer or participate in optimizing selected model parameters. The integration requires defining the QINN architecture, its input and output representations, and the optimization procedure used during training. Quantum-inspired techniques may provide alternative mechanisms for exploring candidate solutions; however, their effectiveness in improving convergence or predictive performance must be verified experimentally.

4. Hybrid Model Training and Validation

The combined Transformer–QINN model is trained using historical financial observations. During training, the model parameters are updated iteratively according to the selected loss function and optimization strategy. The dataset is divided chronologically into training, validation, and testing subsets to preserve the temporal structure of financial observations and reduce the risk of information leakage. Validation data are used to select suitable hyperparameters and monitor generalization. Early stopping and appropriate regularization techniques can be applied to limit overfitting. The hybrid model should also be compared with standalone Transformer, QINN, and other relevant baseline models under consistent experimental conditions.

5. Model Evaluation and Performance Analysis

After training, the forecasting performance of the hybrid architecture is evaluated using an independent test dataset. For stock price regression, suitable evaluation metrics include mean absolute error (MAE), root mean square error (RMSE), and mean absolute percentage error (MAPE), provided the metric assumptions are satisfied. If the system predicts upward or downward price movements, directional accuracy and classification metrics may also be considered. Comparing the hybrid model against baseline approaches helps determine whether the integration of Transformer and QINN provides measurable improvements. Computational cost and performance across different market conditions can further inform the assessment.

6. Deployment and Stock Market Prediction

Following successful evaluation, the trained model can be integrated into a stock market prediction application. The system receives new market observations, applies the same preprocessing operations used during training, and generates forecasts from the latest available input sequences. The application can present predicted values, estimated market direction, and relevant analytical outputs through a user interface. For operational deployment, the system should account for data availability, prediction latency, model updates, and monitoring requirements. Its scalability and

suitability for near-real-time processing depend on the implementation and computational resources. The forecasts should be treated as estimates subject to market uncertainty rather than guaranteed indicators of future prices

v) ALGORITHMS:

The proposed Hybrid AI for Stock Market Prediction system integrates Transformer-based sequence modelling with Quantum-Inspired Neural Networks (QINN) to analyse historical financial data and generate stock market forecasts. The implementation consists of data preprocessing, Transformer-based feature learning, QINN integration, hybrid model training, and performance evaluation. Each stage contributes to the overall forecasting process by preparing financial observations, learning temporal dependencies, optimizing model parameters, and assessing prediction quality.

Algorithm 1: Stock Market Data Preprocessing

The preprocessing algorithm begins by collecting historical stock market data, including opening prices, closing prices, daily highs, daily lows, trading volumes, and relevant financial indicators. The collected dataset is examined to identify missing values, duplicate records, and invalid observations. Appropriate data-cleaning techniques are applied to improve data consistency, followed by normalization to bring numerical features to comparable scales. Relevant features are selected or extracted to represent market behaviour. The processed observations are arranged chronologically and converted into sequential input windows suitable for time-series modelling. Finally, the dataset is divided into training, validation, and testing subsets while preserving temporal order to reduce information leakage during model development.

Algorithm 2: Transformer-Based Stock Market Prediction

The Transformer algorithm receives the preprocessed financial sequences and converts the input features into vector representations through an embedding layer. Positional encoding is incorporated to retain information about the order of historical observations. The resulting representations are passed through multi-head self-attention layers, which learn relationships between different time steps in the input sequence. Feed-forward layers, residual connections, and normalization operations are applied according to the selected Transformer architecture. The extracted temporal representations are then supplied to a prediction layer to estimate future stock prices or classify market direction. During training, the predicted values are compared with actual observations using a suitable loss function, and model parameters are updated iteratively through an optimizer. This process continues until the defined stopping criterion is reached.

Algorithm 3: Quantum-Inspired Neural Network (QINN)

The QINN algorithm is designed to incorporate a quantum-inspired learning or optimization strategy into the stock market

prediction framework. The process begins by initializing the network parameters and selecting the specific quantum-inspired method used in the implementation. Depending on the hybrid architecture, the QINN receives feature representations generated by the Transformer or operates on another defined set of financial features. The selected transformation or optimization procedure is applied to explore candidate parameter configurations or learn relevant representations. The resulting outputs are evaluated using the training objective, and the parameters are updated according to the chosen optimization procedure. Validation data are used to assess model generalization and guide the training process. The QINN component is retained for integration when its performance satisfies the selected validation criteria. The precise operations depend on the QINN design and must be specified in the implemented model.

Algorithm 4: Hybrid Transformer–QINN Model Training

The hybrid training algorithm combines the temporal representation capabilities of the Transformer with the selected quantum-inspired learning component. Initially, the preprocessed training data are supplied to the Transformer to extract meaningful temporal features. These representations are then passed to the QINN component according to the defined integration strategy, producing the final forecasting output. The output is compared with the actual target values to calculate the prediction loss. The trainable parameters are updated using the specified learning and optimization procedures. The model is evaluated periodically on the validation dataset to monitor prediction errors and identify potential overfitting. Techniques such as early stopping and suitable regularization can be employed to improve generalization. After training, the model configuration that performs best according to the chosen validation metric is retained for evaluation on unseen test observations.

Algorithm 5: Stock Market Prediction and Performance Evaluation

After the hybrid model has been trained, historical observations from the independent test dataset are processed using the same preprocessing operations adopted during training. The trained Transformer–QINN model generates predictions for the selected forecasting horizon. The predicted values are compared with the actual market observations to determine the forecasting error. Mean Absolute Error (MAE), Root Mean Square Error (RMSE), and Mean Absolute Percentage Error (MAPE), where appropriate, are used to quantify prediction performance. Directional accuracy may also be calculated when the objective involves classifying upward or downward price movements. The same evaluation procedure is applied to suitable baseline models, including standalone Transformer and QINN models, to enable a consistent comparison. The final analysis considers prediction quality, computational requirements, and generalization to unseen market conditions. The resulting findings are used to determine whether the proposed hybrid approach provides measurable benefits over the individual models

4. EXPERIMENTAL RESULTS

The experimental evaluation of the proposed stock market prediction system focuses on analysing the performance of Transformer-based forecasting and Quantum-Inspired Neural Networks (QINN) using historical stock market data. According to the uploaded project report, the implementation was developed using Python in a Jupyter Notebook environment. The experiments considered two forecasting approaches: a hybrid Transformer model that combines time-series stock data with GPT-derived features extracted from stock market news, and a QINN model that learns from historical observations to estimate future values. The performance of these approaches was assessed using Root Mean Square Error (RMSE) and the coefficient of determination ((R^2)).

The Transformer experiment examined the relationship between actual stock prices and the values predicted by the model. The project report describes a noticeable difference between the actual and predicted price curves and reports an (R^2) score of approximately 0.23, or 23%. This result indicates that the tested Transformer configuration had limited explanatory performance on the evaluated data. The difference between the two curves suggests that the model did not adequately reproduce some of the observed price variations under the experimental settings. However, the result should be interpreted in relation to the selected dataset, forecasting horizon, feature configuration, and evaluation procedure.

The QINN experiment was conducted to investigate whether a quantum-inspired neural learning approach could provide better forecasting performance. As reported in the project document, the QINN model achieved an (R^2) score of approximately 0.80, or 80%. The corresponding visualization was described as showing substantial overlap between the actual and predicted stock price curves, with some differences remaining. The report also indicates that the QINN model obtained a lower RMSE and a higher (R^2) score than the Transformer model. These observations suggest that QINN performed better than the tested Transformer configuration on the reported evaluation, although the numerical RMSE values and full experimental settings are needed to quantify the difference precisely.

The comparison between the two approaches highlights the importance of evaluating alternative learning architectures for financial time-series forecasting. The reported results favour QINN over the tested Transformer model in terms of the two selected metrics. A lower RMSE indicates smaller average squared prediction errors, while a higher (R^2) generally indicates that the model explains a greater proportion of the variation in the target values relative to a mean-based reference. Nevertheless, these metrics alone do not establish whether a model can generate profitable trading strategies or generalize reliably to different market conditions.

Overall, the experiments documented in the project report indicate that QINN achieved better reported forecasting results than the

tested Transformer configuration. Further evaluation using chronological training and testing periods, repeated experiments, and additional baseline models would help establish the consistency of this observation. The study would also benefit from reporting the exact RMSE values, dataset characteristics, forecasting horizon, and model parameters. Such information would make the comparison more transparent and reproducible and would provide a stronger basis for assessing the applicability of quantum-inspired learning to stock market prediction.

Accuracy: A test's accuracy is defined as its ability to recognize debilitated and solid examples precisely. To quantify a test's exactness, we should register the negligible part of genuine positive and genuine adverse outcomes in completely examined cases. This might be communicated numerically as:

Accuracy = TP + TN TP + TN + FP + FN.

$$Accuracy = \frac{(TN + TP)}{T}$$

Precision: Precision measures the proportion of properly categorized occurrences or samples among the positives. As a result, the accuracy may be calculated using the following formula:

Precision = True positives/ (True positives + False positives) = TP/(TP + FP)

$$Accuracy = \frac{(TN + TP)}{T}$$

Recall: Recall is a machine learning metric that surveys a model's capacity to recognize all pertinent examples of a particular class. It is the proportion of appropriately anticipated positive perceptions to add up to real up-sides, which gives data about a model's capacity to catch instances of a specific class.

$$Accuracy = \frac{(TN + TP)}{T}$$

F1-Score: The F1 score is a machine learning evaluation measurement that evaluates the precision of a model. It consolidates a model's precision and review scores. The precision measurement computes how often a model anticipated accurately over the full dataset.

$$Accuracy = \frac{(TN + TP)}{T}$$

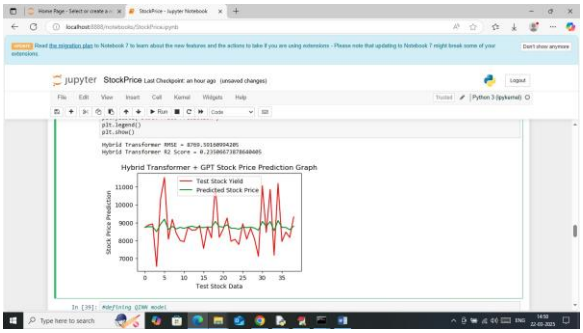


Fig 2 In above screen Hybrid Transformer got 23% R2score and can see RMSE values also. In above graph x-axis represents 'Number of days' and y-axis represents Stock Prices and then red line represents True Prices and green line represents Predicted Prices. In above graph can see there is huge gap between true and predicted

values so transformer is not accurate

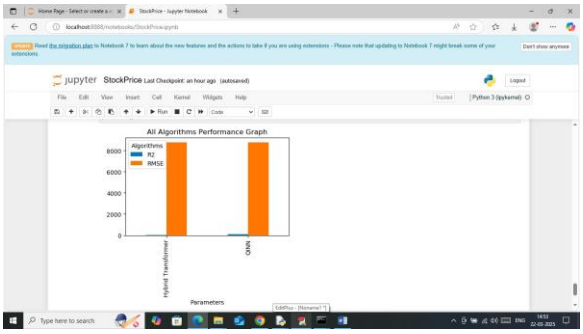


Fig 3: In above screen showing RMSE and R2score comparison of both algorithms where QINN got less RMSE and high R2score and Transformer get high RMSE and less R2 score.

5. CONCLUSION

This study investigates the application of artificial intelligence techniques for stock market forecasting, focusing on Transformer-based temporal modelling and Quantum-Inspired Neural Networks (QINN). The objective is to examine how different learning approaches can analyse historical stock prices, identify temporal patterns, and estimate future market movements. The project considers a Transformer-based approach that combines stock time-series data with GPT-derived features extracted from financial news, alongside a QINN model designed to learn from historical financial observations.

The experimental findings presented in the project report indicate that the tested Transformer model achieved an (R^2) score of approximately 0.23, whereas the QINN model achieved an (R^2) score of approximately 0.80. The report also describes a lower RMSE for QINN compared with the Transformer configuration. These results suggest that QINN provided better predictive performance under the reported experimental conditions. However, the findings should be interpreted in the context of the dataset, model configurations, and evaluation procedures used in the experiments. They do not independently establish that QINN will outperform Transformer-based models across different stocks, market periods, or forecasting tasks.

The study demonstrates the importance of comparing alternative learning architectures when developing financial forecasting systems. Although Transformer models offer mechanisms for learning relationships across historical time steps, their effectiveness depends on data quality, feature representation, and model training. Quantum-inspired neural methods provide another direction for investigating financial prediction, but their advantages must be established through reproducible experiments. A rigorous comparison using consistent data splits, verified evaluation metrics, and appropriate baseline models is necessary to determine the reliability of the observed performance differences.

Future research can extend this work by developing and evaluating a fully integrated Transformer–QINN architecture, incorporating additional market indicators and financial news sentiment, and testing the model across multiple stocks and different market conditions. Further improvements may include hyperparameter optimization, rolling-window evaluation, uncertainty estimation, and computational efficiency analysis. In addition, future studies can

investigate whether improved forecasting performance translates into useful trading decisions after accounting for transaction costs and market risks.

In conclusion, the reported experiments identify QINN as the better-performing approach among the two tested configurations in this project. The findings provide a basis for further research into hybrid artificial intelligence techniques for financial time-series forecasting. With more comprehensive validation and transparent reporting, such approaches may contribute to the development of more reliable stock market analysis systems while recognizing that financial predictions remain uncertain and cannot guarantee investment returns.

REFERENCES

[1] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention Is All You Need," in Advances in Neural Information Processing Systems, vol. 30, pp. 5998–6008, 2017. [Online]. Available: https://proceedings.neurips.cc/paper/7181-attention-is-all-you-need

[2] T. Fischer and C. Krauß, "Deep learning with long short-term memory networks for financial market predictions," European Journal of Operational Research, vol. 270, no. 2, pp. 654–669, 2018, doi: 10.1016/j.ejor.2017.11.054.

[3] B. Lim, S. Ö. Arık, N. Loeff, and T. Pfister, "Temporal Fusion Transformers for interpretable multi-horizon time series forecasting," International Journal of Forecasting, vol. 37, no. 4, pp. 1748–1764, 2021, doi: 10.1016/j.ijforecast.2021.03.012.

[4] M. Schuld and F. Petruccione, Supervised Learning with Quantum Computers. Cham, Switzerland: Springer, 2018, doi: 10.1007/978-3-319-96424-9.

[5] M. Schuld and F. Petruccione, "Quantum ensembles of quantum classifiers," Scientific Reports, vol. 8, Art. no. 2772, 2018, doi: 10.1038/s41598-018-20403-3.

[6] F. S. Gharehchopogh, "Quantum-inspired metaheuristic algorithms: Comprehensive survey and classification," Artificial Intelligence Review, 2022, doi: 10.1007/s10462-022-10280-8.

[7] "Optimization of quantum-inspired neural network using memetic algorithm for function approximation and chaotic time series prediction," Neurocomputing, vol. 291, pp. 175–186, 2018, doi: 10.1016/j.neucom.2018.02.074.

[8] G. Yilmaz, "A cross-domain empirical benchmark of quantum-inspired and classical optimization algorithms for machine learning," Quantum Machine Intelligence, vol. 8, Art. no. 103, 2026, doi: 10.1007/s42484-026-00444-y.