Prediction of Stocks Index Price Using Quantum GANs

Sangram Deshpande[1,2], Gopal Ramesh Dahale[1*], Sai Nandan Morapakula[1,3], Uday Wad[1]

[1]Qkrishi Quantum Pvt. Ltd., Gurgaon, Haryana, India
[2]Electrical and Computer Engineering, NC State University, Raleigh, USA
[3]Physics Department, University of Massachusetts, Boston, MA, USA
ssdesh24@ncsu.edu, dahalegopal27@gmail.com, sainandan.m@umb.edu, uday.wad@gmail.com

Abstract

This paper investigates the application of Quantum Generative Adversarial Networks (QGANs) for stock price prediction. Financial markets are inherently complex, marked by high volatility and intricate patterns that traditional models often fail to capture. QGANs, leveraging the power of quantum computing, offer a novel approach by combining the strengths of generative models with quantum machine learning techniques. We implement a QGAN model tailored for stock price prediction and evaluate its performance using historical market data. Results demonstrate that QGANs can generate synthetic data closely resembling actual market behavior, leading to enhanced prediction accuracy. The experiment was conducted using stock index price data and the AWS Braket SV1 simulator for training QGAN circuits. The quantum-enhanced model outperforms classical LSTM and GAN models in both convergence speed and prediction accuracy. This research marks a key step toward integrating quantum computing in financial forecasting, offering potential advantages in speed and precision over traditional methods. These findings hold promising implications for traders, financial analysts, and researchers.

Introduction

Accurate price prediction can aid in determining risk exposure, setting margin limits, and issuing margin calls, among other things. Nevertheless, the volatile nature of markets and the influence of multiple factors, such as policy changes, interest rate shifts, and currency fluctuations, make this process complex. Stock price prediction is essential for investors, financial analysts, and traders as it provides insights into future market trends. Accurate predictions enable investors to make informed decisions regarding buying, holding, or selling stocks, which can significantly affect their financial returns. Given the inherent volatility and complexity of financial markets, effective prediction models are crucial for risk management, portfolio optimization, and strategic planning. Reliable stock price forecasts help in minimizing losses and maximizing profits by identifying potential opportunities and threats in the market. For individual investors, it means better investment strategies and higher returns. For financial institutions, it enables the development of sophisticated trading algorithms, enhancing the efficiency and profitability of their trading operations. Companies can use these predictions to make strategic business decisions, such as timing for issuing new shares or buybacks. Additionally, accurate forecasts contribute to market stability by reducing the likelihood of large, unexpected price swings. This predictability can also foster investor confidence, attracting more capital into the markets and supporting overall economic growth. (Herman 2023)

Fully Quantum Generative Adversarial Networks (QGANs) are advanced machine learning models that leverage the principles of quantum computing to enhance generative modeling capabilities. Traditional GANs consist of a generator and a discriminator, where the generator creates synthetic data, and the discriminator evaluates its authenticity.(Orús, Mugel, and Lizaso 2019) QGANs incorporate quantum algorithms into this framework, potentially providing exponential speed-ups and improved accuracy due to quantum parallelism and entanglement. In stock price prediction, QGANs can process complex patterns and correlations in historical market data more efficiently than classical models, generating realistic synthetic data that closely mirrors actual market behaviors. This capability allows QGANs to produce more accurate and reliable stock price forecasts, aiding in better decision-making for investors and financial analysts. (Zhang, Huang, and Li 2022a)

Classical Generative Adversarial Networks (GANs) in finance are primarily used for generating synthetic financial time-series data, which is valuable for backtesting trading strategies, risk assessment, and stress testing financial models. They can capture complex patterns in financial data, such as volatility and correlations. Additionally, GANs are used for anomaly detection in financial data, identifying potential fraud, market manipulation, or sudden price changes. By learning normal patterns, they can pinpoint suspicious deviations. Classical GANs are also employed to enhance financial forecasting by creating more realistic data for training predictive models, thereby improving the accuracy of stock price forecasts. They aid in risk management by generating various market scenarios, helping financial institutions better understand their risk exposure.

[*]These authors contributed equally.

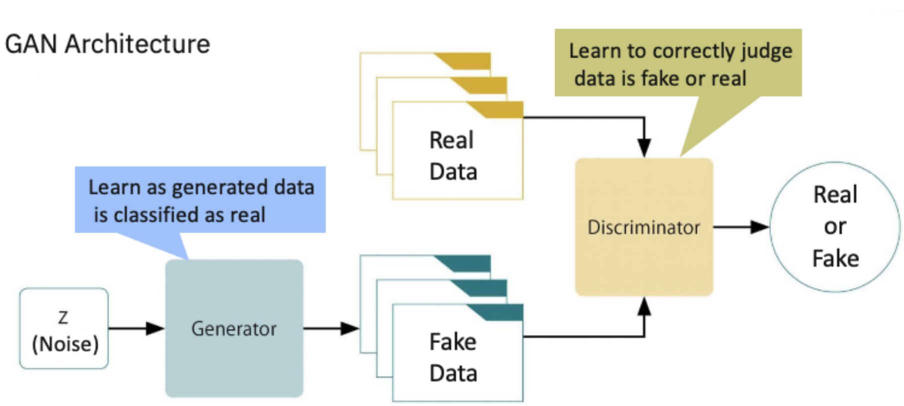

Figure 1: GAN Architecture from paper (Nakamura 2023)

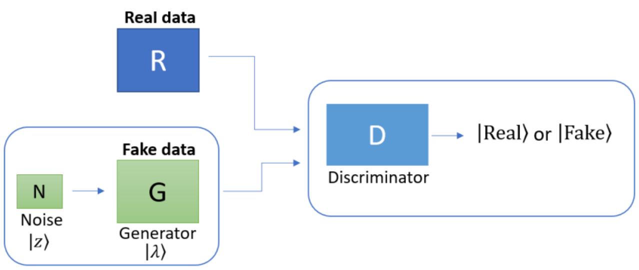

Figure 2: QGAN Architecture from paper (PennyLane 2019)

Problem Formulation: GANs To QGANs

We propose a generalized Quantum Generative Adversarial Network that takes advantage of the classical GAN architecture integrated with quantum-enhanced prediction power. (Lin et al. 2021) GANs are a popular artificial intelligence model that consists of two neural networks, as shown in Figure 1:

- Generator: Replicates a real-world dataset by producing synthetic data samples (e.g., text, audio, and/or photos).
- Discriminator: Assists in differentiating between authentic data and the generator's synthetic examples.

To "fake" realistic historical and future stock price data, mimicking real market dynamics and capturing complex relationships, the two competing entities, Generator and Discriminator, collaborate when using GANs for stock price prediction. The other part of the system acts as a critic, attempting to discern fake data from the real historical price data. As a result of their ongoing competition, the discriminator improves at spotting departures from actual market trends, while the generator gradually grows more skilled at producing realistic and representative data. The idea behind this is that after going through this training process, the generator could be able to provide "future" market data that almost exactly matches current prices, which could aid in prediction.

Training Loop and Feedback

- The generator is trained based on input from the discriminator, which pushes it to produce more realistic data that can trick the discriminator. The generator and discriminator are simultaneously improved by this iterative adversarial process. (Gonzalez and Miikkulainen 2020)
- The model's performance is determined by calculating the difference between the discriminator's predicted probability and the true label (real or created) using the loss function.
- The discriminator's weights and biases are updated using the loss using backpropagation, a method of modifying model parameters in response to prediction mistakes. The goal of this procedure is to increase the discriminator's capacity to differentiate between produced and actual data in later rounds.

The iterative adversarial process improves both the generator and discriminator simultaneously. Using a loss function, model performance is determined by evaluating the difference between predicted probabilities and true labels. Backpropagation updates the discriminator's weights and biases, enhancing its ability to differentiate between produced and actual data in subsequent rounds.

In the realm of QGANs, the discriminator employs a supervised learning technique, often a Quantum Deep Neural Network (DNN). (Amzhao 2022) Maintaining a delicate balance between the discriminator's ability to identify fakes and the generator's capacity to produce realistic data ensures continuous improvement throughout training.

Like in classical techniques, to categorise data as genuine or produced, the discriminator in QGAN usually uses a supervised learning technique, most often a quantum deep neural network (DNN). Unstable QGAN training and less-than-ideal outcomes might arise from a discriminator with inadequate training. The discriminator's capacity to identify fakes and the generator's capacity to generate realistic data must be balanced. This equilibrium makes sure that throughout training, both components become better repeatedly.

Another vital aspect is data encoding, converting processed data into a format compatible with the QGAN architecture, typically involving numerical vectorization or tensor conversion. Finally, data partitioning divides the dataset into training, testing sets to facilitate model training, evaluation, and validation. This partitioning ensures that the model's performance is assessed on unseen data, guarding against overfitting and enhancing generalization to new market conditions. In summary, meticulous data preparation establishes the groundwork for developing robust predictive models, enabling the effective utilization of QGANs for stock price prediction.

$$\lceil \log_2 b \rceil + \lceil \log_2 f \rceil + 1 \tag{1}$$

Overcoming the Limitation To overcome the limitation of the normalization factor, we use a simple strategy with which we can obtain the normalization factor for the predictions from the training dataset. We refer to the FQGAN, which uses this strategy as Invertible FQGAN. We explain the strategy with the help of an example.

Consider a past window of size 16 and a future window of size 8, which means that we need to predict 8 values. Instead of predicting 8 values, we try to predict 16 values in the future, i.e. we predict the 8 unknowns as well as the 8 knowns that are in the past window of size 16. Once trained on this type of data, given a past window of size 16 as input to the Quantum Generator, we can generate 16 values in the future out of which the first 8 overlap with the input data (Fig 3). Using classical optimization, we try to minimize the following objective function to find the normalization factor.

$$\sum_{i=1}^{8} (a_i - f\hat{y_i})^2 \tag{2}$$

Where a_i is the min-max scaled price of the overlapping input data, $\hat{y_i}$ is the normalized predicted price of the overlapping part, and f is the normalization factor.

Once f is found, we can multiply it with the $\hat{y_i}$ and then perform the inverse min-max transform to obtain the prediction prices.

Results

The results presented in this section encompass the outcomes of our experiments, offering a comprehensive analysis of the performance of each model—Classical GAN, Hybrid Quantum-Classical GAN, and Fully Quantum GAN—based on key evaluation metrics: Root Mean Square Error (RMSE), Mean Absolute Error (MAE), and the coefficient of determination (R^2 score). This comparative analysis not only highlights the advantages and limitations of quantum computing in financial forecasting but also provides a foundation for future research avenues in this emerging field. By contributing to the ongoing discourse on the feasibility of quantum-enhanced predictive models, this study paves the way for significant advancements in stock market prediction.

One Step Forecast Results

We present results from experiments conducted with classical, hybrid, and fully quantum Generative Adversarial Networks (GANs), using a past window size of 3 days to predict the stock index for the next day. This window size strikes a balance between capturing short-term market trends and limiting computational complexity.

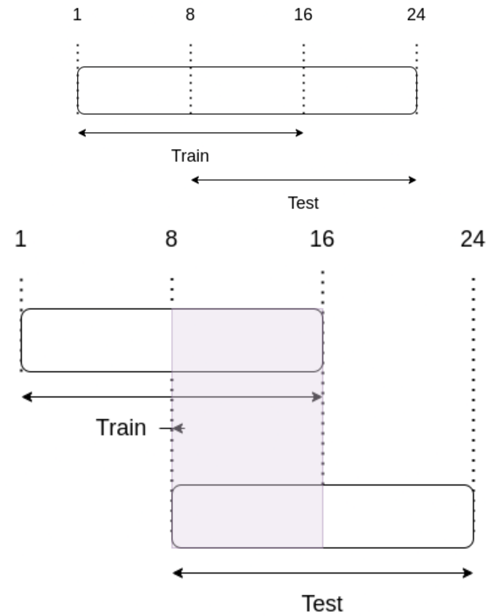

Figure 3: Data preparation for Invertible FQGAN. We use overlapped train and test data to train the model(left). Once trained, the Quantum Generator can generate 16 values in the future, which have an overlap with the input data and therefore the normalization factor can be obtained (right).

Results with Classical GAN We calculated the most popular technical indicators for classical GAN, including 7-day and 21-day moving averages, exponential moving average, and momentum. Along with this, we created Fourier transforms to extract long-term and short-term trends in the stock prices.

The Classical GAN model was used as a baseline for comparing the performance of quantum-enhanced models. The study involved two variants: one that incorporated technical indicators (referred to as Classical GAN with TI), and another that used only adjusted closing prices (referred to as Simple GAN).

Performance Metrics: The Classical GAN with TI model underwent training for 150 epochs with a learning rate of 0.00016 and a batch size of 128, utilizing the Adam optimizer. The performance was primarily evaluated using RMSE, providing a quantitative measure of the deviation between predicted and actual stock prices. Training and test performance results are presented in Fig. 4.

It should be noted that these classical results are taken as standards, and we compare our other method results with these.

Results with Hybrid Quantum GAN The Hybrid Quantum-Classical GAN marks an important step in integrating quantum computation with classical models. In this approach, the generator employs a quantum circuit while the discriminator is based on classical architecture. Despite similar hyperparameters to the classical model, the Hybrid Quantum GAN demonstrated distinct performance characteristics.

Performance Metrics: Trained over 150 epochs with the same learning rate and batch size as the Classical GAN, the Hybrid Quantum GAN achieved a training RMSE of 49.88 and a higher test RMSE of 84.27, indicating that the quantum component introduced significant variability in model performance. These results highlight the challenges associated with combining quantum and classical systems in stock price prediction. Fig. 5

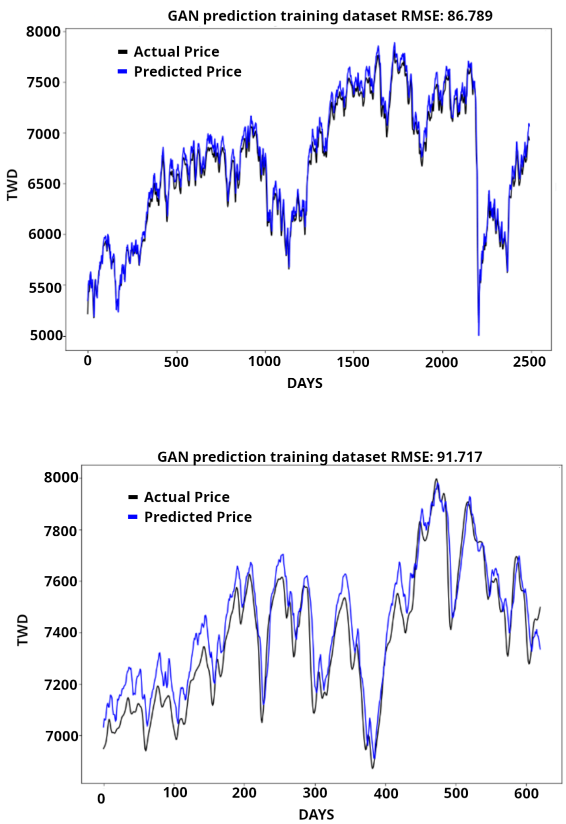

Figure 4: Classical GAN with Technical Indicators. Train (Top) and Test (Bottom).

Results with Full Quantum GAN The Fully Quantum GAN (FQGAN) leverages both the generator and discriminator as quantum circuits. This model seeks to fully exploit quantum computing's advantages, such as parallelism and entanglement, to capture complex stock price patterns.

Performance Metrics: Trained for just 5 epochs on the AWS SV1 quantum simulator with a learning rate of 0.016, the FQGAN model showed considerable promise despite its short training duration. The model achieved a training RMSE of 571.36 and a test RMSE of 251.89, suggesting a competitive performance, particularly when considering the limited quantum resources and training epochs. Fig. 6.

Results Across Varying Window Sizes We also explored the effect of varying past window sizes on the performance

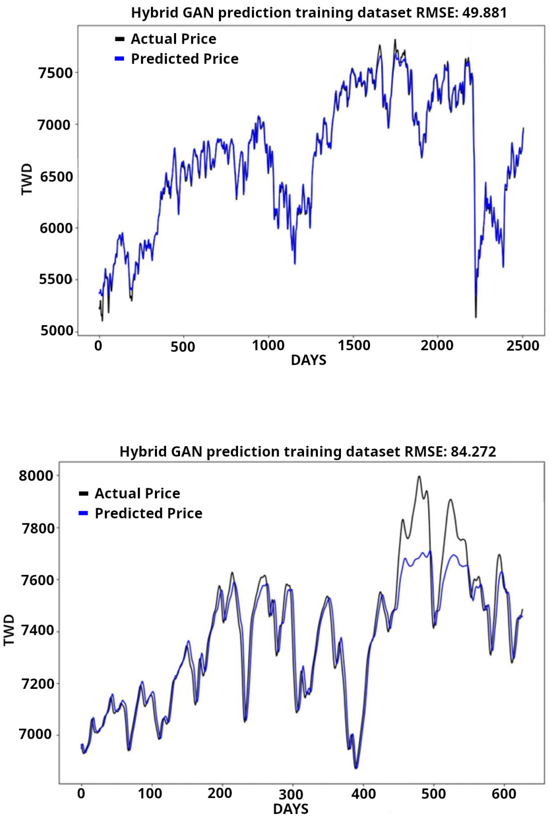

Figure 5: Hybrid Quantum GAN. Train (Top) and Test (Bottom).

of the four models. The results indicate that the Hybrid Quantum GAN exhibits a decline in performance as the window size increases, while the FQGAN's RMSE improves up to a window size of 5, before deteriorating at a window size of 10. Fig. 8

Results on AWS using SV1 Simulator While running the experiments on the SV1 quantum simulator, we faced challenges due to the unavailability of real hardware devices and extended queue times. Each epoch took two hours, and with a minimum of 5 epochs required to obtain significant results, the total simulation time for each model stretched beyond 10 hours. Due to these limitations, we were unable to perform tests on real hardware.

The Classical GAN achieved a testing RMSE of 91.71, which serves as the standard for comparing the hybrid and fully quantum models. Notably, the Fully Quantum GAN performed poorly in comparison, which was expected due to its nascent stage of development. The hyperparameters, particularly the learning rate, played a significant role in model performance, with lower learning rates improving test accuracy. Fig. 7

Quantum Circuit Resource Usage From Eq. 1, we can calculate the number of qubits required by the FQGAN cir-

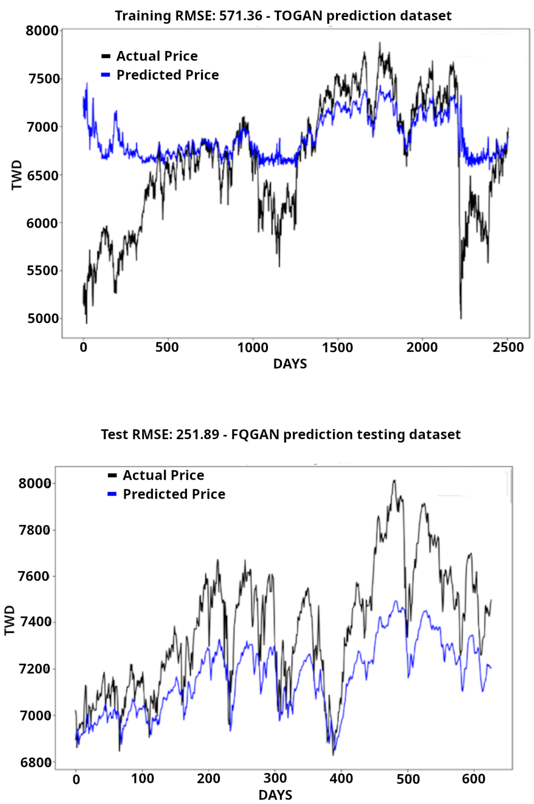

Figure 6: Full Quantum GAN. Train (Top) and test (Bottom).

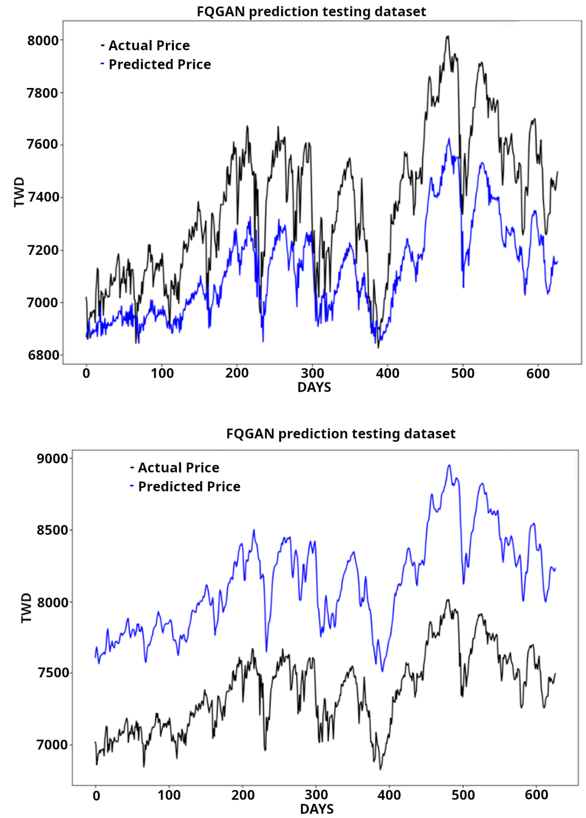

Figure 7: From AWS SV1: 2 epochs (Top) and 3 epochs (Bottom).

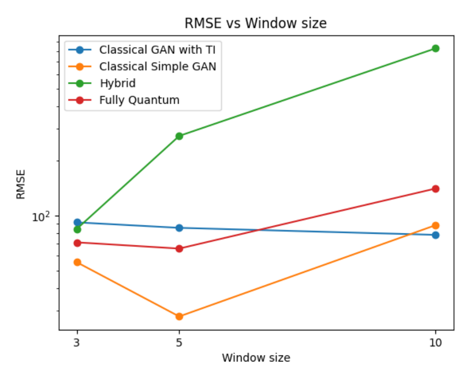

Figure 8: RMSE for the models across varying past window sizes

cuit. The depth of the circuit scales exponentially because we use Amplitude Embedding for encoding the data. One way to reduce the depth of the circuit is to use efficient data embedding schemes. (Zhang, Huang, and Li 2022b) proposes a parallel amplitude embedding method which can reduce the depth of the amplitude embedding circuit by 25% for more than 10 qubits (pennylane 2022). Fig. 9

Advantages Over Other Approaches

Traditional methods for stock price forecasting, such as AutoRegressive Integrated Moving Average (ARIMA), Long Short-Term Memory (LSTM) networks, and multivariate regression, have been widely used to analyze and predict stock price movements. ARIMA models rely on historical data trends, while LSTM networks, a class of neural networks designed for time-series data, excel at capturing complex patterns and long-term dependencies in stock prices. Multivariate regression, on the other hand, seeks to model relationships between stock prices and various external variables. (Fan et al. 2021)

Despite their effectiveness, these classical approaches face significant challenges when dealing with large, dynamic, and high-dimensional financial datasets. The computational limitations of classical algorithms often result in slow processing times, which in turn hinder their ability to facilitate real-time decision-making. This limitation becomes particularly pronounced in the context of financial markets, where the need for quick and accurate predictions is paramount.

Quantum computing presents a promising alternative by exploiting the principles of parallelism, superposition, and entanglement to solve complex optimization problems more efficiently than classical methods. These quantum advantages hold the potential to significantly improve the speed and accuracy of stock price forecasting, particularly in sce-

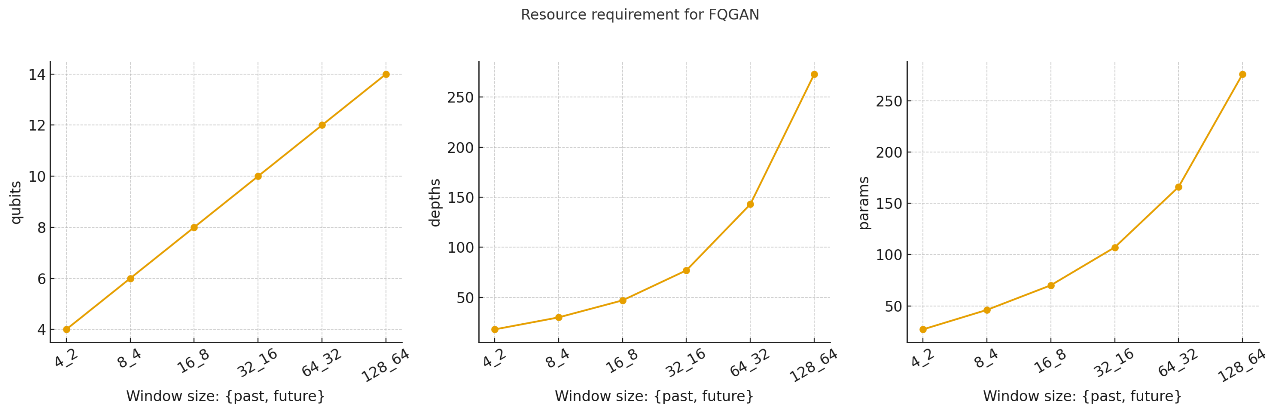

Figure 9: Number of qubits (left), circuit depth (middle) and number of trainable parameters (right) for fully quantum GAN. The x-axis has the different window sizes.

narios where traditional algorithms struggle. (Srivastava et al. 2023)

To address these challenges, we propose the use of Quantum Generative Adversarial Networks (QGANs) as an innovative solution. QGANs overcome some of the inherent limitations of classical Generative Adversarial Networks (GANs), particularly the issue of saddle point problems, which can slow down convergence and limit model robustness. By leveraging quantum properties, QGANs enable faster convergence, enhancing their ability to learn from complex patterns in stock price data. Additionally, QGANs are well-suited to handle large and diverse financial datasets, making them a versatile tool for financial modeling. One of the key advantages of QGANs is their ability to operate effectively with limited training data, facilitating efficient learning from smaller datasets, a critical factor in the financial domain where data availability may be constrained.

Business Use Case: Forecasting the NIFTY Index Prices

The potential of quantum computing within the financial services industry is substantial. Quantum algorithms hold the potential to enhance predictive models utilized for stock price forecasting. Through the processing of extensive volumes of market data, encompassing historical prices, trading volumes, and economic indicators, Quantum Machine Learning (QML) models could discern patterns that might elude classical machine learning algorithms.

Our solution of employing QGANs for stock price prediction marks a significant advancement in the field. By leveraging quantum computing, we have achieved a quantum advantage over classical methods, notably by accurately predicting prices with a smaller dataset. Our idea of utilising a quantum generator and quantum discriminator has played a pivotal role in improving the effectiveness, resulting in superior outcomes. This breakthrough holds immense promise for fund managers, enabling them to make more informed decisions by providing highly accurate price predictions. With the ability to forecast stock prices more precisely, fund managers can optimize their investment strategies, mitigate risks, and maximize returns.

Our focus has been on leveraging quantum computing to refine prediction accuracies, empowering fund managers to make informed decisions and optimize profits across different time horizons. By integrating quantum-driven price predictions for both indices and individual stocks, we enable investors to construct portfolios that not only mitigate future risks but also yield better returns. Additionally, our assistance in individual stock analyses has bolstered fund management effectiveness, allowing the bank to provide superior investment strategies to its clients.

Our approach not only enhances decision-making but also boosts portfolio performance and risk management by accurately predicting stock prices and facilitating insightful portfolio analyses. With strategic recommendations supported by state-of-the-art quantum computing technology, we empower the securities department to take a forward-looking approach in buying and selling decisions, ensuring optimal returns for their clients.

Improvement

Quantum computing presents a substantial and transformative opportunity for the financial services industry, particularly in predictive modeling. Quantum algorithms can significantly improve the accuracy and efficiency of stock price forecasting by processing vast market data, including historical prices, trading volumes, and economic indicators. Quantum Machine Learning (QML) models excel at identifying complex patterns that classical models often miss.

We propose using Quantum Generative Adversarial Networks (QGANs) for stock price prediction, marking a key advancement in the field. By integrating quantum computing, our approach achieves a clear advantage over traditional methods, especially in generating accurate forecasts from smaller datasets.

Our model leverages both a quantum generator and discriminator to enhance predictive performance, outperforming classical counterparts. This innovation has major implications for fund managers, enabling highly accurate forecasts that improve investment strategies, reduce risk, and maximize returns.

Our work emphasizes the use of quantum computing to refine prediction accuracy, empowering fund managers to make data-driven, informed decisions and optimize portfolio performance across various time horizons. By integrating quantum-powered price predictions for both market indices and individual stocks, we enable investors to construct portfolios that not only minimize potential risks but also generate superior returns. Additionally, our contributions to individual stock analysis have significantly enhanced fund management capabilities, allowing financial institutions to provide clients with more effective and tailored investment strategies.

Beyond improving decision-making processes, our approach also strengthens portfolio performance and risk management strategies. By providing accurate stock price forecasts and enabling insightful portfolio analysis, we facilitate a more strategic investment approach. Leveraging state-of-the-art quantum computing technologies, our methodology empowers securities departments to adopt a forward-looking approach to trading decisions, ensuring optimal returns for clients and enhancing overall financial outcomes. (Zhang, Huang, and Li 2022a)

Conclusion

This research explores the use of Quantum Generative Adversarial Networks (QGANs) for forecasting stock index prices, contributing to the growing field of Quantum Finance. By comparing classical GANs, hybrid quantum-classical GANs, and fully quantum GANs, we demonstrate the superior performance of quantum models in handling complex, high-dimensional financial data.

The advantage of quantum models stems from their unique properties—superposition and entanglement—which enhance data representation and learning capabilities. A key innovation of this study is the introduction of the Invertible Fully Quantum GAN (FQGAN), which addresses normalization issues in quantum systems, improving model stability and prediction accuracy.

Our findings suggest that QGANs can outperform traditional forecasting techniques, paving the way for more accurate financial predictions. The research highlights the practical viability of quantum computing for tasks such as predictive analytics, portfolio optimization, and investment strategy development.

As quantum technologies advance, their integration into financial modeling promises to transform the industry, marking a significant step toward more powerful and intelligent decision-making tools in finance. This work lays a strong foundation for future studies and applications in Quantum Finance. (Zhang, Huang, and Li 2022b)

Acknowledgments

The authors would like to recognize AWS Braket, as all simulations were performed on SV1 simulators. All the authors would also like to thank Prof. Greg Byrd and Prof. Kazuki Ikeda for the meaningful discussions and comments.

References

Amzhao. 2022. Quantum Neural Networks. https://medium.com/mit-6-s089-intro-to-quantum-computing/quantum-neural-networks-7b5bc469d984. Published in MIT 6.s089 Intro to Quantum Computing, Accessed January 15, 2024.

Fan, D.; Sun, H.; Yao, J.; Zhang, K.; Yan, X.; and Sun, Z. 2021. Well production forecasting based on ARIMA-LSTM model considering manual operations. *Energy*, 220: 119708.

Gonzalez, S.; and Miikkulainen, R. 2020. Improved Training Speed, Accuracy, and Data Utilization Through Loss Function Optimization. In *2020 IEEE Congress on Evolutionary Computation (CEC)*.

Herman, P. 2023. The Importance of Price Prediction — Blog Future Processing. Technology & Software Development Blog — Future Processing. Last modified March 7, 2023. urlhttps://www.future-processing.com/blog/the-importance-of-price-prediction/.

Lin, H.; Chen, C.; Huang, G.; and Jafari, A. 2021. Stock price prediction using Generative Adversarial Networks. *Journal of Computer Science*, 17(3): 188–196.

Nakamura, Y. 2023. Efficient Preparation of Arbitrary Probability Distribution with Quantum. urlhttps://medium.com/@yuma.nakamura.private/theory-section-efficient-preparation-of-arbitrary-probability-distribution-with-quantum-963367ba876f. Last modified September 22, 2023.

Orús, R.; Mugel, S.; and Lizaso, E. 2019. Quantum computing for finance: overview and prospects. *Reviews in Physics*, 4: 100028.

PennyLane. 2019. Quantum Generative Adversarial Networks with Cirq + TensorFlow — PennyLane Demos. urlhttps://pennylane.ai/qml/demos/tutorial-QGAN/. Last modified October 11, 2019.

pennylane. 2022. Optimizing the depth of the Mottonen state preparation circuit in PennyLane. urlhttps://github.com/king-p3nguin/qhack2024-amplitude-embedding?tab=readme-ov-file.

Srivastava, N.; Belekar, G.; Shahakar, N.; and Babu H., A. 2023. The Potential of Quantum Techniques for Stock Price Prediction. In *2023 IEEE International Conference on Recent Advances in Systems Science and Engineering (RASSE)*, 1–7. Kerala, India.

Zhang, S.; Huang, K.; and Li, L. 2022a. Automatic Depth-Optimized Quantum Circuit Synthesis for Diagonal Unitary Matrices with Asymptotically Optimal Gate Count. *arXiv*. ArXiv preprint arXiv:2212.01002.

Zhang, S.; Huang, K.; and Li, L. 2022b. Automatic Depth-Optimized Quantum Circuit Synthesis for Diagonal Unitary Matrices with Asymptotically Optimal Gate Count. *arXiv*. ArXiv preprint arXiv:2212.01002.