

Quantum Generative Diffusion for Real-World Time Series

Jack Waller, Filippo Caruso, Dimitrios Makris, Rajagopal Nilavalan,
and Xing Liang

Kingston
University
London



Motivation



- **Financial data is scarce, and non-stationary** - unlike images or text, historical financial records are limited in volume, and market dynamics shift over time such that older data can become statistically irrelevant to current conditions.
- **Data-hungry models meet data-scarce markets** - financial forecasting models perform better with more training data, yet real market data is both scarce and time-sensitive, creating a persistent shortfall.
- **High fidelity synthetic data offers a practical solution** - data generated to reproduce the current statistical properties of real markets can directly augment training sets, improving model performance in a sector where even marginal gains carry substantial competitive advantage.

Prior work and research gap

- Classical diffusion models have shown strong results for classical time series generation [4] but rely on classical neural architectures to learn complex data distributions.
- Quantum generative models offer an alternative approach to learning complex probability distributions using parameterised quantum circuits, with early results showing promise [5]. However, their application to time series generation remains largely unexplored.
- Here we introduce the first quantum generative diffusion model for time series and demonstrate its implementation on quantum hardware using financial data.



Diffusion Models



Diffusion models learn to generate data by reversing a noise process [6].

Forward Process: Gradually corrupt real data with Gaussian noise until only noise remains.


$$x_0 \rightarrow x_1 \rightarrow \cdots \rightarrow x_T \sim \mathcal{N}(0, I)$$



$$x_0 \leftarrow x_1 \leftarrow \cdots \leftarrow x_T$$


Reverse Process: A trainable model learns how to reverse this process, transforming noise back into a realistic sample.

Quantum Neural Networks




Quantum neural networks use a parameterised quantum circuit as a trainable model — gate parameters θ are optimised analogous to weights in a classical neural network [1].


Model Architecture Overview



Diffusion Model Architecture



Diffusion Model Architecture - Decoder



Parameter Reduction

- Replacing each feed forward neural network with quantum neural networks reduces per-component parameters by three orders of magnitude (33,088 vs 36)
- Across the full model this corresponds to a parameter reduction of approximately 20% (589,030 vs 461,254)














Dataset and experimental setup

- Apple and Amazon daily stock data are used, each comprising Open, Close, High, Low prices and Volume across the same time period.
 - Two evaluation tasks are defined; the distributional fidelity of generated synthetic data against real log-return distributions, and the performance on a downstream financial forecasting task.
 - The forecasting task uses a BiLSTM model, benchmarked on real data alone, then with classical and quantum synthetic data augmentation, to isolate the contribution of each generative model.
- 


Distribution Comparisons



(a)




(b)






Forecasting
Task Results

Quantum hardware experimental setup

- **Hybrid execution:** Inference is performed on the **IQM Emerald** processor via Amazon Braket, with the model trained on classical simulator.
- **Hardware efficient ansatz:** Quantum circuit altered to avoid additional SWAP gates, adapting the native connectivity of the qubits.



Quantum Hardware Validation



Conclusions



- The quantum diffusion model achieves superior distributional fidelity over its classical counterpart, with equivalent gains carried through to the downstream forecasting task.
- Quantum components offer a dramatic reduction in model parameters demonstrating a clear path toward more efficient large-scale AI models as quantum hardware matures.
- Performance is maintained despite the noise constraints of current quantum devices.



Future work and open questions

- **Incorporating quantum noise into the diffusion process** — replacing the classical noise schedule with quantum-native noise sources, more deeply integrating quantum mechanics into the generative process beyond the neural network components.
- **Extending the framework to broader time series domains** — the architecture is not inherently finance-specific, and applying it to domains such as climate, energy, or healthcare data would demonstrate the generalisability of the approach.
- **Scaling with advances in quantum hardware** — current results are constrained by the noise and qubit counts of today's devices; as hardware fidelity improves, deeper quantum circuits and larger models become accessible.



Acknowledgements


This project was funded and supported by the UK National Quantum Computing Centre (NQCC) [NQCC200921], as part of the UK National Quantum Technologies Programme (NQTP).



National Quantum
Computing Centre



UK NATIONAL
QUANTUM
TECHNOLOGIES
PROGRAMME



Quantum Generative Diffusion Model for Real-World Time Series

Paper: arxiv.org/abs/2606.27561

Thanks for listening!

References

- [1] Benedetti, M., Lloyd, E., Sack, S. & Fiorentini, M. Parameterized quantum circuits as machine learning models. Quantum Science and Technology 4 (2019).
- [2] . Li, Z. Wang, R. Han, S. Shi, J. Li, R. Shang, H. Zheng, G. Zhong, Y. Gu, Quantum recurrent neural networks for sequential learning, Neural Networks 166 (2023) 148–161.
- [3] Abbas, A. et al. The power of quantum neural networks. Nature Computational Science 1, 403-409 (2021).
- [4] Wiese, M., Knobloch, R., Korn, R. & Kretschmer, P. Quant GANs: deep599 generation of financial time series. Quantitative Finance 20, 1419–1440 (2020).
- [5] Niu, M. Y. et al. Entangling quantum generative adversarial networks. Phys.629 Rev. Lett. 128, 220505 (2022).
- [6] Ho J, Jain A N and Abbeel P 2020 Denoising diffusion probabilistic models Advances in Neural Information Processing Systems 33 (NeurIPS 2020).