

Learning Temporal Patterns in Financial Time Series: A Comparative Study of Different Quantum Algorithms

Gerhard Hellstern

DHBW Stuttgart, Stuttgart, Germany

gerhard.hellstern@dhw-stuttgart.de

Extended Abstract— Within the research project QROM, funded by the German Federal Ministry of Research, Technology and Space, we investigate different approaches to time series prediction for real financial data and compare them with established classical methods. To encode classical data into quantum circuits, we implement not only standard techniques such as amplitude encoding, but also novel encoding schemes developed within the project consortium [1]. These are designed to efficiently handle both small and larger datasets.

As a classical benchmark, we first implemented linear regression, random forests, boosted trees and a neural network using lagged variables as features. Their out-of-sample performance is then compared to several quantum models: a Quantum Neural Network (QNN), Quantum Long Short-Term Memory (QLSTM) and Quantum Reservoir Computing (QRC) architectures [2]. In these models, the quantum components are realized as parameterized quantum circuits; classical optimizers are used to train the readout and, where applicable, the variational circuit parameters. Initial results indicate that, with appropriate lag selection and suitable data encoding, quantum-enhanced architectures match classical baselines in univariate settings and can modestly outperform them in multivariate scenarios with correlated inputs, where more expressive encodings are especially beneficial [2]. However, a systematic and balanced hyperparameter optimization for both classical and quantum models is essential to ensure a fair comparison and to draw robust conclusions about performance differences.

Building on this, we conduct a study comparing a conditional classical Gaussian–Bernoulli restricted Boltzmann machine (CRBM) [3] with three quantum or hybrid variants: a hybrid quantum–classical QCRBM, a fully quantum QQRBM, and a lag-feature-based QFeatureQRBM. All models support conditional sampling, where the conditioning variables are given by the lagged inputs. In contrast to many existing comparisons, our evaluation enforces symmetric hyperparameter optimization: both classical and quantum-specific hyperparameters are subjected to an equally thorough grid search across thirteen structured experiments. We test on two regimes: (i) financial time series from the project use case and (ii) an input-driven NARMA-10 nonlinear benchmark. Each setting comprises three time series and four random seeds ($n = 12$ paired observations), and we apply Holm–Bonferroni correction to control for multiple comparisons.

At the current sample size ($n=12$), the study is not powerful enough to detect small effects, so a definitive picture on quantum advantage (e.g. lower mean squared error) remains premature. The hybrid QCRBM already matches the classical CRBM across all budget settings, and the fully quantum architectures show clear room for improvement with more expressive ansätze. Larger-scale experiments are therefore the natural next step. [4].

Keywords— *Machine Learning, Quantum Computing, Time Series, Data Encoding, Boltzmann machines*

References

- [1] Data Cybernetics, “q-alchemy-sdk-py: Python SDK for the Q-Alchemy API,” GitHub, 2024.
- [2] Maheshwari, D.; Hellstern, G.; Zaefferer, M.; Braun, M., and Döhler, T.; Learning Temporal Patterns in Financial Time Series: A Comparative Study of Quantum LSTM and Quantum Reservoir Computing, <https://doi.org/10.48550/arXiv.2605.02656>
- [3] Taylor, G.W. and Hinton, G.E. Factored Conditional Restricted Boltzmann Machines for Modeling Motion Style. In Proceedings of the 26th International Conference on Machine Learning (ICML), 2009.
- [4] Hellstern, G.; Maheshwari, D.; Zaefferer, M.; Braun, M., and Döhler, T.; Variational Quantum Conditional Boltzmann Machines for Time-Series Forecasting: Architectures, Symmetric Hyperparameter Evaluation, and a Nonlinear Benchmark, in preparation.