

Quantum Machine Learning Architectures for Semi-Supervised Time-Series Anomaly Detection

J. Bonilla^{1,*}, I. Montalban², E. Combarro², J. Ranilla²

¹*Neural Domains, ARQUIMEA Research Center, Spain*

²*Computer Science Department, University of Oviedo, Spain*

*E-mail: jbonilla@arquimea.com

Semi-supervised anomaly detection in multivariate time series requires models that learn normal system dynamics without labelled anomalies while remaining compact enough for resource-constrained applications. We investigate two lightweight quantum machine learning architectures based on complementary paradigms: Quantum Reservoir Computing (QRC) and a quantum-enhanced Transformer, denoted as iQTransformer. Both models are trained exclusively on normal operating data and detect anomalies through forecasting residuals or representation-aware scores.

The QRC architecture employs a fixed quantum reservoir governed by a transverse-field Ising Hamiltonian,

$$H = - \sum_{i < j} J_{ij} Z_i Z_j - \sum_i h_i X_i,$$

which maps sequential multivariate observations into nonlinear, high-dimensional dynamical representations. Since the reservoir parameters remain fixed, the quantum subsystem requires no gradient-based optimization. Different encoding strategies are considered to incorporate temporal and cross-variable information. The resulting states are processed through two alternative lightweight pipelines: a linear readout that predicts future observations and derives anomaly scores from forecasting residuals, and a One-Class Support Vector Machine (OCSVM) that detects deviations directly in the QRC feature space.

The iQTransformer extends the iTransformer architecture [2], in which variables are represented as tokens and self-attention is applied across the variate dimension. Following the quantum forecasting architecture introduced in [1], the linear projections used to generate query, key, and value representations are replaced by independent variational quantum circuits. Classical embeddings are encoded into quantum states, processed through parameterized entangling layers, and measured to obtain the attention representations. This preserves the overall Transformer structure while introducing nonlinear quantum mappings and reducing the number of trainable parameters.

Both approaches are evaluated on the public GECCO dataset using point-wise and event-based metrics. Their normalized anomaly scores are classified using an adaptive Dynamic Streaming Peaks-Over-Threshold (DSPOT) method fitted on normal training data. Among the QRC variants, the OCSVM model clearly outperforms the linear readout, increasing AUC-ROC from 0.7330 to 0.9048 and recall from 0.2255 to 0.5971. The iQTransformer improves AUC-ROC from 0.8831 to 0.9512 and point-wise F1 from 0.4400 to 0.5454 relative to the classical iTransformer, while using approximately 50% fewer trainable parameters.

These results highlight two complementary approaches to efficient quantum-enhanced anomaly detection. QRC generates expressive temporal representations without optimizing the quantum substrate, whereas the iQTransformer integrates trainable quantum mappings into an attention-based architecture, improving detection performance with reduced model complexity.

[1] S. Ranilla-Cortina, D. A. Aranda, J. Ballesteros, J. Bonilla, N. Monrio and E. F. Combarro, *Quantum-enhanced architectures for multivariate time-series forecasting*, The Journal of Supercomputing **82**, 155 (2026).

[2] Y. Liu, T. Hu, H. Zhang, H. Wu, S. Wang, L. Ma and M. Long, *iTransformer: Inverted Transformers Are Effective for Time Series Forecasting*, Proceedings of the International Conference on Learning Representations (ICLR), 2024.