Quantum Architectures for Multivariate Time-Series Forecasting

D. A. Aranda[1], J. Ballesteros[3], J. Bonilla[1], E. F. Combarro[2], N. Monrio[1], S. Ranilla-Cortina[2], J. Ranilla[2]
[1]IAQ Orbital, ARQUIMEA Research Center (Spain)
[2]Department of Computer Science, University of Oviedo (Spain)
[3]Instituto Tecnológico y de Energías Renovables (Spain)

Introduction and Motivation

- **QML** combines quantum computing with classical learning to enhance pattern extraction and data-driven prediction, showing promise in classification, reinforcement learning, and generative modelling.
- While early efforts have addressed univariate time-series prediction, there is still **no clear state-of-the-art methodology for tackling multivariate time-series forecasting** within the quantum paradigm.
- Classical deep learning models, especially **Transformer architectures**, achieve strong results through attention mechanisms that capture inter-variable dependencies, though at high computational cost.
- This work extends QML to the multivariate regime, exploring several quantum and hybrid architectures and introducing a novel **quantum Transformer model** for multivariate forecasting.
- We compare their performance against leading **classical and hybrid baselines** on both synthetic and real-world datasets.

QML for Multivariate Time-Series Forecasting

Models were implemented using the **PennyLane** quantum ML framework.

1. **Baseline VQC with independent channels:**
   - Each input channel is encoded into its own quantum circuit with as many qubits as time steps, using single-parameter RY rotations.
   - No cross-variable correlations are modeled.
2. **Hybrid VQC + Post-Quantum MLP:**
   - Same independent-channel design, followed by a shallow MLP to capture inter-channel correlations from circuit outputs.
3. **Dense embedding VQC:**
   - All channels jointly encoded in one circuit via RZ–RX–RY rotations (`qml.Rot`), up to three channels per qubit.
   - Two variants:
     - · Multi-observable single-qubit: measure Pauli X, Y, Z on one qubit per channel.
     - · Single-observable multi-qubit: measure Pauli Z on three qubits.
4. **Hybrid Encoder–VQC–Decoder:**
   - A classical encoder maps inputs to a lower dimension, processed by one quantum circuit; a classical decoder restores the output size.
5. **Sequential Data Re-uploading:**
   - Each time step is sequentially encoded in the same circuit, alternating encoding and variational layers.

Quantum Transformer Model: iQTransformer

- Variables are embedded as tokens (iTransformer paradigm); attention operates *across variables*, while feed-forward networks act along time per token.

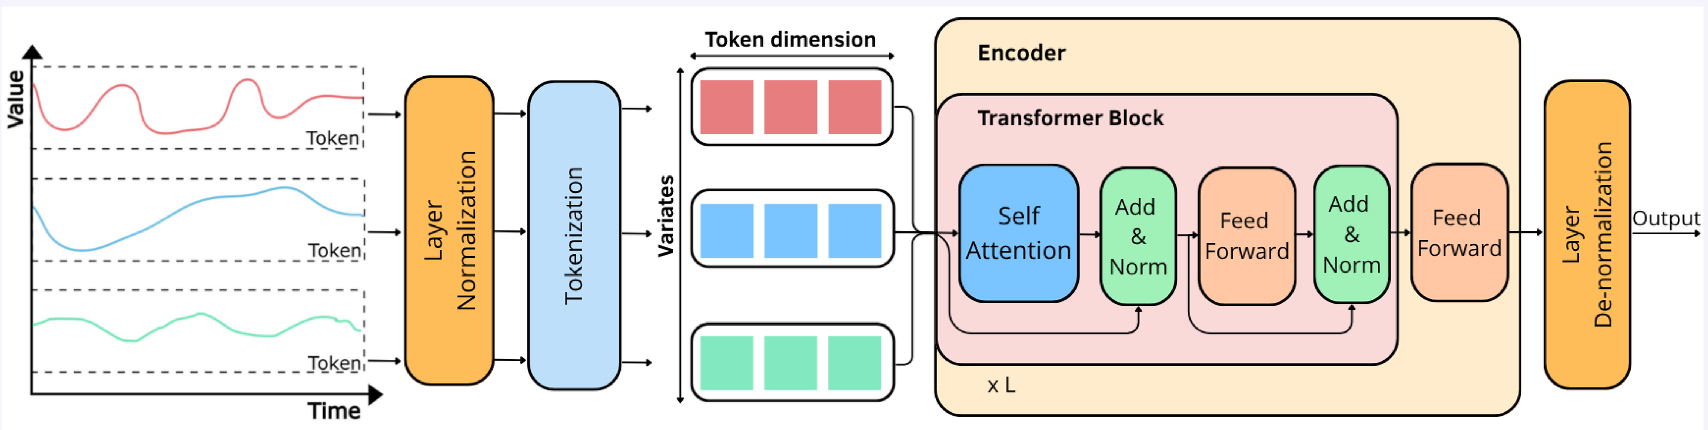

- The classical self-attention is replaced by a **Quantum Self-Attention Layer (QSAL)** based on **VQCs (QSANN)**: queries, keys, and values are produced by separate variational circuits.
- Attention coefficients arise from **Gaussian projections** of query/key measurements; outputs aggregate value measurements across variables via residual connections.
- The QSANN ansatz uses layers of single-qubit rotations with entangling CNOTs, repeated p times; depth and qubit count control expressivity and parameter efficiency.

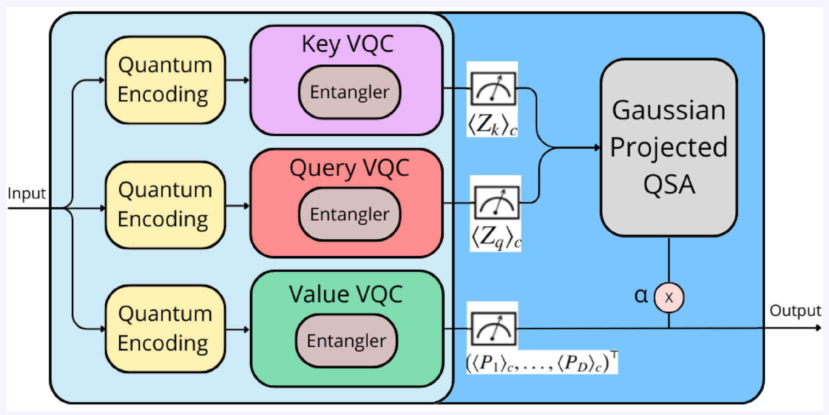

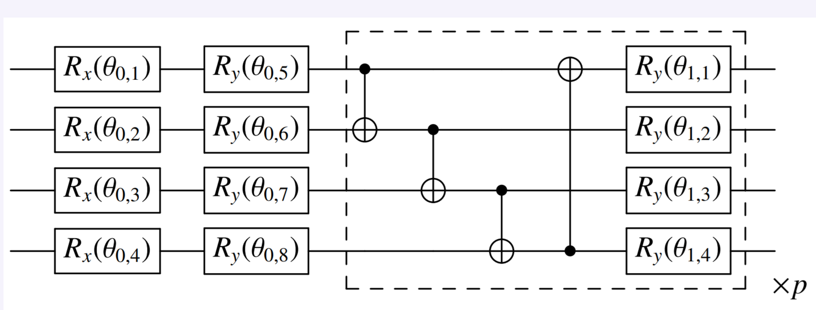

Experimental Setup

- **Synthetic dataset — Lorenz system:** a 3-channel chaotic system defined by the Lorenz equations.
- **Real dataset — ITER:** a 7-channel dataset from a wind turbine in Tenerife, Spain, covering four months of operation in 2024. Variables include total energy demand, renewable generation, normalized power, wind speed/direction, and curtailment setpoint (target variable).
- Two forecasting regimes were considered: **short-term (ST)** and **long-term (LT)**.
  - · ST: next prediction point computed from every 5 ground-truth points.
  - · LT: for Lorenz, the next 5 points are predicted; for ITER, 3.5 days of data are used to forecast the next 6 hours.

Results

- **Validation metrics (MAPE, MAE, RMSE)** are reported for all implemented models, averaged over 10 random initializations. Additionally, state-of-the-art classical (**1D CNN**) and quantum (**QGRU**) baselines are included for comparison.

Model	MAPE		MAE		RMSE	
	Lorenz	ITER	Lorenz	ITER	Lorenz	ITER
Short-Term Forecasting						
VQC (indep.)	0.0353	0.1490	0.0173	0.0905	0.0281	0.1341
VQC + MLP	0.305	0.1037	0.139	0.0522	0.184	0.0963
DE. (obs.)	0.117	n/a	0.0558	n/a	0.0982	n/a
DE. (qubits)	0.116	n/a	0.0563	n/a	0.0985	n/a
Enc.–VQC–Dec.	0.454	0.3800	0.210	0.3435	0.259	0.3574
Data re-upload.	0.0388	0.0518	0.0192	0.0353	0.0308	0.0534
QGRU	0.432	0.1852	0.195	0.1205	0.243	0.1577
1D CNN	0.204	0.0798	0.082	0.0434	0.101	0.0741
iTransformer	**0.00807**	0.0154	**0.00394**	0.0069	**0.00636**	0.0354
iQTransformer	0.00857	**0.0152**	0.00414	**0.0071**	0.00668	**0.0351**
Long-Term Forecasting						
VQC + MLP	0.235	n/a	0.0976	n/a	0.127	n/a
Enc.–VQC–Dec.	0.283	0.0947	0.114	0.0477	0.141	0.1076
1D CNN	0.235	0.1191	0.0954	0.0515	0.122	0.1241
iTransformer	0.0498	0.0874	0.0234	0.0352	0.0371	0.1050
iQTransformer	**0.0490**	**0.0849**	**0.0230**	**0.0340**	**0.0364**	**0.1019**

- **Time-series reconstruction** from the best-performing **iQTransformer** model for the Lorenz (short-term forecasting) and ITER (long-term forecasting) datasets:

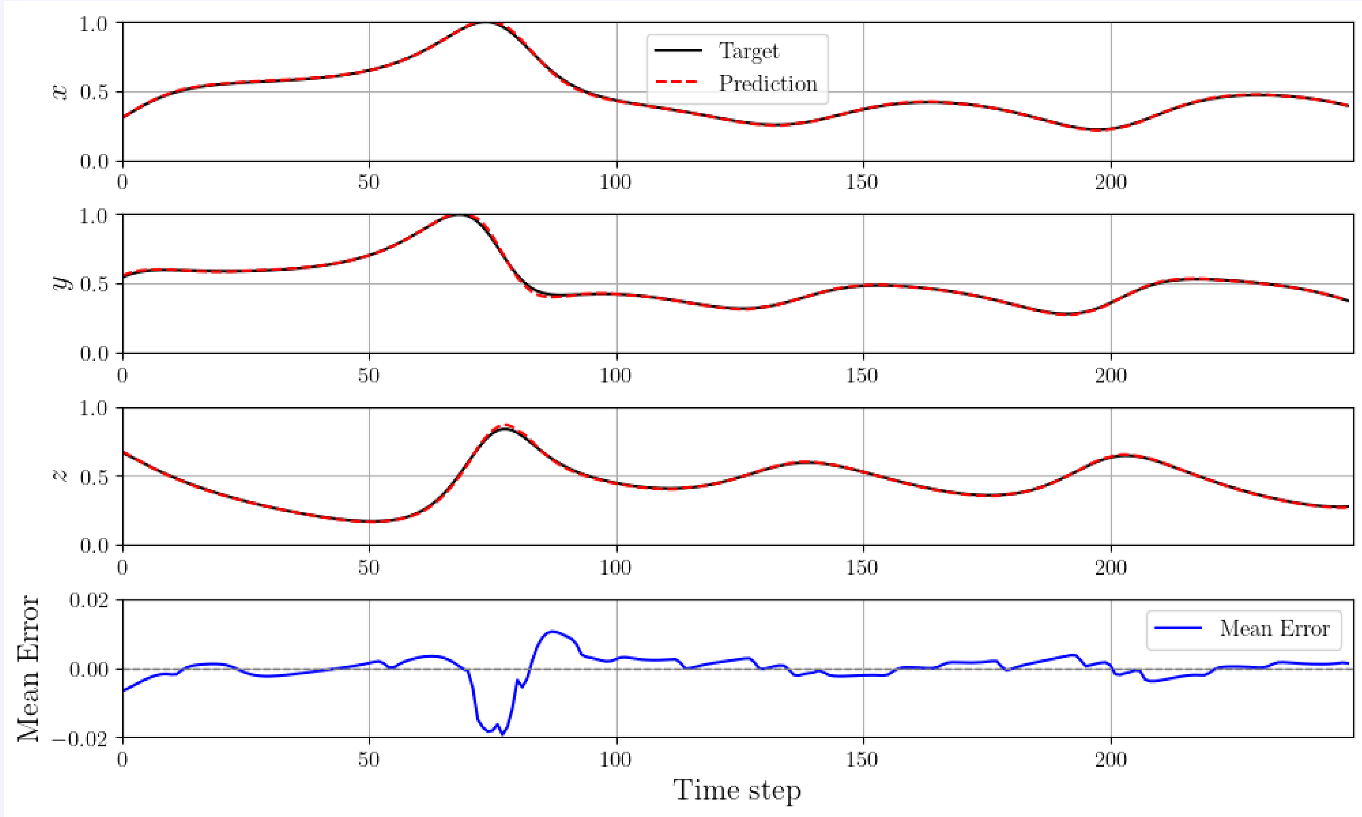

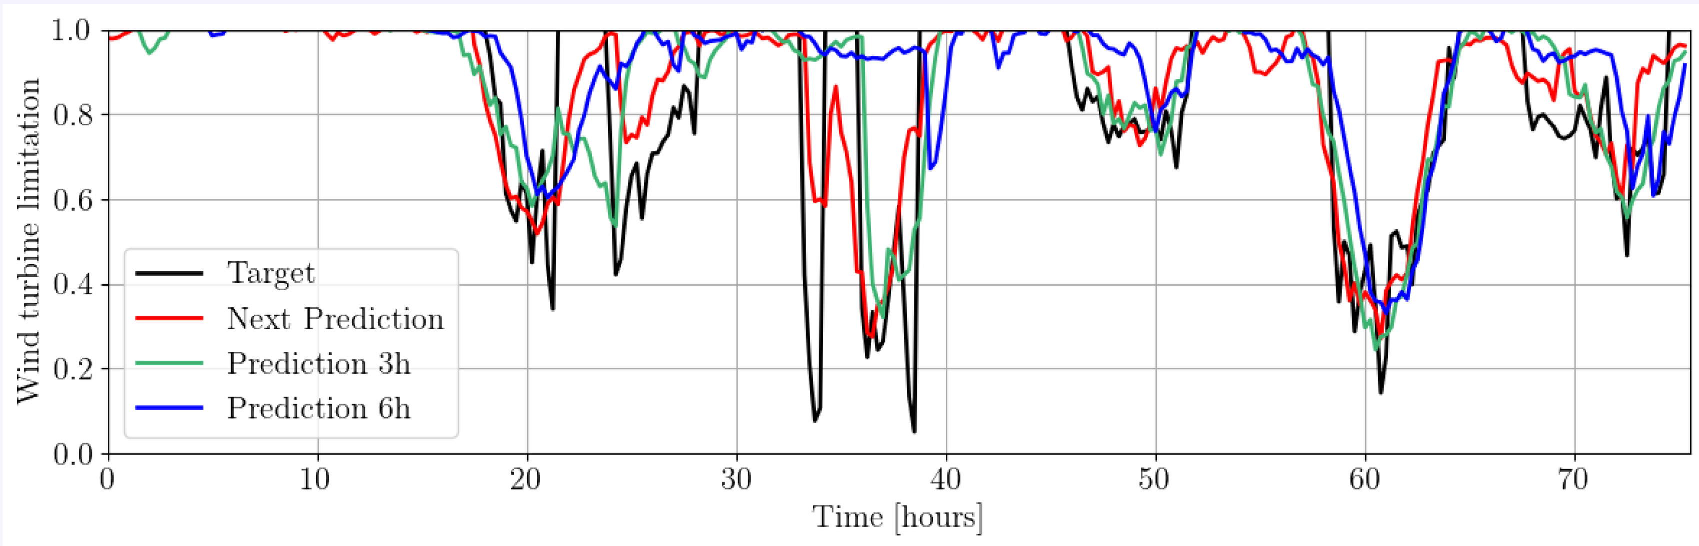

Conclusions

- **QML Adaptations:**
  - · Several **VQC architectures** were extended to multivariate forecasting, enabling modeling of inter-variable dependencies previously unexplored in QML.
- **Quantum Transformer:**
  - · The **iQTransformer** integrates quantum self-attention into the iTransformer backbone, capturing cross-variable correlations with fewer parameters and faster convergence.
- **Benchmarking and Results:**
  - · Benchmarks on **Lorenz** and **ITER** datasets show that hybrid quantum–classical models can **match or surpass** state-of-the-art classical and quantum baselines, achieving an effective balance between accuracy and scalability.

Acknowledgements

This work has been supported by grant PID2023-146520OB-C2, funded by MICIU/AEI/10.13039/501100011033; by the Ministry for Digital Transformation and Civil Service of the Spanish Government through the QUANTUM ENIA project call – Quantum Spain project, and by the European Union through the Recovery, Transformation and Resilience Plan – NextGenerationEU within the framework of the Digital Spain 2026 Agenda; and by ARQUIMEA Research Center and Horizon Europe, Teaming for Excellence, under grant agreement No 101059999, project QCircle. The work of Jesús Bonilla was partially supported by the Spanish Ministry of Science, Innovation and Universities through the Torres Quevedo grant PTQ2023-013228. The work of Jorge Ballesteros is financially supported by Instituto Tecnológico y de Energías Renovables (ITER) and Cabildo de Tenerife.

Universidad de Oviedo
Q·CIRCLE × ARQUIMEA RESEARCH CENTER
CABILDO DE TENERIFE
ITER
Funded by the European Union