Towards Practical Advantage with Quantum Machine Learning

Maida Wang | Supervisor: Peter V. Coveney

2nd Year PhD Review (2024.10-2026.5)

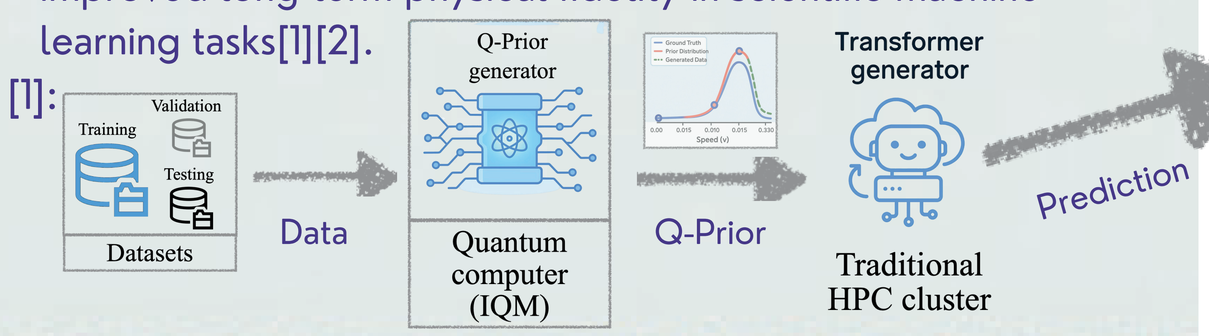
Research Aim

This PhD project aims to identify where quantum machine learning can deliver practical advantage on near-term hardware, moving toward forms of advantage that are scientifically meaningful, testable, and relevant to real computational problems. The work focuses on developing resource-efficient hybrid quantum–classical methods and evaluating whether they offer benefits such as parameter efficiency, compact representation, and memory advantage in challenging learning tasks.

Key Contributions

- Developed QIML (Quantum-Informed Machine Learning), a hybrid quantum–classical framework for long-term prediction of chaotic systems using quantum generator (Q-Prior)[1].
- Developed PEQAD (Parameter-Efficient Quantum Anomaly Detection), a quantum anomaly detection framework[2].
- Established a research direction centred on practical quantum advantage, with emphasis on realistic near-term benefits. Demonstrated that quantum models can be useful through parameter efficiency, compact statistical representation, and improved long-term physical fidelity in scientific machine learning tasks[1][2].

[1]:



Publications

[1] Maida Wang†, Xiao Xue†,..., Peter V. Coveney*. "Quantum-Informed Machine Learning for Predicting Spatiotemporal Chaos with Practical Quantum Advantage," **Science Advances (2026)**.

[2] Maida Wang,..., Peter V. Coveney*. "Parameter-Efficient Quantum Anomaly Detection Method on a Superconducting Quantum Processor," **Physical Review Research (2025)**.

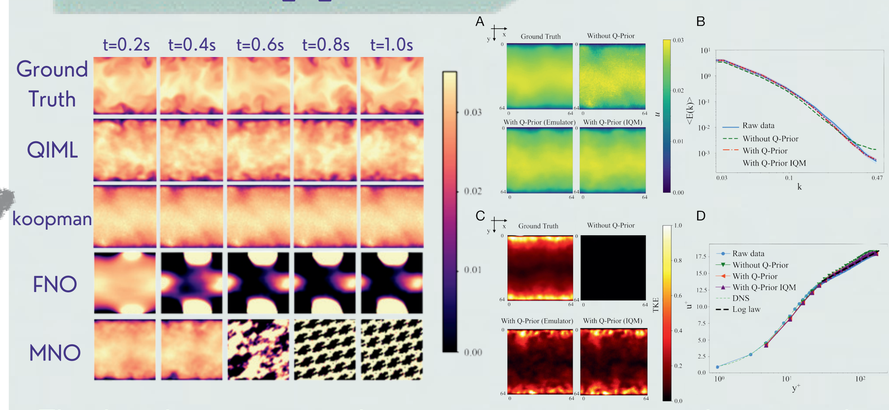
Methods

This research search for practical quantum advantage by using quantum learning models on real scientific tasks (chaos prediction and anomaly detection) where classical methods struggle, and by measuring useful gains in parameter efficiency, memory efficiency, stability, and statistical fidelity. The quantum advantage here aims to overcome key barriers: NISQ noise, limited qubits, costly measurement and data-encoding bottlenecks. For example, the QIML learns a quantum prior (Q-Prior) that guides the representation of small-scale interactions and improves the modelling of fine-scale dynamics[1].



IQM, a European superconducting quantum chip company, has provided us with its most accurate chip: IQM-Garnet, by connecting it to a supercomputing cluster to enable hybrid computing[1].

Results[1]

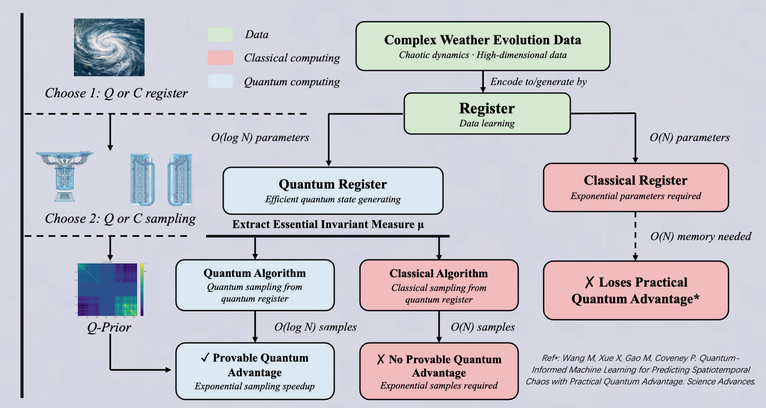


The hardware-trained quantum prior proved essential for stable long-term prediction: without it, forecasts became unstable, whereas QIML produced physically consistent long-term predictions and outperformed strong classical baselines.

System	Qubits	Trainable θ parameters	Raw data (full set)	Q-Prior file (full set)	Compression (ratio)	Device
Kuramoto–Sivashinsky	10	~ 120	300 MB	0.25 MB	$\approx 1.2 \times 10^3 : 1$	Emulator
2D Kolmogorov flow	10	~ 180	400 MB	0.40 MB	$\approx 10^3 : 1$	Emulator
Turbulent channel flow	15	~ 300	500 MB	2.3 MB	$\approx 2.2 \times 10^2 : 1$	Emulator
Turbulent channel flow	10	~ 240	500 MB	2.0 MB	$\approx 2.5 \times 10^2 : 1$	IQM-Garnet

It also demonstrated a memory advantage by compressing multi-megabyte datasets into a kilobyte-scale Q-Prior, highlighting a practical route for integrating quantum resources into scientific modelling.

Future works



Future applications include weather forecasting, multiscale simulation, and drug discovery.

Centre for Computational Science, Department of Chemistry, University College London