Qudit extension of parameterized IQP circuits: A generative quantum machine learning approach to calorimeter data

Robert Banks, Arianna Crippa, Matthias Traube, Josua Unger, Christian Ertler and Wolfgang Lechner

arXiv:2606.28236

Abstract

Simulating electromagnetic particle showers in calorimeters is a key task in HEP. To model these complex, correlated distributions on a quantum circuit, the energy data are reduced to binary or encoded into gate angles. When applied to non-binary distributions, these methods exhibit notable limitations: mapping integer numbers into qubit-compatible binary representations often destroys the original metric structure of the data. Instantaneous Quantum Polynomial (IQP) circuits have been proven useful in quantum generative learning models. Expectation values of IQP circuits can be estimated classically, while sampling remains hard. This means that training can be carried out using classical computing resources, while generation necessitates a quantum device. However, IQP circuits mainly focus on the generation of **binary distributions**. In our work, we extend parameterized IQP circuits to qudits, developing the framework for addressing the non-binary energy deposits from single-particle electron showers. We design a qudit-based quantum circuit and correspondingly adapt the loss function. The method can be extended to other applications that utilize quantum generative machine learning for non-binary data.

The dataset

A centered window of size 25 × 25 × 25 is then extracted around the barycenter [1].

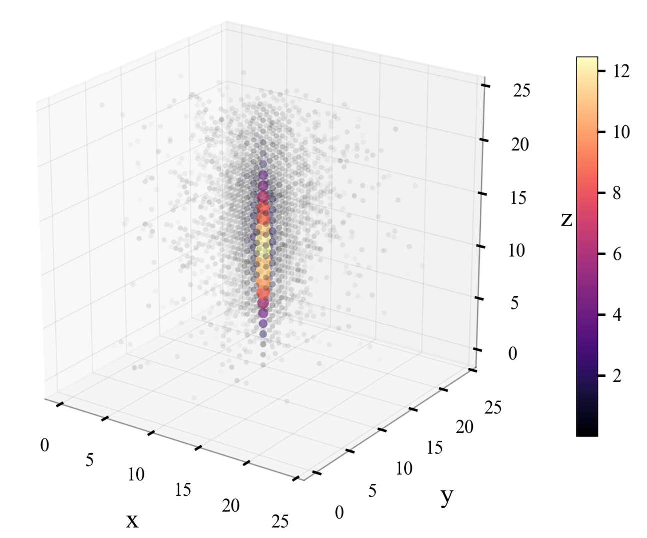

Data between [225, 275] GeV are selected and projected onto the z-axis to form a 1D image.

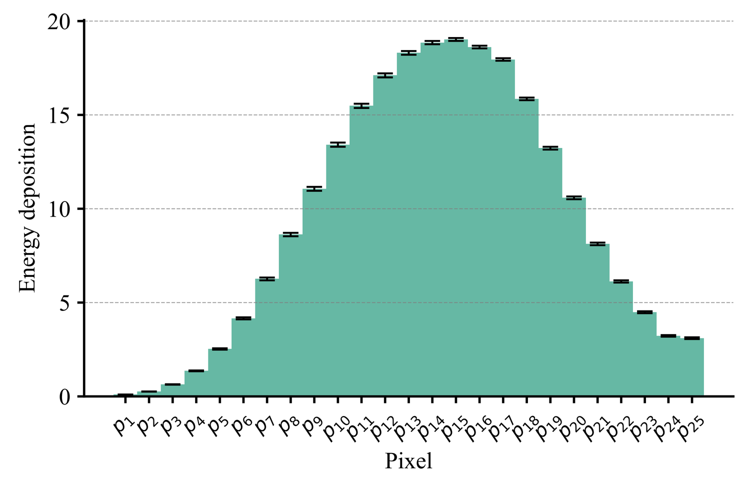

Parameterized IQP circuits

QUBITS:

Each qubit is initialized in the zero state $|0\rangle$. A layer of Hadamards is followed by a sequence of diagonal multi-body Pauli-Z rotations. A final layer of Hadamard gates is then applied.

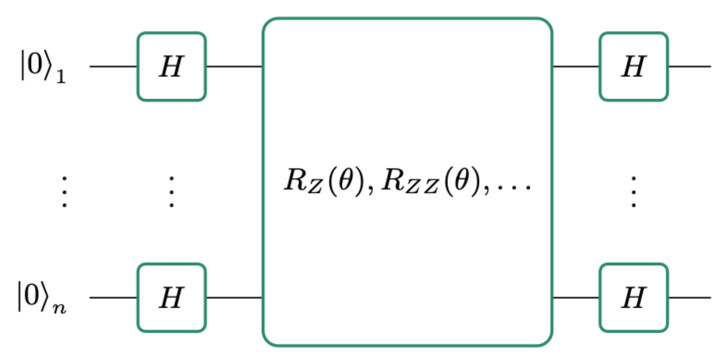

QUDITS:

For each of the n pixels, $\log_2 d$ qubits are initialized in the zero state $|0\rangle$. The dimension of each qudit is therefore d. A layer of QFTs is followed by a sequence of multi-body diagonal rotations. A final layer of inverse QFT gates is then applied.

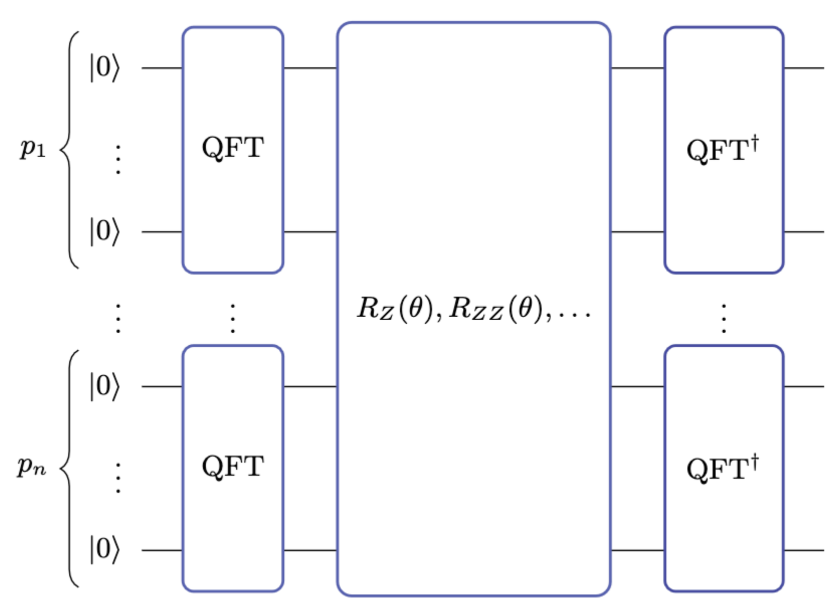

Training

The model is trained classically. The loss-function depends on being able to estimate expectation values from the IQP circuit. The model was developed for binary IQP circuits by Recio-Armengol *et al.* [4].

ESTIMATING EXPECTATION VALUES

- We can classically estimate expectation values of powers of the qudit-Z operator acting simultaneously on different qudits.
- Each expectation value is estimated stochastically by sampling N_s strings from a uniform distribution. Each string is of length n and each element takes integer values between 0 and $d-1$.

EVALUATING THE LOSS FUNCTION

The loss-function compares expectation values from the training data with estimated expectation values from the IQP circuit.

1. Sample N_{ops} operators that consist of tensor products of powers of Z. In this work the operators are drawn from a discretized Gaussian distribution.
2. Stochastically approximate the expectation values of the N_{ops} operators for the IQP circuit.
3. Evaluate the corresponding expectation values on the training data.
4. Compare the discrepancy between the expectation values to give the loss-function.
5. With the samples fixed, the gradient of the loss-function is found with automatic differentiation.

12 pixels, $d = 32$: 60 qubits needed

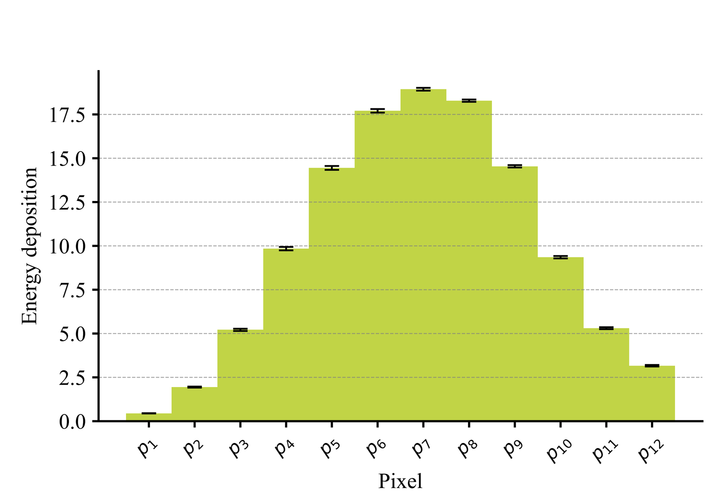

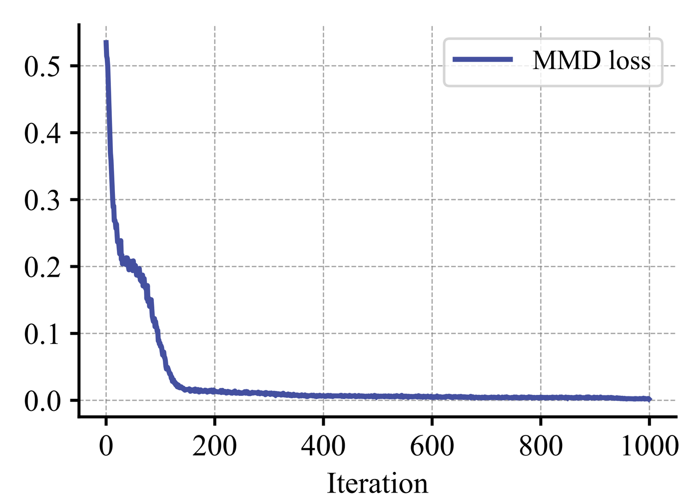

Circuit implementation

The IQP circuit can be efficiently implemented using Parity Twine [2].

For the diagonal gates, we take the ansatz to be all-to-all two-body qubit-Z rotations as well as single-body qubit-Z rotations.

Using Parity Twine the circuit depth of the diagonal layer scales as $\nu n \log_2 d$ (n total number of pixels). Where ν reflects the device connectivity, $\nu = 2$ for all-to-all device connectivity.

The inverse QFT can also be implemented using Parity Twine. The depth scales as $\nu \log_2 d$. Parity Twine QFT on 52 superconducting qubits has already been realized [3].

The initial QFT can be replaced by a layer of Hadamard gates.

A comparison of the QFT process fidelity for different system sizes. The data is obtained from running both methods on the *ibm boston* Heron r3 device. Figure from [3].

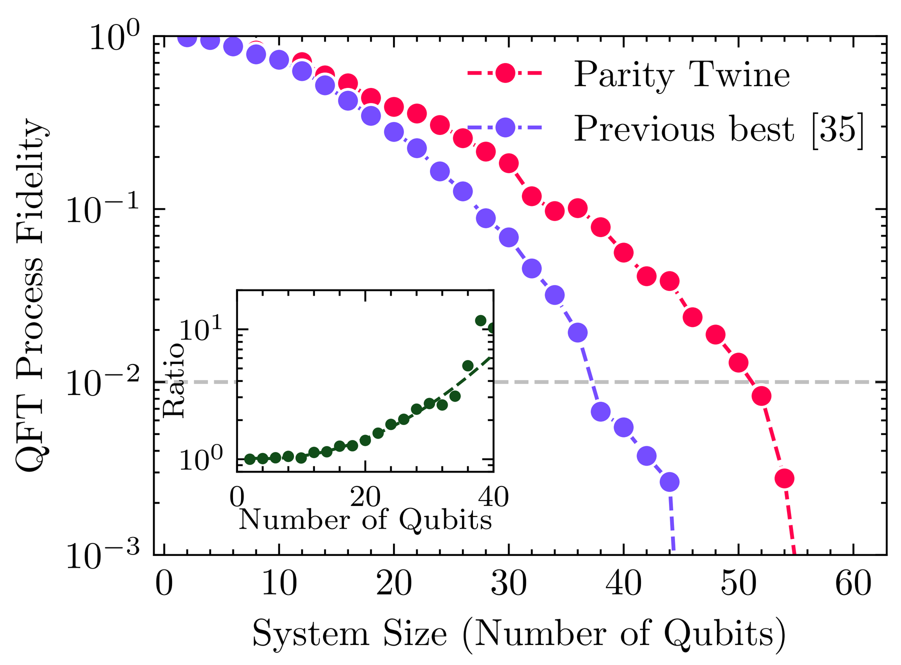

Validation

12 pixels, $d = 32$

(left) Covariance matrix between pixels of the validation dataset.

(right) Estimated covariance matrix from the IQP circuit, obtained with 1000 samples.

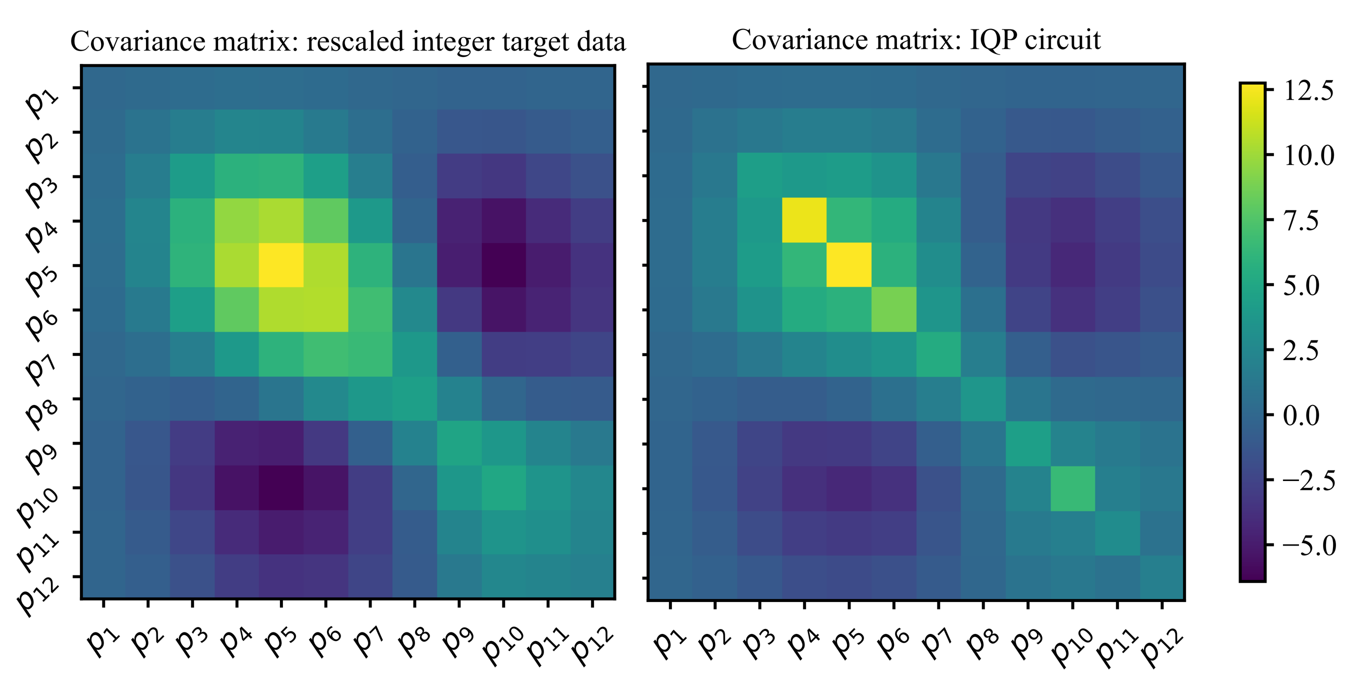

(left) Correlation matrix between pixels of the same set.

(right) Estimated correlation from the IQP circuit, obtained from the covariance calculation.

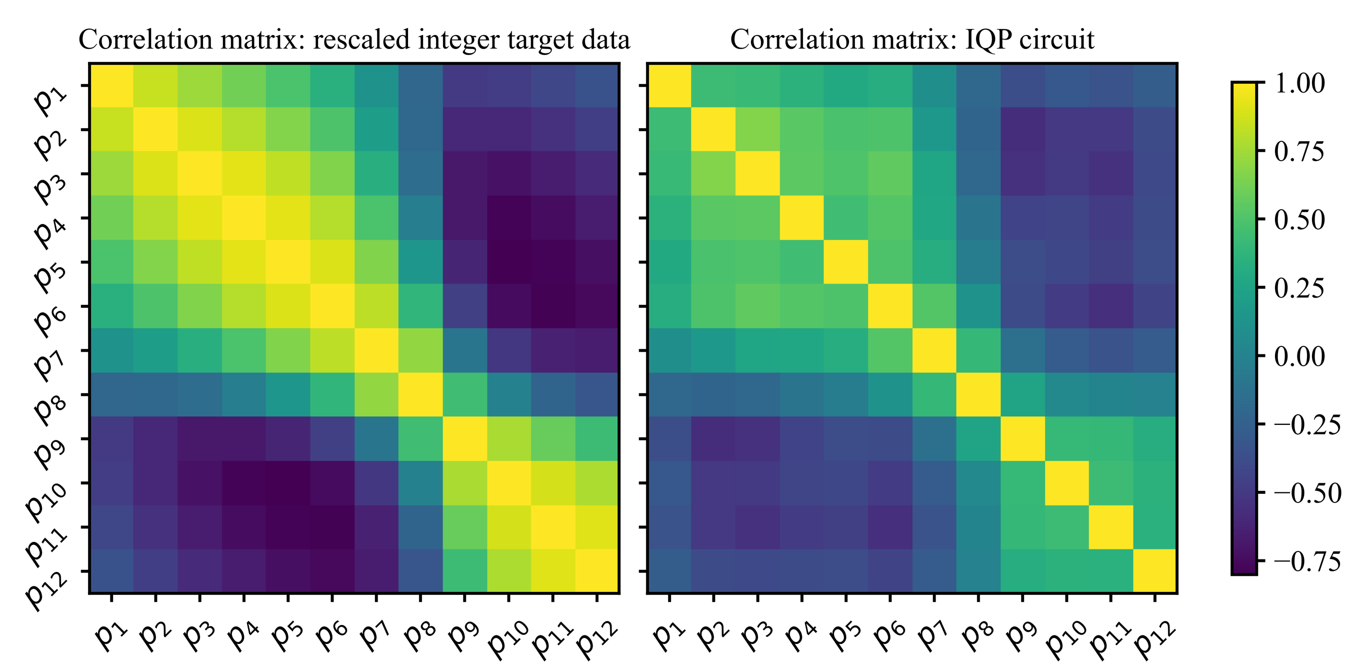

References

[1] Belayneh, Dawit, et al. "Calorimetry with deep learning: particle simulation and reconstruction for collider physics." *The European Physical Journal C* 80.7 (2020): 1-31.

[2] Dreier, Florian, *et al.* "Connectivity-aware synthesis of quantum algorithms." *arXiv preprint arXiv:2501.14020* (2025).

[3] Aumann, Philipp, et al. "Demonstrating Record Fidelity for the Quantum Fourier Transform." *arXiv preprint arXiv:2604.12465* (2026).

[4] Recio-Armengol, Erik, Shahnawaz Ahmed, and Joseph Bowles. "Train on classical, deploy on quantum: scaling generative quantum machine learning to a thousand qubits." *arXiv preprint arXiv:2503.02934* (2025).

This work was developed within the project Hamburg Full Stack Quantum Machine Learning (HQML) in collaboration with consortium partners DESY, eleQtron and the DLR Quantum Computing Initiative (DLR QCI).

This project was funded by IFB Hamburg.