PROCEEDINGS OF SCIENCE

Quantum Invertible Neural Networks for Fast Calorimetry as Quantum Generative Model with High Interpretability at Future Colliders

F. Corchia[a,b] **and M. Franchini**[a,b,*]

[a] *Istituto Nazionale di Fisica Nucleare (INFN), Sezione di Bologna, Viale Berti-Pichat 6/2, 40127 Bologna, Italy*

[b] *Dipartimento di Fisica e Astronomia, Università di Bologna, Viale Berti-Pichat 6/2, 40127 Bologna, Italy*

E-mail: f.corchia@bo.infn.it, matteo.franchini@bo.infn.it

The fast calorimeter simulation is a critical computational bottleneck for the ATLAS experiment at the High-Luminosity LHC and for the design of future colliders. Invertible Neural Networks (INNs) and normalizing flows have demonstrated impressive accuracy as generative models for calorimeter shower simulation, forming the backbone of the emerging FastCaloSim framework. In this work we explore the extension of this classical approach toward quantum-conditioned invertible neural networks (qINN). We present the conceptual motivation, discuss the fundamental tension between the unitarity of quantum circuits and the irreversibility of projective measurement, and outline two complementary solution strategies. As a first feasibility study we implement a hybrid architecture in which a small quantum circuit replaces the classical conditioning sub-network inside an otherwise classical INN, trained on ATLAS Open-Data pion showers from the Fast Calorimeter Challenge dataset. The preliminary results, obtained with a 6-qubit, 2-layer variational ansatz, show encouraging convergence and motivate further optimisation. The work is a *feasibility study*; no definitive performance claims are made.

International Symposium on Grids and Clouds (ISGC2026)
15–20 March 2026
Academia Sinica, Taipei, Taiwan

*Speaker

© Copyright owned by the author(s) under the terms of the Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License (CC BY-NC-ND 4.0) All rights for text and data mining, AI training, and similar technologies for commercial purposes, are reserved. ISSN 1824-8039 . Published by SISSA Medialab.

PoS(ISGC2026)022

1. Introduction

Modern high-energy physics (HEP) experiments rely on Monte Carlo (MC) simulation for virtually every aspect of their scientific programme, from detector calibration to new-physics searches. Among the most computationally demanding components is the simulation of particle interactions in calorimeters, where the development of electromagnetic and hadronic showers must be reproduced with high fidelity. Full simulation based on GEANT4 [1] achieves excellent accuracy but comes at prohibitive CPU cost: at the High-Luminosity LHC (HL-LHC), expected to collect data from $\sim$ 2030 [2], the number of events to simulate will increase by roughly an order of magnitude compared to Run 3, making full simulation alone untenable within projected computing budgets [3].

Generative machine-learning models offer a pathway to orders-of-magnitude speedups while preserving distribution-level accuracy. They are used in HEP for fast detector simulation, calorimeter shower generation, background modelling, anomaly detection, and sampling of complex physics distributions [4, 5]. As the community looks beyond the HL-LHC toward proposed future colliders such as the FCC or CLIC, detector designs will produce even denser, more granular showers, amplifying the computational challenge further.

Invertible Neural Networks (INNs) and normalizing flows have emerged as particularly attractive generative architectures for this domain because they provide exact likelihood evaluation and efficient sampling alongside a tractable Jacobian determinant [6–9]. Several ATLAS-related developments—CALOINN, CALOCFM, and CALODIT—now exploit INN-based approaches for both energy and shower-shape generation.

In parallel, the field of quantum machine learning (QML) has produced encouraging demonstrations of generative quantum neural networks on low-dimensional distributions and reduced calorimeter images [23, 24]. Quantum circuits are described by unitary operators and are invertible by construction, suggesting a natural structural connection to INNs. However, the projective measurement step that extracts classical information from a quantum state destroys unitarity, creating a fundamental tension that must be addressed before the analogy can be exploited in practice.

In this proceedings we report on a preliminary feasibility study of *quantum invertible neural networks* (QINN) for fast calorimeter simulation in the context of the ATLAS FASTCALOSIM programme [11]. The paper is organised as follows. Section 2 reviews generative models for calorimeter simulation. Section 3 describes INNs and normalizing flows as currently deployed in FASTCALOSIM. Section 4 introduces the QINN concept and the measurement problem. Section 5 details the implementation strategy. Section 6 presents preliminary results, and Section 8 summarises conclusions and next steps.

2. Generative Models for Calorimeter Simulation in HEP

Given a training dataset $\mathcal{D} = \{x_i\}$ drawn from an unknown distribution $p_{\text{data}}(x)$, a generative model learns to produce new samples $\hat{x} \sim p_\theta(x) \approx p_{\text{data}}(x)$ using a parametric family with parameters θ. The specific requirements in calorimeter simulation are demanding: the output space can have $O(10^3$–$10^5)$ dimensions (calorimeter cells), the distributions span many orders

PoS(ISGC2026)022

of magnitude in energy, and both local shower shapes and global energy containment must be reproduced faithfully.

Several generative paradigms have been explored for this task. Generative Adversarial Networks (GANs) [16] were among the first to be applied, with FastCaloGAN V2 representing the baseline fast-simulation model currently deployed in the ATLAS experiment [5]. A known limitation of the GAN approach is that shower selection in FastCaloGAN V2 is based on total simulated energy rather than the actual shower shape, leading to systematic mismodelling in certain phase-space regions. Active work within the ATLAS simulation group is addressing this issue.

Variational Autoencoders (VAEs), diffusion models, and, most relevantly here, normalizing flows and INNs have subsequently been explored, with the latter showing particular promise for their exact density evaluation and well-defined inverse [8, 9, 11].

A community-wide benchmarking effort is provided by the Fast Calorimeter Simulation Challenge [12, 13], which supplies public datasets of Geant4-simulated calorimeter showers and standardised metrics, enabling fair comparison of algorithms. This challenge constitutes the primary benchmark environment for the present work.

3. Invertible Neural Networks and Normalizing Flows in FastCaloSim

3.1 Mathematical framework

normalizing flows [14, 15] learn an invertible, differentiable mapping $f_\theta : \mathcal{Z} \to \mathcal{X}$ between a simple base distribution $p_Z(z)$ (typically a standard Gaussian) and the complex data distribution $p_X(x)$. The change-of-variables theorem gives the exact likelihood:

$$p_X(x) = p_Z\left(f_\theta^{-1}(x)\right)\left|\det \frac{\partial f_\theta^{-1}(x)}{\partial x}\right|, \tag{1}$$

enabling training by maximum likelihood:

$$\mathcal{L}(\theta) = -\sum_i \left[\log p_Z\left(f_\theta^{-1}(y_i)\right) + \log\left|\det \frac{\partial f_\theta^{-1}(y_i)}{\partial y_i}\right|\right]. \tag{2}$$

Practical implementations avoid the $O(d^3)$ cost of a general Jacobian determinant by constructing *coupling layers* [6]. A sample $x = (x_1, x_2)$ is split into two halves and transformed as:

$$y_1 = x_1, \tag{3}$$
$$y_2 = x_2 \odot \exp\big(s(x_1)\big) + t(x_1), \tag{4}$$

where $s(\cdot)$ and $t(\cdot)$ are arbitrarily complex (non-invertible) neural networks, and $\odot$ denotes element-wise multiplication. The inverse is analytically available:

$$x_1 = y_1, \tag{5}$$
$$x_2 = \big(y_2 - t(y_1)\big) \odot \exp\big(-s(y_1)\big). \tag{6}$$

PoS(ISGC2026)022

Because s and t are never inverted, the Jacobian $\partial y/\partial x$ is triangular and its determinant reduces to $\exp(\sum_j s(x_1)_j)$, which is computationally cheap. A full normalizing flow is constructed by composing K such coupling layers, $f = f_K \circ \cdots \circ f_1$, whose overall inverse is $f^{-1} = f_1^{-1} \circ \cdots \circ f_K^{-1}$.

The main advantages of this class of models are efficient sampling (one forward pass), exact density evaluation, a useful and inspectable latent representation, and relatively straightforward maximum-likelihood training. Known limitations include difficulty with very noisy data, the cost of the Jacobian determinant in high dimensions, and reduced expressivity for pure discriminative tasks.

3.2 INN-based models in ATLAS FastCaloSim

The ATLAS simulation group has recently moved toward INN and flow-based models to replace and augment the GAN-based approach. In the current framework [11], the simulation pipeline operates on voxelised calorimeter deposits and proceeds in two stages:

1. **CaloINN**: predicts the per-layer energy deposits from the incident particle's kinematics using an INN-based model trained on normalised voxels.

2. **Shape models**: predict the energy distribution within each layer, conditioned on the layer energy $\hat{E}_{\text{layer}}$ from CaloINN. Two parallel approaches are under study:

   - CaloCFM [10]: a Continuous Normalizing Flow model implementing a time-continuous invertible transformation via a neural ODE.
   - CaloDiT-2: a Diffusion Transformer that frames shape generation as a stochastic inversion problem.

Preliminary shower-shape results from these models show good agreement with Geant4 reference distributions for both photon and pion showers across a wide energy range [11], outperforming the uncorrected FastCaloGAN V2 baseline in the modelling of the lateral shower profile. These models are currently undergoing integration into the official ATLAS software release.

4. Toward Quantum Invertible Neural Networks

4.1 Motivation from quantum generative models

Classical neural network architectures are growing rapidly in size and computational cost, approaching the boundaries of practical feasibility for certain tasks. Quantum Generative Neural Networks (Gen-QNNs) offer a potentially complementary paradigm: quantum superposition allows a register of n qubits to represent 2^n states simultaneously, while entanglement enables the representation of complex multi-qubit correlations that have no direct classical analogue. These properties make quantum circuits natural candidates for sampling machines.

Several families of quantum generative architectures have been explored in the literature:

QCBM (Quantum Circuit Born Machines) [17, 18] learn a probability distribution $p_\theta(x)$ directly from the measurement outcomes of a parameterised quantum circuit $U(\theta)|0\rangle$.

Quantum GANs [19, 20] employ a quantum generator, typically with a classical discriminator.

PoS(ISGC2026)022

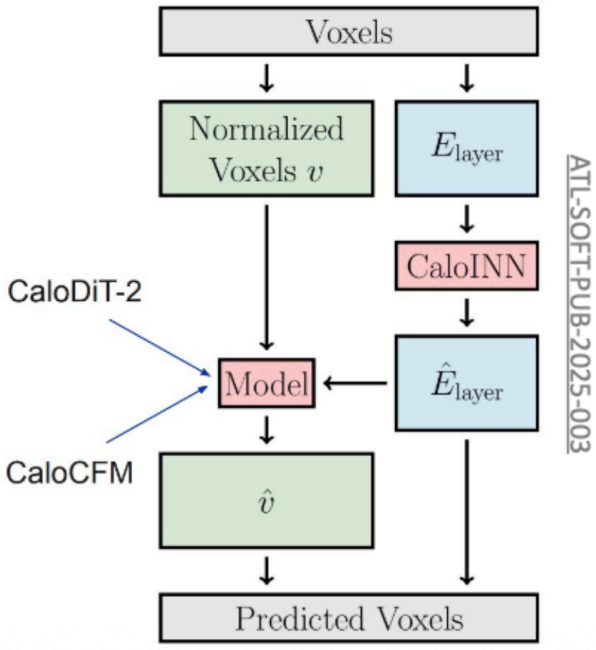

Figure 1: Schematic of the INN-based FastCaloSim pipeline (adapted from [11]). Voxelised calorimeter deposits are processed in two stages: CaloINN predicts layer energies, while shape models (CaloCFM or CaloDiT-2) generate the intra-layer energy distribution of the normalised voxels $\hat{v}$.

Quantum VAEs [21] use quantum circuits in the latent representation or the decoder/generator block.

Quantum-enhanced normalizing flows [22, 23] are hybrid designs in which quantum components enrich the transformations inside an otherwise classical flow.

Encouraging demonstrations exist for low-dimensional distributions and reduced calorimeter images [23, 24], but it remains an open question whether quantum generative models can surpass strong classical baselines in the high-dimensional settings relevant to HEP. A recent study at the CERN IQT workshop [25] identified key open problems: defining feature maps that enable efficient sampling through projective measurements in the computational basis, matching model capacity to the ansatz structure, and finding transformations between bases (e.g., Fourier or Chebyshev spaces) that enable expressive generative modelling.

PoS(ISGC2026)022

4.2 The invertibility tension in quantum circuits

Quantum circuits are described by unitary operators $U(\theta)$ and are therefore invertible by definition:

$$|\psi_{\text{out}}\rangle = U(\theta)|\psi_{\text{in}}\rangle, \quad U(\theta) = U_1(\theta_1)\, U_2(\theta_2) \cdots U_N(\theta_N). \tag{7}$$

This is structurally analogous to the coupling-layer composition of an INN. A natural first attempt is to map the latent–data bijection $z \leftrightarrow x = f_\theta(z)$ onto the unitary transformation $|\psi_\theta\rangle = U(\theta)|z\rangle$, with the Born-rule probability $p_\theta(x) = |\langle x|\psi_\theta\rangle|^2$ playing the role of the model density, and training via

$$\mathcal{L}(\theta) = -\frac{1}{N}\sum_{i=1}^{N} \log p_\theta(x_i). \tag{8}$$

However, the analogy breaks down at the measurement step. *Projective measurement is irreversible*: once a qubit is measured in the computational basis, the quantum state collapses and all information about the pre-measurement superposition is lost. The unitary evolution is invertible; the measurement is not. This fundamental asymmetry prevents the direct application of INN-style inversion in a fully quantum setting and constitutes the central conceptual challenge of QINN design.

4.3 Proposed solution strategies

Two broad strategies can be identified to address the measurement problem:

Algorithmic approach.

1. *Full state-vector retention*: storing the complete quantum state vector avoids the measurement collapse and allows exact inversion. This approach is feasible for small qubit counts in simulation, but scales exponentially in classical memory (2^n complex amplitudes) and requires full state tomography on real hardware, which is prohibitively time-consuming. Future quantum memory devices could make this viable.

2. *Smart projection spaces*: designing feature maps that make projective measurements approximately univocal, for instance via SWAP-test-based similarity metrics. This direction has been explored by Spannowsky et al. [26] but remains technically challenging.

Structural approach (pursued in this work). Rather than attempting a fully quantum INN, one can embed quantum conditioning inside a classical INN architecture. The coupling functions $s(\cdot)$ and $t(\cdot)$ of equation (4) are replaced by quantum neural network (QNN) modules:

$$y_2 = x_2 \odot \exp\big(s_{\text{QNN}}(x_1)\big) + t_{\text{QNN}}(x_1). \tag{9}$$

The overall INN inversion is guaranteed by the classical structure (equations (3)–(4)), and the quantum sub-network is used only in the forward direction. The hope is that the richer feature space accessible to the QNN—through superposition and entanglement—can improve expressivity in difficult regions of phase space and enhance correlation modelling, without sacrificing tractability. The limitations are that the classical core retains the same Jacobian-inversion cost, and improved expressivity over a well-tuned classical sub-network is not guaranteed.

PoS(ISGC2026)022

5. Implementation Strategy

The implementation is carried out within the Fast Calorimeter Challenge (CaloChallenge) [12] framework, which provides standardised datasets and evaluation tools. The dataset used consists of pion (and photon) showers generated at 65 GeV with GEANT4 using ATLAS Open Data detector geometry [13], containing up to 10^6 events.

The software stack is based on PYTORCH [28] for the classical INN components and PENNYLANE [29] for the quantum sub-network, interfaced via PennyLane's `qml.qnn.TorchLayer` wrapper. This setup allows gradient-based training through the quantum circuit using the parameter-shift rule for gradient computation. The starting point was a Keras-based INN implementation from the CaloChallenge baseline, migrated to PyTorch to enable the PennyLane integration.

The model configuration for the preliminary study is as follows:

- Architecture: `BNReLUqINN` with generator layers $(100, 64)$;
- Latent space dimension: 10;
- Input features: 10, processed in batches of 10;
- Training sample: 1 000 events (proof-of-concept scale);
- Quantum sub-network: 1 conditioning QNN module with 6 qubits and 2 variational layers, implemented with linear entanglement;
- Gradient method: `best` (parameter-shift rule);
- Classical device backend: CPU simulation (`default.qubit`).

The code is publicly available at `https://github.com/QC4HEP-BO/FastCaloQ/tree/qINN`.

6. Preliminary Results

The results presented in this section are very preliminary and should be considered a first feasibility study. No definitive performance claims are made.

6.1 Loss convergence

Training was performed using a Wasserstein loss [30], which provides a gradient signal even when the 2 parallel training samples (called generator and discriminator respectively) distributions have disjoint support. Figure 2 shows the Wasserstein loss for two training configurations (different batch sizes) as a function of training iteration (where Iter = epoch × batch size). In both cases the generator loss decreases rapidly in the first few hundred iterations. The convergence is fast (approximately 120 min of training on CPU), but the loss behaviour suggests that the current Wasserstein objective may not fully capture the correct likelihood structure of the flow; replacing it with a normalizing-flow-specific negative log-likelihood objective (equation 2) is planned as a next step.

PoS(ISGC2026)022

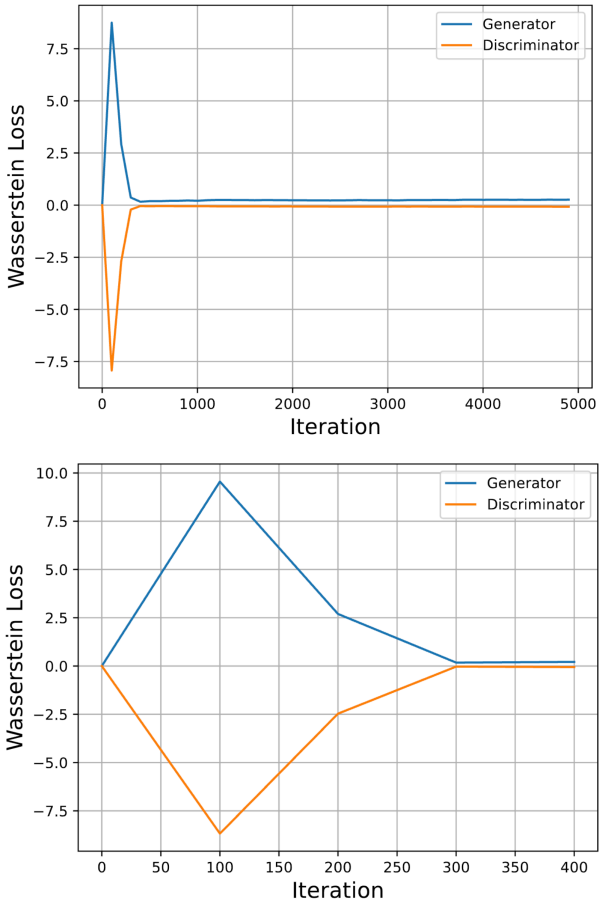

Figure 2: Wasserstein loss as a function of training iteration for the qINN model with two batch-size configurations (left: small batch; right: large batch). Generator (blue) and discriminator (orange) losses are shown. Rapid initial convergence is observed in both cases. *Very preliminary.*

6.2 Energy distribution generation

Figure 3 shows the generated energy distribution for pion showers in the pseudorapidity slice $0.20 < |\eta| < 0.25$ compared to the GEANT4 reference, at several training checkpoints (Iter = 3000, 5000, 7000, and 18000). The χ^2/NDF values decrease from ~ 71 at 3000 iterations to ~ 7 at 18 000 iterations, indicating progressive improvement in the modelling of the energy distribution. The best checkpoint achieves a $\chi^2/\text{NDF} = 200/28 \approx 7.1$, demonstrating that the model is learning non-trivial features of the distribution.

These results must be interpreted with care: the training sample is small (1 000 events), the architecture is minimal, and no hyperparameter optimisation has been performed. Nevertheless, the trend provides a proof-of-concept that the quantum-conditioned INN architecture is trainable and that the quantum conditioning module does not prevent learning.

PoS(ISGC2026)022

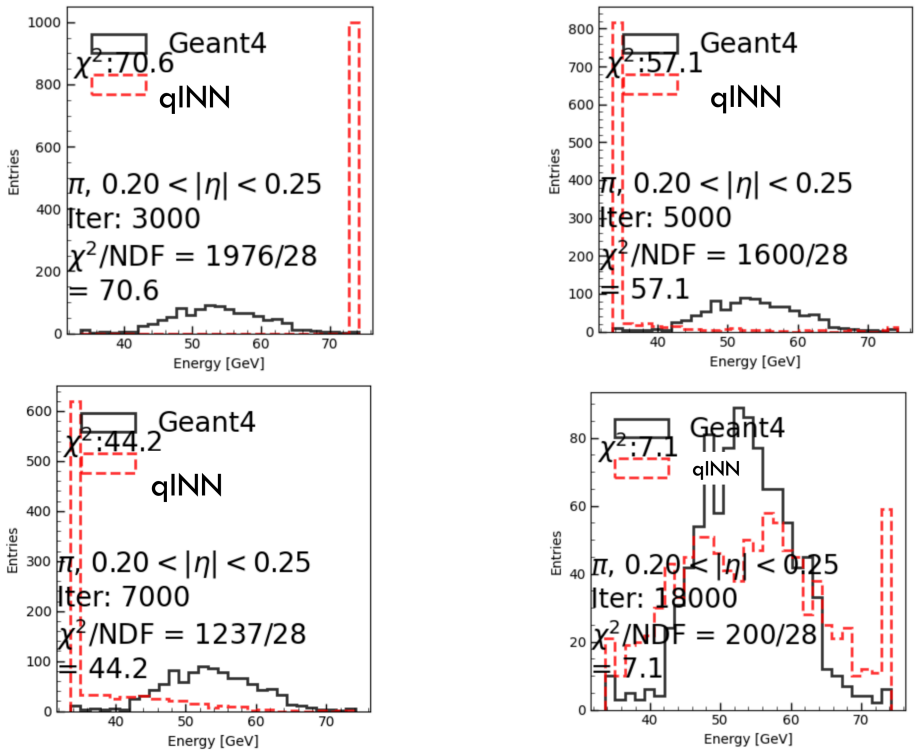

Figure 3: Generated pion energy distributions (dashed red) compared to GEANT4 reference (solid black) for $0.20 < |\eta| < 0.25$, at increasing training iterations. The χ^2/NDF decreases from $\sim$71 at 3000 iterations to $\sim$7 at 18 000 iterations. Training time was approximately 120 minutes on CPU. *Very preliminary.*

7. Discussion and Limitations

The results presented above establish that a hybrid QINN with quantum conditioning is trainable and produces a distribution that progressively approaches the GEANT4 reference during training. However, several important caveats must be noted.

Loss function. The Wasserstein loss drives rapid initial convergence but does not correspond to the exact maximum-likelihood objective of a normalizing flow (equation 2). This may explain the observed discrepancy between fast loss decrease and residual χ^2/NDF values. Adopting a proper negative log-likelihood objective, possibly with a physics-motivated regularisation term, is a necessary step for a rigorous comparison.

Scalability. The proof-of-concept uses 6 qubits, 10 input features, and 1 000 training events. Full-scale calorimeter simulation requires hundreds of input dimensions and $O(10^5$–$10^6)$ training events. Scaling the quantum sub-network is non-trivial: the classical simulation of a 6-qubit circuit is fast, but 20–30 qubits would require exponentially larger memory. On current noisy intermediate-scale quantum (NISQ) hardware, circuit depth and qubit connectivity are further constrained by decoherence. Near-term implementations will likely remain in the hybrid simulation regime.

Barren plateaus. Parameterised quantum circuits with more than a few layers are susceptible to barren plateaus in the loss landscape, where gradients vanish exponentially with system size [27].

PoS(ISGC2026)022

The 2-layer ansatz used here is modest enough to avoid this regime, but deeper circuits will require targeted mitigation strategies.

Quantum advantage. No evidence of quantum advantage has been demonstrated at this stage. The current setup is a feasibility demonstration; whether the quantum conditioning provides an expressivity benefit over a classical sub-network of equivalent parameter count remains to be shown through systematic ablation studies.

Shower shapes. The present implementation targets the energy distribution only. Extending the model to generate full shower shapes (the primary scientific motivation in FastCaloSim) requires significantly higher output dimensionality and is left for future work.

8. Conclusions

We have presented a preliminary feasibility study of quantum-conditioned invertible neural networks (qINN) for fast calorimeter simulation in HEP. The key contributions are:

1. A clear formulation of the invertibility tension in quantum circuits—unitary evolution is invertible, but projective measurement is not—and two complementary strategies to address it: an algorithmic approach (state-vector retention or smart projections) and a structural approach (quantum conditioning inside a classical INN).

2. A first implementation of the structural qINN, integrating a 6-qubit PennyLane variational circuit as the coupling-layer conditioning function within a PyTorch-based INN trained on ATLAS Open Data pion showers from the CaloChallenge framework.

3. Preliminary results showing trainability and progressive improvement in the modelled energy distribution, with χ^2/NDF decreasing from ~ 71 to ~ 7 over 18 000 iterations on a small proof-of-concept dataset.

The idea is confirmed to be feasible. The immediate next steps are: optimising the loss function toward a proper normalizing-flow likelihood; replacing the quantum conditioning module with a deeper QNN acting on both $s(\cdot)$ and $t(\cdot)$; extending the model to generate shower shapes; performing systematic hyperparameter optimisation; and, ultimately, exploring integration as a conditioning module within the full ATLAS FastCaloSim INN framework. A parallel algorithmic direction—implementing fully reversible generative quantum circuits—is already under development.

The HL-LHC will begin operations around 2030, and future colliders and detectors will impose even greater computational demands on simulation. Early investment in quantum-enhanced generative models positions the community to exploit quantum hardware as it matures, while already providing interpretability benefits through the INN structure. The present work, while preliminary, demonstrates that this direction is technically viable and scientifically motivated.

References

[1] S. Agostinelli et al. (Geant4 Collaboration), *Geant4 4—a simulation toolkit*, Nucl. Instrum. Meth. A **506** (2003) 250, 10.1016/S0168-9002(03)01368-8.

PoS(ISGC2026)022

[2] ATLAS Collaboration, *ATLAS Phase-II Upgrade Scoping Document*, CERN/LHCC 2015-020 (2015).

[3] ATLAS Collaboration, *ATLAS Software and Computing HL-LHC Roadmap*, ATL-SOFT-PUB-2020-001 (2020).

[4] M. Paganini, L. de Oliveira and B. Nachman, *CaloGAN: Simulating 3D high energy particle showers in multi-layer electromagnetic calorimeters with generative adversarial networks*, Phys. Rev. D **97** (2018) 014021, 10.1103/PhysRevD.97.014021.

[5] ATLAS Collaboration, *Fast simulation of the ATLAS detector using a Generative Adversarial Network (FastCaloGAN)*, ATL-SOFT-PUB-2020-006 (2020).

[6] L. Dinh, J. Sohl-Dickstein and S. Bengio, *Density estimation using Real-valued Non-Volume Preserving (Real NVP) transformations*, arXiv:1605.08803 (2016).

[7] L. Ardizzone et al., *Analyzing inverse problems with invertible neural networks*, arXiv:1808.04730 (2018).

[8] M. Kofler et al., *CaloINN: Fast calorimeter simulation with invertible neural networks*, arXiv:2104.10170 (2021).

[9] ATLAS Collaboration, *Fast simulation of calorimeter showers with normalizing flows and invertible neural networks in ATLAS*,

[10] ATLAS Collaboration, *CaloCFM: Continuous normalizing flows for calorimeter simulation*,

[11] ATLAS Collaboration, *ATLAS fast calorimeter simulation with INN-based models*, ATL-SOFT-PUB-2025-003 (2025).

[12] C. Krause et al., *CaloChallenge 2022: A community challenge for fast calorimeter simulation*, arXiv:2410.21611 (2024), 10.1088/1361-6633/ae1304.

[13] C. Krause et al., *Fast Calorimeter Simulation Challenge 2022 – dataset 2*, https://calochallenge.github.io/homepage/.

[14] D. J. Rezende and S. Mohamed, *Variational inference with normalizing flows*, arXiv:1505.05770 (2015).

[15] G. Papamakarios et al., *Normalizing flows for probabilistic modeling and inference*, J. Mach. Learn. Res. **22** (2021) 1, arXiv:1912.02762.

[16] I. J. Goodfellow et al., *Generative adversarial nets*, Adv. Neural Inf. Process. Syst. **27** (2014), arXiv:1406.2661.

[17] J.-G. Liu and L. Wang, *Differentiable learning of quantum circuit Born machines*, Phys. Rev. A **98** (2018) 062324, 10.1103/PhysRevA.98.062324.

[18] M. Benedetti et al., *Parameterized quantum circuits as machine learning models*, Quantum Sci. Technol. **4** (2019) 043001, 10.1088/2058-9565/ab4eb5.

PoS(ISGC2026)022

[19] S. Lloyd and C. Weedbrook, *Quantum generative adversarial learning*, Phys. Rev. Lett. **121** (2018) 040502, 10.1103/PhysRevLett.121.040502.

[20] H.-Y. Huang et al., *Quantum advantage in learning from experiments*, Science **376** (2022) 1182, 10.1126/science.abn7293.

[21] A. Khoshaman et al., *Quantum variational autoencoder*, Quantum Sci. Technol. **4** (2019) 014001, 10.1088/2058-9565/aada1f.

[22] M. Benedetti et al., *Adversarial quantum circuit learning for pure state approximation*, New J. Phys. **21** (2019) 043023, 10.1088/1367-2630/ab14b5.

[23] F. Rehm et al., *Quantum machine learning for HEP detector simulations*, arXiv:2109.03430 (2021).

[24] X. Gao, E. R. Anschuetz, S.-T. Wang, J. I. Cirac and M. D. Lukin, *Enhancing generative models via quantum correlations*, Phys. Rev. X **12** (2022) 021037, 10.1103/PhysRevX.12.021037.

[25] M. Grossi, *Quantum generative models for HEP*, Contribution to CERN IQT / QCINFN workshop, 2026, https://agenda.infn.it/event/48148/contributions/283939/attachments/145676/221556/Grossi_CERN_QCINFN26.pdf.

[26] M. Spannowsky et al., *Quantum invertible neural networks*, SciPost Phys. **16** (2024) 146, 10.21468/SciPostPhys.16.6.146.

[27] J. R. McClean, S. Boixo, V. N. Smelyanskiy, R. Babbush and H. Neven, *Barren plateaus in quantum neural network training landscapes*, Nat. Commun. **9** (2018) 4812, 10.1038/s41467-018-07090-4.

[28] A. Paszke et al., *PyTorch: An imperative style, high-performance deep learning library*, Adv. Neural Inf. Process. Syst. **32** (2019), arXiv:1912.01703.

[29] V. Bergholm et al., *PennyLane: Automatic differentiation of hybrid quantum-classical computations*, arXiv:1811.04968 (2018).

[30] M. Arjovsky, S. Chintala and L. Bottou, *Wasserstein GAN*, arXiv:1701.07875 (2017).

[31] F. Corchia and M. Franchini, *Quantum generative INN for fast calorimetry*, PoS(EPS-HEP2025) 639 (2025).

PoS(ISGC2026)022