

One particle - one qubit: Particle physics data encoding for quantum machine learning

Aritra Bal* and Markus Klute
*Institute of Experimental Particle Physics (ETP),
 Karlsruhe Institute of Technology, Karlsruhe, DE*

Benedikt Maier,[†] Melik Oughton, and Eric Pezone
Imperial-X and Department of Physics, Imperial College Of Science, Technology And Medicine, London, UK

Michael Spannowsky[‡]
Institute for Particle Physics Phenomenology (IPPP), Durham University, Durham, UK

We introduce 1P1Q, a novel quantum data encoding scheme for high-energy physics (HEP), where each particle is assigned to an individual qubit, enabling direct representation of collision events on quantum circuits without classical compression. We demonstrate the effectiveness of 1P1Q in quantum machine learning (QML) through two applications: a Quantum Autoencoder (QAE) for unsupervised anomaly detection and a Variational Quantum Circuit (VQC) for supervised classification of top quark jets. Our results show that the QAE successfully distinguishes signal jets from background QCD jets, achieving superior performance compared to a classical autoencoder while utilizing significantly fewer trainable parameters. Similarly, the VQC achieves competitive classification performance, approaching state-of-the-art classical models despite its minimal computational complexity. Furthermore, we validate the QAE on real experimental data from the CMS detector, establishing the robustness of quantum algorithms in practical HEP applications. These results demonstrate that 1P1Q provides an effective and scalable quantum encoding strategy, offering new opportunities for applying quantum computing algorithms in collider data analysis.

Published in: Physical Review D **112** (2025), 076004.

Introduction. The unprecedented collision energies achieved at latest- and next-generation colliders, like the Large Hadron Collider (LHC) and the Future Circular Collider (FCC), produce vast amounts of particle-level data, challenging the limits of conventional data analysis techniques in high energy physics (HEP). With the increasing complexity of these datasets, novel approaches that leverage cutting-edge computational paradigms have become indispensable. Machine Learning (ML) has established itself as an indispensable tool for analysing HEP data in recent years, leading to significant advancements in tasks such as event classification, anomaly detection, and parameter estimation. An emerging frontier in this field is to extend these techniques to the quantum domain, leveraging Quantum Machine Learning (QML) to enhance data analysis capabilities. Quantum computing, with its ability to exploit superposition, entanglement, and interference, offers a promising framework to address some of the most intricate challenges in HEP [1–3].

Any quantum machine learning algorithm consists of three key components: (1) the data encoding, which maps classical data onto quantum states; (2) the quantum model, typically implemented through quantum circuits and quantum operations such as entanglement between qubits; and (3) the loss function, whose optimization is crucial for training the quantum model. While

substantial efforts have been devoted to designing quantum models and optimizing loss functions [4–13], the choice of data encoding remains a critical yet underexplored aspect of QML for HEP applications. In this work, we propose a new encoding scheme tailored to HEP data, which we call 1P1Q (1 Particle - 1 Qubit). This encoding strategy assigns a separate qubit to each particle, enabling an effective representation of collider events within quantum circuits without prior data compression. Establishing a direct correspondence between particles and qubits is warranted to advance the usefulness of quantum computing for HEP in particular upon entering an era in which crucial LHC measurements reveal the quantum mechanical nature of the imprints left behind by the collisions, such as the recently observed entanglement in top quark pair production [14, 15].

We demonstrate the effectiveness of 1P1Q-based QML models in processing HEP data by employing two distinct approaches: a Quantum Autoencoder (QAE) [16] for unsupervised, unlabelled learning and a Variational Quantum Circuit (VQC) [17] for classification tasks. The QAE compresses quantum states by learning a lower-dimensional latent representation, making it well-suited for anomaly detection in collider physics, while the VQC employs a parametrized quantum circuit to discriminate between different class categories based on learned quantum features. Both methods fully exploit the kinematic information of particles encoded via the 1P1Q scheme, using the quantum state representation to retain and process intricate correlations between particles. Such information is at risk of being lost in existing approaches [18–25], which first encode the information of the collider

* aritra.bal@kit.edu

† benedikt.maier@cern.ch

‡ michael.spannowsky@durham.ac.uk

events into an abstract, typically compressed, latent representation using classical machine learning algorithms or domain-inspired high-level features. Instead, 1P1Q guarantees more immediate access to the raw information of collider events and allows for the direct exploitation of their content using quantum variational circuits. 1P1Q therefore also provides a natural way of extending the input space as quantum computers and their simulators can accommodate more and more qubits.

We apply the QML models to one of the most well-established use cases of classical machine learning in HEP: the discrimination of hadronically decaying resonances from the overwhelming background of QCD jets. Identifying such resonances, including those from top quarks, Higgs bosons, or hypothetical new physics states, is crucial for advancing our understanding of fundamental interactions. We find that for this task, VQC and QAE models acting on 1P1Q-encoded particle information result in highly performant QML algorithms, equal to or even better than comparable classical counterparts in the case of the QAE.

For the first time, our study explores QML in HEP using actual experimental data recorded by the CMS detector [26] in 2016. This real-world application provides a critical test of the robustness and feasibility of QML strategies in general, and 1P1Q in particular, beyond simulated datasets. By demonstrating that QML models can be trained on and extract meaningful physics from real collision events, we show that 1P1Q can become a modern, lossless collider data encoding framework as quantum computing hardware advances.

In the following, we present the formulation of the 1P1Q encoding, explore its theoretical underpinnings, and discuss its potential applications to jet physics at the LHC and future colliders with the examples of anomaly detection and a supervised classifier. We demonstrate how this encoding captures the essential kinematic features of jets and showcase its utility in leveraging quantum algorithms to analyze jet substructure, and compare to state-of-the-art classical machine learning algorithms.

1P1Q - Particle Encoding. In collider measurements, reconstructed particles are kinematically fully described by three key parameters: the transverse momentum $p_T = \sqrt{p_x^2 + p_y^2}$, where p_x and p_y are the momentum components in the transverse detector plane, the pseudorapidity η , and the azimuthal angle ϕ . The pseudorapidity η is related to the polar angle ν of the particle's trajectory by $\eta = -\ln \tan(\nu/2)$, which provides an approximation of the rapidity in high-energy regimes. Finally, the azimuthal angle ϕ describes the particle's direction in the transverse plane and is measured relative to a chosen reference axis. Together, these three quantities specify the particle's momentum.

The 1P1Q method directly encodes the kinematics of a particle on a qubit. The pseudorapidity η and the azimuthal angle ϕ of the particle, modulated by the transverse momentum p_T normalized to the p_T of the jet, are used as spherical coordinates on the Bloch sphere to ori-

ent the qubit, enabling a compact and information-rich representation of each particle as a quantum state that is not dependent on the energy scale of the jet, thus facilitating a straightforward transition to the quantum domain. These feature encodings, represented by rotation angles about the Y and X axes respectively, are then additionally scaled by a factor $f = f(w)$ constrained to lie between $[1, 2\pi + 1]$, where w is a trainable parameter. This ensures the particles can spread out across the Bloch sphere instead of clustering too closely around its North Pole. The encoding can then be summarized as:

$$f \cdot \frac{p_T}{p_{T,\text{jet}}} \cdot (\eta - \eta_{\text{jet}}) \rightarrow \theta \quad (1)$$

$$f \cdot \frac{p_T}{p_{T,\text{jet}}} \cdot (\phi - \phi_{\text{jet}}) \rightarrow \varphi \quad (2)$$

$$(p_T, \eta, \phi) \rightarrow |\psi\rangle = R_X(\varphi)R_Y(\theta)|0\rangle \\ = \alpha(\theta, \varphi)|0\rangle + \beta(\theta, \varphi)|1\rangle \quad (3)$$

$$f \rightarrow 1 + \frac{2\pi}{1 + e^{-w}} \quad (4)$$

In the above equations, we take the coordinates η and ϕ relative to the jet axis. While LHC collisions can result in hundreds of final-state particles, we limit our studies to the intrinsic structure of jets rather than entire events due to the limitations of currently available quantum computers and their simulators. Studying the substructure of jets can be a very powerful way to disentangle boosted, hadronically decaying electroweak-scale or beyond-the-standard-model resonances from QCD-induced jets [27]. Furthermore, in order to allow for efficient simulation of our quantum circuits on a simulator, we limit our study by using up to ten hardest jet constituents, which are expected to carry most of the information relevant to jet substructure analysis. Figure 1 shows an example of the encoding for the ten hardest particles in a jet from a top quark decay (red) and in a QCD light-flavor jet (blue). Each particle is encoded on a different qubit. As can be seen, the encoding is rooted in the kinematic structure of jets, mapping angular coordinates and transverse momentum to the Bloch sphere to preserve geometric and energy correlations among constituents. The scaling function f in Eq. 4 increases the angles on the Bloch sphere.

Anomaly detection and classification with 1P1Q. Anomaly detection has emerged as a pivotal tool in model-agnostic searches for new physics, where the only assumption is that deviations from Standard Model signatures are present [28]. Using classical machine learning, unsupervised learning models based on autoencoders are employed as efficient tools for anomaly detection [29–31]. These work by compressing the input data into a lower-dimensional latent space and attempting to reconstruct the input, highlighting anomalous events through elevated reconstruction errors.

Quantum autoencoders extend this paradigm to quantum devices [16]. A QAE consists of an encoder and a




FIG. 1: Bloch Sphere: Effect of the input scaling described in Eq. 4 when applied to the ten hardest particles of a QCD and top jet with comparable jet p_T . The value of w converges to a scaling factor $f = 7.268$ for the VQC.

decoder implemented via variational quantum circuits. The encoder compresses the input quantum state into a smaller latent representation, discarding qubits as necessary. The decoder then reconstructs the input from this latent state using the Hermitian conjugate of the encoding operators. To achieve dimensionality reduction in the latent space, QAEs replace a certain number of qubits (hereafter referred to as N_{trash} trash states) with $N_{\text{ref}} = N_{\text{trash}}$ reference states that are initialized to $|0\rangle$, thus creating an information bottleneck. The unitary transformations within the QAE consist of a combination of parameterized single-qubit rotations and multi-qubit entangling gates, such as the Controlled NOT (CNOT) gate. Using specific examples, QAEs have shown first successes compared to classical AEs regarding training efficiency and performance (e.g., [16]).

To demonstrate the applicability of 1P1Q in a QAE, we use Eqs. 1 and 2 to encode information into a QAE with an architecture following the left circuit of Fig. 2. For illustration purposes, the figure shows a 4-particle-input example, while the QAEs trained are using 6, 8, and 10 particles as input, respectively. To allow the network to learn higher-order non-linear terms, we entangle all possible pairs of qubits using two-qubit CNOT gates. To ensure the network learns an optimal representation of the input space, we apply three parameterized rotations, one along each axis of each qubit. The entanglement and rotation operations are summarized in Eq. 5.

$$U(\Theta) = \left(\bigotimes_{i=1}^N R_X(\phi_i) R_Y(\theta_i) R_Z(\omega_i) \right) \otimes \left(\bigotimes_{1 \leq i < j \leq N} C_{ij} \right). \quad (5)$$

The QAE is trained to reconstruct jets initiated by light quarks or gluons. We expect this reconstruction quality, defined as the fidelity between the trash and reference states, to differ between signal and background, thus pro-

viding a degree of separation.

In addition, we demonstrate that the 1P1Q encoding can just as well be used for classification within a supervised learning paradigm. For this task, we use a VQC [17] that learns to separate signal and background classes, in this case jets from hadronically decaying top quarks versus jets initiated by light quarks or gluons.

The VQC architecture is similar to that of the QAE, with features being first encoded into the circuit using the procedure outlined in Eqs. 1-4, followed by entanglement operations using two-qubit CNOT gates to allow for the introduction of non-linear terms into the network. Only adjacent qubits are entangled for the VQC, which we find gives the best performance. Three trainable rotation gates are applied to each qubit, as summarized in Eq. 6. Unlike the QAE, which is optimized by maximizing the fidelity between the trash and reference states, the VQC is optimized by performing a measurement on the first qubit which is bound in $[-1, 1]$ and comparing it to the ground truth label. The VQC circuit is shown in Fig. 2 for an input of $N = 8$ particles.

$$U(\Theta) = \left(\bigotimes_{i=1}^N R_X(\phi_i) R_Y(\theta_i) R_Z(\omega_i) \right) \otimes \left(\bigotimes_{i=1}^N C_{i, (i+1) \bmod N} \right) \quad (6)$$

We highlight that for an input space comprising of the N hardest input particles per jet, the QAE requires only $3N + 1$ trainable parameters to learn a suitable reconstruction of the inputs. Likewise, the VQC requires only $3N + 2$ parameters (the extra parameter of the VQC being a trainable bias term that is added to the measurement).

Results and Benchmarking. We train the QAE on two different sets: simulated QCD jets from the Jet-




FIG. 2: QAE Circuit (left) used for anomaly detection. VQC (right) used for supervised classification. Example circuits with 4 and 8 input particles, respectively.




FIG. 3: Quantum fidelity distributions for a system with 10 input qubits and a latent space of 2 qubits.

Class dataset [32], and real jet data recorded with the CMS Detector in 2016 [33], which is >99% pure in QCD jets. For the first time, to our knowledge, Fig. 3 shows that a QAE trained on real experimental data is comparable in performance to one trained on simulated events, when applied to the task of differentiating between simulated QCD events and simulated signal models. This suggests that a QAE using 1P1Q encoding, despite not using high-level information, can extract the underlying physics of the jet based on the individual particles encoded onto qubits and is neither corrupted nor biased by the detector response.

Figure 4 illustrates the AUC scores as a function of the quantum fidelity metric $\langle 1 - \text{Fidelity} \rangle_{\text{QCD}}$ for different QAE configurations trained on either simulated JetClass QCD jets (dashed lines) or real CMS data (dot-

ted lines). The three panels correspond to different input sizes, namely 6, 8, and 10 jet constituents, respectively, while different marker shapes indicate varying latent space sizes. Across the various configurations, we observe a trend that higher AUC scores correlate with increased $\langle 1 - \text{Fidelity} \rangle_{\text{QCD}}$, indicating that the fidelity loss serves as a useful proxy for anomaly detection performance. The discrimination is particularly strong for top quark jets ($t \rightarrow bq\bar{q}'$, red), which consistently achieve the highest AUC scores, followed by Higgs decays to bottom quarks ($H \rightarrow b\bar{b}$, blue). The similarity in trends between models trained on real and simulated data highlights the robustness of the QAE in learning fundamental jet substructure features independent of dataset origin. Additionally, larger input sizes tend to improve performance, suggesting that incorporating more jet constituents provides richer representations for classification.

To benchmark the performance of the 1P1Q-encoded QAE against a classical counterpart, we trained on simulated QCD jets and considered as anomalies to this background the signals of hadronic W boson, Higgs boson and top quark decays. While the QAE has a simple structure of 10 input qubits, followed by the circuit of Fig. 2 (left) and a latent space of 2, the classical autoencoder (CAE) model is allowed to be significantly larger, consisting of an encoder model containing an input feature vector of size 30, to be able to encode the same number of features encoded on the QAE, followed by dense layers of size $20 - 16 - 12$ and a latent space 6. The decoder of the CAE has an identical, yet inverse, structure.

The QAE achieves superior performance compared to the CAE across all signal types. The QAE maintains this advantage despite having only 31 trainable parameters, in contrast to the CAE's $\sim 2,500$ parameters, demonstrating the potential of quantum machine learning to capture relevant physics with a significantly reduced model complexity efficiently.

For the supervised classification task, we employ the VQC to distinguish jets originating from top quark decays ($t \rightarrow bq\bar{q}'$) from those initiated by light quarks or gluons. Figure 5 presents the Receiver Operating Char-

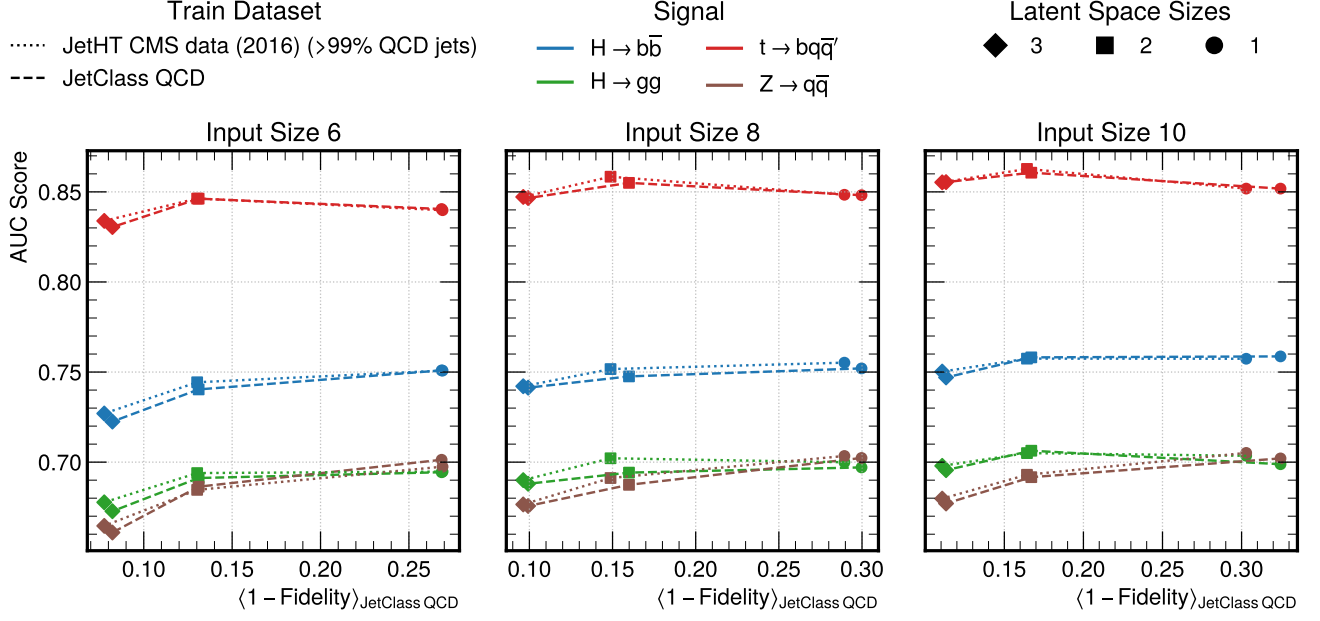


FIG. 4: AUC scores vs. $\langle 1 - \text{Fidelity} \rangle_{\text{QCD}}$ for different QAE configurations. Models trained on simulated (dashed) and real CMS data (dotted) show consistent trends, with larger input sizes and higher fidelity loss correlating with improved anomaly detection performance.

TABLE I: AUC scores for QAE vs CAE, trained on a 10-particle input space.

Algorithm \ Signals	$W \rightarrow q\bar{q}'$	$H \rightarrow b\bar{b}$	$t \rightarrow b\bar{q}q'$
QAE	0.715	0.774	0.872
CAE	0.671	0.739	0.858

acteristic (ROC) curves comparing the VQC trained on the 1P1Q-encoded dataset against the state-of-the-art Particle Transformer (ParT) classifier [34], both trained using the 10 hardest particles per jet, and the same dataset of 1000 jets. The ROC curve demonstrates that the VQC achieves strong classification performance, with an AUC score of 0.885. Although the ParT model achieves a slightly higher AUC of 0.898, the VQC remains competitive despite having significantly fewer trainable parameters (32 vs. over 2 million). Furthermore, for high signal efficiencies ($>80\%$), the VQC yields better background rejection than the Particle Transformer.

This result shows the efficiency and potential of quantum machine learning models in jet classification tasks, especially when using the 1P1Q particle encoding. Such a quantum machine learning method offers competitive performance to classical state-of-the-art frameworks with drastically reduced computational resources.

1P1Q encoding. The expressivity and efficiency of a quantum machine learning approach are jointly defined both by the model (e.g., QAE or VQC) and the encoding step that maps classical data onto a quantum device. Thus, to assess the models' ability to utilize the informa-

tion provided by the 1P1Q data encoding, we successively reduce the features encoded on each qubit and observe the resulting performance degradation.

TABLE II: AUC scores vs input features for the VQC and QAE, for the benchmark signal $t \rightarrow b\bar{q}q'$.

Algorithm \ Inputs	(p_T, η, ϕ)	(p_T, η)	(η, ϕ)	(p_T, ϕ)
VQC	0.886	0.856	0.808	0.857
QAE	0.872	0.825	0.823	0.827

Table II shows the impact of different input feature combinations within the 1P1Q encoding scheme on classification performance. The best AUC scores are achieved when all three features (p_T, η, ϕ) are included, suggesting that the full kinematic variables contribute significantly to the model's performance. When one feature is removed, we observe a performance drop, with the most pronounced reduction occurring when p_T is excluded. This emphasizes the importance of transverse momentum information in jet classification tasks and suggests that quantum models are particularly effective at utilizing correlations between momentum and angular variables. The robustness of these results across different models further validates the suitability of 1P1Q encoding for high-energy physics applications.

Summary. In this work, we introduced the 1P1Q encoding scheme, a novel approach for representing particle physics data on quantum hardware by assigning each

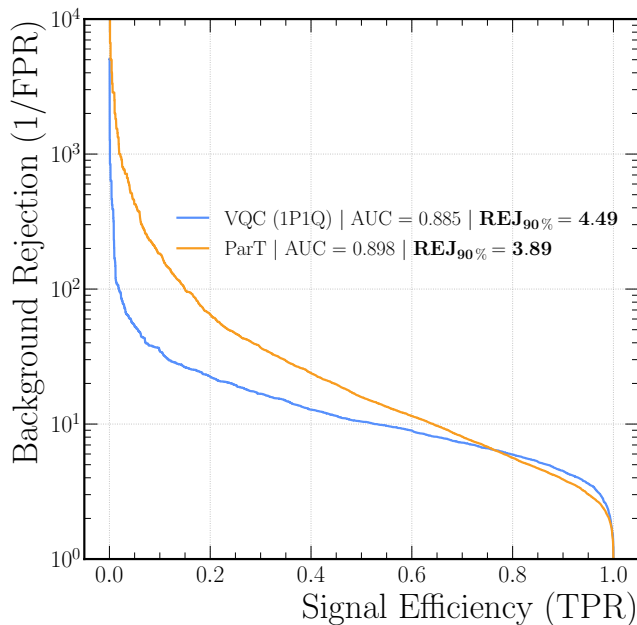


FIG. 5: Comparison of ROC curves for the VQC and Particle Transformer. The VQC trained on the 1P1Q-encoded dataset closely matches the performance of the state-of-the-art Particle Transformer, despite using significantly fewer trainable parameters, with both being trained on the same number of events and input size of 10 particles.

particle to a separate qubit. This encoding allows direct utilization of raw collision event data without classical compression, enhancing the potential expressivity and efficiency of QML models. We demonstrated the effectiveness of this approach using two quantum machine learning models: a QAE for anomaly detection and a VQC for supervised classification.

Our results show that the QAE successfully differentiates signal jets from background QCD jets, achieving superior performance compared to its classical counterpart while requiring significantly fewer trainable parameters. Furthermore, the VQC exhibits strong classification capability, approaching the performance of state-of-the-art classical models despite its minimal parameter count. By systematically reducing the encoded features, we established that the 1P1Q encoding effectively captures jet substructure information, and performance degradation with reduced feature input underscores the importance of a comprehensive quantum data representation.

For the first time, we validate a quantum machine learning model trained on real experimental data from the CMS detector, demonstrating that quantum approaches can extract meaningful physics insights in a real-world setting. These robust results establish the 1P1Q encoding as a viable and scalable data representation framework for quantum computing applications in particle physics. As quantum hardware continues to advance, the efficiency of the 1P1Q approach is tailored for

more intricate and large-scale QML applications in high-energy physics, offering new perspectives for jet classification, anomaly detection, and measurements revealing the quantum mechanical nature of particle production.

Acknowledgements. The authors acknowledge the support of Schmidt Sciences, the Alexander von Humboldt Foundation, and IPPP, as well as the usage of computing resources on the TOPAS GPU cluster at the Scientific Computing Centre (SCC) KIT, and the RCS Cluster at Imperial College London. The authors also thank the NHR Center Karlsruhe for providing access to computing resources on the high-performance computer HoreKa, which is jointly supported by the German Federal Ministry of Education and Research (BMBF) and the state governments.

APPENDIX

Training setup

The quantum circuits presented in this study are simulated and optimized using the Quantum Machine Learning library `pennylane` [35] with the `lightning.gpu` and `lightning.kokkos` devices. The trainable circuit parameters are optimized using the classical Adam Optimizer [36], with a scheduler that periodically decays the learning rate.

To train the QAE, we seek to maximize the fidelity between the output and input states in the subspace relevant for reconstruction:

$$F = \sum_{i=1}^{N_{\text{ref}}} \langle T_i | R_i \rangle \quad \forall \quad T_i \in H^T, R_i \in H^{\text{ref}} \quad (7)$$

Specifically, following the approach first introduced in [37], we minimize the cost function defined as the negative of the fidelity. This measurement is performed using a SWAP test [38], which requires an ancillary qubit initialized to $|0\rangle$.

In the case of the VQC, we use the expectation value of the Pauli Z observable on the target, which is the first qubit in the circuit. The final VQC prediction (Eq. 8) is arrived at by adding a classical, trainable bias term b to the expectation value:

$$f(x) = \langle q(x) | Z | q(x) \rangle + b \quad (8)$$

$$(9)$$

The circuit optimization is performed using the Mean Squared Error (MSE) between the prediction and the truth label as the loss function.

Datasets

The JETCLASS dataset [34] is more recent and far larger in scope and size compared to any of its predecessor

datasets, and is therefore used to train and benchmark the 1P1Q approach. It contains a total of 100M jets divided into 10 classes. The background jets are those initiated by light quarks or gluons, while jets arising from decays such as, but not limited to, $t \rightarrow bq\bar{q}'$ and $W/Z \rightarrow q\bar{q}'$ are treated as signal jets. The production and subsequent decay of top quarks and the W, Z, and Higgs bosons are simulated with `Madgraph5_aMC@NLO` at next-to-leading order precision [39]. Parton showering and hadronisation processes are simulated in `PYTHIA8` [40].

To ensure the simulated jets closely resemble those reconstructed by the CMS detector, the detector effects are simulated with `DELPHES` [41] using a simplified CMS detector configuration. Jets are clustered utilizing the anti- k_T algorithm [42] with a distance parameter $R = 0.8$. As additional criteria, only jets with $p_T \in [500, 1000]$ GeV and pseudorapidity $|\eta| < 2$ are stored. For signal jets, additional quality requirements are imposed to ensure they fully contain the decay products of the initial particles.

The Aspen Open Jets (AOJ) dataset [33] is derived from the CMS 2016 JetHT datasets [43, 44] and presented in a structured format specifically optimised for machine learning applications. Although an extensive

description of CMS data acquisition and processing falls beyond the scope of this letter, it is important to note that this dataset predominantly comprises jets initiated by light quarks or gluons, with contamination from alternative decay processes such as those involving W/Z bosons or top quarks constituting less than 1% of the total sample size. In its entirety, the dataset contains approximately 180M jets recorded in 2016 with the CMS detector.

For training the QAE, we use a total of 10,000 events for training and 2,500 events for validation. Inference is performed on 10,000 events of background and of each of the following class of events: $H \rightarrow b\bar{b}$, $H \rightarrow c\bar{c}$, $H \rightarrow g\bar{g}$, $W \rightarrow q\bar{q}'$, $Z \rightarrow q\bar{q}$ and $t \rightarrow bq\bar{q}'$. The VQC is trained using 1,000 events and validated on 500 events, with an inference dataset of 10,000 events equally distributed between signal and background classes.

For training the QAE and VQC respectively, we first sample jets such that each class has a flat distribution in jet p_T , in the range $[500, 1000]$ GeV, so as to not bias the training towards the scale of the jet, allowing us to purely focus on jet substructure. The effect of this sampling can be seen in Figure 6. The jet samples used in the inference are also sampled to have the same flat distribution.

-
- [1] W. Guan, G. Perdue, A. Pesah, M. Schuld, K. Terashi, S. Vallecorsa, and J.-R. Vlimant, Quantum Machine Learning in High Energy Physics, *Mach. Learn. Sci. Tech.* **2**, 011003 (2021), arXiv:2005.08582 [quant-ph].
 - [2] A. Di Meglio *et al.*, Quantum Computing for High-Energy Physics: State of the Art and Challenges, *PRX Quantum* **5**, 037001 (2024), arXiv:2307.03236 [quant-ph].
 - [3] C. Brown, M. Spannowsky, A. Tapper, S. Williams, and I. Xiotidis, Quantum pathways for charged track finding in high-energy collisions, *Front. Artif. Intell.* **7**, 1339785 (2024), arXiv:2311.00766 [hep-ph].
 - [4] S. Abel, J. C. Criado, and M. Spannowsky, Completely quantum neural networks, *Phys. Rev. A* **106**, 022601 (2022), arXiv:2202.11727 [quant-ph].
 - [5] S. Abel, A. Blance, and M. Spannowsky, Quantum optimization of complex systems with a quantum annealer, *Phys. Rev. A* **106**, 042607 (2022), arXiv:2105.13945 [quant-ph].
 - [6] A. Delgado and K. E. Hamilton, Unsupervised quantum circuit learning in high energy physics, *Phys. Rev. D* **106**, 096006 (2022), arXiv:2203.03578 [quant-ph].
 - [7] S. Alvi, C. W. Bauer, and B. Nachman, Quantum anomaly detection for collider physics, *JHEP* **02**, 220, arXiv:2206.08391 [hep-ph].
 - [8] J. Y. Araz and M. Spannowsky, Classical versus quantum: Comparing tensor-network-based quantum circuits on Large Hadron Collider data, *Phys. Rev. A* **106**, 062423 (2022), arXiv:2202.10471 [quant-ph].
 - [9] A. Hammad, K. Kong, M. Park, and S. Shim, Quantum Metric Learning for New Physics Searches at the LHC, arXiv preprint (2023), arXiv:2311.16866 [hep-ph].
 - [10] J. Schuhmacher, L. Boggia, V. Belis, E. Puljak, M. Grossi, M. Pierini, S. Vallecorsa, F. Tacchino,




FIG. 6: Effect of the p_T sampling on the input datasets used for training

- P. Barkoutsos, and I. Tavernelli, Unravelling physics beyond the standard model with classical and quantum anomaly detection, *Mach. Learn. Sci. Tech.* **4**, 045031 (2023), arXiv:2301.10787 [hep-ex].
- [11] Y.-A. Chen and K.-F. Chen, Jet discrimination with a quantum complete graph neural network, *Phys. Rev. D* **111**, 016020 (2025), arXiv:2403.04990 [hep-ph].
- [12] J.-C. Yang, S. Zhang, and C.-X. Yue, A novel quantum machine learning classifier to search for new physics, arXiv preprint (2024), arXiv:2410.18847 [hep-ph].

- [13] J. L. Scott, Z. Dong, T. Kim, K. Kong, and M. Park, Hybrid quantum-classical approach for combinatorial problems at hadron colliders, arXiv preprint (2024), arXiv:2410.22417 [hep-ph].
- [14] G. Aad *et al.* (ATLAS), Observation of quantum entanglement with top quarks at the ATLAS detector, *Nature* **633**, 542 (2024), arXiv:2311.07288 [hep-ex].
- [15] A. Hayrapetyan *et al.* (CMS), Observation of quantum entanglement in top quark pair production in proton–proton collisions at $\sqrt{s} = 13$ TeV, *Rept. Prog. Phys.* **87**, 117801 (2024), arXiv:2406.03976 [hep-ex].
- [16] V. S. Ngairangbam, M. Spannowsky, and M. Takeuchi, Anomaly detection in high-energy physics using a quantum autoencoder, *Phys. Rev. D* **105**, 095004 (2022), arXiv:2112.04958 [hep-ph].
- [17] A. Blance and M. Spannowsky, Quantum Machine Learning for Particle Physics using a Variational Quantum Classifier, *JHEP* **02**, 212, arXiv:2010.07335 [hep-ph].
- [18] A. Blance and M. Spannowsky, Unsupervised event classification with graphs on classical and photonic quantum computers, *JHEP* **21**, 170, arXiv:2103.03897 [hep-ph].
- [19] S. Y.-C. Chen, T.-C. Wei, C. Zhang, H. Yu, and S. Yoo, Hybrid Quantum-Classical Graph Convolutional Network, arXiv preprint (2021), arXiv:2101.06189 [cs.LG].
- [20] A. Gianelle, P. Koppenburg, D. Lucchesi, D. Nicotra, E. Rodrigues, L. Sestini, J. de Vries, and D. Zuliani, Quantum Machine Learning for b-jet charge identification, *JHEP* **08**, 014, arXiv:2202.13943 [hep-ex].
- [21] M. C. Peixoto, N. F. Castro, M. Crispim Romão, M. G. J. a. Oliveira, and I. Ochoa, Fitting a collider in a quantum computer: tackling the challenges of quantum machine learning for big datasets, *Front. Artif. Intell.* **6**, 1268852 (2023), arXiv:2211.03233 [hep-ph].
- [22] V. Belis, K. A. Woźniak, E. Puljak, P. Barkoutsos, G. Dissertori, M. Grossi, M. Pierini, F. Reiter, I. Tavernelli, and S. Vallecorsa, Quantum anomaly detection in the latent space of proton collision events at the LHC, *Commun. Phys.* **7**, 334 (2024), arXiv:2301.10780 [quant-ph].
- [23] C. Duffy, M. Hassanshah, M. Jastrzebski, and S. Malik, Unsupervised Beyond-Standard-Model Event Discovery at the LHC with a Novel Quantum Autoencoder (2024), arXiv:2407.07961 [quant-ph].
- [24] A. Hammad, M. M. Nojiri, and M. Yamazaki, Quantum similarity learning for anomaly detection (2024), arXiv:2411.09927 [hep-ph].
- [25] V. Belis, P. Odagiu, M. Grossi, F. Reiter, G. Dissertori, and S. Vallecorsa, Guided quantum compression for high dimensional data classification, *Mach. Learn. Sci. Tech.* **5**, 035010 (2024), arXiv:2402.09524 [quant-ph].
- [26] S. Chatrchyan *et al.* (CMS), The CMS Experiment at the CERN LHC, *JINST* **3**, S08004.
- [27] S. Marzani, G. Soyez, and M. Spannowsky, *Looking inside jets: an introduction to jet substructure and boosted-object phenomenology*, Vol. 958 (Springer, 2019) arXiv:1901.10342 [hep-ph].
- [28] V. Belis, P. Odagiu, and T. K. Aarrestad, Machine learning for anomaly detection in particle physics, *Rev. Phys.* **12**, 100091 (2024), arXiv:2312.14190 [physics.data-an].
- [29] G. Aad *et al.* (ATLAS), Anomaly detection search for new resonances decaying into a Higgs boson and a generic new particle X in hadronic final states using $\sqrt{s} = 13$ TeV pp collisions with the ATLAS detector, *Phys. Rev. D* **108**, 052009 (2023), arXiv:2306.03637 [hep-ex].
- [30] G. Aad *et al.* (ATLAS), Search for New Phenomena in Two-Body Invariant Mass Distributions Using Unsupervised Machine Learning for Anomaly Detection at $s=13$ TeV with the ATLAS Detector, *Phys. Rev. Lett.* **132**, 081801 (2024), arXiv:2307.01612 [hep-ex].
- [31] The CMS Collaboration, Model-agnostic search for di-jet resonances with anomalous jet substructure in proton–proton collisions at $\sqrt{s} = 13$ TeV, *Reports on Progress in Physics* **88**, 067802 (2025).
- [32] H. Qu, C. Li, and S. Qian, Jetclass: A large-scale dataset for deep learning in jet physics, 10.5281/zenodo.6619768 (2022).
- [33] O. Amram, L. Anzalone, J. Birk, D. A. Faroughy, A. Hallin, G. Kasieczka, M. Krämer, I. Pang, H. Reyes-Gonzalez, and D. Shih, Aspen open jets: unlocking lhc data for foundation models in particle physics, *Machine Learning: Science and Technology* **6**, 030601 (2025).
- [34] H. Qu, C. Li, and S. Qian, Particle transformer for jet tagging, in *Proceedings of the 39th International Conference on Machine Learning*, Proceedings of Machine Learning Research, Vol. 162, edited by K. Chaudhuri, S. Jegelka, L. Song, C. Szepesvari, G. Niu, and S. Sabato (PMLR, 2022) pp. 18281–18292.
- [35] V. Bergholm, J. Izaac, M. Schuld, C. Gogolin, S. Ahmed, V. Ajith, M. S. Alam, G. Alonso-Linaje, B. Akash-Narayanan, A. Asadi, J. M. Arrazola, U. Azad, S. Banning, C. Blank, T. R. Bromley, B. A. Cordier, J. Ceroni, A. Delgado, O. D. Matteo, A. Dusko, T. Garg, D. Guala, A. Hayes, R. Hill, A. Ijaz, T. Isacsson, D. Ittah, S. Jahangiri, P. Jain, E. Jiang, A. Khandelwal, K. Kottmann, R. A. Lang, C. Lee, T. Loke, A. Lowe, K. McKiernan, J. J. Meyer, J. A. Montañez-Barrera, R. Moyard, Z. Niu, L. J. O’Riordan, S. Oud, A. Panigrahi, C.-Y. Park, D. Polatajko, N. Quesada, C. Roberts, N. Sá, I. Schoch, B. Shi, S. Shu, S. Sim, A. Singh, I. Strandberg, J. Soni, A. Száva, S. Thabet, R. A. Vargas-Hernández, T. Vincent, N. Vitucci, M. Weber, D. Wierichs, R. Wiersema, M. Willmann, V. Wong, S. Zhang, and N. Killoran, PennyLane: Automatic differentiation of hybrid quantum-classical computations (2022), arXiv:1811.04968 [quant-ph].
- [36] D. P. Kingma and J. Ba, Adam: A method for stochastic optimization, in *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, edited by Y. Bengio and Y. LeCun (2015).
- [37] J. Romero, J. P. Olson, and A. Aspuru-Guzik, Quantum autoencoders for efficient compression of quantum data, *Quantum Science and Technology* **2**, 045001 (2017).
- [38] H. Buhrman, R. Cleve, J. Watrous, and R. de Wolf, Quantum fingerprinting, *Phys. Rev. Lett.* **87**, 167902 (2001).
- [39] J. Alwall, R. Frederix, S. Frixione, V. Hirschi, F. Maltoni, O. Mattelaer, H.-S. Shao, T. Stelzer, P. Torrielli, and M. Zaro, The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations, *Journal of High Energy Physics* **2014**, 10.1007/jhep07(2014)079 (2014).
- [40] T. Sjöstrand, S. Ask, J. R. Christiansen, R. Corke, N. Desai, P. Ilten, S. Mrenna, S. Prestel, C. Rasmussen, and P. Z. Skands, An introduction to PYTHIA 8.2, *Comput. Phys. Commun.* **191** (2014).
- [41] S. Ovin, X. Rouby, and V. Lemaitre, Delphes, a framework for fast simulation of a generic collider experiment

- (2010), arXiv:0903.2225 [hep-ph].
- [42] M. Cacciari, G. P. Salam, and G. Soyez, The anti- k_t jet clustering algorithm, *Journal of High Energy Physics* **04**, 063 (2008).
 - [43] CMS Collaboration, JetHT primary dataset in MINIAOD format from Run G of 2016 (/JetHT/Run2016G-UL2016-MiniAODv2-v2/MINIAOD), CERN Open Data Portal (2024), cite as: CMS Collaboration (2024).
 - [44] CMS Collaboration, JetHT primary dataset in MINIAOD format from Run H of 2016 (/JetHT/Run2016H-UL2016-MiniAODv2-v2/MINIAOD), CERN Open Data Portal (2024), cite as: CMS Collaboration (2024).