



PAPER

Guided quantum compression for high dimensional data classification

OPEN ACCESS

RECEIVED

16 February 2024

REVISED

12 June 2024

ACCEPTED FOR PUBLICATION

5 July 2024

PUBLISHED

16 July 2024

Original Content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](#).

Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

Vasilis Belis^{1,*} , Patrick Odagiu¹ , Michele Grossi² , Florentin Reiter³, Günther Dissertori¹ and Sofia Vallecorsa² ¹ Institute for Particle Physics and Astrophysics, ETH Zurich, 8093 Zurich, Switzerland² European Organization for Nuclear Research (CERN), CH-1211 Geneva, Switzerland³ Institute for Quantum Electronics, ETH Zurich, 8093 Zurich, Switzerland

* Author to whom any correspondence should be addressed.

E-mail: vbelis@phys.ethz.ch**Keywords:** quantum machine learning, deep learning, particle physics, feature extraction, quantum computing**Abstract**

Quantum machine learning provides a fundamentally different approach to analyzing data. However, many interesting datasets are too complex for currently available quantum computers. Present quantum machine learning applications usually diminish this complexity by reducing the dimensionality of the data, e.g. via auto-encoders, before passing it through the quantum models. Here, we design a classical-quantum paradigm that unifies the dimensionality reduction task with a quantum classification model into a single architecture: the guided quantum compression model. We exemplify how this architecture outperforms conventional quantum machine learning approaches on a challenging binary classification problem: identifying the Higgs boson in proton-proton collisions at the LHC. Furthermore, the guided quantum compression model shows better performance compared to the deep learning benchmark when using solely the kinematic variables in our dataset.

1. Introduction

Machine Learning (ML) is established as an invaluable tool for analysing data and assisting many physics analyses at the Large Hadron Collider (LHC) [1–4]. Meanwhile, quantum computing is a fundamentally different paradigm for information processing, that is known to provide computational speed-ups over classical methods for a large class of problems [5–11]. Furthermore, Quantum Machine Learning (QML) has the potential to enhance traditional ML methods [12–17] and yields various advantages in specific learning tasks [18–24]. Recent studies have highlighted guarantees regarding the expressivity, generalisation power, and trainability of quantum models [25–30]. Moreover, the efficacy of applying QML models to High Energy Physics (HEP) data analysis is exemplified in studies for classification [31–36], reconstruction [37–39], anomaly detection [40–43], and Monte Carlo integration [44, 45]. A summary of advancements in QML applied to HEP is found in [46].

However, for most realistic applications, the dimensionality of the dataset is usually too large to be directly processed by commonly available quantum computers. Consequently, dimensionality reduction techniques are typically employed and treated as a preprocessing step before the data is loaded into the QML algorithm. Previous studies use manual feature selection, informed by prior knowledge about the given problem [33, 34], feature extraction techniques, such as the popular Principal Component Analysis [31, 32, 41], or more recently, dimensionality reduction using deep learning models, e.g. simple auto-encoders [34, 43, 47]. However, there is no guarantee or even incentive for the lower-dimensional representation produced by these methods to preserve original data structures that are relevant for the QML task, e.g. binary classification.

Crucial information required in discriminating between the classes can be lost in the dimensionality reduction. In extreme cases, the dimensionality reduction performed as a preprocessing step can render impossible the Higgs classification task to be solved by the QML model. To address this challenge,

we draw inspiration from the multi-task learning literature [48, 49] and design a model architecture that generates lower dimensional representations which are suitable for classification by QML models. We call this approach *guided quantum compression*.

In contrast with conventional techniques, the dimensionality reduction performed within the guided quantum compression framework is not treated as a separate preprocessing step. For example, the usual approach is to train an auto-encoder on the data, obtain a lower dimensional representation from its latent space, and then train a QML algorithm on this latter representation [34]. Instead, guided quantum compression *simultaneously* accomplishes the dimensionality reduction and classification tasks with a single architecture. This way, the performance of the QML is not limited by the arbitrary choice of the reduction method, be it manual, classical feature extraction, or deep learning.

We demonstrate this claim by considering a realistic and complex classification task: identifying the Higgs boson in the $t\bar{t}H(b\bar{b})$ semi-leptonic channel for simulated proton collision data at the LHC. On this data, the conventional reduction methods mentioned earlier fail relative to the guided quantum compression method, i.e. independently compressing the dataset before training the QML model leads to poor classification performance. In contrast, the guided quantum compression method is able to solve the classification problem, reaching competitive accuracy with state-of-the-art classical methods [50]. Furthermore, we observe an improved performance of our algorithm compared to the classical benchmark when using only the particle kinematics, suggesting that for QML one should use features representative of the quantum process that generated the data [41, 43, 51, 52].

2. Models

The guided quantum compression network shown in figure 1 is comprised of an auto-encoder (figure 1(a)) and a variational quantum circuit (VQC) (figure 1(b)) that are coupled. In the following, we describe these two elements and how they are simultaneously trained.

2.1. The auto-encoder

The auto-encoder (AE) is a machine learning model that has been used for decades across the historical landscape of neural networks [47, 53, 54]. The most commonly used type of AE consists of two feed-forward neural networks: the encoder \mathcal{E}_ω and the decoder \mathcal{D}_ρ . The encoder \mathcal{E}_ω maps the input feature space x to a *latent space* z of lower dimension ℓ . Conversely, the goal of the decoder \mathcal{D}_ρ is to reconstruct the input x from z . A schematic of an AE neural network is shown in figure 1(a). The objective of the AE training is to minimise the difference between the input data and the reconstructed data; this difference can be quantified by various functions. There exist many types of AEs, based on how this difference is quantified or on architectural extensions [55–58]. The mean squared error (MSE) is the standard function used to quantify the difference between the input data x and its reconstructed counterpart:

$$\mathcal{L}_R = \frac{1}{M} \sum_{m=1}^M [x_m - \mathcal{D}_\rho \circ \mathcal{E}_\omega(x_m)]^2, \quad (1)$$

where \mathcal{D}_ρ is the decoder network with weights ρ , \mathcal{E}_ω is the encoder network with weights ω , and M is the size of the training dataset. The conventional AE learns as any other feed-forward neural network, with one subtlety: the reconstruction loss is not only propagated through the decoder network, but through the encoder as well. Therefore, the latent space and the reconstructed data evolve simultaneously as the AE model is learning.

2.2. The VQC classifier

The VQC [14, 59] is a QML model based on parametrized quantum circuits that are trained variationally to undertake tasks such as classification [32–35], regression [60], and generative modelling [61–64]. The classifier output, in the VQC implementation, is the expectation of an observable \mathcal{M} . This expectation is interpreted as the likelihood of the input sample to belong in a certain data class, e.g. the sample contains a Higgs boson or not, and is extracted from the quantum computer via doing measurements. Specifically, the model output $f_\vartheta(z)$, for a given input data vector z and gate parameters ϑ , is defined as

$$f_\vartheta(z) = \langle 0 | \mathcal{U}_\vartheta^\dagger(z) \mathcal{M} \mathcal{U}_\vartheta(z) | 0 \rangle = \langle \mathcal{M}_\vartheta(z) \rangle, \quad (2)$$

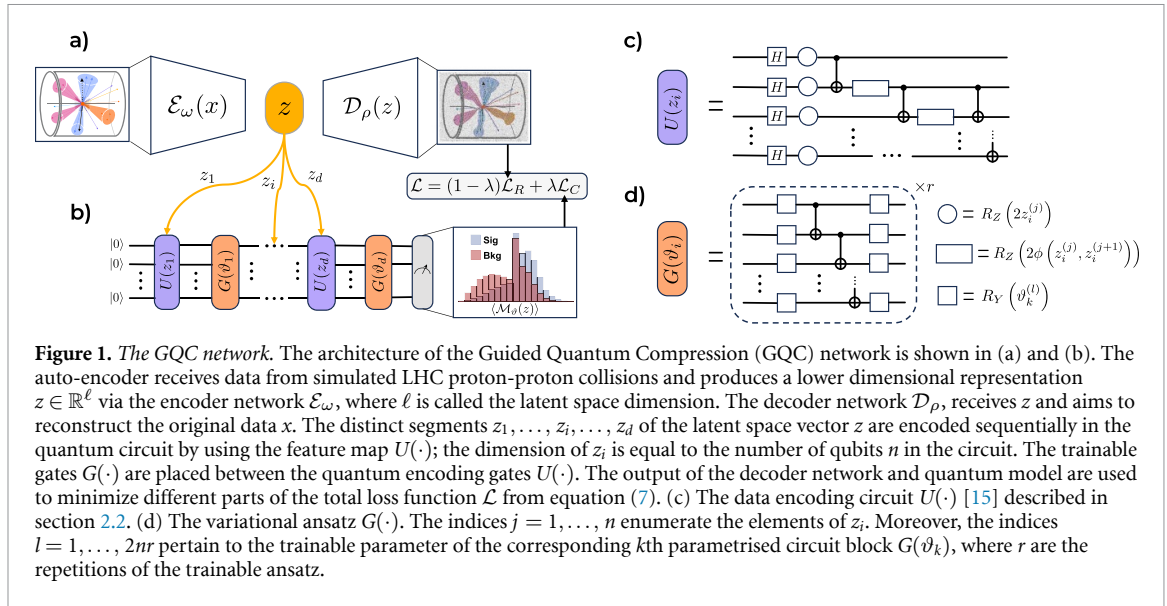


Figure 1. The GQC network. The architecture of the Guided Quantum Compression (GQC) network is shown in (a) and (b). The auto-encoder receives data from simulated LHC proton-proton collisions and produces a lower dimensional representation $z \in \mathbb{R}^\ell$ via the encoder network \mathcal{E}_ω , where ℓ is called the latent space dimension. The decoder network \mathcal{D}_ρ , receives z and aims to reconstruct the original data x . The distinct segments $z_1, \dots, z_i, \dots, z_d$ of the latent space vector z are encoded sequentially in the quantum circuit by using the feature map $U(\cdot)$; the dimension of z_i is equal to the number of qubits n in the circuit. The trainable gates $G(\cdot)$ are placed between the quantum encoding gates $U(\cdot)$. The output of the decoder network and quantum model are used to minimize different parts of the total loss function \mathcal{L} from equation (7). (c) The data encoding circuit $U(\cdot)$ [15] described in section 2.2. (d) The variational ansatz $G(\cdot)$. The indices $j = 1, \dots, n$ enumerate the elements of z_i . Moreover, the indices $l = 1, \dots, 2nr$ pertain to the trainable parameter of the corresponding k th parametrised circuit block $G(\vartheta_k)$, where r are the repetitions of the trainable ansatz.

where $\mathcal{U}_\vartheta(z)$ is the whole quantum circuit of the model, \mathcal{M} is the observable whose expectation value we measure, and $|0\rangle = |0\rangle^{\otimes n}$ the initial n qubit state. Furthermore, the label predicted by the VQC, $\hat{y} \in \{0, 1\}$, is given by

$$\hat{y} = \frac{\text{sign}[\langle \mathcal{M}_\vartheta(z) \rangle] + 1}{2}, \quad (3)$$

where one can assume that $\langle \mathcal{M}_\vartheta(z) \rangle \in [-1, 1]$ without any loss of generality.

We design a VQC architecture as shown in figure 1(b). The input data vector z is split into d sub-vectors $z = (z_1, z_2, \dots, z_d)$, each z_i being of dimensionality n . Furthermore, each sub-vector is encoded into the quantum circuit sequentially by using the feature map $U(\cdot)$; between each $U(\cdot)$ there is a set of trainable gates $G(\cdot)$. Hence, the whole quantum circuit of the model is

$$\mathcal{U}_\vartheta(z) = \prod_{k=1}^d G(\vartheta_k) U(z_k). \quad (4)$$

The proposed architecture design is theoretically motivated by how the VQC model expressivity increases with the circuit depth [25, 65]. Furthermore, through our encoding strategy, we are able to tune the number of qubits and the number of gate operations in the VQC circuit. This way, we ensure $d = \ell/n$, where d is the number of segments of the latent vector z and ℓ is the dimensionality of the latent space produced by the auto-encoder.

Here, we set the observable $\mathcal{M} \equiv \sigma_z$, where σ_z is the Pauli-Z operator acting on the first qubit. The data encoding $U(\cdot)$ used herein is the feature map from [15]. This consists of R_z rotation gates, which encode one data feature per qubit, and nearest-neighbor entanglement between these qubits, as presented in figure 1(c). The data encoding blocks include interactions between the features:

$$\phi(z_i^{(j)}, z_i^{(j+1)}) = \prod_{j=1}^n (\pi - z_i^{(j)}) (\pi - z_i^{(j+1)}), \quad (5)$$

where the symbols are described in figures 1(c) and (d). A single repetition of the variational ansatz applies R_Y rotations on each qubit and nearest-neighbor entanglement via the CNOT (CX) gates, as in figure 1(d).

We choose these general purpose data encoding and trainable circuits to highlight the effectiveness of the GQC network in solving classification tasks that conventional methods would struggle with. Hence, we do not focus on an extensive search for a quantum circuit that would yield the best classifier performance; our results do not depend on a specific VQC architecture choice.

All classifiers in this work are trained to minimise the binary cross entropy loss function:

$$\mathcal{L}_C = -\frac{1}{M} \sum_{m=1}^M [y_m \log(\hat{y}_m) + (1 - y_m) \log(1 - \hat{y}_m)], \quad (6)$$

where M is the number of data points, e.g. in a batch, y_m is the true label of the data point m , and \hat{y}_m is its predicted label, as in equation (3). This includes the VQC algorithm presented in section 3.2, any traditional methods, defined in section 3.1, and the VQC part of the guided quantum compression paradigm, introduced in section 3.3. All of these are precisely defined in the following section. This scoring rule dates back to 1952 [66], albeit introduced in a different context, and is the most popular loss function used for ML classification tasks.

3. Training paradigms

Different training methods are investigated to solve the binary classification problem of the $t\bar{t}H(b\bar{b})$ dataset [67], in order to contrast and exemplify the dimensionality reduction advantages of the guided quantum compression algorithm introduced in this work. Specifically, we study two processes occurring in proton-proton collision events at the LHC: the *signal* $t\bar{t}H(b\bar{b})$ process, in which a Higgs boson is produced, and the *background* $t\bar{t}(b\bar{b})$ process, where it is not produced. These two processes lead to the same final-state particles, generating similar high-dimensional signatures in the detector. The task of distinguishing these two processes is challenging [34, 67]. However, accurate classification of such events is crucial for enhancing the sensitivity of LHC experiments, enabling study Higgs boson processes with high precision. Furthermore, we assume that this dataset is representative of common high-dimensional datasets in which the probability distributions of the features for the two classes have substantial overlap. For these reasons, this dataset and the corresponding classification task is well-suited for ML and QML methods. It enables us to demonstrate that the proposed guided quantum compression paradigm generates a better latent space compared to conventional techniques in a challenging real-world problem, important to fundamental physics research.

To this end, we use three distinct training paradigms. First, the so-called classical approach, which serves as our baseline benchmark. Second, the 2Step method, which represents the usual way to perform dimensionality reduction when using QML for classifying complex data. Finally, the GQC paradigm developed in this work, which performs the dimensionality reduction and classification objectives at the same time.

Note, the three aforementioned strategies pertain to different computational requirements: it is less computationally intensive to train and robustly optimise the hyperparameters of a classical algorithm compared to a quantum one. Thus, while we are able to find the optimal values for the parameters of the classical network, this is not technically achievable for the quantum models presented in this work. The amount of compute time required to find the optimal hyperparameters for the quantum circuit is simply unfeasible. Furthermore, a vast amount of computational resources (cf appendix A) are also required to find the best quantum hyperparameter combination, including the best number of ansatz repetitions, the best number of qubits, or the best learning rate.

We present the classical and 2Step methods since they contrast and highlight the advantages our method brings in performing the needed dimensionality reduction. The classical baseline is used to disambiguate the performance of the quantum algorithm and classical component of the guided quantum compression paradigm. Meanwhile, the 2Step method is used to contrast the guided quantum compression with the conventional way of applying QML to our high-dimensional data.

3.1. Classical

A fully-connected feed-forward network is our classical benchmark. This network is trained to minimise the binary cross entropy loss, in equation (6), via stochastic gradient descent. In essence, the classical paradigm represents the most conventional way to address our classification task.

The hyperparameters of the feed-forward network, such as the learning rate, are optimized via a grid search. This is the only type of model and paradigm for which an exhaustive hyperparameter optimisation is feasible; see appendix A for more details.

3.2. 2Step

In the 2Step paradigm, the dimensionality reduction algorithm, namely the AE from section 2.1, is trained independently from the VQC classifier presented in section 2.2. As the name suggests, the classification task is performed in two separate steps. First, the hyperparameters of the AE are optimised and the resulting architecture is trained with the goal of minimising the MSE loss between the input data and the output of the decoder. Secondly, the VQC is trained by using the latent space representation z of each sample that the AE produces. This way, the VQC acts on a lower dimensional input and consequently its quantum circuit can be smaller. As mentioned in section 1, the dimensionality reduction allows us to employ a reasonable amount of resources in simulating quantum circuits in classical devices and in implementing our quantum models in currently available quantum computers. Hitherto, we outlined the most common way in which AEs are used

for dimensionality reduction in the context of a classification task [68]; in the following, a paradigm which integrates the two steps is introduced.

3.3. Guided quantum compression

Our strategy aims to address the problem of generating a low-dimensional representation that ensures the discrimination of the classes by the quantum model. The GQC network implements a trainable data encoding map $\mathcal{U}_\theta(\mathcal{E}_\omega(x))$. The lower-dimensional representation learned by the encoder network is *guided* by the quantum classification algorithm.

The guidance is performed during the training of the GQC network by coupling the auto-encoder and VQC models through the following loss function:

$$\mathcal{L} = (1 - \lambda)\mathcal{L}_R + \lambda\mathcal{L}_C, \quad (7)$$

where \mathcal{L}_R is the MSE loss as defined in equation (1), \mathcal{L}_C is the loss of the VQC classifier defined in equation (6), and $\lambda \in (0, 1)$ is the hyperparameter that specifies the coupling between the reconstruction and classification optimisation tasks. The λ is fine-tuned separately to maximise the classification accuracy on the validation dataset, see appendix A.

The simultaneous learning of the two objectives improves the generalisation power of the classifier [69, 70]. This synergy arises from the regularisation imposed through equation (7) of the quantum classifier. Namely, the classifier performance is enhanced by the additional task of reconstructing data from the latent space [70].

In our algorithm, we train the GQC network via stochastic gradient descent with the Adam optimizer [71]. The gradients of the classical parts of the model are computed using backpropagation, while the quantum circuit gradients are computed via the adjoint differentiation method [72], for efficient training of the model on classical processors used to simulate the quantum software. The proposed GQC architecture can also be trained on quantum hardware using the parameter-shift rule [73] instead of the adjoint differentiation method used here. The hyperparameters of the GQC network are tuned using a sequential grid search for each hyperparameter while keeping all the rest fixed. For more details on the hyperparameter tuning procedure, see appendix A.

4. Results

The classification performance of the three training paradigms described in section 3 is benchmarked using the same simulated dataset consisting of $\bar{t}\bar{t}H(b\bar{b})$ (signal) and $\bar{t}\bar{t}(b\bar{b})$ (background) events [67, 74, 75]. For all models except the AE used in the 2Step paradigm, the training data consists of 20 000 samples; meanwhile, the AE uses 1.44×10^6 samples for training. The test dataset consists of 5 k-folds, with 20 000 samples for each fold. The number of samples is balanced across the two classes for all datasets. Additionally, a validation dataset of 1500 samples is used during training to monitor for overfitting and to fix the hyperparameters of the models. Specifically, at the end of each training epoch, the validation loss is computed on the the validation dataset for each model: the classification loss in equation (6) for the classifier models, the reconstruction loss in equation (1) for the AE, or the combined loss in equation (7) for GQC model. In each case, the model that yields the minimum validation loss is selected after the training.

The initial dimensionality of each data sample is 67 (60), if high-level, classically preprocessed, features pertaining to the collision event are included (excluded). This is always reduced to a final dimensionality of 16 by the AE, except in the classical training workflow from section 3.1, which does not use dimensionality reduction. The initial 60 features are the kinematic variables of the particles included in the collision event, e.g. their energy, angle of incidence on the detector, and so on [67]. These kinematic variables describe the quantum interaction of colliding particles and can be calculated from first-principles within the framework of quantum field theory. Furthermore, the additional 7 features that introduce high-level properties of the collision event are the so called *btag* variables. These determine the likelihood of a certain particle, the *b* quark, being produced in the event. These likelihoods are conventionally obtained using classical ML models when producing this dataset.

The dataset is preprocessed using physical criteria before the training and testing of the investigated models. The exact requirements can be found in [34, 67]. These preprocessing steps accommodate for the geometric acceptance of the detector used to record the proton collision data and ensure that each collision event has the desired characteristics for the study, e.g. at least one detected lepton. Subsequently, we normalise all data features to the range [0, 1].

The lower dimensional latent spaces produced through the conventional 2Step method and the GQC paradigm are shown in figure 2. As clearly visible in this figure, the GQC network learns a better separation

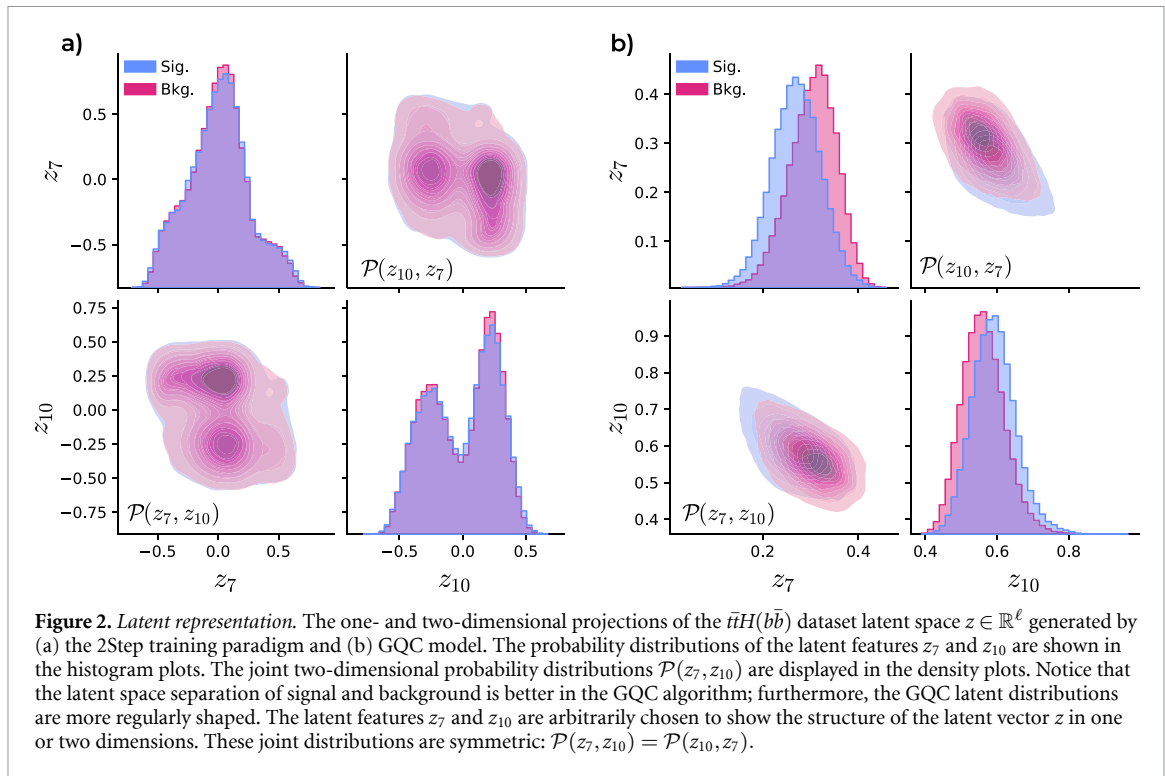


Table 1. Model performances. The first two columns present the Area Under the Curve (AUC) values of the three training paradigms with and without the b-tag variable in the training data, respectively. The largest difference between the GQC and conventional methods is highlighted in bold. The last two columns show the inverse FPR at TPR = 0.8.

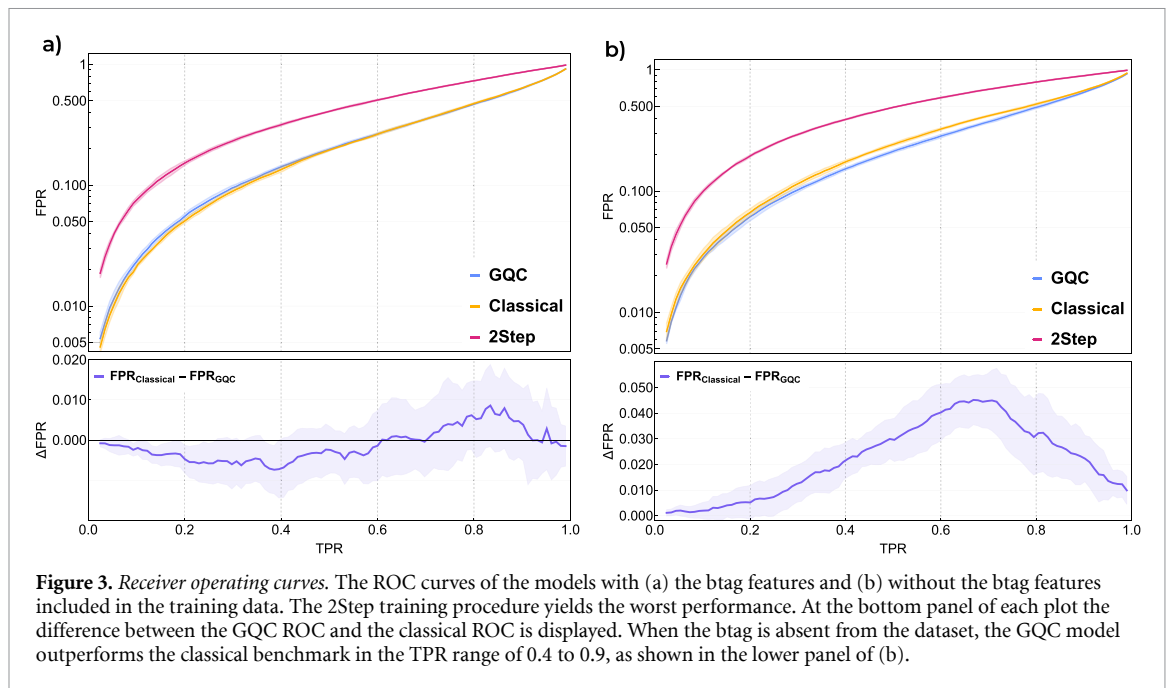
Model	AUC w/ b-tag	AUC w/o b-tag	FPR ⁻¹ w/ b-tag	FPR ⁻¹ w/o b-tag
GQC	0.733 ± 0.003	0.720 ± 0.005	2.134 ± 0.028	2.045 ± 0.049
2Step	0.561 ± 0.003	0.508 ± 0.002	1.263 ± 0.004	1.368 ± 0.007
Classical	0.734 ± 0.002	0.699 ± 0.004	2.107 ± 0.029	1.921 ± 0.035

between the two data classes in its latent space. Hence, it is favorable to use the GQC paradigm when performing dimensionality reduction on the input data. In appendix B, we quantify the class separation and the GQC improvement using feature-wise Kullback–Leibler divergences (KLD) [76].

The receiver operating curve (ROC) produced through each paradigm from section 3 is shown in figure 3. The 5 k-folds of the test data are used to compute these ROC curves and their corresponding uncertainties. The uncertainty is represented as error bands with a width of one standard deviation around the ROC curves. The conventional 2Step method performs the worst. Meanwhile, the classical and the GQC methods yield a similar classification accuracy when the btag is included in the dataset, cf figure 3(b); excluding the btag, i.e. keeping only the kinematic variables of the particles, leads to the GQC significantly outperforming the classical approach. The advantage in classification performance appears in the relevant range of True Positive Rate between 0.4 and 0.9, which is a typical choice for physics analyses at the LHC. The results in figure 3 are condensed in table 1, where summary performance metrics derived from these ROCs are shown. The GQC latent space allows for a competitive performance with the classical benchmark.

5. Conclusions

The choice of the compression method can have a significant impact on the classifier performance. We show that for the $\bar{t}tH(b\bar{b})$ classification task, applying dimensionality reduction as a preprocessing step renders the problem impenetrable for a currently applicable VQC classifier algorithm, as highlighted in figure 3 and table 1. In contrast, when an identical VQC is used as part of the GQC training paradigm, its classification performance drastically increases. Furthermore, the GQC model achieves a consistent enhanced classification accuracy without an increase in its training time complexity compared to the 2Step approach; it inherits its training time complexity from its components: the autoencoder and the VQC. Thus, integrating the dimensionality reduction and classification tasks is shown to provide a better low dimensional latent space for our problem.



Additionally, the hyperparameter optimisation procedure employed for each network in this work leads to a potential bias towards the classical paradigm when comparing the model performances: it is computationally feasible to perform an extensive hyperparameter optimisation for the classical algorithm, while this is currently not the case for the GQC model. Thus, the hyperparameters of the presented networks that include a QML element are only approximately optimal while the hyperparameters of the classical algorithm are the best they can possibly be. Nevertheless, the conclusions we draw from our results are not affected by these limitations.

For simple datasets, dimensionality reduction and classification tasks can potentially be decoupled, i.e. they can be treated as independent problems [31, 32, 34, 41]. However, for realistic datasets it is possible that such a separate treatment yields worse results. The dimensionality reduction algorithm can obfuscate the class structure of the original problem, apparent in our results. The GQC encoding strategy ensures that the VQC is both expressive and flexible for choosing the desired circuit width and depth, while retaining its overall accuracy. The training paradigms presented in this work show the benefits of carefully constructing hybrid QML models. These advantages stem from aligning the goals of the dimensionality reduction algorithm, typically chosen *a priori* and arbitrarily, with that of the quantum model. Hence, the use of classical methods for dimensionality reduction can facilitate current quantum computing applicability in realistic settings.

As seen in figure 3(a), in the worst case the performance of the GQC network is equivalent to that of the classical benchmark. The GQC network outperforms both the classical model and the quantum model in the 2Step approach when the highly processed bt tag feature is not used in the training, shown in figure 3(b). Moreover, it is competitive with state-of-the-art results on this data subset [50]. This suggests that for improved QML performance on particle physics data, features that are representative of the quantum process that generated it [51] are preferred, e.g. the angular distribution of particles [41, 43]. Investigating this in more detail is left for future work.

Data availability statement

The data that support the findings of this study are openly available at the following DOI: <https://doi.org/10.5281/zenodo.7267942>.

Code availability

The code developed for this paper is available publicly in the GitHub repository: <https://github.com/CERN-IT-INNOVATION/gqc>. A combination of PyTorch and PennyLane [77] were used to implement and train the GQC, 2Step, and the classical model architectures. The website in [78] was used for figure 1(a) to distort the reconstructed high-energy physics event schematic.

Acknowledgments

V B is supported by an ETH Research Grant (Grant No. ETH C-04 21-2). P O is supported by Swiss National Science Foundation Grant No. PZ00P2_201594. The authors would like to thank Elías F Combarro for the helpful discussions. M G and S V are supported by CERN through the Quantum Technology Initiative. F R acknowledges financial support by the Swiss National Science Foundation (Ambizione Grant No. PZ00P2_186040).

Appendix A. Model hyperparameter selection

The classical autoencoder used in this study for the 2Step paradigm has an encoder network consisting of six layers with nodes [67, 64, 44, 32, 24, 16], respectively. The decoder network mirrors the architecture of the encoder: [16, 24, 32, 44, 64, 67]. The ReLU activation function is applied to all layers except for the layer corresponding to the latent space, i.e. the one with 16 nodes, and the output layer, i.e. the last layer of the decoder with 67 nodes. For the latent and output layers both the sigmoid and tanh activation functions have been assessed providing similar results.

The hyperparameters of all the classical models are found through exhaustive grid search. Specifically, we repeat the training and k-fold testing of each model, as described in section 4, for each hyperparameter combination. The considered hyperparameter combination is {32, 64, 128, 256, 512, 1024, 2048} for the batch size and $\{10^{-3}, 10^{-2}, 10^{-1}\}$ for the learning rate.

For the AE, we firstly identify that a batch size of 128 and a learning rate of 10^{-3} yields the best performance. With the batch size fixed at 128, we subsequently perform a finer hyperparameter tuning for the learning rate in the neighborhood of $[10^{-3}, 10^{-2}]$, using the Optuna hyperparameter optimisation framework [79]. Following this procedure, we find an optimal learning rate of 0.0012.

For the classical benchmark, we evaluate the performance of a fully connected feed-forward network. This network has the same architecture as the encoder network of the AE model, with the addition of a single-node layer with a sigmoid activation function serving as the output of the classifier. For studies excluding the btag variable, the model architectures stay the same, except that the input layer is changed to 60 nodes. Additionally, for the AE, the output layer is also set to 60. Furthermore, we investigate a shallower architecture with three layers in total [67, 16, 1] that yielded similar classification performance. Hence, for simplicity, we choose the shallower architecture. To maintain consistency with the classical benchmark, the classical components of the hybrid GQC network (cf figure 1), \mathcal{E}_ω and \mathcal{D}_ρ , also have the shallow architectures of [67, 16] and [16, 67], respectively.

In the following, the sequential grid search we employed to tune the hyperparameters of the GQC network is described; in bold we present the hyperparameter values found to be optimal following our procedure. For our simulations and numerical studies, we used specialized nodes at the CERN and Paul Scherrer Institut PSI (PSI) computing clusters. At each step of the sequential grid search the models are retrained and evaluated using k-fold testing, as discussed in section 4. Firstly, we choose the best performing batch size out of {128, 256, 512, **1024**, 2048}, afterwards we do a scan on the learning rate $\{10^{-3}, \mathbf{10}^{-2}, 10^{-1}\}$, then we optimize for the repetitions of the trainable gates (cf figure 1) of the VQC with $r \in \mathbf{2}, \mathbf{4}, \mathbf{8}$. Lastly, we fix all the above hyperparameters and optimize the coupling λ (see equation (7)) in the range [0.3, 0.9] with a step of 0.1 finding the best performing value at $\lambda = \mathbf{0.7}$. The λ is fine-tuned to maximise the classification accuracy on the validation dataset.

An exhaustive grid search, as the one performed for the classical model (cf section 4), would require assessing all possible combinations of the above hyperparameters leading to a drastic increase of compute time. Specifically, it would increase the current quantum simulation runtime, which is at the order of days, to multiple weeks.

Appendix B. Quantifying the class separation in the latent distributions

The KLD, also called relative entropy, is a measure that quantifies the difference between a probability distribution P and a reference probability distribution Q [76]. Specifically, given a continuous random variable $x \in \mathcal{X}$, over a sampling space \mathcal{X} , the KLD is defined as

$$D_{\text{KL}}(P||Q) = \int_{\mathcal{X}} P(x) \log \left(\frac{P(x)}{Q(x)} \right) dx. \quad (\text{B1})$$

The KLD approaches zero as the distributions become more similar. In general, one does not have access to an analytical expression for P and Q but only to finite samples from these distributions. In such cases, the

Table 2. Quantifying the separation between the background $\mathcal{B}(z_i)$ and signal $\mathcal{S}(z_i)$ distributions in the latent space produced by the 2Step and GQC methods, respectively, using the KL divergence D_{KL} . The presented values correspond to each latent space feature z_i , where $i = 0, 1, \dots, 15$.

	z_0	z_1	z_2	z_3	z_4	z_5	z_6	z_7	z_8	z_9	z_{10}	z_{11}	z_{12}	z_{13}	z_{14}	z_{15}
$D_{\text{KL}}^{(2\text{Step})}(\mathcal{S}(z_i) \mathcal{B}(z_i))$	0.16	0.09	0.31	0.16	0.34	0.23	0.12	0.05	0.07	0.04	0.07	0.10	0.03	1.81	0.08	1.42
$D_{\text{KL}}^{(\text{GQC})}(\mathcal{S}(z_i) \mathcal{B}(z_i))$	3.53	3.84	17.60	6.25	2.17	2.77	3.12	23.40	7.06	6.06	7.23	5.89	4.94	2.73	0.41	2.29

KLD can be computed by constructing histograms from the P and Q samples,

$$D_{\text{KL}}(P || Q) = \sum_{i=k}^N P_k \log \left(\frac{P_k}{Q_k} \right), \quad (\text{B2})$$

where N is the number of bins, and P_k and Q_k represent the probabilities of observing values that fall into the k th bin.

In our work, we are interested in quantifying the separation between the background \mathcal{B} and signal \mathcal{S} distributions in the latent space produced by the conventional 2Step method and the proposed GQC network, as seen in figure 2 and discussed in section 3. Specifically, using equation (B2), we compute the difference between the background $\mathcal{B}(z_i)$ and signal $\mathcal{S}(z_i)$ latent distributions for each latent feature z_i , where $i = 0, 1, \dots, 15$ enumerates here the latent features. The obtained values are presented in table 2. By computing the average ratio of KL divergences $D_{\text{KL}}^{(\text{GQC})}$ and $D_{\text{KL}}^{(2\text{Step})}$, obtained from the GQC network and 2Step latent distributions, respectively, over all features we arrive at

$$\mathcal{R} = \frac{1}{16} \sum_{i=0}^{15} \frac{D_{\text{KL}}^{(\text{GQC})}(\mathcal{S}(z_i) || \mathcal{B}(z_i))}{D_{\text{KL}}^{(2\text{Step})}(\mathcal{S}(z_i) || \mathcal{B}(z_i))} \approx 79.58. \quad (\text{B3})$$

Hence, we observe a significant improvement by a factor of 79.58 with the GQC over the 2Step method in terms of signal and background separation in the latent representations, as quantified by the KL divergence. This in turn also leads to higher classification accuracy for the GQC network, as presented in section 4.

ORCID iDs

Vasilis Belis  <https://orcid.org/0000-0001-5920-8998>

Patrick Odagiu  <https://orcid.org/0009-0005-0342-0128>

Michele Grossi  <https://orcid.org/0000-0003-1718-1314>

Sofia Vallecorsa  <https://orcid.org/0000-0002-7003-5765>

References

- [1] Plehn T, Butter A, Dillon B and Krause C 2022 Modern machine learning for LHC physicists (arXiv:2211.01421 [hep-ph])
- [2] Guest D, Cranmer K and Whiteson D 2018 Deep learning and its application to LHC physics *Annu. Rev. Nucl. Part. Sci.* **68** 161–81
- [3] Karagiorgi G, Kasieczka G, Kravitz S, Nachman B and Shih D 2021 Machine learning in the search for new fundamental physics (arXiv:2112.03769 [hep-ph])
- [4] Belis V, Odagiu P and Aarrestad T K 2024 Machine learning for anomaly detection in particle physics *Rev. Phys.* **12** 100091
- [5] Grover L K 1996 A fast quantum mechanical algorithm for database search (arXiv:quant-ph/9605043 [quant-ph])
- [6] Shor P W 1997 Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer *SIAM J. Comput.* **26** 1484–509
- [7] Harrow A W, Hassidim A and Lloyd S 2009 Quantum algorithm for linear systems of equations *Phys. Rev. Lett.* **103** 150502
- [8] Arute F et al 2019 Quantum supremacy using a programmable superconducting processor *Nature* **574** 505–10
- [9] Zhong H-S et al 2020 Quantum computational advantage using photons *Science* **370** 1460–3
- [10] Madsen L S et al 2022 Quantum computational advantage with a programmable photonic processor *Nature* **606** 75–81
- [11] Babbush R, Berry D W, Kothari R, Somma R D and Wiebe N 2023 Exponential quantum speedup in simulating coupled classical oscillators (arXiv:2303.13012 [quant-ph])
- [12] Biamonte J, Wittek P, Pancotti N, Rebentrost P, Wiebe N and Lloyd S 2017 Quantum machine learning *Nature* **549** 195–202
- [13] Schuld M and Petruccione F 2018 *Supervised Learning With Quantum Computers* (Springer)
- [14] Benedetti M, Lloyd E, Sack S and Fiorentini M 2019 Parameterized quantum circuits as machine learning models *Quantum Sci. Technol.* **4** 043001
- [15] Havlíček V, Córcoles A D, Temme K, Harrow A W, Kandala A, Chow J M and Gambetta J M 2019 Supervised learning with quantum-enhanced feature spaces *Nature* **567** 209–12
- [16] Schuld M and Killoran N 2019 Quantum machine learning in feature Hilbert spaces *Phys. Rev. Lett.* **122** 040504
- [17] Schuld M M, Bocharov A, Svore K M and Wiebe N 2020 Circuit-centric quantum classifiers *Phys. Rev. A* **101** 032308
- [18] Rebentrost P, Mohseni M and Lloyd S 2014 Quantum support vector machine for big data classification *Phys. Rev. Lett.* **113** 130503
- [19] Liu Y, Arunachalam S and Temme K 2021 A rigorous and robust quantum speed-up in supervised machine learning *Nat. Phys.* **17** 1013–7

- [20] Huang H Y, Broughton M, Mohseni M, Babbush R, Boixo S, Neven H and McClean J R 2021 Power of data in quantum machine learning *Nat. Commun.* **12** 2631
- [21] Huang H Y et al 2022 Quantum advantage in learning from experiments *Science* **376** 1182–6
- [22] Muser T, Zapusek E, Belis V and Reiter F 2023 Provable advantages of kernel-based quantum learners and quantum preprocessing based on Grover's algorithm (arXiv:2309.14406 [quant-ph])
- [23] Pirnay N, Sweke R, Eisert J and Seifert J-P 2023 A super-polynomial quantum-classical separation for density modelling *Phys. Rev. A* **107** 042416
- [24] Gyurik C and Dunjko V 2022 On establishing learning separations between classical and quantum machine learning with classical data (arXiv:2208.06339)
- [25] Pérez-Salinas A, Cervera-Lierta A, Gil-Fuster E and Latorre J I 2020 Data re-uploading for a universal quantum classifier *Quantum* **4** 226
- [26] Schuld M 2021 Supervised quantum machine learning models are kernel methods (arXiv:2101.11020)
- [27] Goto T, Tran Q H and Nakajima K 2021 Universal approximation property of quantum machine learning models in quantum-enhanced feature spaces *Phys. Rev. Lett.* **127** 090506
- [28] Abbas A, Sutter D, Zoufal C, Lucchi A, Figalli A and Woerner S 2021 The power of quantum neural networks *Nat. Comput. Sci.* **1** 403–9
- [29] Caro M C, Huang H-Y, Cerezo M, Sharma K, Sornborger A, Cincio L and Coles P J 2022 Generalization in quantum machine learning from few training data *Nat. Commun.* **13** 4919
- [30] Jerbi S, Fiderer L J, Poulsen Nautrup H, Kübler J M, Briegel H J and Dunjko V 2023 Quantum machine learning beyond kernel methods *Nat. Commun.* **14** 517
- [31] Wu S L et al 2021 Application of quantum machine learning using the quantum kernel algorithm on high energy physics analysis at the LHC *Phys. Rev. Res.* **3** 033221
- [32] Wu S L et al 2021 Application of quantum machine learning using the quantum variational classifier method to high energy physics analysis at the LHC on IBM quantum computer simulator and hardware with 10 qubits *J. Phys. G: Nucl. Part. Phys.* **48** 125003
- [33] Terashi K, Kaneda M, Kishimoto T, Saito M, Sawada R and Tanaka J 2021 Event classification with quantum machine learning in high-energy physics *Comput. Softw. Big Sci.* **5** 2
- [34] Belis V, González-Castillo S, Reissel C, Vallecorsa S, Combarro E F, Dissertori G and Reiter F 2021 Higgs analysis with quantum classifiers *EPJ Web Conf.* **251** 03070
- [35] Blance A and Spannowsky M 2020 Quantum machine learning for particle physics using a variational quantum classifier (arXiv:2010.07335 [hep-ph])
- [36] Gianelle A, Koppenburg P, Lucchesi D, Nicotra D, Rodrigues E, Sestini L, de Vries J and Zuliani D 2022 Quantum machine learning for b -jet charge identification *J. High Energy Phys.* **JHEP08(2022)014**
- [37] de Lejarza J J M, Cieri L and Rodrigo G 2022 Quantum clustering and jet reconstruction at the LHC *Phys. Rev. D* **106** 036021
- [38] Tüysüz C, Carminati F, Demirköz B, Dobos D, Fracas F, Novotny K, Potamianos K, Vallecorsa S and Vlimant J-R 2020 Particle track reconstruction with quantum algorithms *EPJ Web Conf.* **245** 09013
- [39] Magano D et al 2022 Quantum speedup for track reconstruction in particle accelerators *Phys. Rev. D* **105** 076012
- [40] Ngairangbam V S, Spannowsky M and Takeuchi M 2022 Anomaly detection in high-energy physics using a quantum autoencoder *Phys. Rev. D* **105** 095004
- [41] Schuhmacher J, Boggia L, Belis V, Puljak E, Grossi M, Pierini M, Vallecorsa S, Tacchino F, Barkoutsos P and Tavernelli I 2023 Unravelling physics beyond the standard model with classical and quantum anomaly detection *Mach. Learn.: Sci. Technol.* **4** 045031
- [42] Alvi S, Bauer C and Nachman B 2023 Quantum anomaly detection for collider physics *J. High Energy Phys.* **JHEP02(2023)220**
- [43] Woźniak K A, Belis V, Puljak E, Barkoutsos P, Dissertori G, Grossi M, Pierini M, Reiter F, Tavernelli I and Vallecorsa S 2023 Quantum anomaly detection in the latent space of proton collision events at the LHC (arXiv:2301.10780 [quant-ph])
- [44] de Lejarza J J M, Cieri L, Grossi M, Vallecorsa S and Rodrigo G 2024 Loop Feynman integration on a quantum computer (arXiv:2401.03023 [hep-ph])
- [45] Agliardi G, Grossi M, Pellen M and Prati E 2022 Quantum integration of elementary particle processes *Phys. Lett. B* **832** 137228
- [46] Meglio A Di et al 2023 Quantum computing for high-energy physics: state of the art and challenges. summary of the QC4HEP working group (arXiv:2307.03236 [quant-ph])
- [47] Ballard D H 1987 Modular learning in neural networks *Proc. 6th National Conf. on Artificial Intelligence* vol 1 pp 279–84
- [48] Caruana R 1997 Multitask learning *Mach. Learn.* **28** 41–75
- [49] Ruder S 2017 An overview of multi-task learning in deep neural networks (arXiv:1706.05098)
- [50] Donega M, Kasieczka G, Lösche T, Musella P, Reissel C and Wallny R 2021 Attention and dynamic graph convolution neural network in the context of classifying $t\bar{t}H(b\bar{b})$ vs. $t\bar{t}(b\bar{b})$ in the semi-leptonic top quark pair decay channel (available at: <https://indico.cern.ch/event/980214/contributions/4413711/>)
- [51] CMS Collaboration 2019 Measurement of the top quark polarization and $t\bar{t}$ spin correlations using dilepton final states in proton-proton collisions at $\sqrt{s} = 13$ TeV *Phys. Rev. D* **100** 072002
- [52] Kübler J et al 2021 The inductive bias of quantum kernels *Advances in Neural Information Processing Systems* vol 34, ed M Ranzato, A Beygelzimer, Y Dauphin, P S Liang and J W Vaughan (Curran Associates, Inc.) pp 12661–73
- [53] LeCun Y 1987 (Universite P. et M. Curie (Paris 6)) Modeles connexionnistes de l'apprentissage (connectionist learning models) *PhD Thesis*
- [54] Hinton G E and Zemel R 1993 Autoencoders, minimum description length and helmholtz free energy *Advances in Neural Information Processing Systems* vol 6
- [55] Kingma D P and Welling M 2013 Auto-encoding variational Bayes (arXiv:1312.6114 [stat.ML])
- [56] Deja K, Dubiński J, Nowak P, Wenzel S, Spurek P and Trzciński T 2021 End-to-end Sinkhorn autoencoder with noise generator *IEEE Access* **9** 7211–9
- [57] Joo W, Lee W, Park S and Moon I-C 2020 Dirichlet variational autoencoder *Pattern Recognit.* **107** 107514
- [58] Pan S, Hu R, Long G, Jiang J, Yao L and Zhang C 2018 Adversarially regularized graph autoencoder for graph embedding (arXiv:1802.04407)
- [59] Mitarai K, Negoro M, Kitagawa M and Fujii K 2018 Quantum circuit learning *Phys. Rev. A* **98** 032309
- [60] Pérez-Salinas A, Cruz-Martinez J, Alhajri A A and Carrazza S 2021 Determining the proton content with a quantum computer *Phys. Rev. D* **103** 034027
- [61] Chang S Y, Herbert S, Vallecorsa S, Combarro E F and Duncan R 2021 Dual-parameterized quantum circuit GAN model in high energy physics *EPJ Web Conf.* **251** 03050

- [62] Kiss O, Grossi M, Kajomovitz E and Vallecorsa S 2022 Conditional Born machine for Monte Carlo event generation *Phys. Rev. A* **106** 022612
- [63] Delgado A and Hamilton K E 2022 Unsupervised quantum circuit learning in high energy physics *Phys. Rev. D* **106** 096006
- [64] Bravo-Prieto C, Baglio J, Cé M, Francis A, Grabowska D M and Carrazza S 2022 Style-based quantum generative adversarial networks for Monte Carlo events *Quantum* **6** 777
- [65] Schuld M, Sweke R and Meyer J J 2021 The effect of data encoding on the expressive power of variational quantum machine learning models *Phys. Rev. A* **103** 032430
- [66] Good I J 1952 Rational decisions *J. R. Stat. Soc. B* **14** 107–14
- [67] Reissel C and Belis V 2024 $\bar{t}tH(bb)$ dataset in the semi-leptonic decay channel Zenodo <http://dx.doi.org/10.5281/zenodo.7267942>
- [68] Wang Y, Yao H and Zhao S 2016 Auto-encoder based dimensionality reduction *Neurocomputing* **184** 232–42
- [69] Le L, Patterson A and White M 2018 Supervised autoencoders: improving generalization performance with unsupervised regularizers *Advances in Neural Information Processing Systems* vol 31
- [70] Vafaieikia P, Namdar K and Khalvati F 2020 A brief review of deep multi-task learning and auxiliary task learning (arXiv:2007.01126 [cs.LG])
- [71] Kingma D P and Ba J 2017 Adam: a method for stochastic optimization (arXiv:1412.6980 [cs.LG])
- [72] Jones T and Gacon J 2020 Efficient calculation of gradients in classical simulations of variational quantum algorithms (arXiv:2009.02823 [quant-ph])
- [73] Schuld M, Bergholm V, Gogolin C, Izaac J and Killoran N 2019 Evaluating analytic gradients on quantum hardware *Phys. Rev. A* **99** 032331
- [74] CMS collaboration 2019 Search for $\bar{t}tH$ production in the $H \rightarrow b\bar{b}$ decay channel with $\bar{t}t$ leptonic decays in proton-proton collisions at $\sqrt{s} = 13$ TeV *J. High Energy Phys.* **JHEP03(2019)026**
- [75] The ATLAS Collaboration 2022 Measurement of Higgs boson decay into b-quarks in associated production with a top-quark pair in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector *J. High Energy Phys.* **JHEP06(2022)097**
- [76] Kullback S and Leibler R A 1951 On information and sufficiency *Ann. Math. Stat.* **22** 79–86
- [77] Bergholm V et al 2022 PennyLane: automatic differentiation of hybrid quantum-classical computations (arXiv:1811.04968 [quant-ph])
- [78] Narkevich D deepfriedmemes website (available at: <https://github.com/efskap/deepfriedmemes.com>) (Accessed 2 January 2024)
- [79] Akiba T, Sano S, Yanase T, Ohta T and Koyama M 2019 Optuna: a next-generation hyperparameter optimization framework *CoRR* (arXiv:1907.10902)