



OPEN ACCESS

EDITED BY
Matt LeBlanc,
The University of Manchester, United Kingdom

REVIEWED BY
Jack Y. Araz,
Jefferson Lab (DOE), United States
Stefano Scali,
University of Exeter, United Kingdom

*CORRESPONDENCE
Nuno Filipe Castro
✉ nuno.castro@fisica.uminho.pt

RECEIVED 28 July 2023
ACCEPTED 20 November 2023
PUBLISHED 15 December 2023

CITATION
Peixoto MC, Castro NF, Crispim Romão M,
Oliveira MGJ and Ochoa I (2023) Fitting a
collider in a quantum computer: tackling the
challenges of quantum machine learning for
big datasets. *Front. Artif. Intell.* 6:1268852.
doi: 10.3389/frai.2023.1268852

COPYRIGHT
© 2023 Peixoto, Castro, Crispim Romão,
Oliveira and Ochoa. This is an open-access
article distributed under the terms of the
[Creative Commons Attribution License \(CC BY\)](https://creativecommons.org/licenses/by/4.0/).
The use, distribution or reproduction in other
forums is permitted, provided the original
author(s) and the copyright owner(s) are
credited and that the original publication in this
journal is cited, in accordance with accepted
academic practice. No use, distribution or
reproduction is permitted which does not
comply with these terms.

Fitting a collider in a quantum computer: tackling the challenges of quantum machine learning for big datasets

Miguel Caçador Peixoto¹, Nuno Filipe Castro^{1,2*},
Miguel Crispim Romão^{1,3}, Maria Gabriela Jordão Oliveira¹ and
Inês Ochoa⁴

¹LIP—Laboratório de Instrumentação e Física Experimental de Partículas, Escola de Ciências, Universidade do Minho, Braga, Portugal, ²Departamento de Física, Escola de Ciências, Universidade do Minho, Braga, Portugal, ³Department of Physics and Astronomy, University of Southampton, Southampton, United Kingdom, ⁴LIP—Laboratório de Instrumentação e Física Experimental de Partículas, Lisbon, Portugal

Current quantum systems have significant limitations affecting the processing of large datasets with high dimensionality, typical of high energy physics. In the present paper, feature and data prototype selection techniques were studied to tackle this challenge. A grid search was performed and quantum machine learning models were trained and benchmarked against classical shallow machine learning methods, trained both in the reduced and the complete datasets. The performance of the quantum algorithms was found to be comparable to the classical ones, even when using large datasets. Sequential Backward Selection and Principal Component Analysis techniques were used for feature's selection and while the former can produce the better quantum machine learning models in specific cases, it is more unstable. Additionally, we show that such variability in the results is caused by the use of discrete variables, highlighting the suitability of Principal Component analysis transformed data for quantum machine learning applications in the high energy physics context.

KEYWORDS

high energy physics, quantum computing, quantum machine learning, K-means, principal component analysis, data reduction

1 Introduction

The Standard Model of Particle Physics (SM) provides a remarkable description of the fundamental constituents of matter and their interactions, being in excellent agreement with the collider data accumulated so far. Nonetheless, there are still important open questions, unaddressed by the SM, such as gravity, dark matter, dark energy, or the matter-antimatter asymmetry in the universe (Ellis, 2012), motivating a comprehensive search program for new physics phenomena beyond the SM (BSM) at the Large Hadron Collider (LHC) at CERN.

The search for BSM phenomena at colliders poses specific challenges in data processing and analysis, given the extremely large datasets involved and the low signal to background ratios expected. In this context, the analysis of the collision data obtained by the LHC experiments often relies on machine learning (ML), a field in computer science that can harness large amounts of data to train generalizable algorithms for a variety of applications (Guest et al., 2018; Feickert and Nachman, 2021), such as classification tasks. These techniques have shown an outstanding ability to find correlations in high-dimensional parameter spaces to discriminate between potential signal and background processes. They

are known to scale with data, and usually rely on a large number of learnable parameters to achieve their remarkable performance.

In order to train these large models, classical¹ machine learning (CML) takes advantage of hardware accelerators, such as graphics processing units (GPUs), for efficient, parallel, and fast matrix multiplications. On the other hand, a new class of hardware is becoming available, with the advent of noisy intermediate-scale quantum (NISQ) computing devices. This accelerated the development of new quantum algorithms targeted at exploiting the capacity and feasibility of this new technology for ML applications.

Quantum machine learning (QML) is an emerging research field aiming to use quantum circuits to tackle ML tasks. One of the motivations for using this new technology in high energy physics (HEP) relates to the intrinsic properties of quantum computations, namely representing the data in a Hilbert space where the data can be in a superposition of states or in entangled states, which can allow to explore additional information in data analysis and, eventually, contribute to better classification of HEP events, namely in the context of the search for BSM phenomena. Recently, this new technology has been applied to various HEP problems (Guan et al., 2021). Namely, in event reconstruction (Das et al., 2019; Shapoval and Calafiura, 2019; Bapst et al., 2020; Tüysüz et al., 2020; Wei et al., 2020; Zlokapa et al., 2021a; Funcke et al., 2022), classification tasks (Mott et al., 2017; Belis et al., 2021; Blance and Spannowsky, 2021; Terashi et al., 2021; Wu et al., 2021; Zlokapa et al., 2021b; Araz and Spannowsky, 2022; Chen et al., 2022; Gianelle et al., 2022), data generation (Chang et al., 2021a,b; Delgado and Hamilton, 2022; Borrás et al., 2023; Rehm et al., 2023), and anomaly detection problems (Ngairangbam et al., 2022; Alvi et al., 2023; Schuhmacher et al., 2023; Woźniak et al., 2023).

Despite the promising potential of quantum computation, NISQ processors have important limitations, such as the qubit quality (i.e., the accuracy with which it is possible to execute quantum gates), the qubit lifetime and the limited depth of quantum circuits, since for large circuits the noise overwhelms the signal (Li et al., 2018; Preskill, 2018). This necessarily limits the complexity of the circuits and the size of the datasets used to train them.

In this paper, we perform a systematic comparison of the performance of QML and shallow CML algorithms in HEP. The choice to focus on shallow methods rather than state-of-the-art architectures based on deep neural networks is to provide a fair comparison between methodologies, since neural networks are known to require large datasets (both in terms of sample size and dimension) to achieve good performance, something that is not feasible with current quantum computers. By choosing CML algorithms suited for smaller datasets, we will add to the on-going discussion regarding potential advantages of quantum computing by comparing QML and CML in the same footing.

The use of QML algorithms in this context is studied by targeting a common binary classification task in HEP: classifying a BSM signal against SM background. A benchmark BSM signal leading to the Zt final state is considered, in events with multiple leptons and b -tagged jets, which can be used to achieve a reasonable

signal to background ratio. Variational quantum classifiers (VQC) are trained and optimized via a grid search. The use of reduced data is explored, considering both the number of features and the number of events, via different strategies: ranking of features, data transformations aiming for a richer reduced set of features, use of random samples, and choice of representative data samples.

2 Quantum machine learning

The QML algorithms are implemented using a quantum circuit, i.e., a collection of quantum gates applied to an n -qubit quantum state, followed by a measurement (or multiple measurements) that represent the output of the circuit. In order to implement a learning algorithm, the quantum circuit can be parameterized with parameters that can be learned by confronting the measurement to a loss function.

QML is effectively an extension of CML techniques to the Hilbert space, where instead of representing data as vectors in a high-dimensional real space, we encode it in state vectors of a Hilbert space. A QML algorithm, such as a quantum neural network, can be implemented using the quantum equivalent of a perceptron, one of the building blocks of CML. A problem arises from the realization that the activation functions used in CML can not be expressed using a linear operation, which is inherently required from the quantum evolution of a state. Ideas have been proposed to imitate an activation function in the quantum space (Gupta and Zia, 2001; Schuld et al., 2015), but, in the current paper, only variational quantum classifiers (Farhi and Neven, 2018; Schuld et al., 2020) are used for binary classification.

A VQC is a parameterized quantum circuit, a circuit type containing adjustable gates with tunable parameters. These gates are a universal set of quantum gates and, in the current study, rotation $[R_X(w), R_Y(w), R_Z(w)]$ and CNOT gates are used.²

The considered VQC pipeline used has the following components:

- **Data Embedding:** the numerical vector X representing the classical information is converted to the quantum space with the preparation of an initial quantum state, $|\psi_X\rangle$, which represents a HEP event.
- **Model circuit:** a unitary transformation $U(w)$, parameterized by a set of free parameters w , is applied to the initial quantum state $|\psi_X\rangle$. This produces the final state $|\psi'_X\rangle = U(w)|\psi_X\rangle$.
- **Measurement:** a measurement of an observable is performed in one of the qubits of the state $|\psi'_X\rangle$, which will give the prediction of the model for the task at hand. The training of a VQC aims to find the best set of parameters w to match the event labels to the prediction.

Throughout this work, the *PennyLane* package (Bergholm et al., 2018) was used as a basis for the hybrid quantum-classical machine learning applications. Leveraging *PennyLane*'s `default.qubit` quantum simulator, a straightforward tool

¹ *Classical* is used throughout the paper as opposed to *quantum* machine learning.

² Even if, in general, the phase shift gate $P(w)$ should be included, this gate does not change the final outcome (i.e., it does not impact probabilities), so it can be discarded.

for quantum circuit simulations, we trained and assessed the performance of various QML algorithms. Subsequently, the performance of the algorithms trained on IBM's quantum computers was gauged by integrating *PennyLane* with IBM's quantum computing framework, *Qiskit* (Anis et al., 2023).

2.1 Data embedding

Before passing the data through the VQC, the preparation of the initial quantum state $|\psi_X\rangle$ is required. This is called data embedding, and there are a number of proposals to perform this step (LaRose and Coyle, 2020). Among the different possible embeddings, it was chosen to test amplitude embedding against angle embedding. The preliminary results have shown that angle embedding leads to a better performance than the former, as previously reported in a different context (Gianelle et al., 2022). In this paper angle embedding was, therefore, the adopted choice. Further studies on possible embeddings is left for future works.

For an N -dimensional vector of classical information, $X = (x_1, x_2, \dots, x_N)$, the state entering the VQC will be defined via a state preparation circuit applied to the initial state of $|0\rangle^{\otimes N}$. The information contained in X is embedded as angles: these are the values used in rotation gates applied to each qubit, thus requiring N qubits for embedding N features from the original dataset.

In the current study, the embedding is done using rotations around the x -axis on the Bloch sphere, thus defining the quantum state embedded with the classical information as:

$$|\psi_X\rangle = \bigotimes_{i=1}^N R_X(x_i) = \bigotimes_{i=1}^N \left[\cos\left(\frac{x_i}{2}\right) |0\rangle - i \sin\left(\frac{x_i}{2}\right) |1\rangle \right], \quad (1)$$

where $R_X(x) = e^{-ix\hat{\sigma}_x}$ and $\hat{\sigma}_x$ is a Pauli operator. In this embedding each of the considered features of the original dataset is required to be bound between $[-\pi, \pi]$.

2.2 Model circuit

The model circuit is the key component of the VQC and includes the learnable set of parameters. It is defined by a parameterized unitary circuit $U(w)$, with w being the set of tunable parameters, which will evolve a quantum state embedded with classical information ψ_X into the final state ψ'_X .

Analogously to the architecture of a classical neural network, the model circuit is formed by layers. Each layer is composed of an assemblage of rotation gates applied to each qubit in the system, followed by a set of CNOT gates.

A rotation gate, R , is designed to be applied to one single qubit and rotate its state. It is composed by three learnable parameters: ϕ, θ, ω , which enables the gate to rotate any arbitrary state to any location on the Bloch sphere.

$$R(\phi, \theta, \omega) = RZ(\omega) RY(\theta) RZ(\phi) = \begin{bmatrix} e^{-i(\phi+\omega)/2} \cos(\theta/2) & -e^{i(\phi-\omega)/2} \sin(\theta/2) \\ e^{-i(\phi-\omega)/2} \sin(\theta/2) & e^{i(\phi+\omega)/2} \cos(\theta/2) \end{bmatrix} \quad (2)$$

Since all the learnable parameters of the VQC are contained inside the rotation gates, and each gate has three parameters, the shape of the weight vector is $w \in \mathbb{R}^{n \times l \times 3}$, where n is the number of qubits of the current system and l is the number of layers in the network. As mentioned in the previous section, n will depend on the number of features in the data and l is a hyper-parameter (HP) to be tuned.

After rotating the qubits' state, a collection of CNOT gates will be applied to entangle the qubits. The CNOT gate is a 2-qubit gate with no learnable parameters. It will flip the state of the so-called target-qubit, based on the value of the control-qubit, and it is usually represented by having two inputs as such: CNOT(control-qubit, target-qubit). Given the number of qubits, the CNOT arrangement is implemented as detailed in Algorithm 1.

```

Require:  $n \geq 2$ ,  $n$  being the number of qubits.
if  $n == 2$  then
    CNOT(1,0)
else
    for qubit  $\leftarrow 0$  to  $n-1$  do
        if qubit ==  $n-1$  then
            CNOT(qubit,0)
        else
            CNOT(qubit, qubit + 1)
        end if
    end for
end if

```

Algorithm 1. CNOT arrangement.

2.3 Measurement

The output of the model is obtained by measuring the expectation value of the Pauli $\hat{\sigma}_z$ operator in one of the qubits of the final state ψ'_X . An example of the implementations of a VQC is represented in Figure 1.

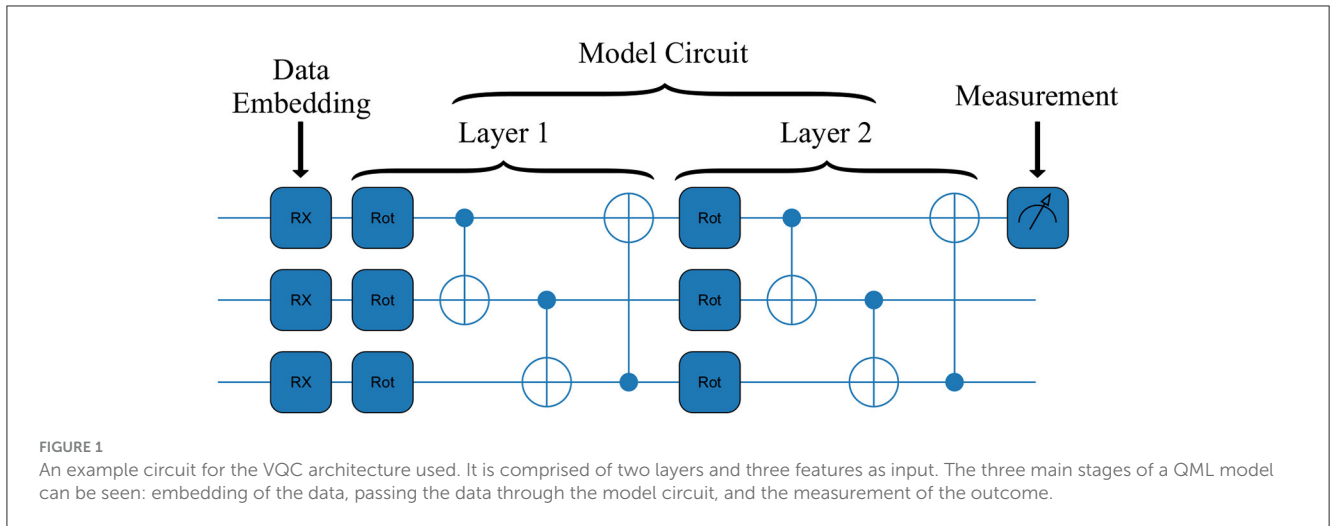
3 Classical machine learning methods

Shallow CML methods are used to provide a baseline comparison to the QML models. The specific methods chosen for the comparison are Logistic Regression (LR) and Support Vector Machines (SVM), with these algorithms being trained with the same data as the QML algorithms.

All the classical methods were implemented using *scikit-learn* (Pedregosa et al., 2011) library and, if not specified otherwise, the default parameters were used.

3.1 Logistic regression

Logistic Regression is one of the simplest ML models and can be formulated as one of the basic building blocks of a neural network—a single-layer perceptron. The goal is to find the best set of weights



w that fit the data x :

$$\hat{y}(w, b, x) = \sigma(w \cdot x + b), \quad (3)$$

where \hat{y} is the probability of an event to belong to class 1, w , and b are learnable parameters, and σ is the sigmoid function.

The learning process is guided by minimizing the loss function, which in our case is the binary cross-entropy:

$$L = -\mathbb{E}_x[y \log(\hat{y}) + (1 - y) \log(1 - \hat{y})], \quad (4)$$

where y is the binary label of whether the event is of the class signal or not, and \mathbb{E}_x is the expectation value over the training data, obtained using the event weights corresponding to each signal and background process.

3.2 Support vector machine

An SVM classifier is trained by finding the hyperplane that best separates two classes of data in the hyperspace of features. It does so by using support vectors, which are the data points from the two classes closer to the hyperplane, influencing the position and orientation of the hyperplane.

The loss function of an SVM revolves around the goal of maximizing the margin, i.e., the distance between the hyperplane and the nearest data point from either class. In other words, the goal is to find the hyperplane with the greatest possible margin between itself and any point within the training set, giving a greater chance of new data being classified correctly.

Just like the Logistic Regression, the base SVM classifier can only learn a linear decision boundary. However, classification problems are rarely simple enough for it to be separable using a hyperplane, thus usually requiring a non-linear separation. SVM can do this by transforming the data using a non-linear function, named kernel, after which it can be split by a hyperplane. For this implementation, the radial-basis function (RBF) was used as kernel. This endows the SVM with a non-linear mapping where it better separates the two classes using a hyperplane.

4 Dataset

The dataset used in this work (Crispim Romão et al., 2021) is comprised of simulated events of pp collisions at 13 TeV, in final states with two leptons, at least 1 b -jet, at least 1 large- R jet, and large scalar sum of transverse³ momentum (p_T) of all reconstructed particles in the event ($H_T > 500$ GeV). Such basic selection corresponds to a topology commonly used in different searches for BSM events at the LHC (Crispim Romão et al., 2021). The dominant SM background for this topology, $Zb\bar{b}$, and the BSM signal corresponding to $t\bar{t}$ production with one of the top-quarks decaying via a flavor changing neutral current decay $t \rightarrow qZ$ ($q = c, u$) (Durieux et al., 2015), were considered. Such signal was chosen given the kinematic similitude to the background, thus providing a good benchmark for the present study.

Both samples were generated with MADGRAPH5 2.6.5 (Alwall et al., 2014) and PYTHIA 8.2 (Sjöstrand et al., 2015), and the detector was simulated using DELPHES 3 (Selvaggi, 2014) with the default CMS card. Jets were clustered using the anti- k_t algorithm (Cacciari et al., 2008), implemented via FASTJET (Cacciari et al., 2012), with R -parameters of 0.5 and 0.8 (the latter for the large- R jets).

The following features were used for training of both the classical and quantum machine learning algorithms:

- $(\eta, \phi, p_T, m, b\text{-tag})$ of the five leading jets, ordered by decreasing p_T , with $b\text{-tag}$ being a Boolean variable indicating if the jet is identified as originating from a b -quark by the b -tagging algorithm emulated by DELPHES;
- (η, ϕ, p_T, m) of the leading large- R jet;
- N -subjettiness of leading large- R jet, τ_n with $n = 1, \dots, 5$ (Thaler and Van Tilburg, 2011).
- (η, ϕ, p_T) of the two leading leptons (electrons or muons);
- transverse momentum (\cancel{E}_T) and ϕ of the missing transverse energy;
- multiplicity of jets, large- R jets, electrons, and muons;

³ The transverse plane is defined with respect to the proton colliding beams.

TABLE 1 Top five features ranked by their AUC score on the training dataset.

Feature	AUC
\cancel{E}_T	0.817
Lepton ₁ p_T	0.692
Lepton ₂ p_T	0.649
large- R jet m	0.609
large- R jet τ_1	0.576

- H_T .

The proportion of signal and background events was kept the same as the original simulated data during training, being 13 and 87%, respectively. Additionally, the Monte-Carlo weights, corresponding to the theoretical prediction for each process at target luminosity of 150 fb^{-1} , were taken into account in the evaluation of all the considered metrics and loss functions.

5 Feature selection

As described in the previous section, a total of 47 features are available for training. Considering the type of data embedding chosen, 47 qubits would be needed to train a VQC using all the dataset features. Such number of qubits is impractical given the currently available quantum computers and thus it is not feasible to train a VQC using all the features in our dataset. For the purposes of the current study, quantum computers with only five qubits were considered and two methods for feature selection were implemented: principal component analysis (PCA) and sequential feature selection (SFS).

A relative comparison of the best five features⁴ is shown in Table 1 while the best performance obtained with state-of-the-art CML methods without any features or data points restrictions can be seen in Figure 2.

5.1 Sequential feature selection

SFS algorithms are a widely used family of greedy search algorithms used for automatically selecting a subset of features that is most relevant to the problem. This is achievable by removing or adding one feature at a time based on the classifier performance until a feature subset of the desired size, k , is reached.

There are different variations of SFS algorithms but for the current paper, the Sequential Backward Selection (SBS) algorithm was chosen. This algorithm starts with the full set of features ($n = 47$) and, at each iteration, it generates all possible feature subsets of size $n - 1$ and trains a ML model for each one of the subsets. The performance is subsequently evaluated and the feature that is absent from the subset of features with the highest performance metric is

⁴ The area under the curve (AUC) of the receiver operating characteristic curve (ROC) is considered as metric for these comparisons.

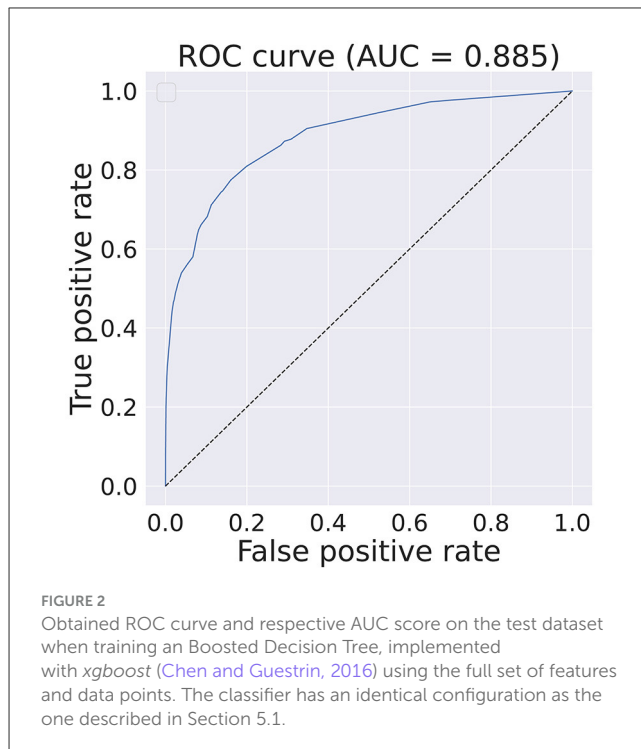


FIGURE 2 Obtained ROC curve and respective AUC score on the test dataset when training an Boosted Decision Tree, implemented with *xgboost* (Chen and Guestrin, 2016) using the full set of features and data points. The classifier has an identical configuration as the one described in Section 5.1.

TABLE 2 List of the features selected by the SBS algorithm for $k = 1, \dots, 5$.

k	Selected features
1	\cancel{E}_T
2	\cancel{E}_T , Number of muons
3	\cancel{E}_T , Number of muons, Jet ₁ b -tag
4	\cancel{E}_T , Number of muons, Jet ₁ b -tag, Jet ₂ p_T
5	\cancel{E}_T , Number of muons, Jet ₁ b -tag, Jet ₂ p_T , large- R τ_3

removed. This process is iterated until the feature subset contains k features.

This technique was used to find subsets of 1–5 features. The ML model assisting the SBS was a boosted decision tree (BDT) with a maximum number of estimators set at 100 and a learning rate of 1×10^{-5} . The considered loss function was a logistic regression for binary classification and the AUC score was used as evaluation metric. The BDT was implemented using *xgboost* (Chen and Guestrin, 2016) and the SBS algorithm using *mlxtend* (Raschka, 2018). The selected features for the different values of k is shown in Table 2 and the AUC scores for each feature in Table 3. It should be noted that Table 2 shows the features selected with the SBS algorithm and Table 3 shows the AUC value of each one of these features. The latter is ordered by descending AUC value.

5.2 Principal component analysis

The PCA transforms a highly correlated, high-dimensional dataset and into a new one with reduced dimensionality and uncorrelated features, by rotating the dataset in the direction of the

TABLE 3 Features selected by the SBS algorithm and their respective AUC score on the training dataset.

Feature	AUC
\not{E}_T	0.817
Number of muons	0.534
Jet ₁ b -tag	0.418
large-R jet τ_3	0.316
Jet ₂ p_T	0.313

TABLE 4 Top five PCA components obtained with the training dataset, ranked by their AUC.

PCA component	AUC
Component 3	0.726
Component 14	0.606
Component 5	0.565
Component 41	0.563
Component 43	0.560

eigenvectors of the dataset covariance matrix. In the present paper, the PCA was performed only to remove the correlation between the features, maintaining the same dimensionality as the original data. The PCA transformation was learned from the training dataset and then applied to all datasets. When training a VQC for a specific number of features, the PCA components were ranked by AUC score and thus selected from the highest to the lowest. This is done by introducing a priority queue, i.e., if training a model using two features is desired, the two top-ranked PCA components will be selected. The *scikit-learn* PCA implementation was used and the obtained five better components are shown in Table 4.

6 Dataset size reduction

The present paper addresses the use of reduced datasets to overcome the limitation of NISQ processors while minimizing the loss of information and thus avoiding a performance loss of the QML algorithms in the HEP context. The primary method used for this purpose in the current study is *KMeans*, where the k th most representative points, i.e., a set of *centroids*, is obtained from the original dataset. Although these *centroids* are the most representative data points, they are not necessarily contained in the original dataset and, consequently, a resampling process, allowing to choose points of the original dataset (centrus), is required.

A study of the performance of the proposed dataset reduction method will be done by training a logistic regression model with the original dataset and comparing the results with those obtained when *Kmeans* and randomly undersampled datasets are used.

6.1 KMeans algorithm

Considering a clustering algorithm, *Kmeans* iteratively tries to separate data into independent groups (MacKay, 2003). This separation is done using the *Lloyd's algorithm* (Wilkin and Xiuzhen,

2008), based on the minimal variability of samples within each cluster. The *KMeans* algorithm requires the specification of the desired number of clusters (k) a priori. The following steps were used:

1. **Initialization of the centroids:** using the *scikit-learn* implementation, it is possible to do it in two different ways, *random* and *k-means++* (Vouros et al., 2021):

- **Random:** k random samples of the original dataset are chosen.
- **K-means++:** k samples of the original dataset are chosen based on a probabilistic approach, leading to the centroids being initialized far away from each other.

Assuming there is enough time, the algorithm will always converge, although the convergence to an absolute minimum is not guaranteed. The *K-means++* initialization helps to address this issue. Furthermore, for both initializations, the algorithm, by default, runs several times with different centroid seeds, with the best result being the output.

2. **Assignment:** Each data point x_i is addressed to a cluster c_k , in such a way that the *inertia* is minimized:

$$k' = \underset{k}{\operatorname{argmin}} \left\{ \sum_{j=0}^{F-1} (x_{ij} - \mu_{kj})^2 \right\}, \quad (5)$$

where F is the dimensionality, i.e., the number of features, μ_k is the centroid of the cluster c_k and j stands for the $(j+1)$ th feature.

3. **Update of the centroids' position:** The new centroids are just the means positions of each cluster, i.e.,

$$\mu_k = \frac{\sum_{i=0, x_i \in c_k}^{n_k-1} x_i}{n_k}, \quad (6)$$

with n_k being the number of samples addressed to c_k . It should be noted that if $n_k = 0$ the centroid $\mu_{k'}$ doesn't change.

4. **Iteration:** Steps 2 and 3 are repeated until the maximum number of iterations is reached or until the result converges, i.e., the centroids don't change.

The *KMeans* algorithm was used separately for the signal and background samples, with the corresponding weights being used.

6.2 Dataset resampling

As previously mentioned, although centroids are the most representative points, they are not necessarily contained in the original dataset. Hence, it was chosen to consider 10 neighbors of each centroid to determine each centrus, i.e., the 10 nearest points of the original dataset.

The position of each centrus was determined using the weighted mean of the position of the neighbors,

$$W = \frac{\sum_{i=0}^9 x_i \times w_i}{\sum_{i=0}^9 w_i}, \quad (7)$$

where W is the mean position, x_i is the $(i+1)$ th nearest point and w_i the weight of the sample.

TABLE 5 List of scanned hyperparameters.

Variable HP	Possible values
Feature selection	[PCA, SBS]
Number of data points	[100, 500, 1 k, 5 k]
Number of features	[1, 2, 3, 4, 5]
Number of layers	[1, 2, 3, 4, 5]
Fixed HP	Fixed values
Max epochs	500
Batch size	Size of the dataset
Learning rate (LR)	0.03

The LR parameter is used only for the VQC optimized by Adam while the number of layers is only used by the VQCs.

The sample weight of each centrus was calculated based on the number of samples of the same label (i.e., signal or background) on the original dataset:

$$w_i = \frac{1}{n}, \quad (8)$$

with w_i being the weight of the $(i + 1)$ th centrus and n the number of samples in the original dataset with the same label of this centrus.

7 Quantum and classical machine learning training

The training of the QML algorithms used in the current paper requires the use of optimizers. Two different ones were considered: *Adam* (Kingma and Ba, 2014) and tree-structured Parzen estimator sampler (*TPE*; Bergstra et al., 2011, 2013).

The *Adam* optimizer uses an extension of stochastic gradient descent, leveraging techniques such as adaptive moment estimation, being extensively used in optimization problems, namely in the context of machine learning. Nonetheless, since there is no reason to expect, *a priori*, that it will work equally well in the context of QML, where specific challenges are expected, the *TPE* optimizer was also tested.

The *TPE* is a Bayesian optimization algorithm first developed for HP tuning in the context of machine learning. In the current study, it will be used to optimize VQC weights in a way very similar to what is typically done for HP tuning. *TPE* is implemented using *Optuna* (Akiba et al., 2019), a library focused on HP optimization for machine learning models. *TPE* works by choosing a parameter candidate that maximizes the likelihood ratio between a Gaussian Mixture Model (GMM) fitted to the set of parameters associated with the best objective values, with another GMM being fitted to the remaining parameter values. In the context of HEP, *TPE* has also been used to explore parameter spaces of BSM models (de Souza et al., 2022).

Different machine learning methods were optimized, namely a LR, a SVM, and a VQC. The corresponding training was done for the set of HP summarized in Table 5, where the scanned values are also listed. For each set of HP, 5 models were trained on 5 different subsets of the initial dataset (random sampling).

For both optimizers, the considered cost function used is the squared error, with the individual Monte Carlo samples being properly weighted. During the training of VQCs, the inference was done on the validation dataset at five epoch intervals, the AUC computed and only the best-performing model, according to the previously mentioned metric, was considered.

7.1 Adam implementation details

The training starts with the initialization of the weight vector. This is done randomly with an order of magnitude of 10^{-2} , which is followed by training iterations until a maximum number of epochs is reached. At each iteration, the model is inferred with the training dataset, the cost function calculated and the model parameters updated via the *Adam* optimizer. A summary of *Adam*-optimized VQC training is shown in Algorithm 2.

```

params ← params_initialization()
for epoch ← 1 to max_epochs do
    loss ← cost(params)
    params ← optimizer.step()
    if epoch_number%5 == 0 or epoch_number == max_epochs
    then
        validation_step()
    end if
end for

```

Algorithm 2. Adam training.

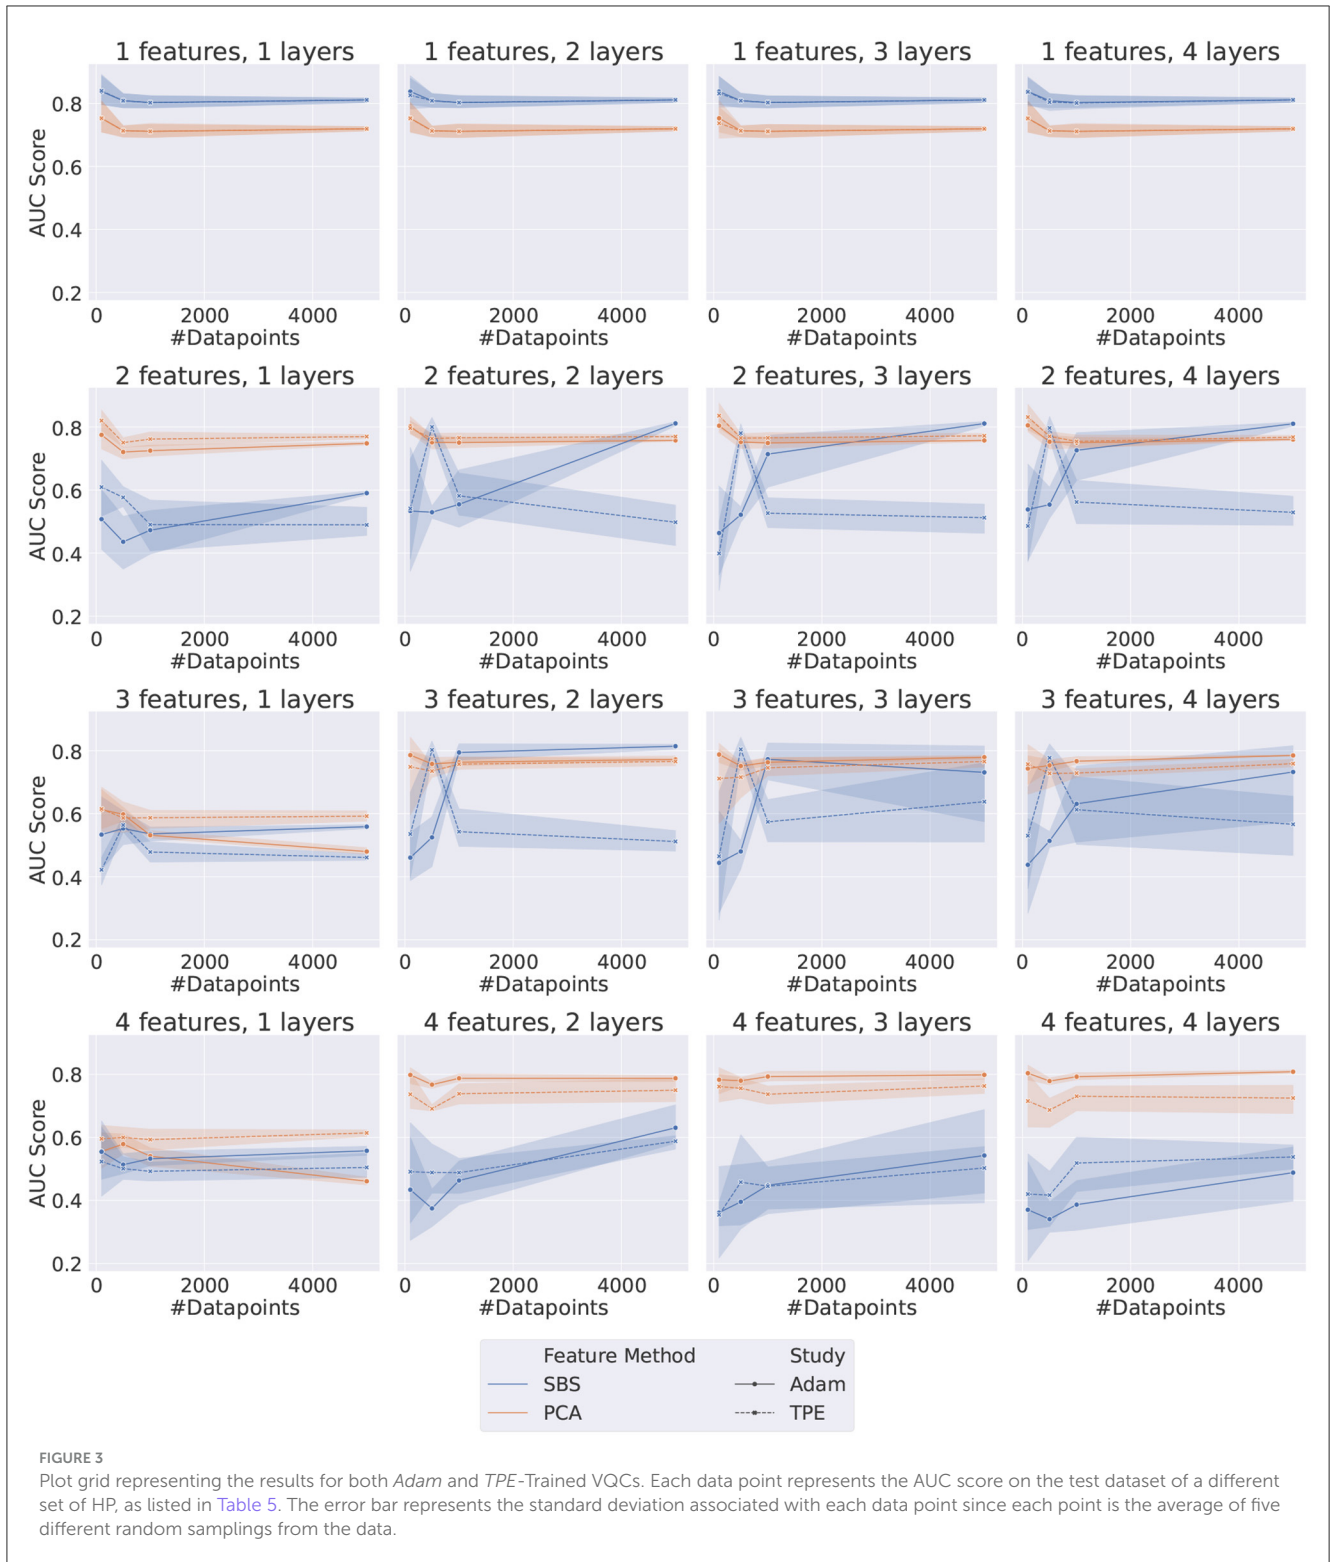
7.2 TPE implementation details

We use the *Optuna* implementation of the *TPE* sampler. Being a Bayesian optimization algorithm, *TPE* works very differently to *Adam*, which is a gradient descent algorithm. In *TPE*, for every training iteration, each parameter is replaced by a new value acquired sampling from a Gaussian Mixture Model of good points, which is then used to compute the loss function. At each epoch, the algorithm computes new values for the model parameters. With the value of the loss function of the suggested parameters, *TPE* will update its internal Gaussian Mixture Models of good and bad points, which will allow it to learn what are good suggestions as more parameter values are sampled. Since *TPE* is a Bayesian algorithm, it does not need to compute derivatives of the loss function, as *Adam* does, which might allow for a light workload when running trainings on quantum computers.

8 Simulation results

8.1 Feature reduction

The results indicate that QML circuits trained with SBS data are generally unstable and very susceptible to fluctuations in the randomly sampled data, as can be seen in Figure 3. Specifically, it is



evident that using PCA-originated data produces significantly more stable results.

The performance of both optimizers, *Adam* and *TPE*, is usually saturated with only two layers. This effect is most noticeable when the number of features is greater or equal to 3. When considering only the PCA-obtained results, the

two optimizers are compatible for most of the configurations tested. Exceptions occur when using a high number of features (≥ 4) and only one layer, where *TPE* outperforms *Adam*, and when using a high number of features (≥ 4) and more than one layer, where the opposite happens and *Adam* outperforms *TPE*.

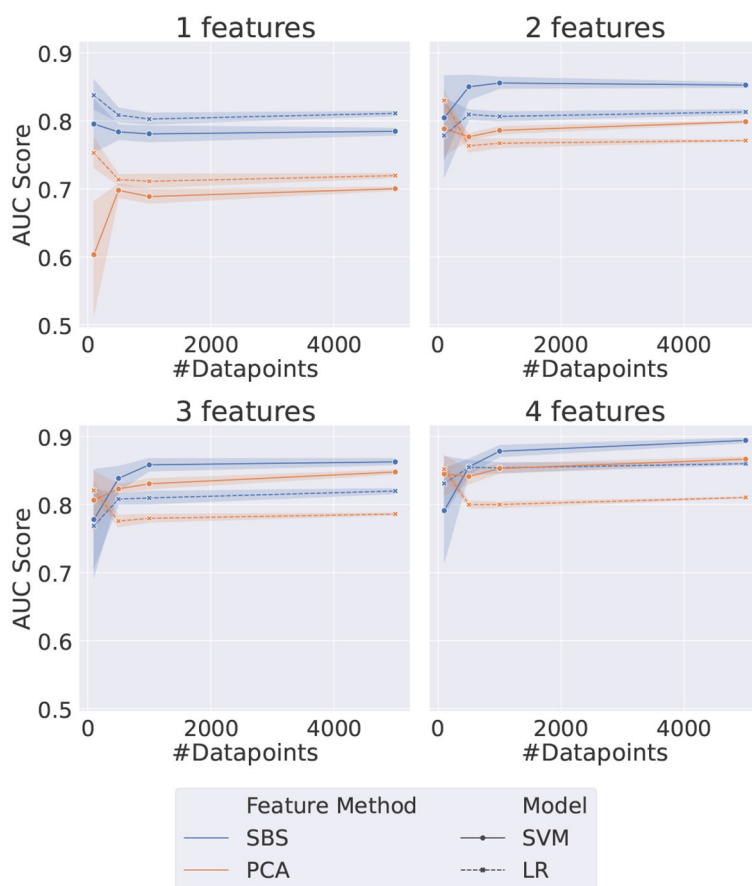


FIGURE 4 Plot grid representing the results for the considered shallow methods. Each data point represents the AUC score on the test dataset of a different set of HP, as listed in Table 5. The error bar represents the standard deviation associated with each data point since each point is the average of five different random samplings from the data.

The shallow ML methods trained on the same data as the VQCs are shown in Figure 4. The AUC scores obtained in this case are more stable for both the PCA and SBS datasets. The performance in both cases is saturated when using two features and the models trained with SBS data outperform the PCA-trained models, contrary to what was observed for the QML case. It should also be noted that the SVM outperforms LR in all cases except when only one feature is used, which is not surprising since SVMs are more sophisticated classifiers.

For the best set of HP, VQCs trained using *TPE* and *Adam* have performed similarly to the shallow ML methods (c.f. with Figure 5, Figure A1, respectively). It was also observed that there are no cases where QML outperforms any of the shallow methods tested. The *TPE* optimizer regime produced the best performance for QML, achieving an AUC score of 0.841 ± 0.051 , as shown in Figure 5.

The reduction algorithms studied come with an additional computational cost compared to using the original dataset directly. In particular, the SBS algorithm added an overhead of 1 h for running the XGBoost algorithm and selecting the features with more classification power. On the other hand, the PCA algorithm took a sub-minute negligible time to complete. However, since these algorithms only need to be run once, before the training, and given that the grid search for the VQC, SVM, and LR algorithms took over 200 CPU hours on a dual-Intel(R) Xeon(R)

Gold 6,348 machine, in the end the computational cost of the classical reduction algorithms is negligible.

8.1.1 VQC’s robustness to discrete features

In the previous section it was noted that there was a significant variability in the final score of QML models, especially when training with SBS data. In fact, VQCs, being variational algorithms, are highly susceptible to small fluctuations in the data, which can have a correspondingly significant impact on the computed AUC. Additionally, numerical instabilities caused by computational floating point accuracy were observed during the validation step, leading to considerable fluctuations in the computed AUC in this regime.

To further investigate this behavior, which was not observed at the same level on the PCA-trained circuits, we looked at the AUC distributions produced by QML models as a function of the number of features. As shown in Figure 6, it is clear that the instability in SBS results occurs when more than two features are used. The biggest difference in the AUC mean is found for four features, where the value for SBS is 0.471 ± 0.129 and for PCA is 0.719 ± 0.096 . The smallest difference is found for 1 feature, where the value for SBS is 0.814 ± 0.035 and for PCA is 0.724 ± 0.037 .

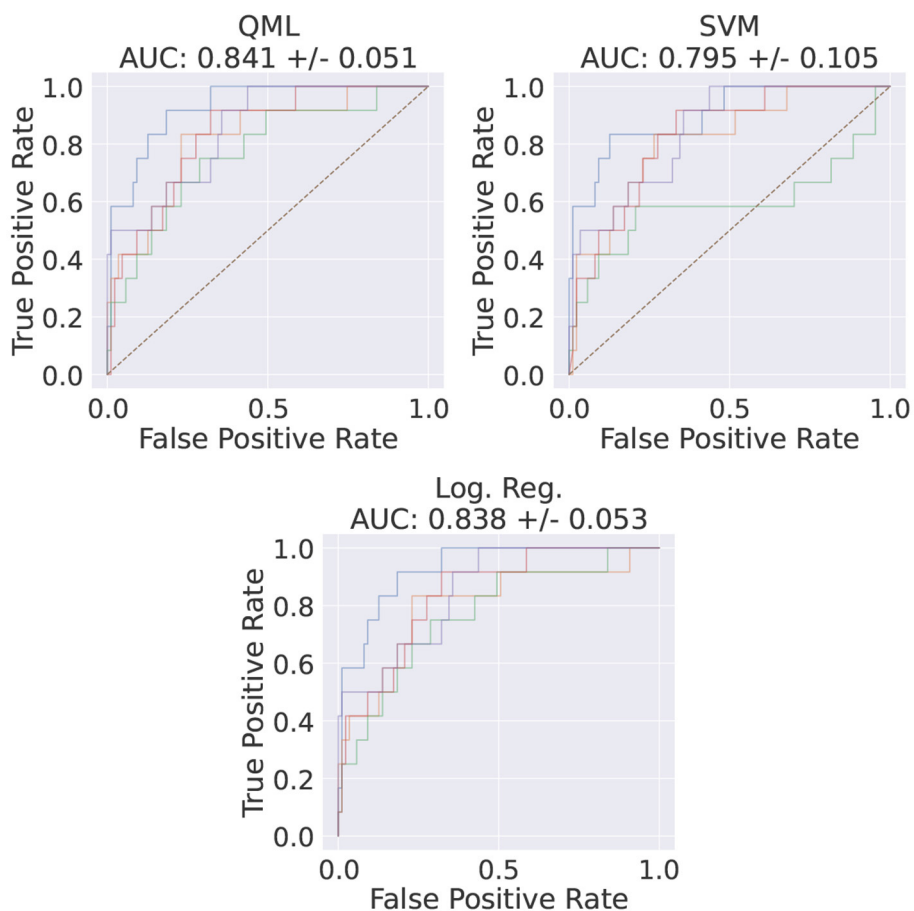


FIGURE 5 ROC of the best HP set, using *TPE*'s QML model average AUC score as a metric and the corresponding shallow methods ROCs for the same data. The HP for this run are SBS for feature method, 100 data points, 1 feature, and 5 VQC layers. The different colors indicate the different random samplings of the data.

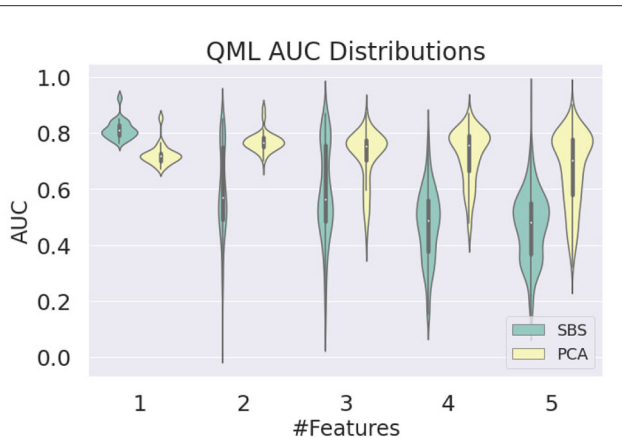


FIGURE 6 Distribution of the AUC values obtained for the QML model as a function of the number of features used in training, evaluated on the test dataset, for SBS and PCA inputs.

Additionally, we produced visualizations of the decision regions of the models trained using both feature selection methods. We focused on runs that used two features, as this is where

the problem originated. Figures 7, 8 show the decision regions obtained with each model for one representative run, illustrating the sensitivity of each boundary to variations in the data, for SBS or PCA. The SBS features used are listed in Table 2, where the second feature, the number of muons in the event, is a discrete variable.⁵

While LR and SVMs are robust in the presence of discrete variables, they may pose a challenge for continuous learning algorithms such as VQCs. It is therefore possible that the variability observed when using different sub-samples of SBS data could be attributed to the use of this discrete variable. To investigate this, we conducted the SBS feature selection once again, this time excluding all discrete variables—yielding Table 6. The VQC circuits were once again trained using this modified list of inputs in a limited study of two features only, as illustrated in Figure 9.

Using the discrete-free SBS version to train the VQC led to significantly better AUC scores, outperforming PCA-trained QML models with an average AUC score of around 0.85, although still

⁵ In ML literature this is called a categorical variable. However, we note that even though it is categorical, it is still ordinal. As there are no non-ordinal or non-binary categorical variables in our dataset, we will refer to these variables as discrete instead of categorical for the rest of this work.

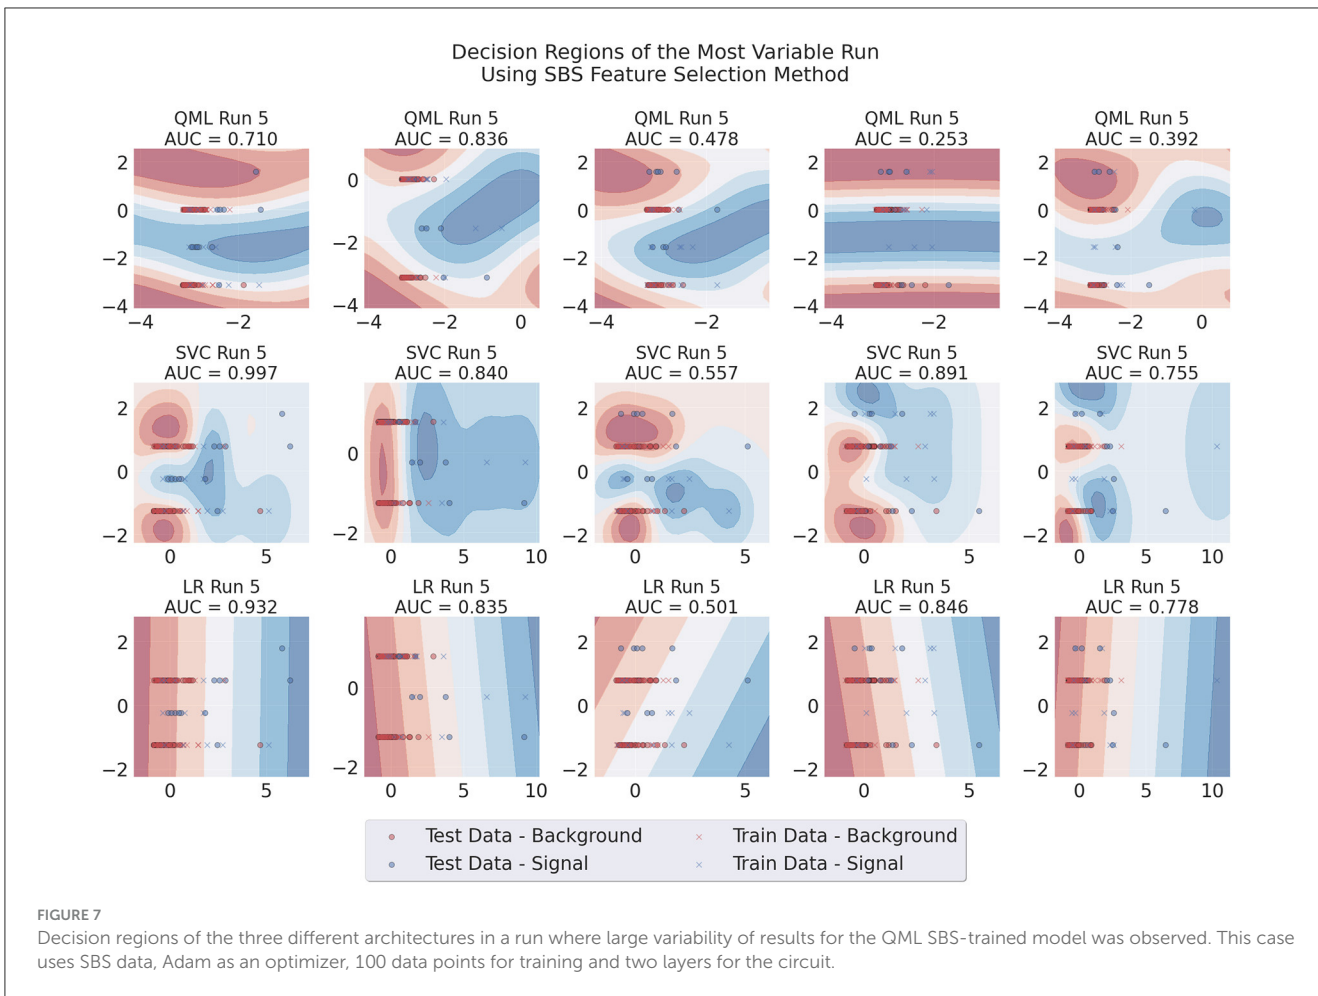


FIGURE 7 Decision regions of the three different architectures in a run where large variability of results for the QML SBS-trained model was observed. This case uses SBS data, Adam as an optimizer, 100 data points for training and two layers for the circuit.

with larger variability than that of the PCA-trained VQCs. This is a notable departure from our previous observations in Section 8.1, where including discrete features in the SBS feature selection methodology resulted in erratic performance with no instance of outperforming PCA (except in cases where only one continuous feature was used). Therefore, we found that excluding discrete variables during feature selection led to better performance for VQC circuits in a limited study of two features, compared to when discrete variables were included. This indicates that the choice of input features is crucial for achieving high accuracy in quantum machine learning, and future studies should consider the impact of discrete variables on VQC performance. The findings may inform future choices in selecting input features for VQC circuits to optimize model performance.

8.2 Dataset reduction

8.2.1 Implementation of *KMeans*

The performance of the *KMeans* algorithm was tested initially by training LR models with 10 reduced datasets and selecting a different number of k features ($k \in [1, 2, 3, 4, 5]$) obtained with the SBS algorithm. The *KMeans* algorithm considers the sample weight and, in order to have an equal number of signal and background centroids, it was separately applied to the signal and background data. Since state-of-the-art quantum computing requires small

datasets, the data reduction studies were done for datasets with 100, 500, 1,000, and 5,000 data points and the number of features previously mentioned.

Two configurations were studied: the framework presented in Section 6.1 was applied to the training and test datasets; and only to the training datasets (with test datasets obtained through random undersampling).⁶

The mean AUC score and respective standard deviation found using *KMeans* for train and test datasets are summarized in Figure 10. The results obtained using the *KMeans* algorithm for the training dataset and random undersampling in the test signal and background samples are presented in Figure 11. In order to provide a benchmark point for comparison with the performance of the reduced datasets, a LR model was trained on the full original dataset, with results shown in both figures.

It can be seen in Figures 10, 11 that using the *KMeans* algorithm to reduce the training dataset results in AUC scores that are compatible with the performance obtained using the full original dataset.

⁶ Throughout this article random undersampling refers to the random selection of data points from the original dataset. In the ML subfield of imbalanced learning, the proper methodology is to use resampling algorithms only during training, but not during validation or test. In this section we present results of these two cases as a comparison, but later we will restrict to random undersampling during validation and testing.

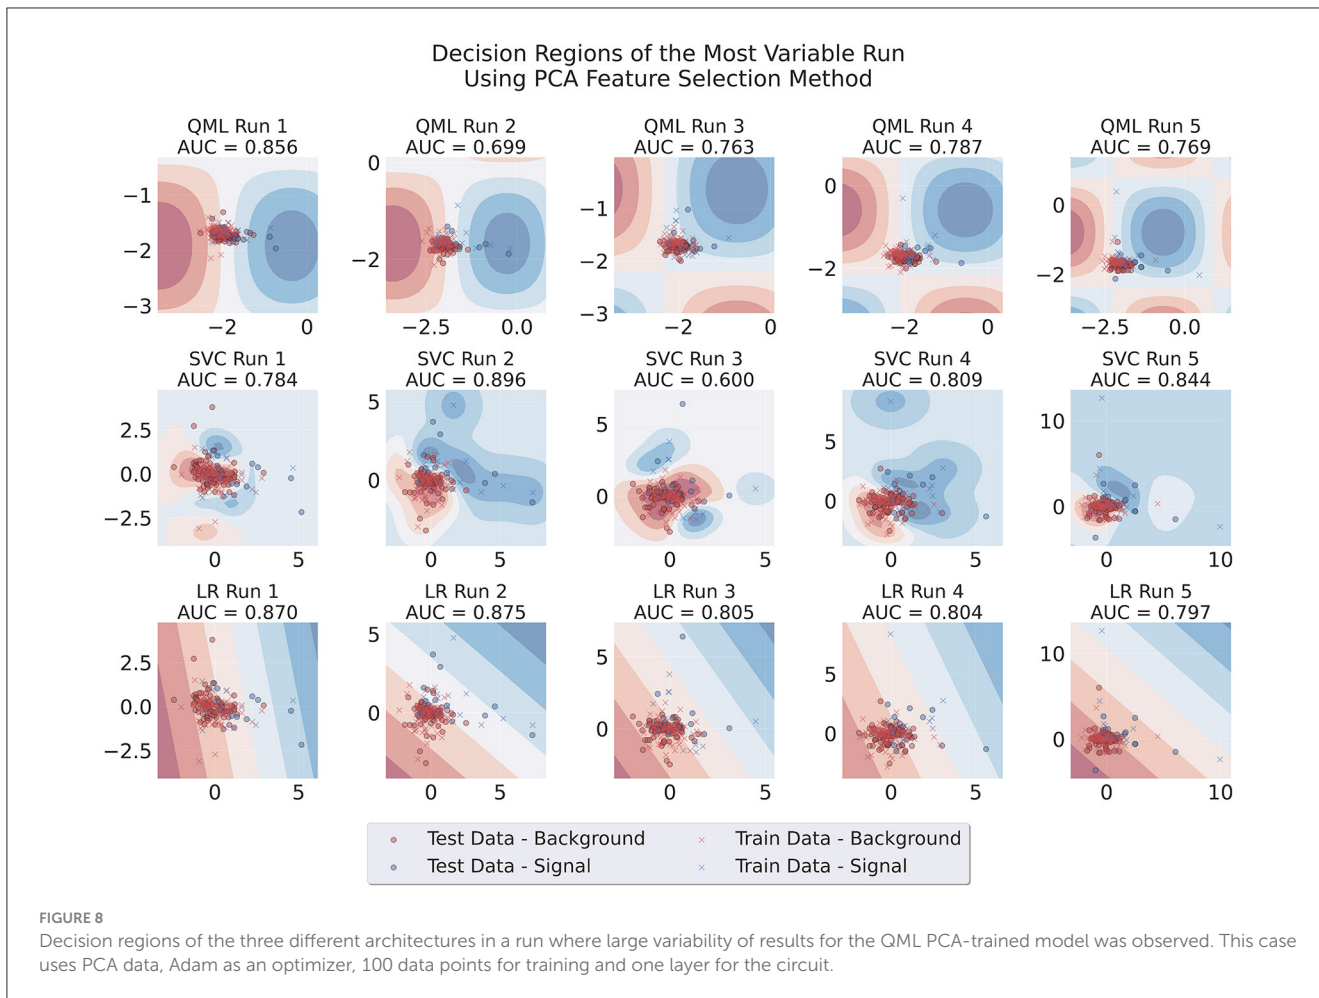


FIGURE 8 Decision regions of the three different architectures in a run where large variability of results for the QML PCA-trained model was observed. This case uses PCA data, Adam as an optimizer, 100 data points for training and one layer for the circuit.

TABLE 6 Features selected by the SBS algorithm and their respective AUC Score on the training dataset with all the discrete features removed.

Feature	AUC
\not{E}_T	0.817
large-R jet τ_1	0.576
large-R jet τ_3	0.316
Jet ₂ p_T	0.313
Jet ₁ p_T	0.292

This study shows that although *KMeans* is a more sophisticated algorithm for data reduction than random undersampling, in the HEP case under consideration no significant deterioration of the performance is observed when using it, suggesting that in this study the dataset composed of prototypes is a good representative of the whole dataset in the small dataset regime, which is explored in this work.

8.2.2 Application to QML

The QML, SVM, and LR models were trained using *KMeans* reduced datasets as well as random undersampling, for different dataset sizes. In this comparison, the HP for the VQC are the ones previously found to be the best, i.e., one feature chosen with the SBS

method and five VQC layers for the architecture. The metric used to compare all models is the AUC score average of five different runs.

For all cases, the test and validation sets were reduced using random undersampling, hence, for each dataset size there are one train, five validation and five test datasets. The choice to keep random sampling for the test dataset, rather than *KMeans* reduction, is to ensure that our methodology represents the test samples as close to the original dataset as possible, ensuring that sophisticated resampling techniques do not significantly modify the data.

The obtained results are shown in Figure 12. It can be seen that the performance for the *KMeans* reduced dataset is compatible with the one obtained using the dataset reduced through random undersampling, for QML and CML models. Furthermore, the performance achieved by the simulated VQCs is identical within the statistical uncertainties to the performances by the SVM and LR, in agreement to what was observed in Section 8.1.

Nonetheless, it should be emphasized that the model trained with random undersampling needs to be trained several times for achieving these average scores, as many times as the number of reduced datasets used. On the other hand, the models using the *KMeans* reduced dataset need to be trained only once. This can be relevant in the context of quantum computers, where access is often subject to long queues and thus the number of accesses can be a limiting factor. While the *KMeans* reduction technique brought

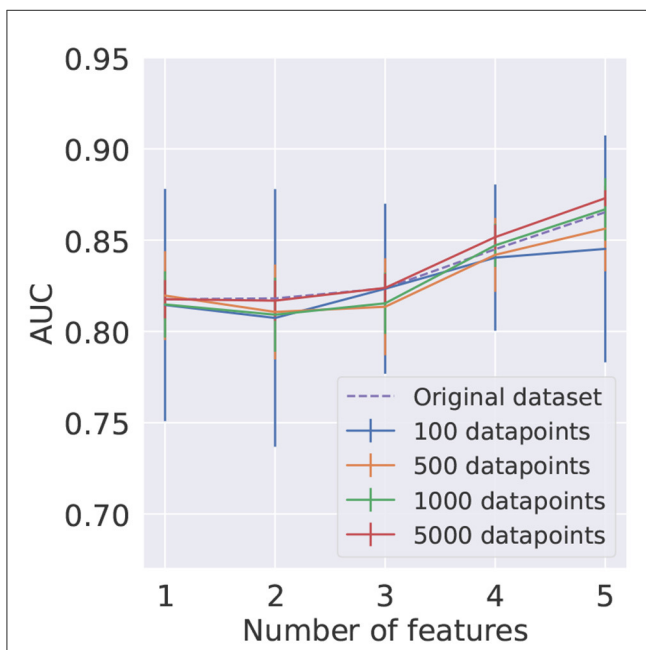
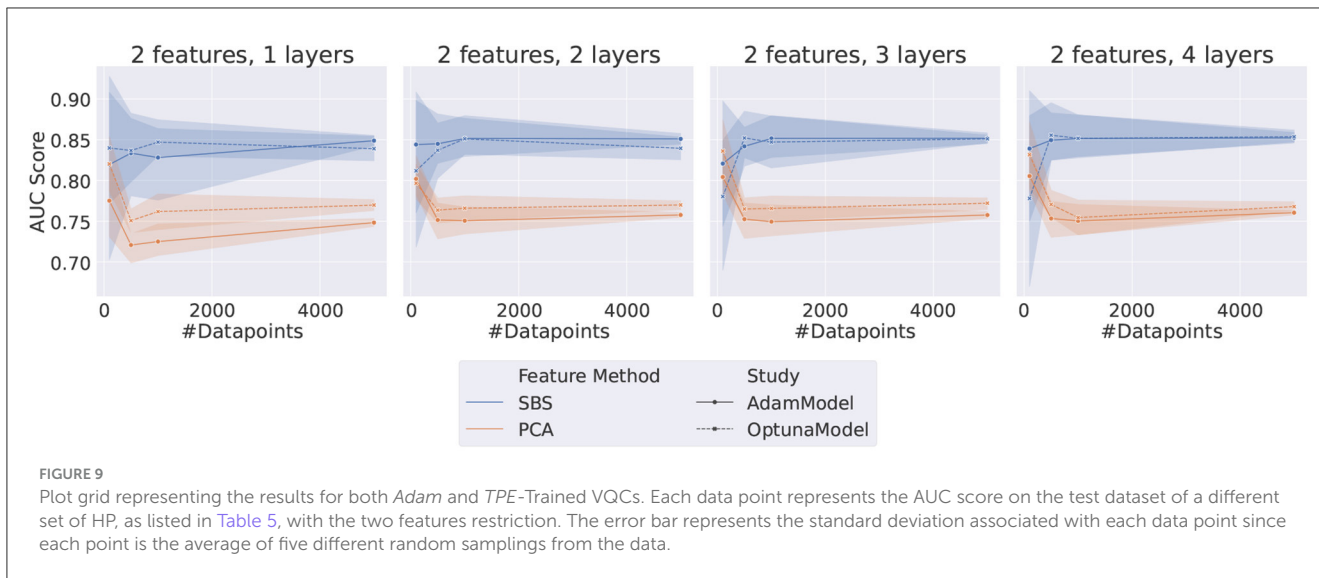


FIGURE 10
Average AUC score and corresponding standard deviation, represented as uncertainty bands, for different numbers of clusters as a function of the number of features. The training and testing datasets were reduced using the *KMeans* algorithm. In each case, 10 different reduced test datasets were used.

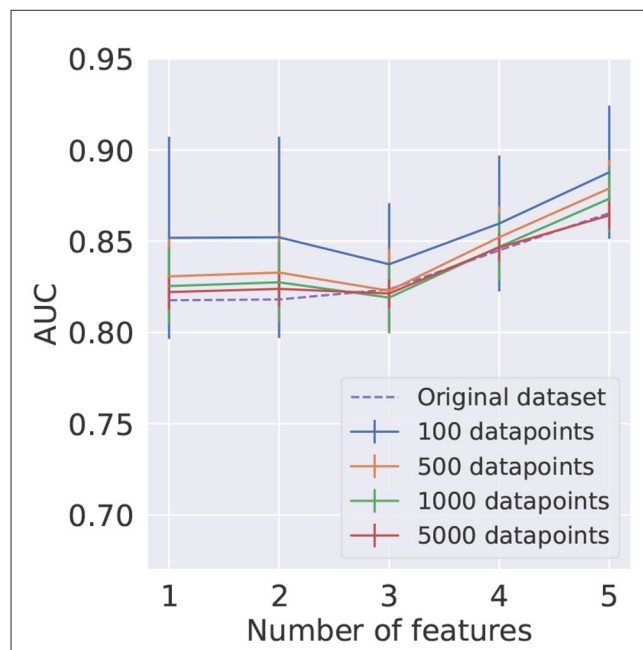


FIGURE 11
Average AUC score and corresponding standard deviation, represented as uncertainty bands, for different numbers of clusters as a function of the number of features. The training dataset was reduced using the *KMeans* algorithm. In each case, 10 different randomly undersampled test datasets were used.

an overall increase in time of around 1%, this change is negligible taking into account the reduction in number of accesses.

9 Real quantum computers results

Until this point, only simulated quantum environments were used. In order to test the performance in real quantum computers, and thus validate the simulation results, the *PennyLane* framework was used as the integration layer with *Qiskit*, which works

effectively as a direct API to the quantum computers provided by IBM.

In this study, only the best performing model HP-set was used, i.e., the *TPE*-trained VQC. This VQC was implemented and its test set was inferred on six different quantum systems with identical architectures, all freely available. Evaluating our model in multiple identical quantum systems allows us to get an idea of the scale of the associated systematic uncertainty via the variability of the observed results. Since the implemented circuits are small, no

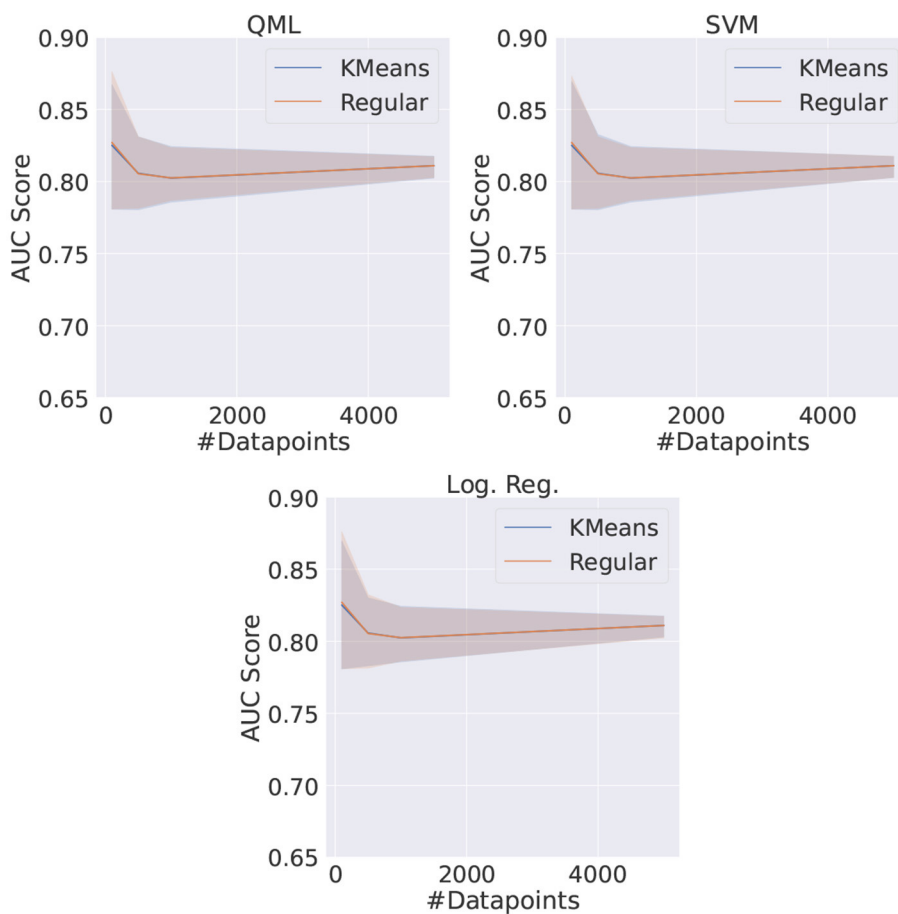


FIGURE 12 Comparison between the QML, SVM, and LR models when trained with the *TPE* and the best set of HP for different dataset sizes for both random undersampling (regular) and *KMeans* reduced datasets.

error mitigation techniques were implemented. IBM’s transpiler optimization level was set to 3⁷ (Anis et al., 2023) and, for each event, the final expectation value was computed by averaging 20k shots on the quantum computer. The obtained results, shown in Figure 13, are compatible with the simulated ones (Figure 5).

10 Conclusion

In this paper, we assessed the feasibility of using large datasets in QML applications by exploring data reduction techniques. To allow for a fair comparison between CML and QML models, we opted to use shallow classical methods as opposed to deep methods, which require large datasets that are not viable given the limitations of the current quantum computers. Our results indicate that there is comparable performance between CML and QML when tested on the same small dataset regime.

To achieve this, our study first compared feature selection techniques, showing that while SBS can produce

the best performant QML model, it generally yielded worse and more unstable results than PCA. Additionally, we found this was produced by using discrete variables in VQCs, highlighting the suitability of PCA-transformed data for QML applications in the HEP context, where discrete variables are commonly used to describe collider events.

Our grid search over different HP combinations of VQC ran in simulation provided no evidence of quantum advantage in our study. We confirmed the results by running the best performing configuration on real-world quantum systems, obtaining compatible performances and therefore validating our conclusions. We compared the performance of TPE and Adam optimizers in QML and found that TPE achieves competitive results. Being a gradient-free optimizer, TPE offers the advantage that it can lead to faster training with a smaller memory usage when compared to Adam, which in principle can further facilitate the application of QML in current quantum computers.

We then explored data reduction techniques, finding that reducing the dataset size with the KMeans algorithm produces results that are similar to those obtained

7 The level 3 of optimization corresponds to the heaviest optimization inherently implemented.

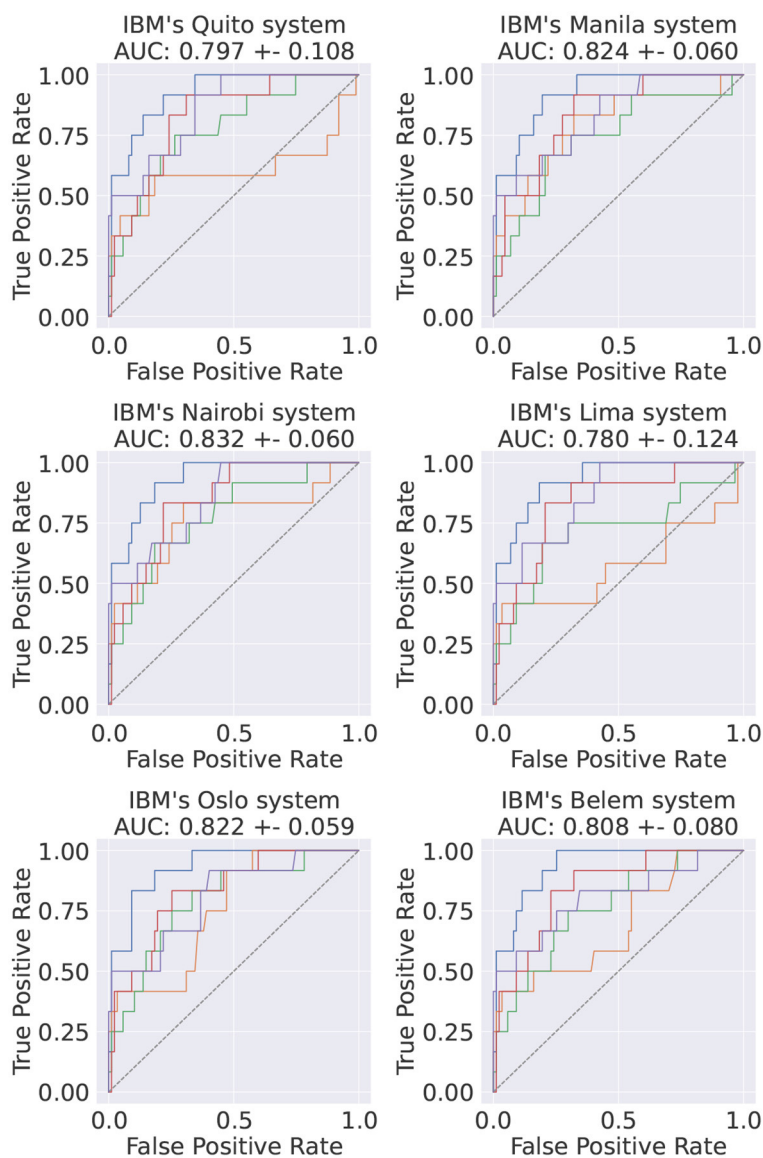


FIGURE 13 Final ROC curve of the best-performing model when inferred on the test dataset in six different IBM systems. The average AUC scores and the corresponding standard deviations are also shown. The colors in each subplot stand for different runs of the same circuits in the same QC.

from random undersampling. This finding is significant in that it means that the model can achieve similar performance with fewer accesses to a quantum computer during training, which is a considerable bottleneck in current QML.

In conclusion, while our study found no evidence of quantum advantage in the current state of QML within the context of large HEP datasets, the performance of QML models was comparable to that of classical machine learning models when restricted to small dataset regimes. Our findings suggest that using dataset reduction techniques enables us to use large datasets more efficiently to train VQCs, facilitating the usage of current quantum computers in large datasets often found in HEP.

Data availability statement

The dataset used in the studies reported in the current paper can be found in <https://zenodo.org/doi/10.5281/zenodo.5126746>, while the computational code used to obtain the present results is publicly available via <https://github.com/mcpeixoto/QML-HEP>.

Author contributions

MP: Conceptualization, Data curation, Formal analysis, Investigation, Project administration, Software, Validation, Visualization, Writing—original draft, Writing—review & editing. NC: Conceptualization, Data curation, Funding

acquisition, Investigation, Methodology, Resources, Supervision, Writing—original draft, Writing—review & editing. MR: Conceptualization, Data curation, Formal analysis, Investigation, Methodology, Software, Supervision, Validation, Writing—original draft, Funding acquisition, Writing—review & editing. MO: Conceptualization, Formal analysis, Investigation, Project administration, Software, Validation, Visualization, Writing—original draft, Writing—review & editing. IO: Conceptualization, Formal analysis, Funding acquisition, Investigation, Methodology, Resources, Supervision, Writing—review & editing.

Funding

The author(s) declare financial support was received for the research, authorship, and/or publication of this article. This work was supported by Fundação para a Ciência e a Tecnologia, Portugal, through project CERN/FIS-COM/0004/2021 (“Exploring quantum machine learning as a tool for present and future high energy physics colliders”). IO was supported by the fellowship LCF/BQ/PI20/11760025 from La Caixa Foundation (ID 100010434) and by the European Union Horizon 2020 research and innovation program under the Marie Skłodowska-Curie grant agreement No 847648.

Acknowledgments

We acknowledge the use of IBM Quantum services for this work. We thank Declan Millar, Nuno Peres, and Tiago Antão for the very useful discussions and Ricardo Ribeiro for kindly providing access to some of the computing systems used in this work. We also thank Henrique Carvalho for the help in producing (Figure 1).

References

- Akiba, T., Sano, S., Yanase, T., Ohta, T., and Koyama, M. (2019). “Optuna: a next-generation hyperparameter optimization framework,” in *Proceedings of the 25rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (Anchorage, AK). doi: 10.1145/3292500.3330701
- Alvi, S., Bauer, C. W., and Nachman, B. (2023). Quantum anomaly detection for collider physics. *J. High Energy Phys.* 2023:220. doi: 10.1007/JHEP02(2023)220
- Alwall, J., Frederix, R., Frixione, S., Hirschi, V., Maltoni, F., Mattelaer, O., et al. (2014). The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations. *J. High Energy Phys.* 7:79. doi: 10.1007/JHEP07(2014)079
- Anis, M. S., Abby-Mitchell, Abraham, H., AduOffei, Agarwal, R., Agliardi, G., Aharoni, M., et al. (2023). *QISKIT: An Open-Source Framework for Quantum Computing*. doi: 10.5281/zenodo.2573505
- Araz, J. Y., and Spannowsky, M. (2022). Classical versus quantum: Comparing tensor-network-based quantum circuits on large hadron collider data. *Phys. Rev. A* 106:62423. doi: 10.1103/PhysRevA.106.062423
- Bapst, F., Bhimji, W., Calafiura, P., Gray, H., Lavrijsen, W., Linder, L., et al. (2020). A pattern recognition algorithm for quantum annealers. *Comput. Softw. Big Sci.* 4, 1–7. doi: 10.1007/s41781-019-0032-5
- Belis, V., González-Castillo, S., Reissel, C., Vallecorsa, S., Combarro, E. F., Dissertori, G., et al. (2021). “HIGGS analysis with quantum classifiers,” in *EPJ Web of Conferences*, Vol. 251 (EDP Sciences), 03070. doi: 10.1051/epjconf/202125103070
- Bergholm, V., Izaac, J., Schuld, M., Gogolin, C., Ahmed, S., Ajith, V., et al. (2018). PennyLane: automatic differentiation of hybrid quantum-classical computations. *arXiv:1811.04968v4*. doi: 10.48550/arXiv.1811.04968
- Bergstra, J., Yamins, D., and Cox, D. (2013). “Making a science of model search: hyperparameter optimization in hundreds of dimensions for vision architectures,” in *International Conference on Machine Learning* (San Francisco), 115–123. Available online at: <https://proceedings.mlr.press/v28/bergstra13.html>
- Bergstra, J., Bardenet, R., Bengio, Y., and Kégl, B. (2011). “Algorithms for hyperparameter optimization,” in *Advances in Neural Information Processing Systems*, Vol. 24.
- Blance, A., and Spannowsky, M. (2021). Quantum machine learning for particle physics using a variational quantum classifier. *J. High Energy Phys.* 2021, 1–20. doi: 10.1007/JHEP02(2021)212
- Borras, K., Chang, S. Y., Funcke, L., Grossi, M., and Hartung, T. (2023). Impact of quantum noise on the training of quantum generative adversarial networks. *J. Phys.* 2438:012093. doi: 10.1088/1742-6596/2438/1/012093
- Cacciari, M., Salam, G. P., and Soyez, G. (2008). The anti- k_t jet clustering algorithm. *J. High Energy Phys.* 4:63. doi: 10.1088/1126-6708/2008/04/063
- Cacciari, M., Salam, G. P., and Soyez, G. (2012). Fastjet user manual. *Eur. Phys. J. C* 72:1986. doi: 10.1140/epjc/s10052-012-1896-2
- Chang, S. Y., Herbert, S., Vallecorsa, S., Combarro, E. F., and Duncan, R. (2021a). Dual-parameterized quantum circuit gan model in high energy physics. *EPJ Web Conf.* 251:03050. doi: 10.1051/epjconf/202125103050

Conflict of interest

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

The handling editor ML declared a past co-authorship/collaboration with authors NC and IO.

The author(s) declared that they were an editorial board member of Frontiers, at the time of submission. This had no impact on the peer review process and the final decision.

Publisher’s note

All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

Author disclaimer

The views expressed are those of the authors, and do not reflect the official policy or position of IBM or the IBM Quantum team.

Supplementary material

The Supplementary Material for this article can be found online at: <https://www.frontiersin.org/articles/10.3389/frai.2023.1268852/full#supplementary-material>

- Chang, S. Y., Vallecorsa, S., Combarro, E. F., and Carminati, F. (2021b). Quantum generative adversarial networks in a continuous-variable architecture to simulate high energy physics detectors. *arXiv:2101.11132*. doi: 10.48550/arXiv.2101.11132
- Chen, S. Y.-C., Wei, T.-C., Zhang, C., Yu, H., and Yoo, S. (2022). Quantum convolutional neural networks for high energy physics data analysis. *Phys. Rev. Res.* 4:013231. doi: 10.1103/PhysRevResearch.4.013231
- Chen, T., and Guestrin, C. (2016). "XGBoost: a scalable tree boosting system," in *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '16* (New York, NY: ACM), 785–794. doi: 10.1145/2939672.2939785
- Crispim Romão, M., Castro, N. F., and Pedro, R. (2021). Finding new physics without learning about it: anomaly detection as a tool for searches at colliders. *Eur. Phys. J. C* 81:27. doi: 10.1140/epjc/s10052-020-08807-w
- Crispim Romão, M., Castro, N. F., and Pedro, R. (2021). *Simulated PP Collisions at 13 TeV With 2 Leptons + 1 b Jet Final State and Selected Benchmark Beyond the Standard Model Signals*. Dataset on Zenodo. doi: 10.5281/zenodo.5126746
- Das, S., Wildridge, A. J., Vaidya, S. B., and Jung, A. (2019). Track clustering with a quantum annealer for primary vertex reconstruction at hadron colliders. *arXiv:1903.08879*. doi: 10.48550/arXiv.1903.08879
- de Souza, F. A., Crispim Romão, M., Castro, N. F., Nikjou, M., and Porod, W. (2022). Exploring parameter spaces with artificial intelligence and machine learning black-box optimisation algorithms. *Phys. Rev. D* 107:035004. doi: 10.1103/PhysRevD.107.035004
- Delgado, A., and Hamilton, K. E. (2022). Unsupervised quantum circuit learning in high energy physics. *Phys. Rev. D* 106:096006. doi: 10.1103/PhysRevD.106.096006
- Durieux, G., Maltoni, F., and Zhang, C. (2015). Global approach to top-quark flavor-changing interactions. *Phys. Rev. D* 91:074017. doi: 10.1103/PhysRevD.91.074017
- Ellis, J. (2012). Outstanding questions: physics beyond the standard model. *Philos. Trans. R. Soc. Lond. A* 370, 818–830. doi: 10.1098/rsta.2011.0452
- Farhi, E., and Neven, H. (2018). Classification with quantum neural networks on near term processors. Available online at: <https://arxiv.org/abs/1802.06002>
- Feickert, M., and Nachman, B. (2021). A living review of machine learning for particle physics. Available online at: <https://arxiv.org/abs/2102.02770>
- Funcke, L., Hartung, T., Heinemann, B., Jansen, K., Kropf, A., Kühn, S., et al. (2022). Studying quantum algorithms for particle track reconstruction in the LUXE experiment. *J. Phys.* 2438:12127. doi: 10.1088/1742-6596/2438/1/012127
- Gianelle, A., Koppenburg, P., Lucchesi, D., Nicotra, D., Rodrigues, E., Sestini, L., et al. (2022). Quantum machine learning for b-jet charge identification. *J. High Energy Phys.* 08:014. doi: 10.1007/JHEP08(2022)014
- Guan, W., Perdue, G., Pesah, A., Schulz, M., Terashi, K., Vallecorsa, S., et al. (2021). Quantum machine learning in high energy physics. *Mach. Learn. Sci. Technol.* 2:011003. doi: 10.1088/2632-2153/abc17d
- Guest, D., Cranmer, K., and Whiteson, D. (2018). Deep learning and its application to LHC physics. *Ann. Rev. Nucl. Part. Sci.* 68, 161–181. doi: 10.1146/annurev-nucl-101917-021019
- Gupta, S., and Zia, R. (2001). Quantum neural networks. *J. Comput. Syst. Sci.* 63, 355–383. doi: 10.1006/jcss.2001.1769
- Kingma, D. P., and Ba, J. (2014). Adam: a method for stochastic optimization. Available online at: <https://arxiv.org/abs/1412.6980>
- LaRose, R., and Coyle, B. (2020). Robust data encodings for quantum classifiers. *Phys. Rev. A* 102:032420. doi: 10.1103/PhysRevA.102.032420
- Li, G., Ding, Y., and Xie, Y. (2019). "Tackling the Qubit Mapping Problem for NISQ-Era Quantum Devices," in *ASPLOS '19: Proceedings of the Twenty-Fourth International Conference on Architectural Support for Programming Languages and Operating Systems*.
- MacKay, D. (2003). *Information Theory, Inference and Learning Algorithms*. Cambridge University Press.
- Mott, A., Job, J., Vlimant, J.-R., Lidar, D., and Spiropulu, M. (2017). Solving a higgs optimization problem with quantum annealing for machine learning. *Nature* 550, 375–379. doi: 10.1038/nature24047
- Ngairangbam, V. S., Spannowsky, M., and Takeuchi, M. (2022). Anomaly detection in high-energy physics using a quantum autoencoder. *Phys. Rev. D* 105:095004. doi: 10.1103/PhysRevD.105.095004
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., et al. (2011). Scikit-learn: machine learning in python. *J. Mach. Learn. Res.* 12, 2825–2830. Available online at: <https://www.jmlr.org/papers/v12/pedregosa11a.html>
- Preskill, J. (2018). Quantum computing in the NISQ era and beyond. *Quantum* 2:79. doi: 10.22331/q-2018-08-06-79
- Raschka, S. (2018). Mlxtend: Providing machine learning and data science utilities and extensions to python's scientific computing stack. *J. Open Source Softw.* 3:24. doi: 10.21105/joss.00638
- Rehm, F., Vallecorsa, S., Grossi, M., Borrás, K., and Krücker, D. (2023). A full quantum generative adversarial network model for high energy physics simulations. *arXiv:2305.07284*. doi: 10.48550/arXiv.2305.07284
- Schuhmacher, J., Boggia, L., Belis, V., Puljak, E., Grossi, M., Pierini, M., et al. (2023). Unravelling physics beyond the standard model with classical and quantum anomaly detection. doi: 10.1088/2632-2153/ad0f7f Available online at: <https://arxiv.org/abs/2301.10787>
- Schuld, M., Bocharov, A., Svore, K. M., and Wiebe, N. (2020). Circuit-centric quantum classifiers. *Phys. Rev. A* 101:32308. doi: 10.1103/PhysRevA.101.032308
- Schuld, M., Sinayskiy, I., and Petruccione, F. (2015). Simulating a perceptron on a quantum computer. *Phys. Lett. A* 379, 660–663. doi: 10.1016/j.physleta.2014.11.061
- Selvaggi, M. (2014). DELPHES 3: a modular framework for fast-simulation of generic collider experiments. *J. Phys. Conf. Ser.* 523:012033. doi: 10.1088/1742-6596/523/1/012033
- Shapoval, I., and Calafiura, P. (2019). "Quantum associative memory in HEP track pattern recognition," in *EPJ Web of Conferences*, Vol. 214, 01012. doi: 10.1051/epjconf/201921401012
- Sjöstrand, T., Ask, S., Christiansen, J. R., Corke, R., Desai, N., Ilten, P., et al. (2015). An introduction to PYTHIA 8.2. *Comput. Phys. Commun.* 191, 159–177. doi: 10.1016/j.cpc.2015.01.024
- 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, 1–11. doi: 10.1007/s41781-020-00047-7
- Thaler, J., and Van Tilburg, K. (2011). Identifying boosted objects with n-subjettiness. *J. High Energy Phys.* 2011, 1–28. doi: 10.1007/JHEP03(2011)015
- Tüysüz, C., Carminati, F., Demirköz, B., Dobos, D., Fracas, F., Novotny, K., et al. (2020). Particle track reconstruction with quantum algorithms. *EPJ Web Conf.* 245:09013. doi: 10.1051/epjconf/202024509013
- Vouros, A., Langdell, S., Croucher, M., and Vasilaki, E. (2021). An empirical comparison between stochastic and deterministic centroid initialisation for k-means variations. *Mach. Learn.* 110, 1975–2003. doi: 10.1007/s10994-021-06021-7
- Wei, A. Y., Naik, P., Harrow, A. W., and Thaler, J. (2020). Quantum algorithms for jet clustering. *Phys. Rev. D* 101:094015. doi: 10.1103/PhysRevD.101.094015
- Wilkin, G. A., and Xiuzhen, H. (2008). A practical comparison of two k-means clustering algorithms. *BMC Bioinformatics* 9(Suppl. 6):S19. doi: 10.1186/1471-2105-9-S6-S19
- Woźniak, K. A., Belis, V., Puljak, E., Barkoutsos, P., Dissertori, G., Grossi, M., et al. (2023). Quantum anomaly detection in the latent space of proton collision events at the LHC. *arXiv:2301.10780*. doi: 10.48550/arXiv.2301.10780
- Wu, S. L., Sun, S., Guan, W., Zhou, C., Chan, J., Cheng, C. 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. doi: 10.1103/PhysRevResearch.3.033221
- Zlokapa, A., Anand, A., Vlimant, J.-R., Duarte, J. M., Job, J., Lidar, D., et al. (2021a). Charged particle tracking with quantum annealing optimization. *Quant. Mach. Intell.* 3, 1–11. doi: 10.1007/s42484-021-00054-w
- Zlokapa, A., Mott, A., Job, J., Vlimant, J.-R., Lidar, D., and Spiropulu, M. (2021b). Quantum adiabatic machine learning with zooming. *Bull. Am. Phys. Soc.* 66:62405. doi: 10.1103/PhysRevA.102.062405