# Quantum Generative Adversarial Networks For Anomaly Detection In High Energy Physics

Elie Bermot,[1,2] Christa Zoufal,[1] Michele Grossi,[3] Julian Schuhmacher,[1]
Francesco Tacchino,[1] Sofia Vallecorsa,[3] and Ivano Tavernelli[1]

[1] *IBM Quantum, IBM Research Europe – Zurich, Rueschlikon 8803, Switzerland*
[2] *ETH Zurich, Zurich 8092, Switzerland*
[3] *European Organization for Nuclear Research (CERN), Geneva 1211, Switzerland*

(Dated: May 1, 2023)

The standard model (SM) of particle physics represents a theoretical paradigm for the description of the fundamental forces of nature. Despite its broad applicability, the SM does not enable the description of all physically possible events. The detection of events that cannot be described by the SM, which are typically referred to as anomalous, and the related potential discovery of exotic physical phenomena is a non-trivial task. The challenge becomes even greater with next-generation colliders that will produce even more events with additional levels of complexity. The additional data complexity motivates the search for unsupervised anomaly detection methods that do not require prior knowledge about the underlying models. In this work, we develop such a technique. More explicitly, we employ a quantum generative adversarial network to identify anomalous events. The method learns the background distribution from SM data and, then, determines whether a given event is characteristic for the learned background distribution. The proposed quantum-powered anomaly detection strategy is tested on proof-of-principle examples using numerical simulations and IBM Quantum processors. We find that the quantum generative techniques using ten times fewer training data samples can yield comparable accuracy to the classical counterpart for the detection of the Graviton and Higgs particles. Additionally, we empirically compute the capacity of the quantum model and observe an improved expressivity compared to its classical counterpart.

## I. INTRODUCTION

In data science, an anomaly corresponds to a data point that does not fit into the considered data distribution [1, 2]. Anomaly detection corresponds to the identification of these outliers. Being able to detect these data points is crucial in domains such as network- and cyber-security [3, 4], fraud detection, cancer screening [5], and many more.

Another interesting application of anomaly detection is in High Energy Physics (HEP). Currently, the Large Hadron Collider (LHC) generates data sets of $O(1)$ MB per sample. Efficiently storing and analyzing the data will become more challenging when the LHC will be updated to the High Luminosity Large Hadron Collider (HL-LHC) in 2029, since more collisions will be generated at a time [6]. There are many open computational challenges in particle physics concerning data generation, data processing and data analysis [7–11], e.g., the detection of events that deviate from the Standard Model (SM) [12], the currently most general description of the basic building blocks of the universe. While the SM is a successful theory that is able to account for many physical phenomena, it is not compatible with, e.g., gravity. In order to fully understand nature, events that cannot be described by the SM need to be investigated in detail. In fact, currently one of the main goals of data analysis in particle collider experiments is the machine learning (ML)-based detection of rare collisions corresponding to Beyond-the-Standard Model (BSM) particles.

There are two main approaches for ML-based anomaly detection, supervised and unsupervised. To train an anomaly detection algorithm in a supervised manner, we require labeled data sets representing the SM and BSM events. In HEP, the reference data sets are usually generated through numerical simulations of the SM or an BSM theory, respectively. Even though this approach has already been applied to detect several rare events in experimental data [13–16], the main drawback is that it does not allow for the search of particles which are not described by an existing BSM theory. Therefore, recent efforts have been focused on the development of unsupervised approaches [17–24]. In the unsupervised settings, a model is trained on unlabeled data, learning its underlying structure with the aim of identifying events that do not conform with the reference data as outliers. If all current knowledge about particle physics is included in the numerical generation of the training data, any outlier identified in a set of measured collisions could then be a sign of new physics. One family of promising unsupervised ML approaches for anomaly detection are generative methods [22, 25], specifically the one based on Generative Adversarial Networks (GAN)s [26–28]. Such approaches employ an adversarial training of two classical neural networks, a generator and a discriminator [29, 30]. Based on the two networks it is then possible to derive an anomaly score that allows to identify a data point as an anomaly or not.

In this work, we explore and investigate the feasibility of an unsupervised quantum machine learning (QML) approach based on a quantum implementation of a GAN [31, 32], i.e., a Quantum Generative Adversarial Networks (qGAN), that we shall refer to as *Anomaly qGAN* (A-qGAN). We demonstrate how a qGAN may be used to detect BSM particles. A qGAN is trained on

an embedded SM data set and then used to evaluate an anomaly score by calculating a distance measure between embedded data samples and quantum states generated by the trained qGAN. Furthermore, we investigate the practical performance of anomaly detection for BSM data with quantum simulation and quantum hardware experiments. These experiments reveal that an A-qGAN can achieve the same anomaly detection accuracy as its classical counterpart using ten times fewer training data points. Additionally, we study the expressive power of qGANs in the context of anomaly detection with respect to a capacity measure – the so-called *effective dimension* [33–36] – and find that the quantum models can have advantageous capacity properties.

This paper is organized as follows: In Section II, we present both the classical and quantum GAN models. We then describe the classical and quantum algorithms to perform anomaly detection in Section III. In Section IV, we present the results of empirical anomaly detection experiments with compressed features of an artificial HEP data set using numerical simulation and actual quantum hardware. Finally, we conclude with a general discussion of anomaly detection with A-qGANs in Section V.

## II. THEORY

### A. Generative Modelling

The experimental setup considered here corresponds to a black-box setting, where we only have access to the input and output of an experiment. A powerful algorithm for approximating the underlying process is given by generative learning, which has already been investigated in different scientific domains, including chemistry [37] and biology [38]. In the following, we will focus our study on GANs. The given task is to learn a representation of the probability distribution underlying a data set $P_{\text{data}}$ using a parameterized ansatz. After a generative model is trained, it can be used to generate new synthetic data that are aligned with the generation process of the training set or to investigate the learned approximation to the model distribution $P_{\text{data}}$.

#### 1. *Classical GAN*

GANs [26, 39] consist of two competing components, the generator $\mathcal{G}^{\theta_G}$ and the discriminator $\mathcal{D}^{\theta_D}$, each represented by a differentiable, deterministic neural network (see Fig. 1). The generator $\mathcal{G}^{\theta_G}$ aims at generating samples that could be mistaken for being drawn from the data distribution $P_{\text{data}}$ while the discriminator $\mathcal{D}^{\theta_D}$ tries to distinguish between the data coming from the generator $\mathcal{G}^{\theta_G}$ and samples from the training data set. The generator takes as input a random sample $z$ drawn from a fixed prior distribution $P_z$ for enabling the generation of multiple outputs.

Multiple loss functions can be used to train this machine learning architecture. In this paper, we focus on a non-saturating loss function [27]. The generator minimizes the following loss function

$$\mathrm{C}_G(\theta_G,\theta_D) = -\frac{1}{2}\mathbb{E}_{z\sim P_z}\log \mathcal{D}^{\theta_D}(\mathcal{G}^{\theta_G}(z))\,, \tag{1}$$

while the discriminator tries to maximize the following loss function

$$\begin{aligned}\mathrm{C}_D(\theta_G,\theta_D) = &-\frac{1}{2}\mathbb{E}_{z\sim P_z}\log(1-\mathcal{D}^{\theta_D}(\mathcal{G}^{\theta_G}(z)))\\ &-\frac{1}{2}\mathbb{E}_{x\sim P_{\text{data}}}\log \mathcal{D}^{\theta_D}(x)\,.\end{aligned} \tag{2}$$

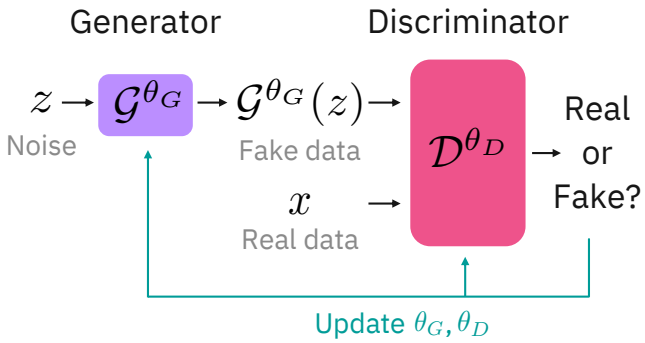


FIG. 1. Working principle of the GAN: a generator $\mathcal{G}^{\theta_G}$ samples $k$ data points, and the discriminator $\mathcal{D}^{\theta_D}$ receives $k$ data samples at a time from either the real data set and from the generated samples and independently classifies them.

For the optimization of the parameters $\theta_G$ and $\theta_D$, standard optimizers can be employed, such as stochastic gradient descent [40], ADAM [41] or AMSGRAD [42], with alternating updates for generator and discriminator parameters.

#### 2. *Quantum GAN*

Several approaches are possible to translate the GAN framework into a quantum machine learning context [7, 31, 32, 43]. In this work, we realize the generator and the discriminator with a parameterized quantum circuit, as is illustrated in Fig. 2. Both, the generator $\mathcal{G}$ and the discriminator $\mathcal{D}$, are defined on $n$-qubit registers with parameterized unitaries $\mathcal{G}(\theta_G)$ and $\mathcal{D}(\theta_D)$. The generator aims at generating a quantum state that represents the distribution underlying the training data set. The discriminator classifies data points as real/generated depending on a single-qubit measurement on the discriminator's qubit register.

To enable a compatibility of this quantum model with classical training data, we need to map each classical training data point $\{\mathbf{x}_1,\ldots,\mathbf{x}_M\}$ to a quantum state $\{|\mathbf{x}_1\rangle,\ldots,|\mathbf{x}_M\rangle\}$. It should be noted that the chosen

mapping needs to be efficient. The generator circuit outputs a state $|\mathcal{G}\rangle = \mathcal{G}(\theta_G)\,|0\rangle^{\otimes n}$. Generating sampled from our model requires to take measurements from $|\mathcal{G}\rangle$ which are then mapped to the feature space $\{\mathbf{x}_1, \ldots, \mathbf{x}_M\}$. The discriminator takes as input an $n$-qubit quantum state and labels it as real or generated based on the measurement of the Pauli $Z$ observable on the last qubit. If the resulting measurement corresponds to the $-1$ eigenvalue, the input data is classified as real and otherwise as generated.

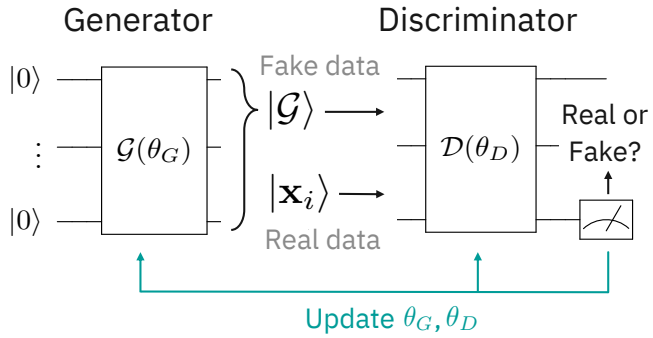


FIG. 2. Illustration of a qGAN model for learning the distribution underlying a training data set $\{|\mathbf{x}_1\rangle, \ldots, |\mathbf{x}_M\rangle\}$. At each training step, the generator generates a state $|\mathcal{G}\rangle$ and the discriminator labels the generated state and a batch of the embedded data points $|\mathbf{x}_i\rangle = \mathcal{R}_i\,|0\rangle$ as real or fake depending on the $Z$-measurement of the last qubit.

To simplify the notation, we introduce the following helper functions:

$$\mathrm{C}^{\mathrm{generated}}(\theta_G, \theta_D) = \frac{1}{2} - \frac{1}{2}\,\langle\mathcal{G}|\mathcal{D}^\dagger(\theta_D)\,\hat{O}\,\mathcal{D}(\theta_D)|\mathcal{G}\rangle\,, \quad (3)$$

$$\mathrm{C}^{\mathrm{data}}(\theta_D) = \frac{1}{2} - \frac{1}{2M}\sum_{i=1}^{M}\langle\mathbf{x}_i|\mathcal{D}^\dagger(\theta_D)\,\hat{O}\,\mathcal{D}(\theta_D)|\mathbf{x}_i\rangle\,, \quad (4)$$

where $\hat{O} = (\mathbb{1}\otimes\mathbb{1}\otimes\ldots\otimes Z)$ and $M$ denotes the number of data samples. Both functions determine the probability of the discriminator labelling a state as real.

The loss function for the generator is then given by

$$\mathrm{C}_G(\theta_G, \theta_D) = -\mathrm{C}^{\mathrm{generated}}(\theta_G, \theta_D). \quad (5)$$

The probability for the generated state $|\mathcal{G}\rangle$ to be labeled as real by the discriminator is maximized by minimizing $\mathrm{C}_G(\theta_G, \theta_D)$ with respect to $\theta_G$. The discriminator $\mathcal{D}(\theta_D)$, on the other hand, tries to discriminate between the generator output and training data by maximizing the following cost function with respect to $\theta_D$:

$$\mathrm{C}_D(\theta_G, \theta_D) = C^{\mathrm{generated}}(\theta_G, \theta_D) - \mathrm{C}^{\mathrm{data}}(\theta_D). \quad (6)$$

During the training the discriminator receives the quantum generator output, i.e. a single quantum state representing the full spectrum of samples, and a batch of $M$ individual samples $\{|\mathbf{x}_i\rangle\}$ from the training data. This may induce a bias in the training that should be monitored carefully and potentially counteracted with multiple optimization steps for the discriminator before performing a single optimization step for the generator.

The loss functions in Eqs. (5) and (6) are optimized in an alternating fashion with, e.g., the ADAM optimizer, by computing analytic quantum gradients [44] using a parameter shift rule [45, 46]. To ensure trainability, i.e. avoid exponentially vanishing gradients, the ansatz should be chosen sufficiently shallow and with a suitable entanglement structure [47, 48]. Furthermore, the method can be expected to be stable against cost-function dependent barren plateaus [49] since only a single qubit measurement is applied in the discriminator.

### B. Anomaly Detection

#### 1. *Classical benchmark*

The GAN framework introduced in Ref. [25] uses the adversarial nature of the two networks to flag anomalous data points. In the training, the generator learns the structure of the non-anomalous data samples. Then, an anomaly score is calculated on the testing data set and/or on new data points to determine whether a new sample aligns with the data distribution. This anomaly score is obtained by optimizing an anomaly loss function based on a combination of the discriminator and generator output. This realizes a combination of all possible information available to the system and can help to compensate for potential instabilities in the generator/discriminator training.

Given a data point $\mathbf{x}$, the minimization of the anomaly loss function aims to find an input noise $z_{\mathrm{opt}}$ corresponding to a generated event $\mathcal{G}(z_{\mathrm{opt}})$ as similar as possible to $\mathbf{x}$ and located on the learned manifold of the GAN. The corresponding anomaly loss function reads

$$\begin{aligned}\mathcal{S}_{\mathcal{C}}(\mathbf{x};\alpha) &= \min_z \mathcal{L}_{\mathcal{C}}(z)\\ &= \min_z (1-\alpha)\|\mathbf{x} - \mathcal{G}(z)\| + \alpha\|\mathcal{D}(\mathbf{x}) - \mathcal{D}(\mathcal{G}(z))\|.\end{aligned} \quad (7)$$

The first term computes the similarity between the generated event $\mathcal{G}(z)$ and the data point $\mathbf{x}$ and the second term measures the similarity of the discriminator outputs $\mathcal{D}(\mathcal{G}(z))$ and $\mathcal{D}(\mathbf{x})$. Furthermore, the parameter $\alpha$ is used to weigh the importance of both terms.

Given $\alpha$ and a well-chosen threshold $\delta > 0$, a data point $\mathbf{x}$ is identified as an anomaly if

$$\mathcal{S}_{\mathcal{C}}(\mathbf{x};\alpha) > \delta. \quad (8)$$

In practice, the threshold $\delta$ is chosen such that it realizes a representative differentiation between non-anomalous and anomalous data. In our case, we may employ training data corresponding to SM and BSM to find a reasonable baseline for this parameter.

### 2. Quantum approach

Next, we introduce our A-qGAN algorithm for the detection of anomalies. We consider a qGAN architecture with a quantum generator and a quantum discriminator as described in Section II A 2. First, each classical data point $\mathbf{x}_j$ is mapped to a quantum state $|\mathbf{x}_j\rangle$. During the training, the qGAN learns a representation for the distribution underlying the generation process of the classical $\{\mathbf{x}_j\}_{j=1}^{M}$. The anomaly score is then evaluated for data points $|\mathbf{x}\rangle$ of the test data set. The loss for the anomaly score is defined as:

$$\mathcal{S}_{\mathcal{Q}}(\mathbf{x};\alpha) = (1-\alpha)|\langle\mathbf{x}|\mathcal{G}\rangle|^2 + \alpha\frac{1+\langle Z\rangle_{\mathcal{D}|\mathbf{x}\rangle}\langle Z\rangle_{\mathcal{D}|\mathcal{G}\rangle}}{2}, \quad (9)$$

with $\langle Z\rangle_{|u\rangle} = \langle u|\mathbb{1}\otimes\mathbb{1}\otimes\ldots\otimes Z|u\rangle$. The first term is a *residual score*. It measures the fidelity between the embedded input point $|\mathbf{x}_j\rangle$ and the trained generator state $|\mathcal{G}\rangle$. It should be noted that the fidelity measure can suffer from exponential concentration effects [50, 51]. Finding scalable distance measures remain an open question for future research. The second term is a *discriminator score* based on the classification of the embedded input point $|\mathbf{x}_j\rangle$ and the trained generator state $|\mathcal{G}\rangle$. Finally, the parameter $\alpha$ weighs the importance of both circuits for the anomaly score. Equivalently to the classical case, an event is labelled as anomalous if

$$\mathcal{S}_{\mathcal{Q}}(\mathbf{x};\alpha) > \delta, \quad (10)$$

for reasonably chosen $\delta$ and $\alpha$. The quantum anomaly score differs from the classical score in the sense that it does not require an additional optimization over the random prior $z$ due to the intrinsic stochasticity of quantum measurements.

## III. METHODS

### A. High Energy Physics Data Set

In this work, we are focusing on an artificial data set generated according to HEP experiments. The training data set consists of a weighted mixture of different SM processes typically observed at 13 TeV with weights given by the production cross section of the corresponding processes [52]. This accounts for the most representative processes in the SM. The anomalous data sets were obtained through the simulation of BSM processes obtained with the PYTHIA8 Monte Carlo simulator [53–55]. In this work, we consider the Higgs boson and the Graviton as anomalies. The SM training data set contains 3'450'279 different events and the testing data set contains 3'450'277 events. Additionally to the SM data set, we also work with a Higgs boson data set [56] containing 139'991 events and a Graviton data set [57] containing 6'910 events. In all data sets, an event is characterized by a list of 23 high-level features [58]. For the quantum approach, we randomly select 100 SM events to build a training set and 100 events for each the testing SM data set, the Graviton data set and the Higgs data set. For the classical benchmark, we randomly select 10 times more, i.e., 1'000 events from each data set including the 100 events used for the quantum approach.

We apply classical pre-processing steps: we efficiently extract a compressed representation of the data sets with the Principal Component Analysis (PCA) method [59] and normalize the training data set to be in $[-\frac{\pi}{2},\frac{\pi}{2}]$. This is necessary due to the periodicity of Pauli rotations that builds the foundation of the data encoding in our ansatz.

The training quantum data set is obtained by mapping each data point $\mathbf{x}\in[-\frac{\pi}{2},\frac{\pi}{2}]^n$ to an $n$-qubit quantum state $|\mathbf{x}\rangle$ obtained with the angle encoding [60–62]:

$$|\mathbf{x}\rangle = \bigotimes_{i=1}^{n} R_y(x_i)\,|0\rangle_i\,, \quad (11)$$

where $x_i$ corresponds to the feature $i$ of the data point $\mathbf{x}$ and $|0\rangle_i$ to the ground state of qubit $i$.

### B. Quantum implementation

We implement the quantum anomaly detection algorithm in Qiskit [63]. The quantum generator circuit $\mathcal{G}(\theta_G)$ consists of an initialization layer and $k_G$ alternating layers of parameterized Pauli-$Y$ single-qubit rotations and blocks of controlled-$Z$ gates (see Fig. 3a) [64, 65]. The initialization layer consists of Hadamard gates $H$ [66] on all qubits. In total, the circuit has $(k_G+1)n$ parameterized Pauli-$Y$ rotations and $k_G n$ controlled-$Z$ gates.

The quantum discriminator is also a parameterized quantum circuit. We use an ansatz containing approximately as many parameters as the generator. The variational form of the discriminator consists of several parts: first, we apply a layer of Hadamard gates on all qubits followed by $k_D$ alternating layers of parameterized single-qubit rotations, Pauli-$Y$($R_Y$) and Pauli-$Z$($R_Z$) rotations on all qubits followed by controlled-$Z$ gates ($CZ$). Then, a series of controlled-$X$ ($CNOT$) [66] is applied with the target on the last qubit. Finally, $R_X$, $R_Y$ and $R_Z$ single-qubit rotations are applied on the last qubit. In total, the circuit has $3(k_D n+1)$ parameterized single-qubit rotations. The quantum circuit of the discriminator is shown in Fig. 3b.

In practice, at each optimization step, the training data is shuffled and split into batches of size 10. The discriminator is trained with samples from the batch while the generator is trained to output a single quantum state. We improve the stability of the algorithm with more optimization iterations for the discriminator than for the generator. The optimizers and respective hyperparameters which are used in the experiments are given in Section IV.

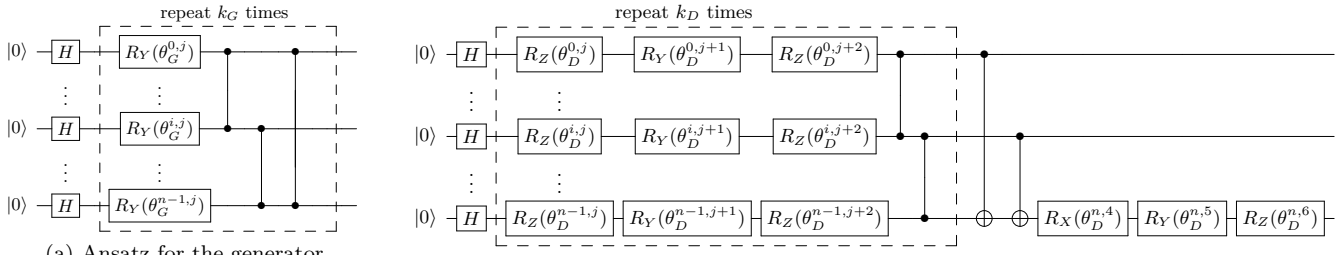


(a) Ansatz for the generator

(b) Ansatz for the discriminator

FIG. 3. Ansatz for the quantum architecture. The quantum generator (a) consists of $R_y$ rotation gates and $CZ$ gates with variational parameters $\theta_G^{i,j}$ which corresponds to the parameter of the $R_Y$ rotation acting on the $i$ qubit on the $j$ layer. The quantum discriminator (b) consists of $R_X$, $R_Y$ and $R_Z$ rotation gates, $CZ$ and $CNOT$ gates with variational parameters $\theta_D^{i,j}$.

### C. Classification scores

We measure the performance of this anomaly detection problem with different metrics, the Receiver Operator Characteristics (ROC) [67, 68], accuracy, F1 score, and precision, which are defined as:

$$\text{Accuracy} = \frac{\text{True Anomalies}}{\text{Total number of events}}, \tag{12}$$

$$\text{Precision} = \frac{\text{True Anomalies}}{\text{True Anomalies} + \text{False Anomalies}}, \tag{13}$$

and

$$\text{F1} = \frac{2}{\text{Recall}^{-1} + \text{Precision}^{-1}} \tag{14}$$

$$\text{with Recall} = \frac{\text{True Anomalies}}{\text{True Anomalies} + \text{False Non-Anomalies}}.$$

An ROC curve corresponds to the plot of the true positive rate versus the false positive rate. It essentially measures the performance of a binary classifier as its cutoff threshold is varied. In practice, the area under the ROC curve, the Area Under Curve (AUC), represents the degree of separability of a classifier. The higher the AUC the better the model is at predicting the correct class of its input.

### D. Anomaly detection algorithm

The anomaly score is computed for each sample in the SM, Graviton and Higgs data sets. The mean of the anomaly scores for the data sets define thresholds which in turn help to identify whether a test sample is an anomaly or not. We apply a grid search over the $\alpha$ parameter, to find the one which results in the best AUC on the test data set. The anomaly detection procedure is explained in Algorithm 1.

**Algorithm 1:** Anomaly detection algorithm

**Data:** Training data set $\mathbb{D} = \{\mathbf{x}_1, \ldots, \mathbf{x}_M\}$ (or $\{|\mathbf{x}_1\rangle, \ldots, |\mathbf{x}_M\rangle\}$), testing data set $\mathcal{T}$, generator optimizer $\text{Opt}_G$, discriminator optimizer $\text{Opt}_D$.
**for** *number of training steps* **do**
  Train generative model $(\mathcal{G}, \mathcal{D})$ based on Section. II A 1 or II A 2 with $\text{Opt}_G$ and $\text{Opt}_D$.
**end**
**for** $\alpha \in [0, 1]$ **do**
  **for** $\mathbf{x} \in \mathbb{D} \cup \mathcal{T}$ **do**
    **if** *classical* **then**
      $\mathcal{S}_\mathcal{C}(\mathbf{x}) \leftarrow \min_z[(1-\alpha)\|\mathbf{x} - \mathcal{G}(z)\| + \alpha\|\mathcal{D}(\mathbf{x}) - \mathcal{D}(\mathcal{G}(z))\|]$;
    **end**
    **if** *quantum* **then**
      $\mathcal{S}_\mathcal{Q}(\mathbf{x}) \leftarrow (1-\alpha)|\langle\mathbf{x}|\mathcal{G}\rangle|^2 + \alpha\frac{1 + \langle Z\rangle_{\mathcal{D}|\mathbf{x}\rangle}\langle Z\rangle_{\mathcal{D}|\mathcal{G}\rangle}}{2}$;
    **end**
  **end**
  Compute ROC Curve and AUC($\alpha$) score;
  Compute $F1(\alpha)$, Accuracy($\alpha$), and Precision($\alpha$) scores;
**end**
Compute $\text{AUC}(\alpha_{\max}) = \max_\alpha \text{AUC}(\alpha)$;
**return** $\text{F1}(\alpha_{\max})$, $\text{Accuracy}(\alpha_{\max})$, and $\text{Precision}(\alpha_{\max})$

### E. Effective dimension study

In the last part of this work, we will discuss the representation power of qGANs and classical equivalents for anomaly detection in HEP. We consider the effective dimension, initially introduced in Ref. [33], which corresponds to a measure of the expressibility of a neural network. The effective dimension is an expressibility measure that can be applied to quantum as well as classical models. More specifically, this measure quantifies what fraction of the model space can actually be covered with the available system parameters. In the case of (q)GANs, the effective dimension measures how the (quantum) generator explores the model space for a given amount of parameters. A higher effective dimension then indicates a higher capacity of the model.

## IV. RESULTS AND DISCUSSION

In the following, we apply the proposed quantum anomaly detection technique to the HEP data sets introduced in Section III A. First, we present the results for the training of the (q)GANs on up to 8 features, and then also demonstrate its feasibility on actual quantum hardware using 3 features. It should be noted that experiments at larger scale are currently infeasible due to the noise present in today's quantum devices. Finally, we study the effective dimension of the applied generators, in order to asses their learning capacity.

### A. Numerical results

In the quantum simulations, we first optimize the qGANs on a noiseless simulator for 500 epochs with the AMSGRAD optimizer [42], applying a learning rate of $10^{-3}$, and $(\beta_1, \beta_2) = (0.7, 0.99)$. Once trained, we compute the anomaly score for the SM testing, the Higgs, and the Graviton data sets. Based on the anomaly score, we then perform the classification between the SM and the anomalies using a varying threshold. Next, we repeat the A-qGAN procedure on a noisy simulator that mimics the behavior of the superconducting IBM Quantum processor *ibmq_belem*. This optimization is performed with the same optimizer but using a learning rate of $10^{-2}$ to handle the noise of the hardware and ten times fewer training epochs. Once trained, we also compute the anomaly score for all data sets and classify the different events.

We benchmark the quantum models against classical neural networks consisting of fully-connected layers with different non-linear activation functions (sigmoid $\sigma$, ReLU, LeakyReLU with parameters $\alpha = 0.2$) or dropout layers of probability 0.25. The architectures of the classical neural networks are chosen, such that they have a number of trainable parameters comparable to the quantum case.

In the following, we use 10 randomly selected instances of the training and testing data sets introduced in Section III A. We perform anomaly detection using between 3 and 8 features, each obtained with PCA-compression. The classification metrics resulting from the anomaly detection using 3 or 7 features are compared in Table I. The data shows that classical and quantum anomaly detection methods achieve comparable classification metrics in both cases. For the 3-feature (7-feature) simulation, we use 3 (9) repetitions of the generator ansatz in Fig. 3a, and 2 (3) repetitions of the discriminator ansatz in Fig. 3b, yielding 6 (70) trainable parameters for the generator and 21 (66) trainable parameters for the discriminator, respectively. Furthermore, in each training epoch, we apply more optimization updates for the discriminator than for the generator, 5 times more for the training on 3 features, and 10 times more for 7 features, respectively.

We plot the evolution of the AUC score for an increasing number of features in Fig. 4. As we increase the number of features, the classical model gets better at detecting the Graviton (see Fig. 4a). The A-qGAN performance with less training data also improves, reaching the performance of the classical algorithm for six PCA-compressed features. After 6 features, we reach a plateau where the performance of the classical and quantum models is comparable. Similarly, the classical Higgs detection performance slightly improves as we add more features before reaching a plateau for 6 or more features, where the average AUC scores are worse than those obtained with less features (see Fig. 4b). The classical algorithm slightly outperforms the quantum one for less than seven features. For seven and eight features, the quantum simulations with less training data are similar in performance to the classical case. Comparing the two figures, we can observe that detecting the Higgs particle is more challenging than detecting the Graviton for both, quantum and classical models, as also noted in Ref. [69]. We can also notice that for both, the Graviton and the Higgs detection, we get similar performance between the noiseless and noisy quantum simulations, which indicates some level of noise robustness of the A-qGAN.

It is worth stressing that the quantum models were trained with 10 times fewer data samples than the classical models, but still achieve comparable accuracy, AUC, precision and F1 scores.

The large standard deviations shown in Fig. 4 can be explained by the intrinsic volatility in the GAN training and by the data set reshufflings.

### B. Hardware experiments

Next, we perform anomaly detection on the superconducting IBM Quantum processor *ibmq_belem*. Fig. 5 shows the topology of the hardware with the CNOT errors and the $T_1$ time for each qubit. For the training of the qGAN, we use the same protocol as before (with the same generator and discriminator optimization hyperparameters as in the noisy simulation). We use $10^4$ shots for each expectation value estimation during the training and the anomaly score computation. Additionally, to ease the hardware training, we reduce the number of training epochs to 5.

In this part, we only consider the anomaly detection of the Graviton particle using 3 features. To ease the computation on the quantum hardware, we train on subsets of 50 events taken from the 10 training data sets used in Section IV A and then compute the anomaly scores for the SM and Graviton events. Fig. 6 presents the average ROC curve of the A-qGAN trained on the noisy quantum hardware (green), compared to the classical method (blue) and the noiseless A-qGAN (orange).

Compared to the classical and noiseless methods, the AUC curve of the A-qGAN evaluated on the quantum hardware is lower. These differences can be explained by the reduced number of epochs, and smaller data sets compared to the noiseless algorithm.

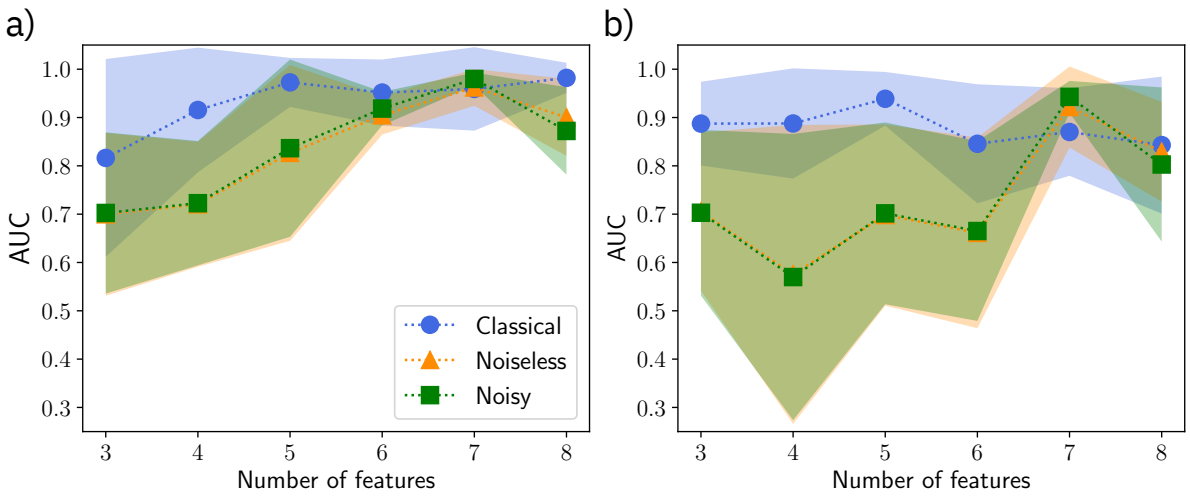


FIG. 4. AUC score for increasing number of features used for the detection of the Graviton (a) and the Higgs (b) particles. For the detection of the Graviton, the AUC score increases with the number of features. For the detection of the Higgs particle, the AUC is constant between 3 and 5 features and slowly decreases after 6 features for the classical model. This is in contrast to the quantum models, which reach their best performance only for 7 features.

| Anomaly | Features | Method | Best $\alpha$ value | F1 score | Accuracy | Precision | ROC-AUC |
|---|---|---|---|---|---|---|---|
| Graviton | 3 | **GAN** | **0.25** | $\mathbf{0.83 \pm 0.127}$ | $\mathbf{0.80 \pm 0.162}$ | $\mathbf{0.91 \pm 0.183}$ | $\mathbf{0.82 \pm 0.205}$ |
| | | qGAN | 0 | $0.77 \pm 0.084$ | $0.71 \pm 0.126$ | $0.90 \pm 0.163$ | $0.70 \pm 0.169$ |
| | 7 | **GAN** | **0.75** | $\mathbf{0.96 \pm 0.059}$ | $\mathbf{0.95 \pm 0.072}$ | $\mathbf{1.0 \pm 0.001}$ | $\mathbf{0.96 \pm 0.086}$ |
| | | qGAN | 0.75 | $0.92 \pm 0.035$ | $0.92 \pm 0.042$ | $1.0 \pm 0.002$ | $0.96 \pm 0.038$ |
| Higgs | 3 | **GAN** | **0.5** | $\mathbf{0.83 \pm 0.088}$ | $\mathbf{0.83 \pm 0.093}$ | $\mathbf{0.98 \pm 0.036}$ | $\mathbf{0.89 \pm 0.087}$ |
| | | qGAN | 0.25 | $0.75 \pm 0.081$ | $0.70 \pm 0.125$ | $0.87 \pm 0.162$ | $0.71 \pm 0.164$ |
| | 7 | GAN | 1 | $0.82 \pm 0.089$ | $0.81 \pm 0.098$ | $1.0 \pm 0.003$ | $0.87 \pm 0.091$ |
| | | **qGAN** | **0.5** | $\mathbf{0.89 \pm 0.063}$ | $\mathbf{0.88 \pm 0.080}$ | $\mathbf{1.0 \pm 0.001}$ | $\mathbf{0.92 \pm 0.084}$ |

TABLE I. Classification scores for the detection of BSM anomalies using three and seven principal components, given for the best $\alpha$ value in the anomaly score. The bold characters correspond to the best results between the classical and quantum methods. In all cases, the quantum calculations are done with 100 data samples while the classical one contains 1000 data samples containing these 100 data samples. For 3 features, the classical model yields better classification scores than for the quantum one for both the Graviton and Higgs detections. For 7 features, the quantum and classical models yield similar classification scores for detection of both the Graviton and Higgs events.

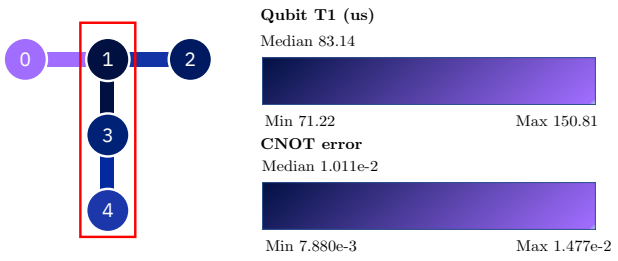


FIG. 5. Topology of the IBM Quantum processor *ibmq_belem*. Qubits are represented as circles and the supported two-qubit gate operations are displayed as edges connecting the qubits. The coloring of the qubits indicates their respective $T_1$ times and the coloring of the edges indicates the errors induced by CNOT gates. The $T_1$ times and CNOT errors were obtained on February 7, 2023. The highlighted qubits are the ones used for the experiments.

### C. Effective dimension study

We now compare the effective dimensions of the quantum and classical generators for different system sizes. It should be noted that we use a similar amount of trainable parameters for the quantum and classical neural networks employed for both the generator and discriminator of the GAN. Hence, their effective dimensions can be directly compared. Fig. 7 shows the effective dimension for increasing number of features. The green (red) dots represent the effective dimension of the trained quantum (classical) model. The effective dimension of the quantum model is always larger than for the classical model which indicates a higher expressibility. It should also be noted that as we increase the dimension of the system, the effective dimension decreases for the chosen generator ansatzes with $\mathcal{O}(n)$ parameters. Furthermore, the effective dimension of the classical model appears to converge towards 0, hinting at

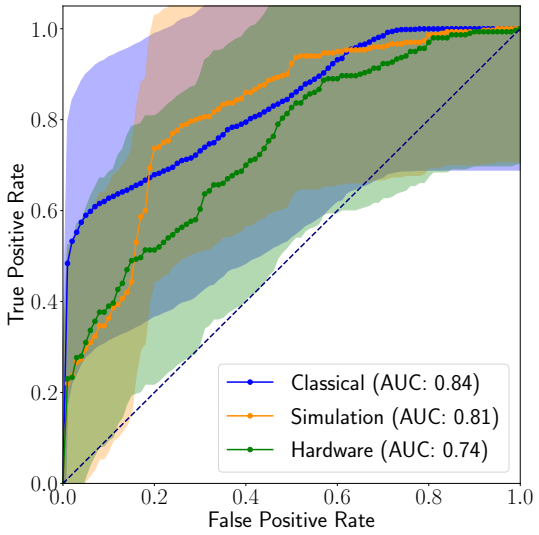


FIG. 6. ROC-AUC curve of the classification between SM and Graviton events for the classical GAN (blue), the noiseless simulation of a qGAN (orange), and the qGAN executed on the IBM Quantum processor *ibmq_belem*. All models are trained on 3 features of the SM data set, and evaluated on SM and Graviton events.

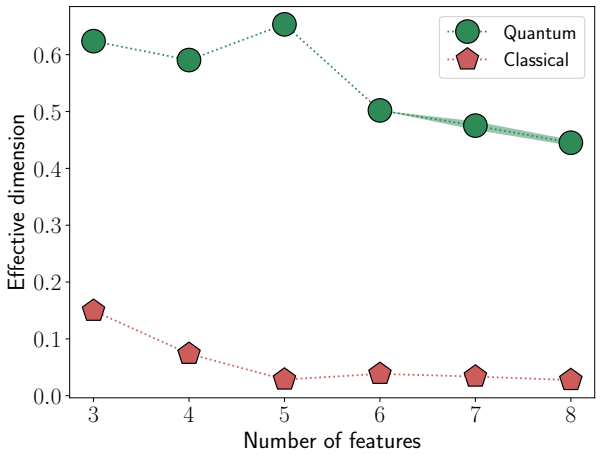


FIG. 7. Effective dimension of the classical (red pentagons) and quantum (green circles) generator for increasing number of features. For each number of features, the classical and quantum generator have the same number of trainable parameters. Therefore, the higher effective dimension of the quantum generator indicates a potential advantage of the quantum approach over the classical one.

a lack of capacity when the classical model is restricted to the same amount of trainable parameters as the quantum model. This may be indicative for a potential advantage of our quantum methodology in larger dimensions as the effective dimension and, hence, expressive capacity of the classical model for $\mathcal{O}(n)$ parameters could become insufficient.

## V. DISCUSSIONS AND CONCLUSIONS

In this work, we proposed and successfully tested an unsupervised learning strategy based on quantum generative modeling techniques for anomaly detection in HEP. In the introduced A-qGAN algorithm, we apply a qGAN to learn the underlying probability distribution from a data set of SM events encoded in quantum states, and then use it to identify HEP processes that do not conform with the SM. The classification is based on an anomaly score, which serves as a distance measure between SM events and other events of interest, not contemplated in the SM. Specifically, we verify the proposed anomaly detection scheme by identifying Higgs and Graviton events. Our quantum anomaly detection scheme is designed to be robust to training instabilities by using an anomaly score that relies on the quantum generator and the quantum discriminator output. Additionally, the proposed generative method can detect new-physics events without modelling a specific BSM scenario.

The presented experiments show that in both cases the anomaly detection performance improves as we increase the number of compressed features with quantum and classical method yielding comparable accuracy. Overall, our results indicate that the proposed quantum approach can detect Graviton and Higgs events using ten times fewer data points compared to a classical benchmark based on classical GANs with the same number of trainable parameters. It should, however, be mentioned that a lack of data is typically not a problem in this type of HEP task. In fact, the amount of training data will even increase in the next years. Nevertheless, the indication that anomaly detection could be performed with a reduced training data set, could set the ground for other applications suffering from data scarcity in other HEP tasks or other fields such as medicine. Moreover, we showed that the quantum model exhibits a larger expressive power, measured by a higher effective dimension, than the classical counterpart. This suggests that our qGAN-based anomaly detection scheme could be better suited to model and detect BSM anomalies in more complex and/or scaled settings.

Future research directions aimed at confirming and extending our preliminary results could address a more extensive collection of data sets and scale up the relevant model and feature dimensions, approaching the regime where our A-qGAN can no longer be simulated classically. As the field of unsupervised anomaly detection in HEP is still relatively young, particularly for what concerns the application of quantum methods, future studies could also investigate alternative approaches, which could be compared or could complement generative models. As an example, it was recently shown that a supervised kernel method based on the generation of a scrambled anomalous data set [69] can achieve performances close to the best known classical equivalents, and an unsupervised technique [70] can potentially outperform its classical counterpart. It is also worth mentioning that the use of classically prepared and pre-processed features might

hide some of the quantum correlations originally present in the data [69, 70]. One could hence envisage the use of the A-qGAN framework for detecting anomalies directly on quantum data sets, obtained for instance through quantum sensing or, e.g., by directly coupling quantum processors to a new generation of quantum detectors.

Lastly, while in this work we mainly focused on detecting anomalies, the proposed qGAN framework could potentially be modified to facilitate the generation of events beyond the SM. In particular, once trained on the SM, one could possibly use the model to generate events on a channel orthogonal to the learned data representation. As a result, our framework could open up novel research avenues at the interface between quantum technologies and HEP.

## VI. ACKNOWLEDGEMENTS

We acknowledge the use of IBM Quantum services for this work. IBM, the IBM logo, and ibm.com are trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. The current list of IBM trademarks is available at https://www.ibm.com/legal/copytrade. MG and SV are supported by CERN through the Quantum Technology Initiative.

[1] Frank E Grubbs. Procedures for detecting outlying observations in samples. *Technometrics*, 11(1):1–21, 1969.

[2] Douglas M Hawkins. *Identification of outliers*, volume 11. Springer, 1980.

[3] Sutharshan Rajasegarar, Christopher Leckie, and Marimuthu Palaniswami. Anomaly detection in wireless sensor networks. *IEEE Wireless Communications*, 15(4):34–40, 2008.

[4] Richard Zuech, Taghi M Khoshgoftaar, and Randall Wald. Intrusion detection and big heterogeneous data: a survey. *Journal of Big Data*, 2(1):1–41, 2015.

[5] Gwenolé Quellec, Guy Cazuguel, Béatrice Cochener, and Mathieu Lamard. Multiple-instance learning for anomaly detection in digital mammography. *Ieee transactions on medical imaging*, 35(7):1604–1614, 2016.

[6] Giorgio Apollinari, I Béjar Alonso, O Brüning, P Fessia, M Lamont, L Rossi, and L Tavian. High-luminosity large hadron collider (hl-lhc). technical design report v. 0.1. Technical report, Fermi National Accelerator Lab.(FNAL), Batavia, IL (United States), 2017.

[7] Gabriele Agliardi, Michele Grossi, Mathieu Pellen, and Enrico Prati. Quantum integration of elementary particle processes. *Physics Letters B*, 832:137228, 2022.

[8] Wen Guan, Gabriel Perdue, Arthur Pesah, Maria Schuld, Koji Terashi, Sofia Vallecorsa, and Jean-Roch Vlimant. Quantum machine learning in high energy physics. *Machine Learning: Science and Technology*, 2(1):011003, 2021.

[9] Cenk Tüysüz, Carla Rieger, Kristiane Novotny, Bilge Demirköz, Daniel Dobos, Karolos Potamianos, Sofia Vallecorsa, Jean-Roch Vlimant, and Richard Forster. Hybrid quantum classical graph neural networks for particle track reconstruction. *Quantum Machine Intelligence*, 3:1–20, 2021.

[10] Cenk Tüysüz, Federico Carminati, Bilge Demirköz, Daniel Dobos, Fabio Fracas, Kristiane Novotny, Karolos Potamianos, Sofia Vallecorsa, and Jean-Roch Vlimant. Particle track reconstruction with quantum algorithms. In *EPJ Web of Conferences*, volume 245, page 09013. EDP Sciences, 2020.

[11] Koji Terashi, Michiru Kaneda, Tomoe Kishimoto, Masahiko Saito, Ryu Sawada, and Junichi Tanaka. Event classification with quantum machine learning in high-energy physics. *Computing and Software for Big Science*, 5:1–11, 2021.

[12] CDF Collaboration, T Aaltonen, S Amerio, D Amidei, A Anastassov, A Annovi, J Antos, G Apollinari, JA Appel, T Arisawa, et al. High-precision measurement of the W boson mass with the CDF II detector. *Science*, 376(6589):170–176, 2022.

[13] Pierre Baldi, Peter Sadowski, and Daniel Whiteson. Searching for exotic particles in high-energy physics with deep learning. *Nature communications*, 5(1):4308, 2014.

[14] Amit Chakraborty, Sung Hak Lim, and Mihoko M Nojiri. Interpretable deep learning for two-prong jet classification with jet spectra. *Journal of High Energy Physics*, 2019(7):1–36, 2019.

[15] Jie Ren, Daohan Wang, Lei Wu, Jin Min Yang, and Mengchao Zhang. Detecting an axion-like particle with machine learning at the LHC. *Journal of High Energy Physics*, 2021(11):1–26, 2021.

[16] Huifang Lv, Daohan Wang, and Lei Wu. Deep learning jet images as a probe of light Higgsino dark matter at the LHC. *Physical Review D*, 106(5):055008, 2022.

[17] Benjamin Nachman. Anomaly detection for physics analysis and less than supervised learning. In *Artificial Intelligence for High Energy Physics*, pages 85–112. World Scientific, 2022.

[18] Georges Aad, Brad Abbott, Dale Charles Abbott, A Abed Abud, Kira Abeling, Deshan Kavishka Abhayasinghe, Syed Haider Abidi, OS AbouZeid, Nadine L Abraham, Halina Abramowicz, et al. Dijet Resonance Search with Weak Supervision Using $\sqrt{s} = 13$ TeV pp Collisions in the ATLAS Detector. *Physical review letters*, 125(13):131801, 2020.

[19] CMS Collaboration, AM Sirunyan, A Tumasyan, W Adam, F Ambrogi, T Bergauer, M Dragicevic, J Erö, A Escalante Del Valle, R Frühwirth, et al. MUSiC: a model-unspecific search for new physics in proton–proton collisions at $\sqrt{s} = 13$ TeV. *The European Physical Journal C*, 81:1–49, 2021.

[20] Gregor Kasieczka, Benjamin Nachman, David Shih, Oz Amram, Anders Andreassen, Kees Benkendorfer, Blaz Bortolato, Gustaaf Brooijmans, Florencia Canelli, Jack H Collins, et al. The LHC Olympics 2020 a community challenge for anomaly detection in high energy physics.

*Reports on progress in physics*, 84(12):124201, 2021.

[21] Thea Aarrestad, Melissa van Beekveld, Marcella Bona, Antonio Boveia, Sascha Caron, Joe Davies, Andrea De Simone, Caterina Doglioni, Javier Duarte, Amir Farbin, et al. The dark machines anomaly score challenge: benchmark data and model independent event classification for the large hadron collider. *SciPost Physics*, 12(1):043, 2022.

[22] Katherine Fraser, Samuel Homiller, Rashmish K Mishra, Bryan Ostdiek, and Matthew D Schwartz. Challenges for unsupervised anomaly detection in particle physics. *Journal of High Energy Physics*, 2022(3):1–31, 2022.

[23] Anders Andreassen, Ilya Feige, Christopher Frye, and Matthew D Schwartz. JUNIPR: a framework for unsupervised machine learning in particle physics. *The European Physical Journal C*, 79:1–24, 2019.

[24] Tommaso Dorigo, Martina Fumanelli, Chiara Maccani, Marija Mojsovska, Giles C Strong, and Bruno Scarpa. RanBox: anomaly detection in the copula space. *Journal of High Energy Physics*, 2023(1):1–46, 2023.

[25] Thomas Schlegl, Philipp Seeböck, Sebastian M Waldstein, Ursula Schmidt-Erfurth, and Georg Langs. Unsupervised anomaly detection with generative adversarial networks to guide marker discovery. In *Information Processing in Medical Imaging: 25th International Conference, IPMI 2017, Boone, NC, USA, June 25-30, 2017, Proceedings*, pages 146–157. Springer, 2017.

[26] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks. *Communications of the ACM*, 63(11):139–144, 2020.

[27] Ian Goodfellow. NIPS 2016 tutorial: Generative adversarial networks. *arXiv preprint arXiv:1701.00160*, 2016.

[28] Thomas Schlegl, Philipp Seeböck, Sebastian M Waldstein, Georg Langs, and Ursula Schmidt-Erfurth. f-AnoGAN: Fast unsupervised anomaly detection with generative adversarial networks. *Medical image analysis*, 54:30–44, 2019.

[29] Emily L Denton, Soumith Chintala, Rob Fergus, et al. Deep generative image models using a laplacian pyramid of adversarial networks. *Advances in neural information processing systems*, 28, 2015.

[30] Alec Radford, Luke Metz, and Soumith Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. *arXiv preprint arXiv:1511.06434*, 2015.

[31] Seth Lloyd and Christian Weedbrook. Quantum generative adversarial learning. *Physical review letters*, 121(4):040502, 2018.

[32] Pierre-Luc Dallaire-Demers and Nathan Killoran. Quantum generative adversarial networks. *Physical Review A*, 98(1):012324, 2018.

[33] Amira Abbas, David Sutter, Christa Zoufal, Aurélien Lucchi, Alessio Figalli, and Stefan Woerner. The power of quantum neural networks. *Nature Computational Science*, 1(6):403–409, 2021.

[34] Oksana Berezniuk, Alessio Figalli, Raffaele Ghigliazza, and Kharen Musaelian. A scale-dependent notion of effective dimension. *arXiv preprint arXiv:2001.10872*, 2020.

[35] Jorma J Rissanen. Fisher information and stochastic complexity. *IEEE transactions on information theory*, 42(1):40–47, 1996.

[36] Thomas M Cover. *Elements of information theory*. John Wiley & Sons, 1999.

[37] Artur Kadurin, Sergey Nikolenko, Kuzma Khrabrov, Alex Aliper, and Alex Zhavoronkov. druGAN: an advanced generative adversarial autoencoder model for de novo generation of new molecules with desired molecular properties in silico. *Molecular pharmaceutics*, 14(9):3098–3104, 2017.

[38] Nathan Killoran, Leo J Lee, Andrew Delong, David Duvenaud, and Brendan J Frey. Generating and designing DNA with deep generative models. *arXiv preprint arXiv:1712.06148*, 2017.

[39] Karol Kurach, Mario Lučić, Xiaohua Zhai, Marcin Michalski, and Sylvain Gelly. A large-scale study on regularization and normalization in GANs. In *International conference on machine learning*, pages 3581–3590. PMLR, 2019.

[40] Herbert Robbins and Sutton Monro. A stochastic approximation method. *The annals of mathematical statistics*, pages 400–407, 1951.

[41] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

[42] Sashank J Reddi, Satyen Kale, and Sanjiv Kumar. On the convergence of adam and beyond. *arXiv preprint arXiv:1904.09237*, 2019.

[43] Christa Zoufal, Aurélien Lucchi, and Stefan Woerner. Quantum generative adversarial networks for learning and loading random distributions. *npj Quantum Information*, 5(1):103, 2019.

[44] Iordanis Kerenidis and Anupam Prakash. Quantum gradient descent for linear systems and least squares. *Physical Review A*, 101(2):022316, 2020.

[45] Maria Schuld, Ville Bergholm, Christian Gogolin, Josh Izaac, and Nathan Killoran. Evaluating analytic gradients on quantum hardware. *Physical Review A*, 99(3):032331, 2019.

[46] Kosuke Mitarai, Makoto Negoro, Masahiro Kitagawa, and Keisuke Fujii. Quantum circuit learning. *Physical Review A*, 98(3):032309, 2018.

[47] Jarrod R McClean, Sergio Boixo, Vadim N Smelyanskiy, Ryan Babbush, and Hartmut Neven. Barren plateaus in quantum neural network training landscapes. *Nature communications*, 9(1):4812, 2018.

[48] Carlos Ortiz Marrero, Mária Kieferová, and Nathan Wiebe. Entanglement-induced barren plateaus. *PRX Quantum*, 2:040316, Oct 2021.

[49] Marco Cerezo, Akira Sone, Tyler Volkoff, Lukasz Cincio, and Patrick J Coles. Cost function dependent barren plateaus in shallow parametrized quantum circuits. *Nature communications*, 12(1):1791, 2021.

[50] Supanut Thanasilp, Samson Wang, Marco Cerezo, and Zoë Holmes. Exponential concentration and untrainability in quantum kernel methods. *arXiv preprint arXiv:2208.11060*, 2022.

[51] Francesco Tacchino, Stefano Mangini, Panagiotis Kl Barkoutsos, Chiara Macchiavello, Dario Gerace, Ivano Tavernelli, and Daniele Bajoni. Variational learning for quantum artificial neural networks. *IEEE Transactions on Quantum Engineering*, 2:1–10, 2021.

[52] Oliver Knapp, Olmo Cerri, Günther Dissertori, Thong Q Nguyen, Maurizio Pierini, and Jean-Roch Vlimant. Adversarially Learned Anomaly Detection on CMS Open Data: re-discovering the top quark. *The European Physical*

*Journal Plus*, 136(2):236, 2021.

[53] Andy Buckley, Jonathan Butterworth, Stefan Gieseke, David Grellscheid, Stefan Höche, Hendrik Hoeth, Frank Krauss, Leif Lönnblad, Emily Nurse, Peter Richardson, et al. General-purpose event generators for LHC physics. *Physics Reports*, 504(5):145–233, 2011.

[54] J De Favereau, Christophe Delaere, Pavel Demin, Andrea Giammanco, Vincent Lemaitre, Alexandre Mertens, and Michele Selvaggi. DELPHES 3: a modular framework for fast simulation of a generic collider experiment. *Journal of High Energy Physics*, 2014(2):1–26, 2014.

[55] Nady Bakhet, Maxim Yu Khlopov, and Tarek Hussein. Phenomenology of large extra dimensions models at hadrons colliders using monte carlo techniques (spin-2 graviton). *arXiv preprint arXiv:1507.03888*, 2015.

[56] CMS collaboration (2016). Simulated dataset VBFHiggs0MToBB_M-125p6_7TeV-JHUGenV4-pythia6-tauola in AODSIM format for 2011 collision data (BSM Higgs). CERN Open Data Portal. DOI:10.7483/OPENDATA.CMS.3R3P.5JYR.

[57] CMS collaboration (2016). Simulated dataset Graviton2BPqqbarToZZTo4L_M-125p6_7TeV-JHUGenV3-pythia6 in AODSIM format for 2011 collision data (BSM Higgs). CERN Open Data Portal. DOI:10.7483/OPENDATA.CMS.SZWT.H9MC.

[58] Olmo Cerri, Thong Q Nguyen, Maurizio Pierini, Maria Spiropulu, and Jean-Roch Vlimant. Variational autoencoders for new physics mining at the large hadron collider. *Journal of High Energy Physics*, 2019(5):1–29, 2019.

[59] Harold Hotelling. Analysis of a complex of statistical variables into principal components. *Journal of educational psychology*, 24(6):417, 1933.

[60] Edward Grant, Marcello Benedetti, Shuxiang Cao, Andrew Hallam, Joshua Lockhart, Vid Stojevic, Andrew G Green, and Simone Severini. Hierarchical quantum classifiers. *npj Quantum Information*, 4(1):65, 2018.

[61] Edwin Stoudenmire and David J Schwab. Supervised learning with tensor networks. *Advances in neural information processing systems*, 29, 2016.

[62] Shuxiang Cao, Leonard Wossnig, Brian Vlastakis, Peter Leek, and Edward Grant. Cost-function embedding and dataset encoding for machine learning with parametrized quantum circuits. *Physical Review A*, 101(5):052309, 2020.

[63] Qiskit contributors. Qiskit: An open-source framework for quantum computing, 2023.

[64] Jarrod R McClean, Jonathan Romero, Ryan Babbush, and Alán Aspuru-Guzik. The theory of variational hybrid quantum-classical algorithms. *New Journal of Physics*, 18(2):023023, 2016.

[65] Marcello Benedetti, Erika Lloyd, Stefan Sack, and Mattia Fiorentini. Parameterized quantum circuits as machine learning models. *Quantum Science and Technology*, 4(4):043001, 2019.

[66] Michael A Nielsen and Isaac Chuang. Quantum computation and quantum information, 2002.

[67] James A Hanley and Barbara J McNeil. The meaning and use of the area under a receiver operating characteristic (ROC) curve. *Radiology*, 143(1):29–36, 1982.

[68] Elizabeth R DeLong, David M DeLong, and Daniel L Clarke-Pearson. Comparing the areas under two or more correlated receiver operating characteristic curves: a nonparametric approach. *Biometrics*, pages 837–845, 1988.

[69] Julian Schuhmacher, Laura Boggia, Vasilis Belis, Ema Puljak, Michele Grossi, Maurizio Pierini, Sofia Vallecorsa, Francesco Tacchino, Panagiotis Barkoutsos, and Ivano Tavernelli. Unravelling physics beyond the standard model with classical and quantum anomaly detection. *arXiv preprint arXiv:2301.10787*, 2023.

[70] Kinga Anna Woźniak, Vasilis Belis, Ema Puljak, Panagiotis Barkoutsos, Günther Dissertori, Michele Grossi, Maurizio Pierini, Florentin Reiter, Ivano Tavernelli, and Sofia Vallecorsa. Quantum anomaly detection in the latent space of proton collision events at the LHC. *arXiv preprint arXiv:2301.10780*, 2023.