



# Improving Variational Autoencoders for New Physics Detection at the LHC With Normalizing Flows

Pratik Jawahar<sup>1\*</sup>, Thea Aarrestad<sup>1</sup>, Nadezda Chernyavskaya<sup>1</sup>, Maurizio Pierini<sup>1</sup>, Kinga A. Wozniak<sup>1,2</sup>, Jennifer Ngadiuba<sup>3,4</sup>, Javier Duarte<sup>5</sup> and Steven Tsan<sup>5</sup>

<sup>1</sup> Experimental Physics Department, European Center for Nuclear Research (CERN), Geneva, Switzerland, <sup>2</sup> Faculty of Computer Science, University of Vienna, Vienna, Austria, <sup>3</sup> Particle Physics Division, Fermi National Accelerator Laboratory (FNAL), Batavia, IL, United States, <sup>4</sup> Lauritsen Laboratory of High Energy Physics, California Institute of Technology, Pasadena, CA, United States, <sup>5</sup> Department of Physics, University of California, San Diego, San Diego, CA, United States

We investigate how to improve new physics detection strategies exploiting variational autoencoders and normalizing flows for anomaly detection at the Large Hadron Collider. As a working example, we consider the DarkMachines challenge dataset. We show how different design choices (e.g., event representations, anomaly score definitions, network architectures) affect the result on specific benchmark new physics models. Once a baseline is established, we discuss how to improve the anomaly detection accuracy by exploiting normalizing flow layers in the latent space of the variational autoencoder.

**Keywords:** anomaly detection (AD), variational auto encoder (VAE), normalizing flow (NF), Large Hadron Collider (LHC), new physics beyond standard model, graph convolutional network (GCN), convolutional neural net

## OPEN ACCESS

### Edited by:

Michela Paganini,  
Facebook, United States

### Reviewed by:

Alexander Radovic,  
Borealis AI, Canada  
Tobias Golling,  
Université de Genève, Switzerland

### \*Correspondence:

Pratik Jawahar  
pjawahar@wpi.edu

### Specialty section:

This article was submitted to  
Big Data and AI in High Energy  
Physics,  
a section of the journal  
Frontiers in Big Data

**Received:** 28 October 2021

**Accepted:** 17 January 2022

**Published:** 28 February 2022

### Citation:

Jawahar P, Aarrestad T,  
Chernyavskaya N, Pierini M,  
Wozniak KA, Ngadiuba J, Duarte J  
and Tsan S (2022) Improving  
Variational Autoencoders for New  
Physics Detection at the LHC With  
Normalizing Flows.  
Front. Big Data 5:803685.  
doi: 10.3389/fdata.2022.803685

## 1. INTRODUCTION

Most searches for new physics at the CERN Large Hadron Collider (LHC) target specific experimental signatures. The underlying assumption of a specific new physics model could enter at various stages in the search design, e.g., when reducing the data rate from 40 M to 1,000 collision events per second in real time (Trocino, 2014; Aad et al., 2020; Sirunyan et al., 2020), when designing the event selection, or when running the final hypothesis testing. When searching for pre-established and theoretically well-motivated particles (e.g., the Higgs boson), this strategy is extremely successful because the underlying assumption can be exploited to maximize the search sensitivity. On the other hand, the lack of a predefined target might turn this strength into a limitation.

To compensate for this potential problem, *model independent* searches are also carried out (Aaltonen et al., 2009; Aaron et al., 2009; D0 Collaboration, 2012; CMS-PAS-EXO-14-016, 2017; Aaboud et al., 2019) at hadron colliders. These searches consist in an extensive set of comparisons between the data distribution and the expectation derived from Monte Carlo simulation. Many comparisons are carried out in parallel for multiple physics-motivated features while applying different event selections. However, when searching for new physics among many channels, the “global” significance of observing a particular discrepancy must take into account the probability of observing such a discrepancy anywhere. This so called look-elsewhere effect can be quantified in terms of a trial factor (Gross and Vitells, 2010). While the large trial factor typically reduces the statistical power of this strategy in terms of significance, model independent searches are valuable tools to identify possible regions of interest and provide data-driven motivations for traditional, more targeted searches to be performed on future data.

Recently, the use of machine learning techniques has been advocated as a mean to reduce the model dependence (Weisser and Williams, 2016; Collins et al., 2018, 2019, 2021; Blance et al., 2019; Cerri et al., 2019; D’Agnolo and Wulzer, 2019; De Simone and Jacques, 2019; Heimel et al., 2019; Andreassen et al., 2020; Cheng et al., 2020; Dillon et al., 2020; Farina et al., 2020; Hajer et al., 2020; Khosa and Sanz, 2020; Nachman, 2020; Nachman and Shih, 2020; Park et al., 2020; Amram and Suarez, 2021; Bortolato et al., 2021; D’Agnolo et al., 2021; Finke et al., 2021; Gonski et al., 2021; Hallin et al., 2021; Ostdiek, 2021). In this context, the particle-physics community engaged in two data challenges: the LHC Olympics 2020 (Kasieczka et al., 2021) and the DarkMachines challenge (Aarrestad et al., 2021), where different approaches were explored to attempt to detect an unknown signal of new physics hidden in simulated data.

As part of our contribution to the DarkMachines challenge, we investigated the use of a particle-based variational autoencoder (VAE) (Kingma and Welling, 2014; Rezende et al., 2014) and the possibility of enhancing its anomaly detection capability by using normalizing flows (NFs) (Papamakarios et al., 2021) in the latent space to optimize the choice of the latent-space prior. In this article, we document those studies and expand that effort, investigating the impact of specific architecture choices (event representation, network architecture, usage of expert features, and definition of the anomaly score). This study is an update of our contribution to the DarkMachine challenge (Aarrestad et al., 2021), which benefits from the lessons learned by the DarkMachines challenge. Taking inspiration from solutions presented by other groups in the challenge (e.g., Caron et al., 2021; Ostdiek, 2021), we evaluate the impact of some of their findings on our specific setup. In some cases (but not always), these solutions translate in an improved performance, quantified using the same metrics presented in Aarrestad et al. (2021). In this way, we establish an improved baseline model, on top of which we evaluate the impact of the normalizing flow layers in the latent space.

## 2. DATA SAMPLES AND EVENT REPRESENTATION

This study is based on the datasets released on the Zenodo platform (DarkMachines Community, 2020) in relation to the Dark Machines Anomaly Score Challenge (Aarrestad et al., 2021). They consist of a set of processes predicted in the standard model (SM) of particle physics, mixed according to their production cross section in proton-proton collisions at 13 TeV center-of-mass energy, and a set of benchmark signal samples. The datasets contains labels, identifying the process that generated each event. Labels are ignored during training and used to evaluate performance metrics.

For each sample, four datasets are provided, corresponding to four different event selections (called *channels*; Aarrestad et al., 2021):

- Channel 1:  $H_T \geq 600$  GeV,  $p_T^{\text{miss}} \geq 200$  GeV, and  $p_T^{\text{miss}}/H_T \geq 0.2$ .

**TABLE 1** | Summary of the available dataset size.

Dataset	Channel 1	Channel 2a	Channel 2b	Channel 3
Training	193,800	13,425	238,450	7,100,934
Validation	10,200	707	12,550	373,733
Bkg. test	10,000	5,868	89,000	1,025,333
Sig. test	38,666	5,868	89,676	1,023,320

- Channel 2a:  $p_T^{\text{miss}} \geq 50$  GeV and at least three light leptons (muons or electrons) with  $p_T > 15$  GeV.
- Channel 2b:  $p_T^{\text{miss}} \geq 50$  GeV,  $H_T \geq 50$  GeV and at least two light leptons (muons or electrons) with  $p_T > 15$  GeV.
- Channel 3:  $H_T \geq 600$  GeV,  $p_T^{\text{miss}} \geq 100$  GeV.

where  $p_T$  is the magnitude of a particle’s transverse momentum,  $H_T$  is the scalar sum of the jet  $p_T$  in the event, and  $\vec{p}_T^{\text{miss}}$  is the vector equal and opposite to the vector sum of the transverse momenta of the reconstructed particles in the event, while  $p_T^{\text{miss}}$  is its magnitude<sup>1</sup>. More details are provided in Aarrestad et al. (2021).

The input consists of the momenta of all the reconstructed physics objects in the event (jets, b jets, electrons e, muons  $\mu$ , and photons), ordered by decreasing  $p_T$ . Each list of objects is zero-padded to force each event into a fixed-length matrix with the same order: up to 15 jets, and up to 4 each of b jets,  $\mu^\pm$ ,  $e^\pm$ , and photons. We pre-process the input by applying the `scikit-learn` (Pedregosa et al., 2011) standard scaling and arranging the list of objects into a matrix of 39 particles times four momentum features ( $E, p_T, \eta, \phi$ ), where  $E$  is the particle energy. For e,  $\mu$ , and photons, the energy is computed assuming zero mass. For jets, the measured jet mass is used. The input matrix is interpreted as an image or an unordered point cloud, depending on the underlying VAE architecture.

The training and validation dataset consists of background events from the SM mixture. The available dataset size is detailed in **Table 1** for each of the channels. The background test samples are combined with the benchmark signal samples listed in **Table 2** to form the labeled test dataset on which performance is evaluated.

## 3. TRAINING SETUP AND EVALUATION METRICS

Variational Autoencoders (Kingma and Welling, 2014, 2019; Rezende et al., 2014) are a class of likelihood-based generative models that maximize the likelihood of the training data according to the generative model  $\prod_{x \in \text{data}} p_\theta(x)$  for the set of observed variables  $x$  in the training data. To achieve this

<sup>1</sup>We use a Cartesian coordinate system with the  $z$  axis oriented along the beam axis, the  $x$  axis on the horizontal plane, and the  $y$  axis oriented upward. The  $x$  and  $y$  axes define the transverse plane, while the  $z$  axis identifies the longitudinal direction. The azimuth angle  $\phi$  is computed with respect to the  $x$  axis. The polar angle  $\theta$  is used to compute the pseudorapidity  $\eta = -\log[\tan(\theta/2)]$ . The transverse momentum ( $p_T$ ) is the projection of the particle momentum on the  $(x, y)$  plane. We fix units such that  $c = \hbar = 1$ .

**TABLE 2** | BSM processes contributing to the signal dataset in each channel.

BSM process	Code	Ch.1	Ch.2a	Ch.2b	Ch.3
Z' + jet	monojet_Zp2000.0_DM_50.0	×	×		×
Z' + W/Z	monoV_Zp2000.0_DM_50.0				×
Z' + t	monotop_200_A	×			×
Z' in LFV U(1) <sub>L<sub>μ</sub>-L<sub>τ</sub></sub>	pp23mt_50		×	×	
	pp24mt_50		×	×	
#-SUSY $\tilde{t}\bar{t}$	stlp_st1000	×		×	×
#-SUSY $\tilde{q}\bar{q}$	sqsq1_sq1400_neut800	×			×
SUSY $\tilde{g}\bar{g}$	g1g11400_neutralino1100	×	×	×	×
	g1g11600_neutralino800	×	×	×	×
SUSY $\tilde{t}\bar{t}$	stop2b1000_neutralino300	×			×
SUSY $\tilde{q}\bar{q}$	sqsq_sq1800_neut800	×			×
SUSY $\tilde{\chi}^\pm\tilde{\chi}^0$	chaneut_cha200_neut50		×	×	
	chaneut_cha250_neut150		×	×	
SUSY $\tilde{\chi}^\pm\tilde{\chi}^\pm$	chacha_cha300_neut140			×	
	chacha_cha400_neut60			×	
	chacha_cha600_neut200			×	

The process code, adopted in this study, is taken from Aarrestad et al. (2021).

in a tractable way, the generative model is augmented by the introduction of a set of latent variables  $z$ , such that the marginal distribution over the observed variables  $p_\theta(x)$ , is given by:  $p_\theta(x) = \int p_\theta(x|z)q_\theta(z)dz$ . In this way,  $q_\theta(z)$  can be a relatively simple distribution, such as a Gaussian, while maintaining high expressivity for the marginal distribution  $p_\theta(x)$  as an infinite mixture of simple distributions controlled by  $z$ . Besides being used as generative models, VAEs have been shown to be effective as anomaly detection algorithms (An and Cho, 2015).

In this work, the VAE models are trained on the training and validation datasets, minimizing the loss function:

$$L_{\text{total}} = \beta D_{\text{KL}} + (1 - \beta)L_C, \quad (1)$$

where  $L_C$  is a reconstruction loss, which is chosen to be an  $L_1$ -type permutation-invariant Chamfer loss (Barrow et al., 1977):

$$L_C = \sum_{\vec{x} \in S_{\text{input}}} \min_{\vec{y} \in S_{\text{output}}} |\vec{x} - \vec{y}| + \sum_{\vec{y} \in S_{\text{output}}} \min_{\vec{x} \in S_{\text{input}}} |\vec{x} - \vec{y}|, \quad (2)$$

similar to the  $L_2$ -type Chamfer distance used in Fan et al. (2017) and Zhang et al. (2020). In Equation (2),  $D_{\text{KL}}$  is the Kullback–Liebler divergence term usually employed to force the data distribution in the latent space to a multidimensional Gaussian with unitary covariance matrix (Rezende and Mohamed, 2015), and  $\beta$  is a parameter that controls the relative importance of the two terms (Higgins et al., 2017).

All of our models are optimized using the Adam minimizer (Kingma and Ba, 2015). A learning rate of  $10^{-4}$  is applied along with a brute force early stopping strategy used on an ad-hoc basis. A batch size of 32 is chosen to train models. All models are implemented with the PyTorch (Paszke et al., 2019) deep learning framework and are hosted on GitHub (Jawahar and Pierini, 2021).

We train and test all our models on the WPI Turing Research Cluster<sup>2</sup>, using 8 CPU nodes and 1 GPU node (NVIDIA Tesla V100 or Tesla P100).

At inference time,  $L_C$  is used as an anomaly detection score, to quantify the distance between the input and the output. By applying a lower-bound threshold on  $L_C$ , we identify every event with an  $L_C$  value larger than the threshold as an anomaly. By comparing this prediction to the ground truth, we can assess the performance of the VAE on specific signal benchmark models.

To evaluate model performance we follow the same strategy and code used in Aarrestad et al. (2021) to enable comparison with other models tested on this dataset. As explained in Aarrestad et al. (2021), we extract four main performance parameters from the receiver operating characteristic (ROC) curves based on the chosen anomaly metric for each model, namely the area under the curve (AUC) and true positive rate (also known as the signal efficiency  $\epsilon_S$ ) at three different, fixed values of the false positive rate (also known as background efficiency  $\epsilon_B$ ). We then combine these scores from all models on all available signal regions across all channels of the dataset to form box-and-whisker plots, using six different combination and comparison strategies namely, the highest mean score method, highest median score method, average rank method, top scorer method, top-5 scorer method, and highest minimum scorer method. A box is drawn spanning the inner half (50% quantile centered at the median) of the data as shown in Figure 1. A line through the box marks the median. Whiskers extend from the box to either the maximum and minimum unless these are further away from the edge of the box than 1.5 box lengths. The outlier points are shown as circles.

For Figure 1 and the other figures, the representative ranking as denoted by the legend corresponds to the performance based on the highest mean score method unless mentioned otherwise. However, to choose the best model for each experiment described in this article, we consider all six comparison methods to arrive at a consensus. The code to perform these comparisons and to generate the corresponding plots is available in Aarrestad et al. (2021).

## 4. BASELINE VAE MODEL

The main goal of this study is to evaluate the impact of normalizing flow layers in the latent space on the anomaly detection capability of a reference VAE model. This and the following sections describe how this reference model is built, starting from the VAE based on convolutional layers (Conv-VAE) presented in Aarrestad et al. (2021) and modifying its architecture based on some of the lessons learned during the DarkMachine challenge.

The encoder of the initial Conv-VAE consists of three convolutional layers, with 32, 16, and 8 kernels of size (3, 4), (5, 1), and (7, 1), respectively. For all layers, the stride is set to 1 and zero padding to “same.” The output of the convolutional layers is flattened and passed to two fully-connected neural network (FCN) layers that output the mean and variance for the latent

<sup>2</sup><https://arc.wpi.edu/computing/hpc-clusters/>

space. The cardinality of the latent space is fixed to 15. The decoder mirrors the encoder architecture, returning an output of the same size as the input.

In order to define the reference model, the architecture of the starting model is modified in different ways, each time evaluating the impact of a given choice on the test dataset. Several possibilities are considered: how to embed the event in the two-dimensional (2D) array (see Section 4.1); how to interpret the array, e.g., as an image or a graph (see Section 4.2); whether to extend the event representation beyond the particle momenta, adding domain-specific high level features as an additional input (see Section 4.3); and which anomaly score to use (see Section 4.4). We study various options for each of these points, following this order. Doing so, we establish a candidate model, which replaces the initial model. We evaluate on this new model the benefit of using normalizing flow layers in the latent space (see Section 5) to improve the anomaly detection accuracy.

## 4.1. Data Representation

By their nature, events consist of a variable number of objects. To some extent, this conflicts with most neural network architectures, which assume a fixed-size input. As a baseline, we adopt the simplest solution, i.e., to zero-pad all events to standardized event sizes for all available samples. To get a better idea of how padding affects results, we study performance across alternative input encodings. We consider two main types of encodings, listed as AllObj and TrdObj in **Figure 1**. The former involves considering the entire event which implies allowing for a large enough padding such that every object per event is taken into consideration across the entire dataset. The latter involves cutting down the padding and the input sequence by considering only up to four leading jets and three objects each of the other types per event.

When using the truncated sequence, the model loses information regarding the number of objects of each type per event, which is implicitly learned when the whole sequence is considered. To compensate for this loss, one can explicitly add this information passing a second input to the model, consisting of a vector containing the multiplicities of each object type. This input is concatenated to the flattened output received from the convolutional layers in the encoder before passing them to the fully connected layers. For the sake of comparison, we also do the same for the AllObj case (labeled as “+Mult” in **Figure 1**).

The results in **Figure 1** show that the truncated sequence does worse than the full sequence. We also see little improvement in performance with the addition of multiplicity information per event in both the AUC as well as performance at lower background efficiencies. As a result, we keep the input encoding that considers the complete sequence per event.

## 4.2. VAE Architecture

The convolutional architecture used for the baseline VAE is not the only option to handle the input considered in this study. The ensemble of reconstructed particles in an event can be represented as a point cloud. Doing so, we can process it with a graph neural network. The main advantage of this choice stands with the permutation invariance of the graph processing,

which pairs that of the loss in Equation (2) and complies with the unordered nature of the input list of particles. Graph-based architectures have already been shown to perform better with sparse, non-Euclidean data representations in general (Bronstein et al., 2017; Zhou et al., 2020) and in particle physics in particular (Duarte and Vlimant, 2020; Shlomi et al., 2020).

To this end, we consider a GCN-VAE model composed of multilayer graph convolutional network layers (GCNs) (Kipf and Welling, 2017) and FCN layers in both the encoder and the decoder. As for the VAE, the input graphs are built from the input list described in Section 2, each particle representing one vertex of the graph in the space identified by five particle features:  $E, p_T, \eta, \phi$ , and object type. The object type is a label-encoded integer that signifies the object type. The input is structured as a fully connected, undirected graph which is passed to the GCN layers of the encoder, defined as (Kipf and Welling, 2017):

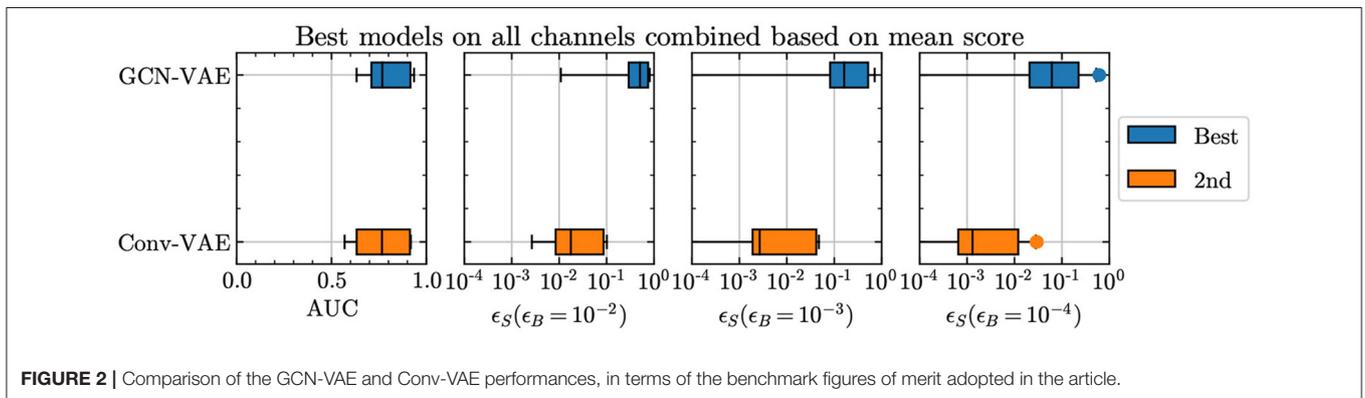
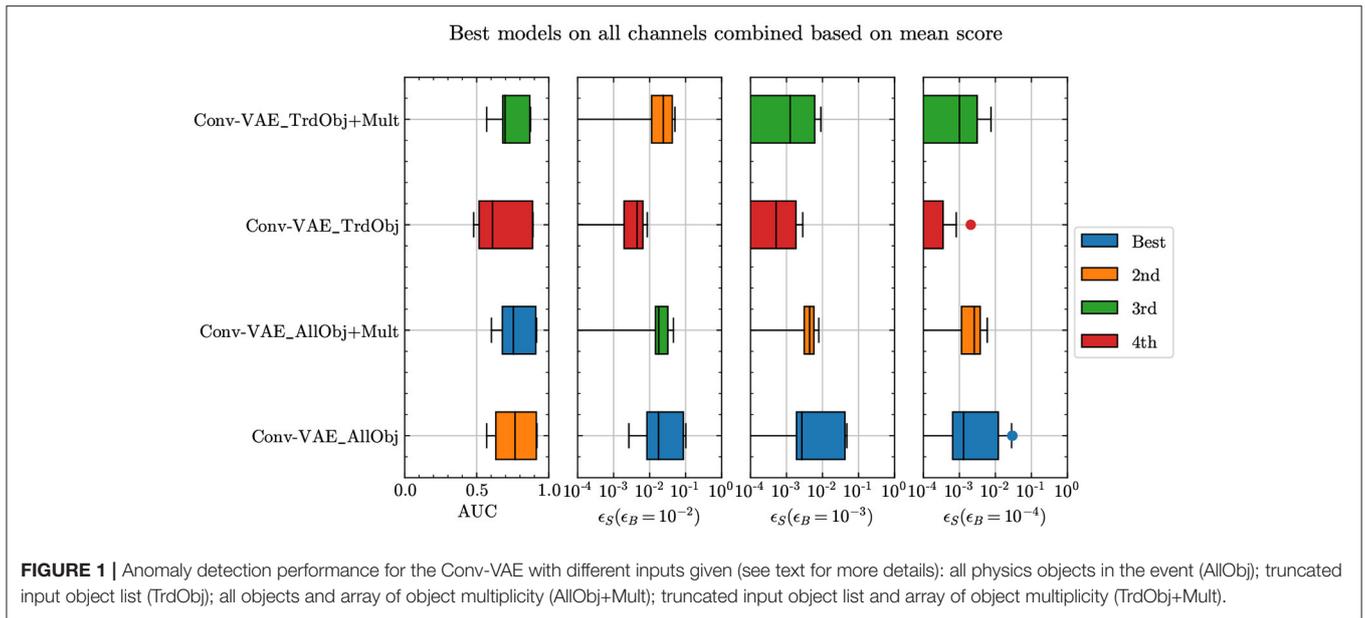
$$H_{(l+1)} = \sigma(\tilde{D}^{-\frac{1}{2}} \tilde{A} \tilde{D}^{-\frac{1}{2}} H_{(l)} W_{(l)}), \quad (3)$$

where  $H_{(l)}$  is the input to the  $(l + 1)$ th GCN layer with  $H_{(0)} = X$  where  $X$  represents the node feature matrix.  $H_{(l+1)}$  is the layer output,  $\tilde{A} = A + I$ , where  $A$  is the adjacency of the graph, with  $I$  being the identity matrix which implies added self connections for each node.  $\tilde{D}_{ii} = \sum_j A_{ij}$  is defined for the normalized adjacency based message passing regime,  $W_{(l)}$  is the layer weights matrix and  $\sigma(\bullet)$  is a suitable nonlinear activation function. The output of the last GCN layer is flattened and passed to an FCN layer which populates the latent space. The encoder has three GCN layers that scale the 5 node features to 32, 16, and 2 respectively, followed by a single FCN layer which generates a 15-dimensional latent space. The decoder has a symmetrically inverted structure with the sampled point being upsampled through an FCN layer first and the resulting output is reshaped and passed to GCN layers that reconstruct the node features.

Considering all comparison metrics along with the representative results shown in **Figure 2**, graph architectures exhibit a definitive improvement in performance compared to the Conv-VAE. The improvement is seen not only in the AUC metric, but more significantly in the  $\epsilon_S$  at low  $\epsilon_B$ . Because of this, the GCN-VAE is used as the reference architecture in the rest of this section and in Section 5.

## 4.3. Physics-Motivated High-Level Features

We also experiment with adding physics-motivated high-level features, as explicit inputs to the model, similar to what was done with object multiplicities in Section 4.1. Doing so, we intend to check if domain knowledge helps in improving anomaly detection capability. We pass event information such as the missing transverse momentum in the event ( $p_T^{\text{miss}}$ ), the scalar sum of the jet  $p_T$  ( $H_T$ ) and  $m_{\text{Eff}} = H_T + p_T^{\text{miss}}$  to the model, by concatenating these with the output of the convolutional layers of the encoder. The concatenated output is then passed to the fully connected layers in the encoder to form the latent space. After the point sampled from the latent space passes through the fully connected layers of the decoder, the reconstructed  $p_T^{\text{miss}}, H_T$ , and



$m_{\text{Eff}}$  are extracted and the rest of the layer output is re-shaped and further passed to the subsequent layers of the decoder.

To include the reconstruction of these features in the loss, we add to Equation (1) a mean-squared error (MSE) term, computed from the reconstructed and input high-level features and weighted by a coefficient. This coefficient is treated as a hyperparameter that is scanned until the best performance is found.

Figure 3 shows that adding high-level features brings no definitive improvement in performance, thereby leading us to conclude that the baseline model with marginally lower number of trainable parameters is a good choice.

#### 4.4. Anomaly Scores

While so far the Chamfer loss has been used as the anomaly score, this is not the only possibility. We consider two alternative metrics: the  $D_{\text{KL}}$  term in Equation (1) and (Aarrestad et al., 2021):

$$R_z = \sum_i \left( \frac{\mu_i}{\sigma_i} \right)^2 \quad (4)$$

where  $\mu$  and  $\sigma$  are the mean and RMS returned by the encoder and the index  $i$  runs across the latent-space dimensions.

The use of different anomaly scores requires a tuning of the  $\beta$  hyperparameter. Since  $\beta$  determines the relative importance of the  $D_{\text{KL}}$  and Chamfer loss terms in the loss, the use of one or the other as anomaly score is certainly related to the choice of the optimal  $\beta$  value. Similarly, the use of  $R_z$  (i.e., anomaly detection in the latent space) might not be optimal when using a  $\beta$  value that was tuned to emphasize the reconstruction accuracy (i.e., the minimization of the Chamfer term in the loss). On the other hand, the study in Aarrestad et al. (2021) shows that an excessive tuning of the hyperparameters affects generalization of performance negatively beyond the available dataset.

In order to address this point, we compare three weights for the  $\beta$  term. The first case ( $\beta = 1$ ) corresponds to training the VAE without the contribution of the reconstruction loss. In the second case ( $\beta = 0.5$ ) the two contributions are equally weighted. The final case ( $\beta = 10^{-6}$ ) corresponds to suppressing the  $D_{\text{KL}}$  term to a negligible level.

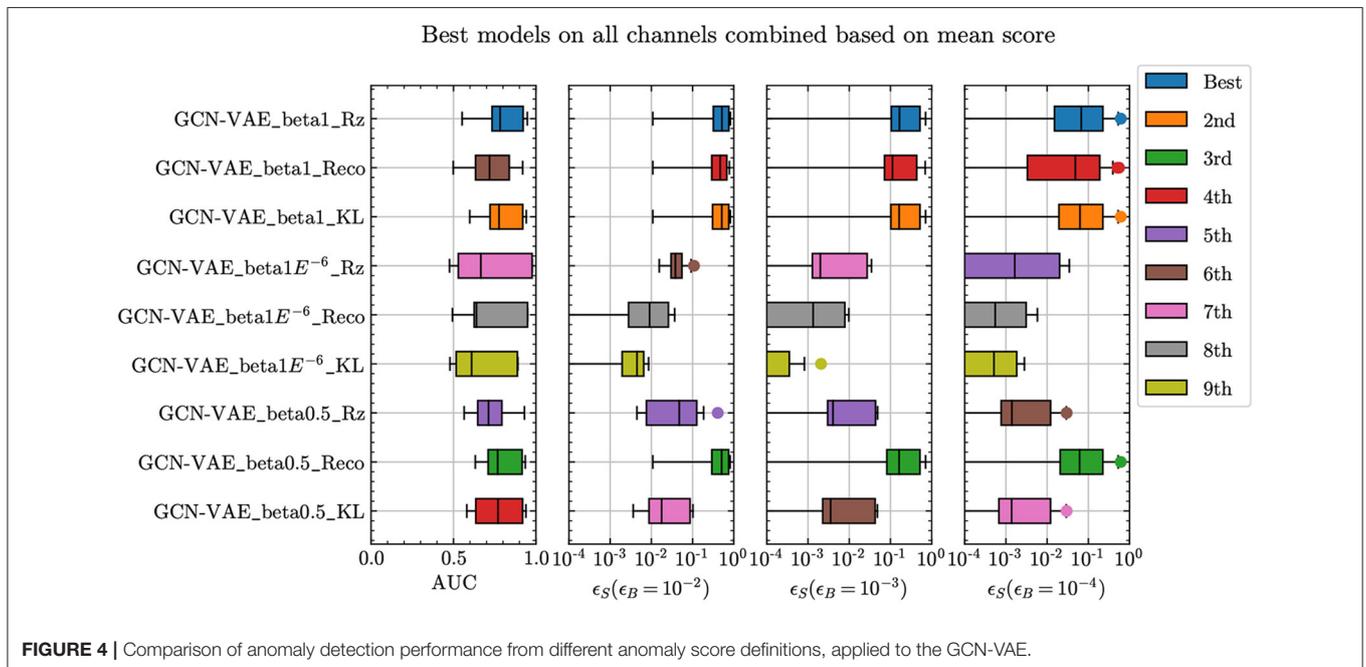
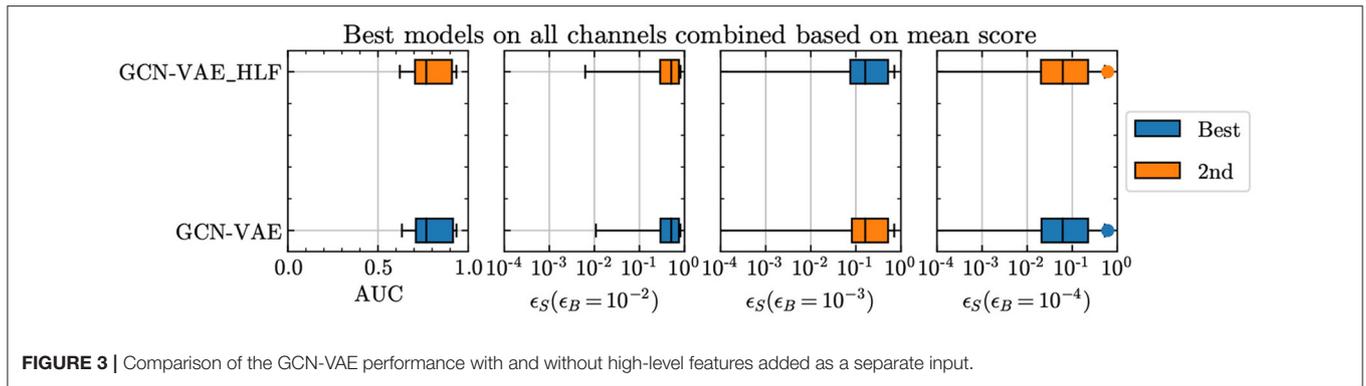


Figure 4 shows that all three anomaly scores underperform in the  $\beta = 10^{-6}$  case. The best performing models overall are the  $\beta = 1$  and  $\beta = 0.5$  cases. Comparing across the three different anomaly scores, we see that the  $\beta = 1$  model that uses  $D_{KL}$  and  $R_z$  metrics, as well as the  $\beta = 0.5$  model that uses the reconstruction metric perform the best. All three cases also show very similar performance across all comparison metrics as well as methods, implying that either model-anomaly score combination is equally suitable. We also find that the  $\beta = 1$   $D_{KL}$  score and the  $\beta = 0.5$  reconstruction score show a similar correlation pattern on signal and background. As a result, we expect that only a limited improvement would be obtained by combining the two, which spares us the cost of introducing a new hyperparameter (the relative weight of the two terms) whose optimal value would be signal-specific, as in the case of Caron et al. (2021).

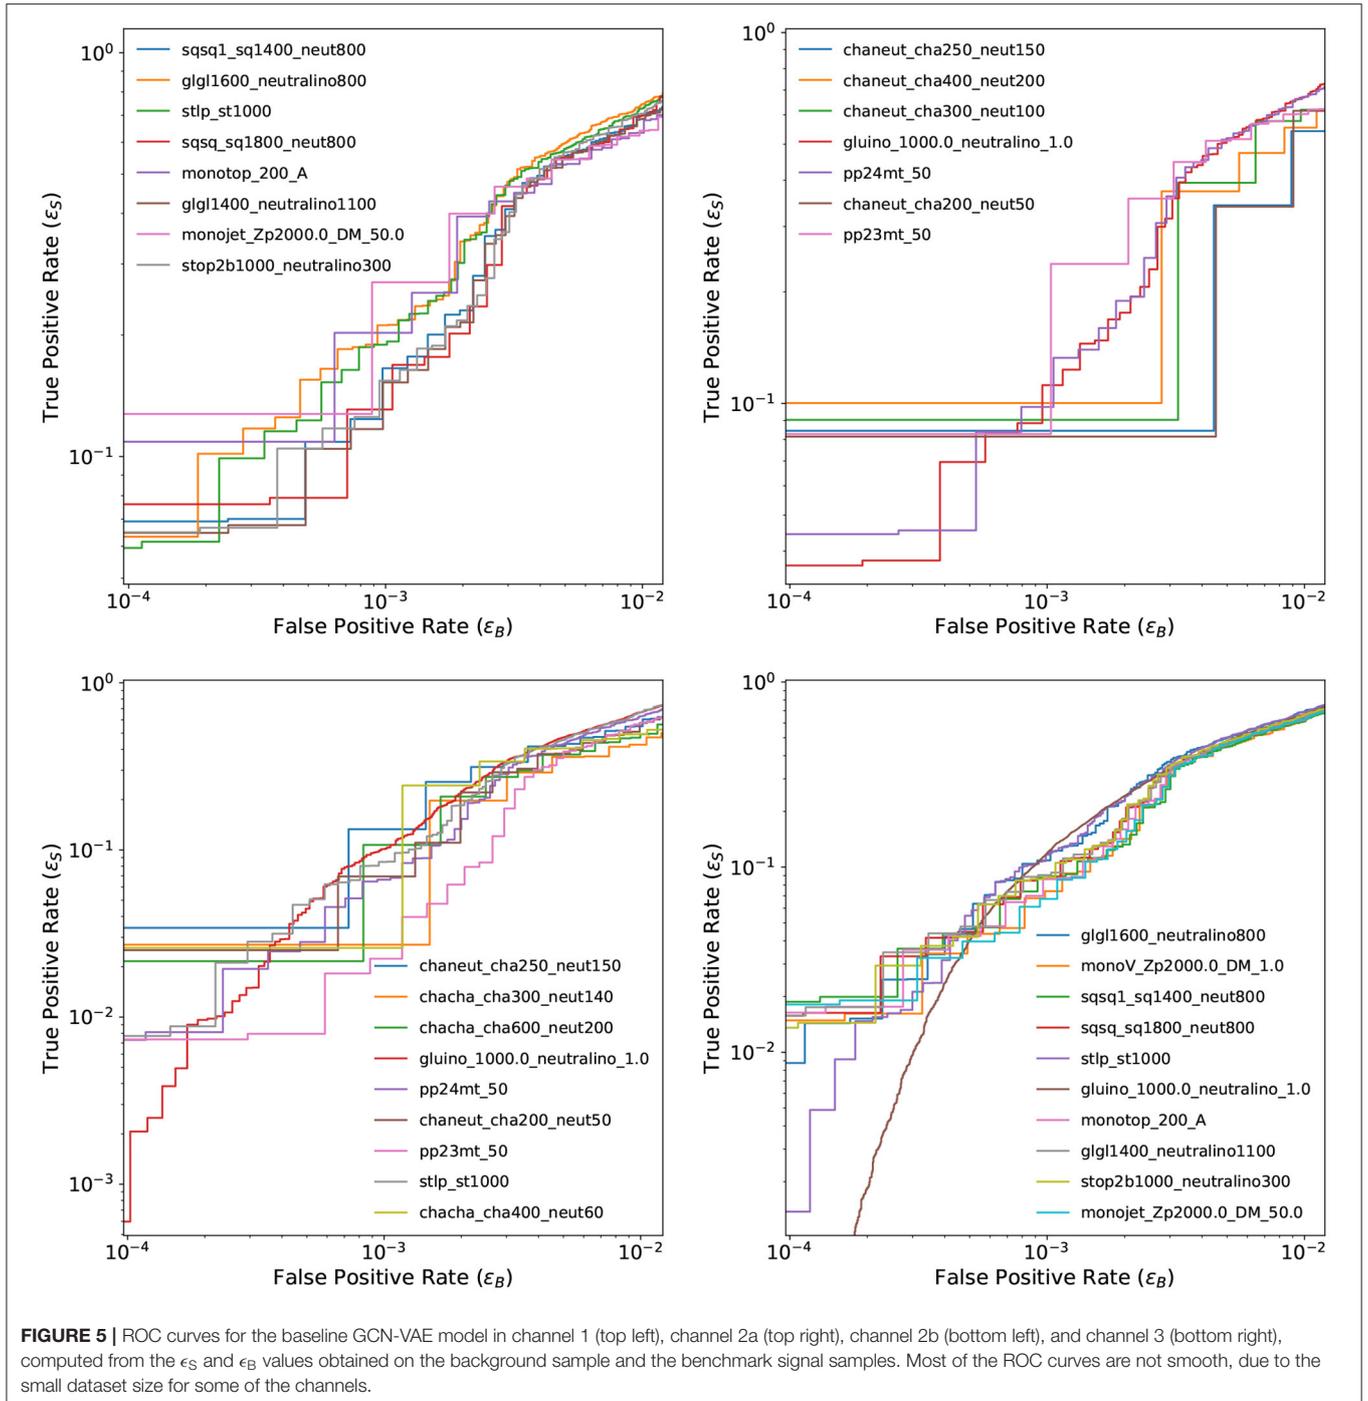
### 4.5. Baseline Discrimination

As a result of the tests presented so far, the baseline VAE model is established as a GCN-VAE taking as input the whole set of

reconstructed physics object but no domain-specific high level features. The Chamfer loss function is used as the anomaly score. The GCN-VAE is trained and tested only with data available within a given channel and the dataset sizes per channel are described in Table 1. Figure 5 shows the ROC curves for the baseline VAE model on benchmark signals in the four channels. It is evident that we suffer from a shortage of events for some signal models at very low  $\epsilon_B$ . We still show ROC curves down to  $\epsilon_B = 10^{-4}$  to allow one to compare our results to those in Aarrestad et al. (2021), where this range was chosen. We see an overall improvement in  $\epsilon_S$  at very low  $\epsilon_B$  for the GCN-VAE compared to our Conv-VAE submission in Aarrestad et al. (2021).

## 5. NORMALIZING FLOWS

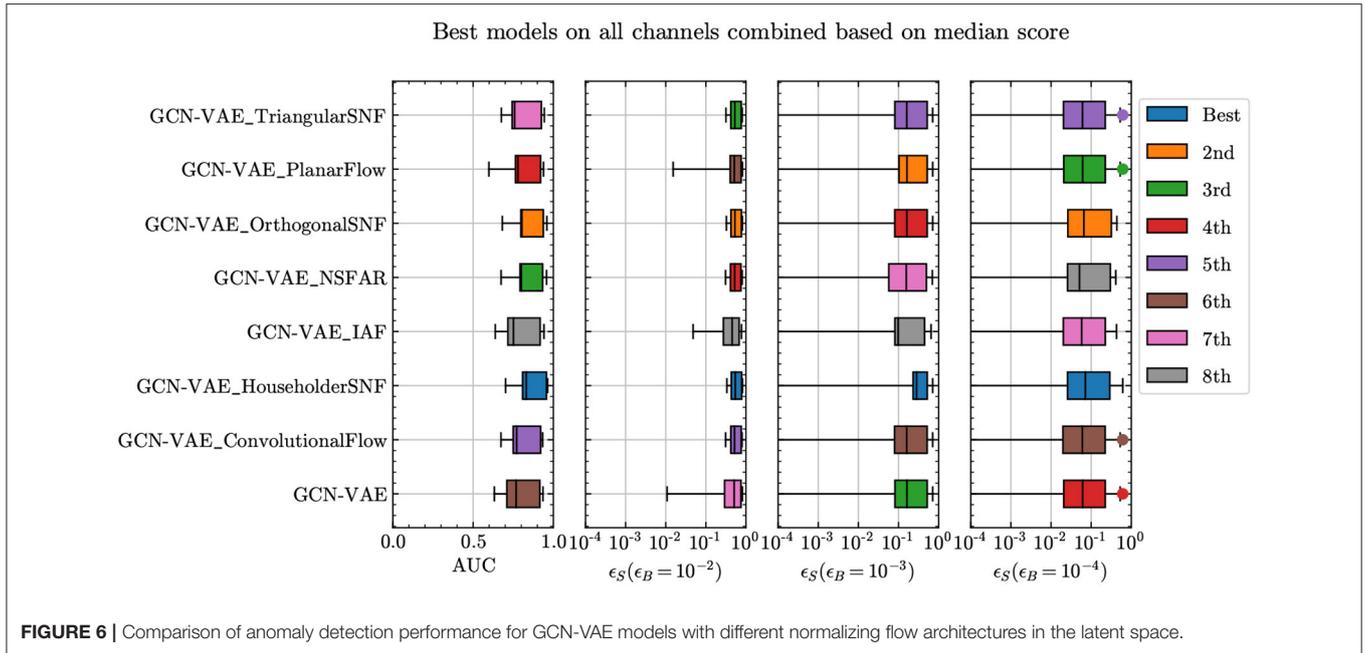
With the GCN-VAE serving as the baseline, we investigate how the use of NFs (Tabak and Vanden-Eijnden, 2010; Tabak and Turner, 2013) impacts the anomaly-detection performance.



Normalizing flow layers are inserted between the Gaussian sampling and the decoder. They provide additional complexity to learn better posterior distributions (Rezende and Mohamed, 2015) by morphing the multivariate prior of the latent space to a more suitable, learned function.

In other words, we use the NF layers to handle the fact that a VAE converging to a good output-to-input matching does not necessarily correspond to a configuration with a Gaussian prior

in the latent space,  $p(z) = \prod G(z)$ . To reach this configuration (e.g., when training a VAE as a generative model), one typically uses a  $\beta$ -VAE with an increased weighting of the  $D_{KL}$  regularizer. This typically results in a degradation of the output-to-input matching. With NFs, we learn a generic prior  $p(z)$  as  $f[G(z)]$ , where  $f$  is the transformation function learned by the NF layers. This is different from the way NFs are traditionally used in VAE training, i.e., to improve the convergence of  $f(z)$  to  $G(z)$  with



a stronger evidence lower bound (ELBO) condition. Because of this, we do not modify the  $D_{KL}$  term in the loss, as done in Rezende and Mohamed (2015). The results obtained following this more traditional training procedure are described in the **Supplementary Material**. Doing so, we observe worse  $\epsilon_S$  for the same  $\epsilon_B$ . This is expected because the ELBO improvement with NFs was introduced in Tomczak and Welling (2017) as a way to improve the VAE generative properties, and it does not imply a better anomaly detection capability.

A NF can be generalized as any invertible, diffeomorphic transformation that can be applied to a given distribution to produce tractable distributions (Kobyzev et al., 2020; Papamakarios et al., 2021). In order to be compatible with variational inference, it is desirable for the transformations to have an efficient mechanism for computing the determinant of the Jacobian, while being invertible (Rezende and Mohamed, 2015). The NFs are trained sequentially, together with the baseline VAE model.

We utilize four major families of flow models:

- **Planar flows** are invertible transformations whose Jacobian determinant can be computed rather efficiently, making them suitable candidates for variational inference (Rezende and Mohamed, 2015). PF transformations are defined as:

$$\mathbf{z}' = \mathbf{z} + u\mathbf{h}(w^T \mathbf{z} + b), \tag{5}$$

where  $u, w \in \mathbb{R}^D$ ,  $b \in \mathbb{R}$  and  $h$  is a suitable smooth activation function.

- **Sylvester normalizing flows** (SNFs) (Berg et al., 2018) build on the planar flow formulation and extend it to be analogous to a multilayer perceptron with one hidden layer of  $M$  units and a residual connection as:

$$\mathbf{z}' = \mathbf{z} + \mathbf{A}\mathbf{h}(\mathbf{B}\mathbf{z} + b), \tag{6}$$

where  $\mathbf{A} \in \mathbb{R}^{D \times M}$ ,  $\mathbf{B} \in \mathbb{R}^{M \times D}$ ,  $b \in \mathbb{R}^M$  and  $M \leq D$ . Computing the Jacobian determinant for such a formulation is made more efficient by utilizing the Sylvester determinant identity (Berg et al., 2018). Depending on the way  $\mathbf{A}$  and  $\mathbf{B}$  are parameterized, we get different types of SNFs. In this article, we use orthogonal, Householder, and triangular SNFs, as described in Berg et al. (2018).

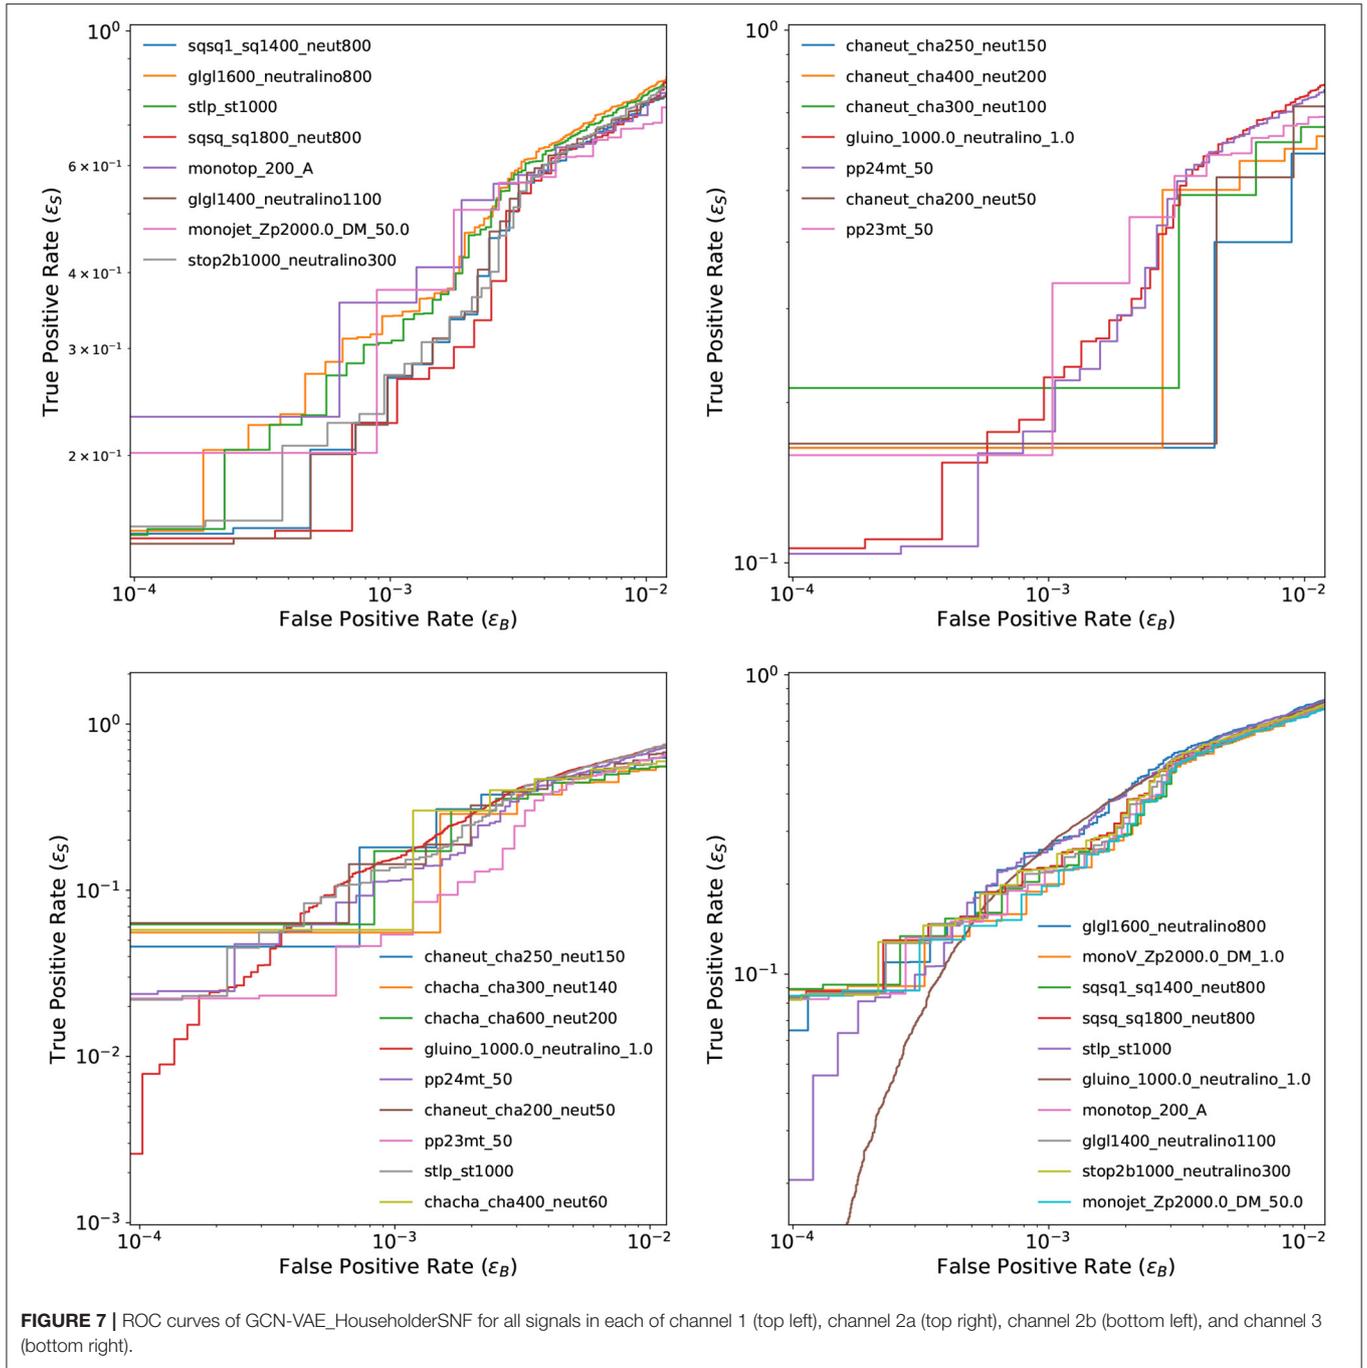
- **Inverse autoregressive flows** (IAFs) (Kingma et al., 2016) are computation-efficient normalizing flows based on autoregressive models. Autoregressive transformations are invertible, making them suitable candidates for our case. However, computing the transformation requires multiple sequential steps (Berg et al., 2018). The inverse transformation however, leads to certain simplifications as described in Berg et al. (2018), allowing more efficient parallel computing, thereby making it a more desirable transformation for our case. We use the IAFs formulated as:

$$z_i^t = \mu_i^t(z_{1:i-1}^{t-1}) + \sigma_i^t(z_{1:i-1}^{t-1}) \cdot z_i^{t-1}, \quad i = 1, 2, \dots, D. \tag{7}$$

Such a formulation allows one to stack multiple transformations to achieve more flexibility in producing target distributions.

- **Convolutional normalizing flows** (ConvolutionalFlows) (Zheng et al., 2018) are an extension of single-hidden-unit planar flows (Kingma et al., 2016) to the case of multiple hidden units, further enhanced by replacing the fully connected network operation with a one-dimensional (1D) convolution, to achieve bijectivity. They are defined by the following transformation:

$$\mathbf{z}' = \mathbf{z} + \mathbf{u} \odot \mathbf{h}[\text{conv}(\mathbf{z}, \mathbf{w})], \tag{8}$$



where  $w \in R^k$  is the parameter of the 1D convolution filter with  $k$ -sized kernel,  $h$  is a monotonic nonlinear activation function and  $\odot$  denotes pointwise multiplication.

- **Autoregressive neural spline flows** (NSFARs) (Durkan et al., 2019) are similar to IAFs, where affine transforms are replaced by monotonic rational-quadratic spline transforms as described in Durkan et al. (2019). They resemble a traditional feed-forward neural network architecture, alternating between linear transformations and elementwise non-linearities, while retaining an exact, analytic inverse.

The hyperparameters for each normalizing flow architecture are chosen arbitrarily to avoid overtuning on the available dataset as learned from Aarrestad et al. (2021). The planar flow model consists of a stack of six flows, each made of three dense layers with 90 neurons each. SNFs are defined by stacking six flows with eight orthogonal, householder, and triangular transformations for each of the respective types of SNF. IAFs are constructed with four masked autoencoder for distribution estimation (MADE) (Germain et al., 2015) layers as described in Kingma et al. (2016), each containing 330 neurons.

ConvolutionalFlows include four flow layers with kernel size  $k = 7$  and applying kernel dilation as described in Zheng et al. (2018). NSFARs are defined by stacking four flow layers each with  $K = 64$  bins and eight hidden features.

**Figure 6** shows the results of all GCN-VAE models combined with all the different types of flows as described in Section 5. Based on results from all data channels combined through all six strategies mentioned in Section 3, and considering variance across trainings from different random seeds (see **Supplementary Material**), it is evident that using normalizing flows improves not only the AUC metric but also the signal efficiencies at low background efficiencies. We find that the Householder variant of SNFs produces the best improvement with respect to the baseline GCN-VAE model. The exercise was also repeated with a Conv-VAE model and similar trends were observed. There, the normalizing flows showed a larger improvement from the baseline Conv-VAE than for the GCN-VAE model but the overall results are less accurate than that of GCN-VAE with normalizing flows.

**Figure 7** shows the ROC curves for the best presented model, GCN-VAE\_HouseholderSNF across all available signal samples in all data channels. For some of the samples, the small dataset size translates in a discontinuous curve and larger uncertainties.

## 6. CONCLUSIONS

We constructed a graph-based anomaly detection model to identify new physics events in the DarkMachines challenge dataset. Inspired by the outcome of this challenge, specific model design choices (data representation, use of physics-motivated high-level features, and anomaly score definition) were further optimized in order to maximize anomaly detection performance. As the case for many other deep learning applications to particle-physics data, we observed that the graph architecture better captures the point-cloud nature of this data, resulting in an enhanced performance.

In this baseline, we investigate the impact of using a stack of normalizing flows in the latent space of the variational autoencoder (VAE), between the Gaussian sampling and the decoding, in order to improve the accuracy of the prior learning process, by morphing the Gaussian prior to a more suitable function, learned during the training.

Testing the trained model on a set of benchmark signal samples, we observe an overall improvement when normalizing

flows are used, with the Householder variant of the Sylvester normalizing flow model giving the best results. With that, we reach a median anomaly identification probability of 72% (34%) for an  $\epsilon_B$  of 1% (0.1%) across all signal samples over all available channels. The median anomaly identification probability increases to 95% (96%) for an  $\epsilon_B$  of 30% (60%).

This work presents an improvement over our Conv-VAE model, submitted to the DarkMachines challenge (Aarrestad et al., 2021).

## DATA AVAILABILITY STATEMENT

The original contributions presented in the study are included in the article/**Supplementary Material**, further inquiries can be directed to the corresponding author/s.

## AUTHOR CONTRIBUTIONS

All authors in equal share developed the baseline VAE model. All authors in equal share took part in writing and editing the manuscript. PJ and MP developed the VAE + normalizing flow models. All authors contributed to the article and approved the submitted version.

## FUNDING

PJ, TA, MP, and KW were supported by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (Grant Agreement No. 772369). JD was supported by the U.S. Department of Energy (DOE), Office of Science, Office of High Energy Physics Early Career Research program under Award No. DE-SC0021187. ST was supported by the University of California San Diego Triton Research and Experiential Learning Scholars (TRELS) program. JN was supported by Fermi Research Alliance, LLC under Contract No. DE-AC02-07CH11359 with the U.S. Department of Energy, Office of Science, Office of High Energy Physics.

## SUPPLEMENTARY MATERIAL

The Supplementary Material for this article can be found online at: <https://www.frontiersin.org/articles/10.3389/fdata.2022.803685/full#supplementary-material>

## REFERENCES

- Aaboud, M., Aad, G., Abbott, B., Abidinov, O., Abeloos, B., Abidi, S. H., et al. (2019). A strategy for a general search for new phenomena using data-derived signal regions and its application within the ATLAS experiment. *Eur. Phys. J. C* 79:120. doi: 10.1140/epjc/s10052-019-6540-y
- Aad, G., Abbott, B., Abbott, D. C., Abed Abud, A., Abeling, K., Abhayasinghe, D. K., et al. (2020). Operation of the ATLAS trigger system in Run 2. *J. Instrum.* 15:P10004. doi: 10.1088/1748-0221/15/10/P10004
- Aaltonen, T., Adelman, J., Akimoto, T., Albrow, M. G., Álvarez González, B., Amerio, S., et al. (2009). Global search for new physics with 2.0 fb<sup>-1</sup> at CDF. *Phys. Rev. D* 79:011101. doi: 10.1103/PhysRevD.79.011101
- Aaron, F. D., Alexa, C., Andreev, V., Antunovic, B., Aplin, S., Asmone, A., et al. (2009). A general search for new phenomena at HERA. *Phys. Lett. B* 674, 257–268. doi: 10.1016/j.physletb.2009.03.034
- Aarrestad, T., van Beekveld, M., Bona, M., Boveia, A., Caron, S., Davies, J., et al. (2021). The dark machines anomaly score challenge: benchmark data and model independent event classification for the large hadron collider. *Sci. Post Phys. arXiv [Preprint]*. arXiv: 2105.14027.

- Amram, O., and Suarez, C. M. (2021). Tag N' Train: a technique to train improved classifiers on unlabeled data. *J. High Energ. Phys.* 1:153. doi: 10.1007/JHEP01(2021)153
- An, J., and Cho, S. (2015). Variational autoencoder based anomaly detection using reconstruction probability. *Spec. Lect. IE 2:1*.
- Andreassen, A., Nachman, B., and Shih, D. (2020). Simulation assisted likelihood-free anomaly detection. *Phys. Rev. D* 101:95004. doi: 10.1103/PhysRevD.101.095004
- Barrow, H. G., Tenenbaum, J. M., Bolles, R. C., and Wolf, H. C. (1977). "Parametric correspondence and Chamfer matching: two new techniques for image matching," in *Proceedings of the 5th International Joint Conference on Artificial Intelligence (IJCAI)*, Vol. 2 (San Francisco, CA: Morgan Kaufmann Publishers Inc.), 659.
- Berg, R. V. d., Hasenclever, L., Tomczak, J. M., and Welling, M. (2018). "Sylvester normalizing flows for variational inference," in *Conference on Uncertainty in Artificial Intelligence (UAI) 2018* (Monterey, CA: UAI). Available online at: <http://auai.org/uai2018/proceedings/papers/156.pdf>
- Blance, A., Spannowsky, M., and Waite, P. (2019). Adversarially-trained autoencoders for robust unsupervised new physics searches. *J. High Energ. Phys.* 10:047. doi: 10.1007/JHEP10(2019)047
- Bortolato, B., Dillon, B., Kamenik, J. F., and Smolkovic, A. (2021). Bump hunting in space. *arXiv [Preprint]*. arXiv: 2103.06595. Available online at: <https://arxiv.org/pdf/2103.06595.pdf> (accessed March 11, 2021).
- Bronstein, M. M., Bruna, J., LeCun, Y., Szlam, A., and Vandergheynst, P. (2017). Geometric deep learning: going beyond Euclidean data. *IEEE Signal Process. Mag.* 34:18. doi: 10.1109/MSP.2017.2693418
- Caron, S., Hendriks, L., and Verheyen, R. (2021). Rare and different: Anomaly scores from a combination of likelihood and out-of-distribution models to detect new physics at the LHC. *arXiv [Preprint]*. arXiv: 2106.10164. Available online at: <https://arxiv.org/pdf/2106.10164.pdf> (accessed December 22, 2021).
- Cerri, O., Nguyen, T. Q., Pierini, M., Spiropulu, M., and Vlimant, J.-R. (2019). Variational autoencoders for new physics mining at the large hadron collider. *J. High Energ. Phys.* 5:36. doi: 10.1007/JHEP05(2019)036
- Cheng, T., Arguin, J. -F., Leissner-Martin, J., Pilette, J., and Golling, T. (2020). Variational autoencoders for anomalous jet tagging. *arXiv [Preprint]*. arXiv: 2007.01850. Available online at: <https://arxiv.org/pdf/2007.01850.pdf> (accessed February 15, 2021).
- CMS-PAS-EXO-14-016 (2017). "MUSIC, a model unspecific search for new physics," in *PP Collisions at textitsqrts = 8 textitmathrmTeV*. Available online at: <https://cds.cern.ch/record/225665/CMS-PAS-EXO-14-016>
- Collins, J. H., Howe, K., and Nachman, B. (2018). Anomaly detection for resonant new physics with machine learning. *Phys. Rev. Lett.* 121:241803. doi: 10.1103/PhysRevLett.121.241803
- Collins, J. H., Howe, K., and Nachman, B. (2019). Extending the search for new resonances with machine learning. *Phys. Rev. D* 99:014038. doi: 10.1103/PhysRevD.99.014038
- Collins, J. H., Martín-Ramiro, P., Nachman, B., and Shih, D. (2021). Comparing weak- and unsupervised methods for resonant anomaly detection. *Eur. Phys. J. C* 81:617. doi: 10.1140/epjc/s10052-021-09389-x
- D0 Collaboration (2012). Model independent search for new phenomena in  $p\bar{p}$  collisions at  $\sqrt{s}=1.96$  TeV. *Phys. Rev. D* 85:092015. doi: 10.1103/PhysRevD.85.092015
- D'Agnolo, R. T., Grosso, G., Pierini, M., Wulzer, A., and Zanetti, M. (2021). Learning multivariate new physics. *Eur. Phys. J. C* 81:89. doi: 10.1140/epjc/s10052-021-08853-y
- D'Agnolo, R. T., and Wulzer, A. (2019). Learning new physics from a machine. *Phys. Rev. D* 99:015014. doi: 10.1103/PhysRevD.99.015014
- DarkMachines Community (2020). *Unsupervised-Hackathon*. DarkMachines Community.
- De Simone, A., and Jacques, T. (2019). Guiding new physics searches with unsupervised learning. *Eur. Phys. J. C* 79:289. doi: 10.1140/epjc/s10052-019-6787-3
- Dillon, B. M., Faroughy, D. A., Kamenik, J. F., and Szwec, M. (2020). Learning the latent structure of collider events. *J. High Energ. Phys.* 10:206. doi: 10.1007/JHEP10(2020)206
- Duarte, J., and Vilmant, J.-R. (2020). "Graph neural networks for particle tracking and reconstruction," in *Artificial Intelligence for High Energy Physics*, eds P. Calafiura, D. Rousseau, and K. Terao (Hackensack, NJ: World Scientific Publishing). doi: 10.1142/12200
- Durkan, C., Bekasov, A., Murray, I., and Papamakarios, G. (2019). Neural spline flows. *Adv. Neural Inform. Process. Syst.* 32, 7511–7522.
- Fan, H., Su, H., and Guibas, L. J. (2017). "A point set generation network for 3D object reconstruction from a single image," in *2017 IEEE Conference on Computer Vision and Pattern Recognition* (Honolulu, HA: IEEE), 2463. doi: 10.1109/CVPR.2017.264
- Farina, M., Nakai, Y., and Shih, D. (2020). Searching for new physics with deep autoencoders. *Phys. Rev. D* 101:075021. doi: 10.1103/PhysRevD.101.075021
- Finke, T., Krämer, M., Morandini, A., Mück, A., and Oleksiyuk, I. (2021). Autoencoders for unsupervised anomaly detection in high energy physics. *J. High Energ. Phys.* 6:161. doi: 10.1007/JHEP06(2021)161
- Germain, M., Gregor, K., Murray, I., and Larochelle, H. (2015). "MADE: masked autoencoder for distribution estimation," in *Proceedings of the 32nd International Conference on Machine Learning*, Vol. 37 of *Proceedings of Machine Learning Research*, eds F. Bach and D. Blei (Lille), 881.
- Gonski, J., Lai, J., Nachman, B., and Ochoa, I. (2021). High-dimensional anomaly detection with radiative return in  $e^+e^-$  collisions. *arXiv [Preprint]*. arXiv: 2108.13451. Available online at: <https://arxiv.org/pdf/2108.13451.pdf> (accessed February 08, 2022).
- Gross, E., and Vitells, O. (2010). Trial factors for the look elsewhere effect in high energy physics. *Eur. Phys. J. C* 70:525. doi: 10.1140/epjc/s10052-010-1470-8
- Hajer, J., Li, Y.-Y., Liu, T., and Wang, H. (2020). Novelty detection meets collider physics. *Phys. Rev. D* 101:076015. doi: 10.1103/PhysRevD.101.076015
- Hallin, A., Isaacson, J., Kasieczka, G., Krause, C., Nachman, B., Quadfasel, T., et al. (2021). Classifying anomalies through outer density estimation (CATHODE). *arXiv [Preprint]*. arXiv: 2109.00546. Available online at: <https://arxiv.org/pdf/2109.00546.pdf> (accessed October 29, 2021).
- Heimel, T., Kasieczka, G., Plehn, T., and Thompson, J. M. (2019). QCD or what? *Sci. Post Phys.* 6:30. doi: 10.21468/SciPostPhys.6.3.030
- Higgins, I., Matthey, L., Pal, A., Burgess, C. P., Glorot, X., Bovinick, M., et al. (2017). "Beta-VAE: Learning basic visual concepts with a constrained variational framework," in *5th International Conference on Learning Representations* (Toulon). Available online at: <https://openreview.net/forum?id=Sy2fzU9gl>
- Jawahar, P., and Pierini, M. (2021). *mpp-hep/DarkFlow repository*. Available online at: <https://github.com/mpp-hep/DarkFlow>
- Kasieczka, G., et al. (2021). The LHC olympics 2020: a community challenge for anomaly detection in high energy physics. *Rep. Prog. Phys.* 84:124201. doi: 10.1088/1361-6633/ac36b9
- Khosa, C. K., and Sanz, V. (2020). Anomaly awareness.
- Kingma, D. P., and Ba, J. (2015). "Adam: A method for stochastic optimization," in *3rd International Conference for Learning Representations* (San Diego, CA). *arXiv [Preprint]*. arXiv: 1412.6980.
- Kingma, D. P., Salimans, T., Jozefowicz, R., Chen, X., Sutskever, I., and Welling, M. (2016). "Improving variational inference with inverse autoregressive flow," in *Advances in Neural Information Processing Systems*, Vol. 29, eds D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, and R. Garnett (Barcelona: Curran Associates, Inc.). *arXiv [Preprint]*. arXiv: 1606.04934. Available online at: <https://proceedings.neurips.cc/paper/2016/file/ddeebdeefdb7e7e7a697e1c3e3d8ef54-Paper.pdf>
- Kingma, D. P., and Welling, M. (2014). "Auto-encoding variational Bayes," in *2nd International Conference on Learning Representations, ICLR 2014* (Banff, AB).
- Kingma, D. P., and Welling, M. (2019). An introduction to variational autoencoders. *Found. Trends Mach. Learn.* 12:307. doi: 10.1561/9781680836233
- Kipf, T. N., and Welling, M. (2017). "Semi-supervised classification with graph convolutional networks," in *5th International Conference on Learning Representations* (Toulon). Available online at: <https://openreview.net/forum?id=SJU4ayYgl>
- Kobyzev, I., Prince, S., and Brubaker, M. (2020). Normalizing flows: an introduction and review of current methods. *IEEE Trans. Pattern Anal. Mach. Intell.* 43, 3964–3979. doi: 10.1109/TPAMI.2020.2992934
- Nachma, B., and Shih, D. (2020). Anomaly detection with density estimation. *Phys. Rev. D* 101, 075042.

- Nachman, B., and Shih, D. (2020). Anomaly detection with density estimation. *Phys. Rev. D* 101:075042. doi: 10.1103/PhysRevD.101.075042
- Ostdiek, B. (2021). Deep set auto encoders for anomaly detection in particle physics. *arXiv [Preprint]*. arXiv: 2109.01695. Available online at: <https://arxiv.org/pdf/2109.01695.pdf> (accessed November 15, 2021).
- Papamakarios, G., Nalisnick, E., Rezende, D. J., Mohamed, S., and Lakshminarayanan, B. (2021). Normalizing flows for probabilistic modeling and inference. *J. Mach. Learn. Res.* 22, 1–64. Available online at: <http://jmlr.org/papers/v22/19-1028.html>
- Park, S. E., Rankin, D., Udrescu, S.-M., Yunus, M., and Harris, P. (2020). Quasi Anomalous Knowledge: Searching for new physics with embedded knowledge. *J. High Energ. Phys.* 21:30. doi: 10.1007/JHEP06(2021)030
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., et al. (2019). “PyTorch: an imperative style, high-performance deep learning library,” in *Advances in Neural Information Processing Systems*, Vol. 32, eds H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett (Vancouver, BC: Curran Associates, Inc.). Available online at: <https://proceedings.neurips.cc/paper/2019/file/bdbca288fee7f92f2bfa9f7012727740-Paper.pdf>
- 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.
- Rezende, D., and Mohamed, S. (2015). “Variational inference with normalizing flows,” in *Proceedings of the 32nd International Conference on Machine Learning*, Vol. 37, eds F. Bach and D. Blei (Lille), 1530.
- Rezende, D. J., Mohamed, S., and Wierstra, D. (2014). “Stochastic backpropagation and approximate inference in deep generative models,” in *Proceedings of the 31st International Conference on Machine Learning*, Vol. 32 of *Proceedings of Machine Learning Research* (Beijing), 1278.
- Shlomi, J., Battaglia, P., and Vlimant, J.-R. (2020). Graph neural networks in particle physics. *Mach. Learn. Sci. Tech.* 2:21001. doi: 10.1088/2632-2153/abbf9a
- Sirunyan, A. M., Tumasyan, A., Adam, W., Ambrogio, F., Arnold, B., Bergauer, H., et al. (2020). Performance of the CMS Level-1 trigger in proton-proton collisions at  $\sqrt{s} = 13$  TeV. *J. Instrum.* 15:P10017. doi: 10.1088/1748-0221/15/10/P10017
- Tabak, E. G., and Turner, C. V. (2013). A family of nonparametric density estimation algorithms. *Commun. Pure Appl. Math.* 66:145. doi: 10.1002/cpa.21423
- Tabak, E. G., and Vanden-Eijnden, E. (2010). Density estimation by dual ascent of the log-likelihood. *Commun. Math. Sci.* 8:217. doi: 10.4310/CMS.2010.v8.n1.a11
- Tomczak, J. M., and Welling, M. (2017). “Improving variational auto-encoders using convex combination linear inverse autoregressive flow,” in *Benelearn 2017* (Eindhoven). Available online at: <http://www.wis.win.tue.nl/~benelearn2017/>
- Trocino, D. (2014). The CMS high level trigger. *J. Phys. Conf. Ser.* 513:012036. doi: 10.1088/1742-6596/513/1/012036
- Weisser, C., and Williams, M. (2016). Machine learning and multivariate goodness of fit. *arXiv [Preprint]*. arXiv: 1612.07186. Available online at: <https://arxiv.org/pdf/1612.07186.pdf> (accessed December 20, 2016).
- Zhang, Y., Jonathon, H., and Adam, P. -B. (2020). “FSPool: Learning set representations with featurewise sort pooling,” in *8th International Conference on Learning Representations* (Addis Ababa). Available online at: <https://openreview.net/forum?id=HJgBA2VYwH>
- Zheng, G., Yang, Y., and Carbonell, J. (2018). “Convolutional normalizing flows,” in *ICML 2018 Workshop on Theoretical Foundations and Applications of Deep Generative Models* (Stockholm). Available online at: [https://drive.google.com/file/d/1\\_40TktDTeKG2eEpbE-D\\_hQcb3oSj5ebQ/view](https://drive.google.com/file/d/1_40TktDTeKG2eEpbE-D_hQcb3oSj5ebQ/view)
- Zhou, J., Cui, G., Hu, S., Zhang, Z., Yang, C., Liu, Z., et al. (2020). Graph neural networks: a review of methods and applications. *AI Open* 1:57. doi: 10.1016/j.aiopen.2021.01.001

**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.

**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.

Copyright © 2022 Jawahar, Arrestad, Chernyavskaya, Pierini, Wozniak, Ngadiuba, Duarte and Tsan. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). 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.