

UC Berkeley

UC Berkeley Previously Published Works

Title

Self-supervised anomaly detection for new physics

Permalink

<https://escholarship.org/uc/item/3wr6h123>

Journal

Physical Review D, 106(5)

ISSN

2470-0010

Authors

Dillon, Barry M

Mastandrea, Radha

Nachman, Benjamin

Publication Date

2022-09-01

DOI

10.1103/physrevd.106.056005

Copyright Information

This work is made available under the terms of a Creative Commons Attribution-NonCommercial License, available at <https://creativecommons.org/licenses/by-nc/4.0/>

Peer reviewed

Self-supervised Anomaly Detection for New Physics

Barry M. Dillon,^{1,*} Radha Mastandrea,^{2,3,†} and Benjamin Nachman^{3,4,‡}

¹*Universität Heidelberg, Heidelberg, Germany*

²*Department of Physics, University of California, Berkeley, CA 94720, USA*

³*Physics Division, Lawrence Berkeley National Laboratory, Berkeley, CA 94720, USA*

⁴*Berkeley Institute for Data Science, University of California, Berkeley, CA 94720, USA*

(Dated: May 24, 2022)

We investigate a method of model-agnostic anomaly detection through studying jets, collimated sprays of particles produced in high-energy collisions. We train a transformer neural network to encode simulated QCD “event space” dijets into a low-dimensional “latent space” representation. We optimize the network using the self-supervised contrastive loss, which encourages the preservation of known physical symmetries of the dijets. We then train a binary classifier to discriminate a BSM resonant dijet signal from a QCD dijet background both in the event space and the latent space representations. We find the classifier performances on the event and latent spaces to be comparable. We finally perform an anomaly detection search using a weakly supervised bump hunt on the latent space dijets, finding again a comparable performance to a search run on the physical space dijets. This opens the door to using low-dimensional latent representations as a computationally efficient space for resonant anomaly detection in generic particle collision events.

I. INTRODUCTION

A central goal of high energy physics is to find the theory that will supersede the Standard Model. A number of competing models exist, with many involving new resonant particles [1, 2] such as supersymmetric partners or weakly interacting massive particles. However, the number of candidate particles is too large to justify a hunt-and-pick procedure of data analysis. Therefore it is advantageous to consider methods of anomaly detection that are agnostic to a particular underlying model of new physics.

There are a number of recent machine learning-based anomaly detection proposals designed to reduce model dependence (see Refs. [3–6] for overviews). Notably, most existing methods (including the first result with data from ATLAS [7]) in the field are best-performing in low dimensional spaces. However, a single event from a particle collision experiment can have on the order of a thousand degrees of freedom.

A resolution to this tension is to reduce the dimensionality of an entire particle collision event while preserving its essential character such that a search for anomalies can be done in the reduced-dimension space. A number of methods exist for carrying out this phase space reduction. For example, one could choose a set of observables (e.g. mass, multiplicity) and perform anomaly detection on this set. However, attempts to reduce dimensionality through selecting a choice of observables implicitly favors a certain class of models.

Dimensionality reduction can also be performed with unsupervised machine learning (ML) techniques, which

ensure a model-agnostic approach. A common tool in the ML literature for this compression is the autoencoder (AE). Autoencoders are a pair of neural networks whereby one function encodes *event space* data into a *latent space* and a second function decodes the latent space back into the event space. No labels are needed because the AE is trained to ensure that the composition of the encoder and decoder is close to the identity. While effective, AEs may not be ideal for anomaly detection. For example, there is nothing in their structure that ensures that anomalies are preserved by the encoder. They are also not well-suited to handle variable length data, as is the case for collider events. The use of variational autoencoders to obtain compressed latent spaces for dijet bump hunting has in fact been explored in [8].

These problems can be circumvented by making use of self-supervised contrastive learning techniques, which use only the inherent symmetries of the physical data to perform the dimensionality reduction. Recent studies have explored this; for example with astronomical images in Ref. [9], and with constituent level jet data in the JetCLR approach of Ref. [10]. In the former a ResNet50 architecture [11] is used to map each astronomical image to a latent representation, while in the latter, a permutation-invariant transformer-encoder architecture maps jet constituents to a latent representation. The networks are trained on the contrastive loss, which ensures that the latent space representations faithfully model the physical symmetries of the original objects. The analysis is then done directly on the latent space representations.

In this paper, we continue the explorations of latent space representations of particle collisions originally carried out in Ref. [10]. We first demonstrate that particle collisions can be well-modeled in a latent space representation with a dimensionality that is an order of magnitude smaller than that of the original events. As part of this work, we extend the per-jet work of Ref. [10] to a per-event structure. We then conduct a low-dimension

* dillon@thphys.uni-heidelberg.de

† rmastand@berkeley.edu

‡ bpnachman@lbl.gov

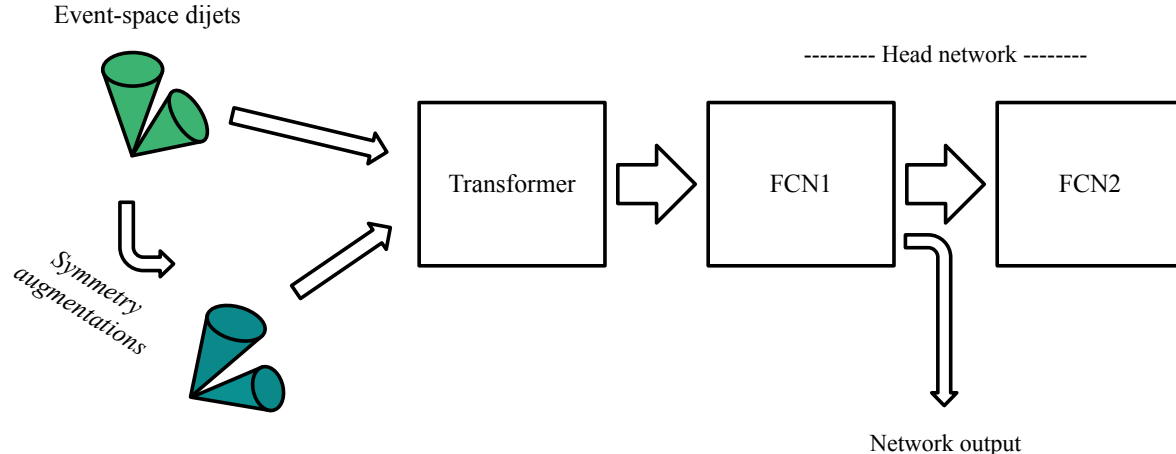


FIG. 1: A schematic of the full transformer encoder network. Event space dijets and their symmetry-augmented versions are fed as input into the transformer network, which creates a mapping into the latent space by training on the contrastive loss function. The output of the transformer network is then passed through a head network, consisting of two fully connected layers. In practice, the representations from the first fully connected layer perform the best in signal versus background classification tasks.

model-agnostic anomaly search in the latent space representations of the particle collision events. For this we use the Classification Without Labels (CWoLa) technique [12–14], which uses deep neural-network classifiers to distinguish between anomaly-enriched events and anomaly-depleted events. We will make use of the dijet resonance events in the LHC Olympics dataset [15].

The structure of this paper is as follows. In Sec. II, we motivate the relevance of contrastive learning to modeling particle collisions and introduce a dataset of dijet events. We further outline a set of *symmetry augmentations* for the contrastive loss function that leave the essential character of dijet events invariant, and how these are used in the contrastive learning approach. Lastly, in this section we outline the CWoLa anomaly detection method where we will use the self-supervised event representations. In Sec. III, we implement the contrastive learning method using a transformer neural network to map the dijet events from the event space into a latent space and evaluate the efficiency of the encoding. In Sec. IV, we use the CWoLa method to perform a relevant, but simplified anomaly bump-hunt analysis using the latent space representations for the dijet events. The paper ends with conclusions and outlook in Sec. V.

II. METHODS

The overarching goal is to optimize a mapping from the event space of particle collision events (i.e. the representation in the space of the individual particles) to a new latent space representation. The mapping between the event and latent space representations of particle collisions should maximally exploit the physical symmetries of particle collision events. In this way, the dimensionality of collision events can be reduced from a few thousand degrees of freedom (a few hundred particles per collision, each with a momentum 4-vector) to a few tens.

Such a mapping can be realised by using a *transformer-encoder* neural network architecture [16]. Event space events are fed into a transformer neural network where they are embedded into a reduced-dimension latent space. A distinguishing feature of the transformer architecture is permutation invariance: the latent space representations are invariant with respect to the order that the event constituents are fed into the network. The transformer network consists of 4 heads each with 2 layers. The output of the transformer encoder network is fed into a two layers of fully connected networks of the same latent space size. These architecture parameters were not heavily optimized.

See Fig. 1 for a schematic of the full network. We find, as do the authors of Ref. [16] and Ref. [10], that the latent space representation of the first head layer output gives a better representation than that of the final output layer, or that of the transformer encoder output.

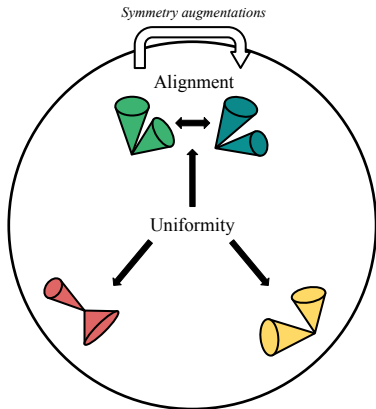


FIG. 2: An illustration of the latent space of the dijet events, built by a transformer network trained on the contrastive loss. The loss function optimizes for both alignment of dijets and their symmetry-augmented versions, and uniformity of physically distinct dijets.

A. Data Selection and Preparation

For this paper, we focus on the LHC 2020 Olympics R&D dataset [3, 15]. The full dataset consists of 1,000,000 background dijet events (Standard Model Quantum Chromodynamic (QCD) dijets) and 100,000 signal dijet events. The signal comes from the process $Z' \rightarrow X(\rightarrow q\bar{q})Y(\rightarrow q\bar{q})$, with three new resonances Z' (3.5 TeV), X (500 GeV), and Y (100 GeV). The only trigger is a single large-radius jet ($R = 1$) trigger with a p_T threshold of 1.2 TeV. The events are generated with Pythia 8.219 [17, 18] and Delphes 3.4.1 [19].

For each event, we cluster the jets using FastJet [20, 21] with a radius $R = 0.8$. We select the two highest-mass jets from each event, and we select the 50 hardest (highest p_T) constituents from each event, zero-padding for any jet with fewer than 50. Note that the average number of constituents per event is 81 with a standard deviation of 16.15. However, we found that including more than 50 constituents per jet did not lead to an appreciable improvement in the performance of the transformer network. The constituents are taken to be approximately massless, and so the relevant degrees of freedom for each constituent are (p_T, η, ϕ) .

For analysis, we select jets from the windows $p_T \in [800, 3000]$ GeV and $\eta \in [-3, 3]$. The cut on p_T was chosen such that invariant dijet mass m_{JJ} has a lower bound at approximately 2 TeV. This cut removes approximately 12% of eligible events from the LHCO dataset and has the benefit of removing a small tail of events with m_{JJ} below 2 TeV that could appear to be artificially anomalous to the transformer network.

B. Contrastive Learning

The contrastive learning method is self-supervised, meaning that it is trained using ‘pseudo-labels’ rather than truth labels. Supervised approaches use truth labels which exactly identify the truth label of the data. Pseudo-labels are artificial labels created from the data alone, without access to the truth labels. This means that the contrastive learning method is also unsupervised and receives no information as to whether the training samples are signal or background. Following JetCLR [10], the pseudo-labels are used to identify jets which are related to each other via some augmentation, for example a symmetry transformation. Using the pseudo-labels this technique aims to construct a latent space representation of events that exploits their physical symmetries.

As an example, consider the transformer-encoder network’s encoding of a dijet event \mathbf{r}_j , and the encoding of an *augmented* version of that event \mathbf{r}'_j . The exact physical symmetries considering in this analysis are outlined in Sec. II C, but for this example, let the augmentation be a rotation of the dijet event about the beam axis by some angle in azimuthal angle, ϕ . These two events represent the same underlying physics, as we expect physical events to be symmetric about the beam axis. Hence, \mathbf{r}_j and \mathbf{r}'_j are often called *positive pairs*. Therefore we would want the transformer-encoder network to map the event and its augmented version into similar regions of the latent space. In contrast, we would expect the transformer to map the jet event \mathbf{r}_j and a different jet event \mathbf{r}_k into different points of the latent space, since we do not expect a high degree of similarity between two arbitrary events. Therefore \mathbf{r}_j and \mathbf{r}_k are often called *negative pairs*. These positive and negative pairs are exactly the pseudo-labels for the contrastive learning method.

These requirements on the transformer-encoder mapping motivate the expression for the contrastive loss:

$$\mathcal{L}(\mathbf{r}_j, \mathbf{r}'_j, \mathbf{r}_k, \mathbf{r}'_k, \tau) = -\log \left(\frac{\exp(\text{sim}(\mathbf{r}_j, \mathbf{r}'_j))}{\sum_{j \neq k} [\exp(\text{sim}(\mathbf{r}_j, \mathbf{r}_k)) + \exp(\text{sim}(\mathbf{r}_j, \mathbf{r}'_k))]} \right). \quad (1)$$

We can interpret this loss function as follows: $\text{sim}(\mathbf{r}_j, \mathbf{r}'_j)$ calculates the similarity between two latent space representations, where

$$\text{sim}(\mathbf{r}_1, \mathbf{r}_2) = \frac{\mathbf{r}_1 \cdot \mathbf{r}_2}{\tau \|\mathbf{r}_1\| \|\mathbf{r}_2\|}. \quad (2)$$

The similarity is parameterized by a temperature τ , which balances the numerator and denominator in the contrastive loss. The numerator of the contrastive loss optimizes for *alignment*, which tries to map jets and their augmented versions to similar regions in the latent space. The denominator of the contrastive loss maximizes the *uniformity*, which tries to use up the entirety of the latent space when creating jet representations (see Fig. 2).

C. Event Augmentations

We now outline the list of symmetry augmentations used to create physically equivalent latent space jets.

We define the following single-jet augmentations:

1. *Rotation*: each jet is randomly (and independently) rotated about its central axis in the $\eta - \phi$ plane. This is not an exact symmetry, but correlations between the radiation patterns of the two jets is negligible.
2. *Distortion*: each jet constituent is randomly shifted in the $\eta - \phi$ plane. The shift is drawn from a gaussian of mean 0 and standard deviation $\sim 1/p_T$, where p_T is the transverse momentum of the constituent being shifted. This shift represents the smearing from detector effects.
3. *Collinear split*: a small number of the jet constituents (“mothers”) are split into two constituents (“daughters”) such that the daughters have η and ϕ equal to that of the mother, and the transverse momenta of the daughters sum to that of the mother.

We define the following event-wide augmentations:

1. η -*shift*: the dijet event is shifted in a random η direction.
2. ϕ -*shift*: the dijet event is shifted in a random ϕ direction.

Augmentations are applied to each training batch of the transformer. Each jet in the dijet event receives all three of the single-jet augmentations. The full event then receives both event-wide augmentations. See Fig. 3 for a visualization of the jet augmentations in the $\eta - \phi$ plane.

The jet augmentations are meant to not modify any of the important physical properties of the jets. As a test, we plot the jet masses of the hardest and second hardest jets from a subset of the LHC Olympics dataset in Fig. 4a, as well as the nsubjettiness variables τ_{21} and τ_{32} in Fig. 4b and Fig. 4c, both before and after receiving the jet augmentations and find no significant change in the distributions.

We also plot m_{JJ} for the dijet system in Fig. 5, again finding good agreement before and after the augmentations are applied. This confirms that our set of jet augmentations can be seen as true symmetry transformations of the dijet events.

D. Training procedure

We train the transformer network on a dataset of 50,000 background dijet events and up to 50,000 signal events. The batch size is set to 400, which is the largest possible given the computing resources available. The

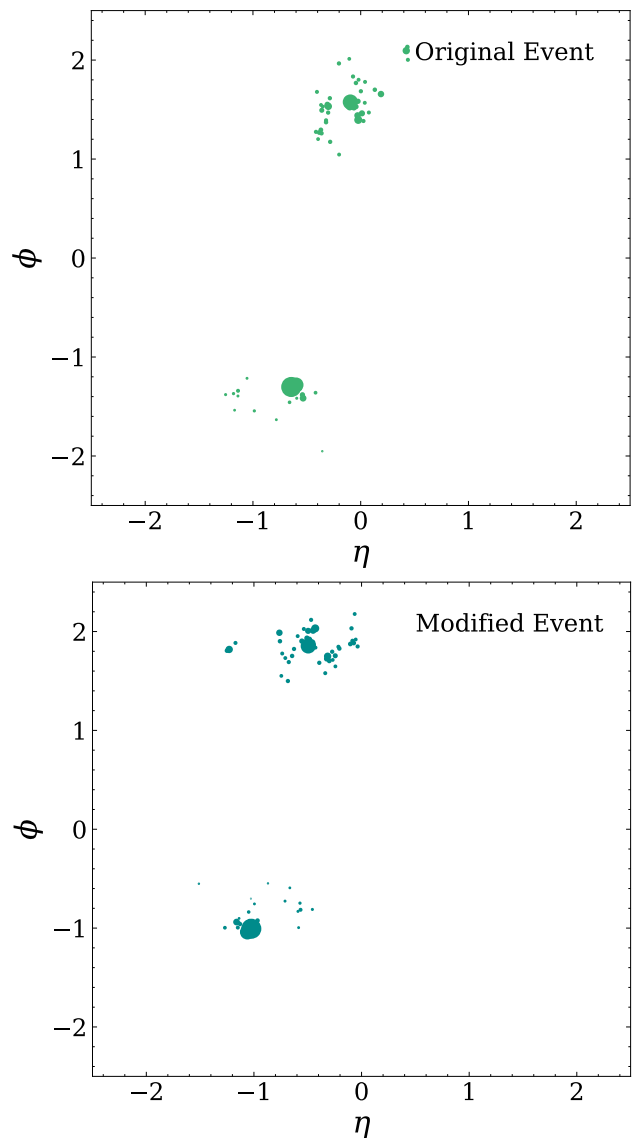


FIG. 3: A dijet event before (top) and after (bottom) receiving the set of single-jet and event-wide augmentations. Note that the upper and lower jets have visibly been rotated about their central axes, and the full event has been shifted in the upper-right direction of the $\eta - \phi$ plane.

network is trained with a learning rate of 0.0001, an early stopping parameter of 20 epochs, and a temperature parameter τ is set to 0.1. All of these hyperparameters were empirically found to deliver the best transformer performance. The transformer network is implemented using Pytorch 1.10.0 [22] and optimized with Adam [23]. Jet augmentations are applied batchwise, with each dijet event receiving a different randomized augmentation.

We also construct a binary classification dataset used to evaluate the latent space jet representations. This dataset consists of 85,000 signal and background dijet events each. We further define a “standard test” dataset

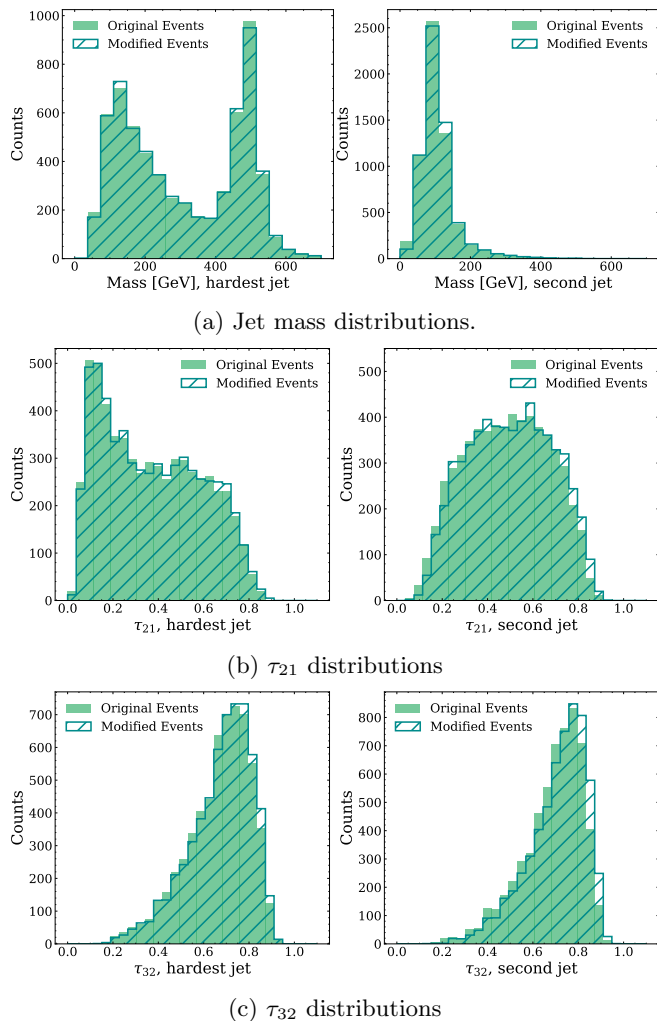


FIG. 4: Jet observable distributions for a sample of dijet events, before and after receiving the symmetry augmentations.

consisting of 10,000 signal and background events each. There is no overlap of events in the standard test dataset with those in the transformer training or binary classifier training sets.

E. Anomaly Detection

The usefulness of the latent space dijet representations is evaluated in a realistic model agnostic anomaly detection search setup. CWOLA (Classification Without Labels) [12] is a weakly supervised training method that allows for signal versus background discrimination in cases where training samples of pure signal and background cannot be provided. Such a scenario might occur in resonant bump-hunting, where it is common to define signal and sideband regions, both of which will have a non-negligible fraction of background events [13, 14]. The authors of Ref. [12] show that a classifier is trained on




FIG. 5: Dijet mass distributions for a sample of dijet events, before and after receiving the symmetry augmentations.

two mixed samples (each with a different signal fraction) is in fact maximally discriminatory for classifying signal from background.

In Sec. IV, we run a CWOLA training procedure on the latent space representations. Our mixtures consist of one background-only sample and one sample with a suppressed signal fraction representing a mixture of background and a rare unknown anomaly. This is the ideal anomaly detection setup in which a sample of pure background can be generated (also the starting point of Refs [24–26]). In practice, this may not be possible, and other methods must be used to obtain this dataset directly from data using sideband information [27–32].

III. EVALUATING THE LATENT SPACE REPRESENTATIONS

We evaluate the ability of our transformer to faithfully translate dijet events into a latent space in the following way: we first train the transformer encoder network to embed event space particle collisions into a reduced-dimension latent space. We then perform a binary classification task on the latent space. We test the sensitivity of this setup to the amount of signal present in the training of the transformer as well as in the training of the classifier. The latter test demonstrates the anomaly detection capability of the approach.

We consider two types of binary classification tasks: fully connected binary classifiers (FCN’s) are implemented in Pytorch and optimized with Adam. Linear classifier tests (LCT’s) are implemented in scikit-learn [33].

A. Exploring the dimensionality of the latent space

We first gauge how the size of the latent space affects the usefulness of representations. The latent space jet representations would ideally be lower in dimension




FIG. 6: Classifier efficiency curves for a fully connected binary classifier (FCN) run on the latent space dijet representations. The performance of the binary classifier increases with the dimension of the latent space. For comparison, we also provide efficiency curves for a FCN and a linear classifier test (LCT) run on the event space dijets.

than the physical space versions so as to save computational processing time by removing nonessential degrees of freedom from a given dataset. However, a latent space embedding with too few dimensions might not contain enough parameters to encode the essential physical symmetries of the jets.

In Fig. 6, we plot the rejection ($1 / \text{false positive rate}$) versus the true positive rate for a FCN trained on the latent space dijet representations. (See Fig. 9 in App. A for curves from a linear classifier.) We scan the latent space size in powers of two from 512 down to 8 dimensions. This classifier has access to signal and background labels (fully supervised).

In general, the performance of the transformer encoder (as quantified by the receiver operators characteristic (ROC) area under the curve (AUC)) improves as the dimension of the latent space increases. We find that a fully supervised classifier trained on latent space jet representations cannot outperform a classifier trained on event space jet representations, but it can outperform a linear classifier trained on event space jet representations. Perhaps more striking is that the linear classifier trained on the compressed latent representations outperforms the linear classifier trained on the full event space data. This indicates that the self-supervised representations are highly expressive despite being compressed, and agrees with the top-tagging results obtained in Ref. [10].

A selection of AUC scores for the FCN’s trained on latent space representations are given in Table I. We additionally include as a performance metric the maximum of the significance improvement characteristic ($\text{max}(\text{SIC})$),

| Training set | Training dim. | Classifier | AUC | $\text{max}(\text{SIC})$ |
|----------------|---------------|------------|-------|--------------------------|
| Particle space | 2100 | FCN | 0.958 | 15.401 |
| | | LCT | 0.883 | 2.277 |
| Latent space | 8 | FCN | 0.904 | 2.542 |
| | | LCT | 0.841 | 1.882 |
| | | Trans+BC | 0.955 | 7.608 |
| | 64 | FCN | 0.915 | 3.163 |
| | | LCT | 0.816 | 1.799 |
| | | Trans+BC | 0.960 | 7.238 |
| | 512 | FCN | 0.945 | 6.396 |
| | | LCT | 0.926 | 4.624 |
| | | Trans+BC | 0.968 | 13.862 |

TABLE I: ROC AUC and $\text{max}(\text{SIC})$ scores for a binary classifier trained to discriminate signal from background on the LHC Olympics dataset. FCN = Fully Connected (Dense) Neural Network; LCT = Linear Classifier Test; Trans+BC = transformer architecture trained on the Binary Cross Entropy loss.

defined as $\text{max}(\frac{\text{true positive rate}}{\sqrt{\text{false positive rate}}})$. The $\text{max}(\text{SIC})$ can be seen as the multiplicative factor by which signal significance improves after performing a well-motivated cut on the dataset.

Table I also contains scores for a binary LCT trained on the latent space representations, as well as a binary classifier constructed from the transformer architecture trained on the Binary Cross Entropy loss with an additional sigmoid function as the final layer (Trans+BC). The Trans+BC network has access to all of the input particles and is not trained post-hoc on the self-supervised latent space, so we expect it to perform the best of all configurations. The hope is that the FCN performance is as close as possible to the performance of the Trans+BC (on the largest latent space).

Note that Table I does show a performance gap between a FCN trained on the latent space and the Trans+BC network. In App. B, we evaluate the performance of both such networks when trained on increasing amounts of data. In both cases, we find that the networks are “data-hungry”; in other words, the classifier performances increase with the amount of training data, and the performances do *not* saturate when trained on the 85,000 signal and 85,000 background dijet events sampled from the LHCO dataset. Therefore the performance of the FCN could likely reach that of the Trans+BC network with a larger training dataset than what was used in this study.

B. Varying the amount of training signal

In practice, we might want to use the transformer network for model-agnostic anomaly detection. In this case, we would not be able to train the transformer on a known signal fraction, as the training data would contain an unknown (and extremely tiny, if any) percentage of BSM signal. It is therefore useful to see if the transformer en-

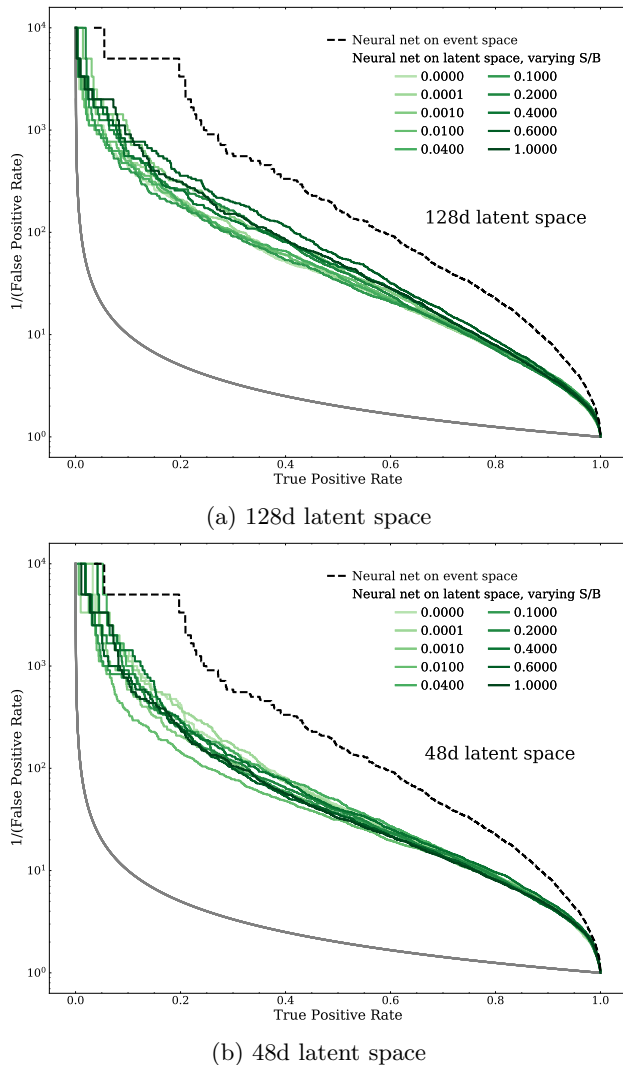


FIG. 7: Performance curves for a latent space FCN classifier trained on varying amounts of signal fraction. The classifier performance is robust with respect to the signal fraction of the training data. This implies that the transformer network can be trained on background alone.

coder network is effective at translating rare events into a latent space. We might expect this to be true if the transformer is learning only generic features about collider events that are true for both signal and background events. This is encouraged by the universality of the symmetry augmentations in the contrastive loss.

In Fig. 7a, we hold the dimension of the transformer latent space fixed at 128, then scan the signal to background ratio S/B down from 1.0 to 0.0. In Fig. 7b, we repeat the previous steps for a transformer latent space of dimension 48. (See Fig. 10 in App. A for curves from a linear classifier.) We find that the classifier performance is robust with respect to the signal to background ratio (as was found in Ref. [10]). This demonstrates that the

transformer network can be trained on background alone and still faithfully model rare signal events.

IV. ANOMALY DETECTION


We now test the usefulness of the latent space jet representations in a more practical setting by performing a CWOLA-style anomaly search. As before, we use a sample of 85,000 signal and 85,000 background events to train the classifier, and we use the “standard test” dataset of 10,000 signal and background events to test.

We first create a baseline against which to compare the CWOLA analyses by carrying out a self-supervised binary classification task on a sample of dijet events in the event space representation.


For the anomaly-detection analysis, we set one CWOLA “mixed sample” to be a set of background-only dijets. The other mixed sample is a mixture of signal and background dijets, with the signal fraction scanned from 0% to 100%. We run the analysis three times, once for the event space dijets and twice for latent space dijets at 128 and 48 dimensions. The evaluation of the performance is always computed with pure signal and background labels. Comparisons of the CWOLA classifier performances are shown in Fig. 8. In Fig. 8a, we use the ROC AUC as a metric for evaluating the CWOLA classifier; in Fig. 8b, we use instead use the $\max(\text{SIC})$ as the metric; in Fig. 8c, we provide the false positive rate at a fixed true positive rate of 50%.

We find that the CWOLA weakly supervised classifier performance of a small dimensional latent space is comparable to (but cannot match) that of the full particle event space, with an improvement in performance for a larger dimension latent space as evaluated by the $\max(\text{SIC})$ metric. The most notable difference between the classifier performance on latent space vs. event space is that in the former case, the classifier performance diminishes to no better than random at a higher signal fraction for the training data. (This is indicated by the ROC AUC dropping to 0.5, the $\max(\text{SIC})$ dropping to 1, and the $\text{FPR} @ \text{TPR} = 0.5$ dropping to 0.5 in Fig. 8.) More specifically, the small-dimension latent space classifier hits random performance at a signal fraction of just below 10^{-2} , while the event space classifier does better than random at all nonzero signal fractions.


Overall, the classifier performances at anomaly-level signal fractions as shown in Fig. 8 are lower than what has been seen in other recent anomaly-detection methods on the LHCO data. In fact, the SIC curves shown in Refs. [27–32] are typically an order of magnitude greater than those in Fig. 8. However, such curves were constructed by training on standard jet observables (e.g. m_J and τ_{12}), and thus represent training methods that are inherently model-dependent. Evidently, this performance decrease from model-dependent anomaly detection methods is the price to pay on this particular signal for using a more widely-applicable, model-agnostic method.



(a) ROC AUC training metric.



(b) max(SIC) training metric.



(c) False positive rate at a true positive rate of 50%.

FIG. 8: Various metrics for evaluating a CWoLa weakly-supervised classifier trained to discriminate a background-only sample from a signal+background sample with a variable signal fraction. The classifier run on the event space representations slightly outperforms that run on the latent space representations, especially at low signal fractions.

There exist a number of avenues for future work to improve on this contrastive-learning trained classifier. For one: we have considered a small set of symmetry augmentations specific to dijet events. However, additional augmentations for dijet events could be added to the contrastive loss. Alternatively, a different selection of augmentations that lead to an even more general event representation could be chosen. As another avenue: we mentioned earlier (and illustrate in App. B) that the transformer network is data-hungry. It would therefore be reasonable to expect an improvement in classifier performance if the training dataset were made larger.

V. CONCLUSIONS

In this paper, we have used transformer neural networks to embed entire collider events into lower- and fixed-dimensional latent spaces. This embedding is constructed using self-supervised learning based on symmetry transformations. Events that are related by symmetry transformations are grouped together in the latent space while other pairs of events are spread out in the latent space.

We have shown that the latent space preserves the salient properties of the events for distinguishing certain BSM events from the SM background. This latent space can then be used for a variety of tasks, including anomaly detection. We have shown that anomalies can still be identified in the reduced representation as long as there is enough signal in the dataset. For the particular signal model studied, the required amount of signal is much higher than reported by other studies using high-level features. This illustrates the tradeoff between signal sensitivity and model specificity. Our reduced latent space knows nothing of particular BSM models and is thus broadly useful but not particularly sensitive. Future work that explores the continuum of approaches by adding more augmentations to the contrastive learning may result in superior performance for particular models in the future.

CODE AVAILABILITY

The code can be found at https://github.com/rmstand/JetCLR_AD.

ACKNOWLEDGMENTS

BMD was supported by a Postdoctoral Research Fellowship from Alexander von Humboldt Foundation. B.N. and R.M. were supported by the Department of Energy, Office of Science under contract number DE-AC02-05CH11231. This material is based upon work supported by the National Science Foundation Graduate Research Fellowship Program under Grant No. DGE 2146752.

Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

Appendix A: Evaluating the latent space with a linear classifier test

In this section, we provide the analogues to Fig. 6 and Fig. 7 with the transformer efficiency curves calculated for a binary linear classifier test run on the latent space representations (rather than a binary neural net). These plots (Fig. 9 and Fig. 10) are not shown in the main text of this report as a realistic anomaly detection analysis would be carried out using neural nets. However, the field-standard way to evaluate representations of jets is through a linear classifier.




FIG. 9: Classifier efficiency curves for a linear classifier test (LCT) run on the latent space dijet representations. For comparison, we also provide efficiency curves for a FCN and a LCT run on the event space dijets.

Appendix B: How data-hungry are the neural networks?

In this section, we provide some plots illustrating the data-hungry nature of the transformer network. The performance of the binary classifier trained on the latent space dijet representations, as shown in Fig. 8 is admittedly low. However, is it likely that the performance could improve if the classifiers were trained on a larger amount of data.

In Fig. 11b, we train a Trans+BC network on a varying fraction of the available dijet dataset (a 100% training fraction makes use of all 85,000 signal and 85,000 background events). In Fig. 11a, we repeat this procedure

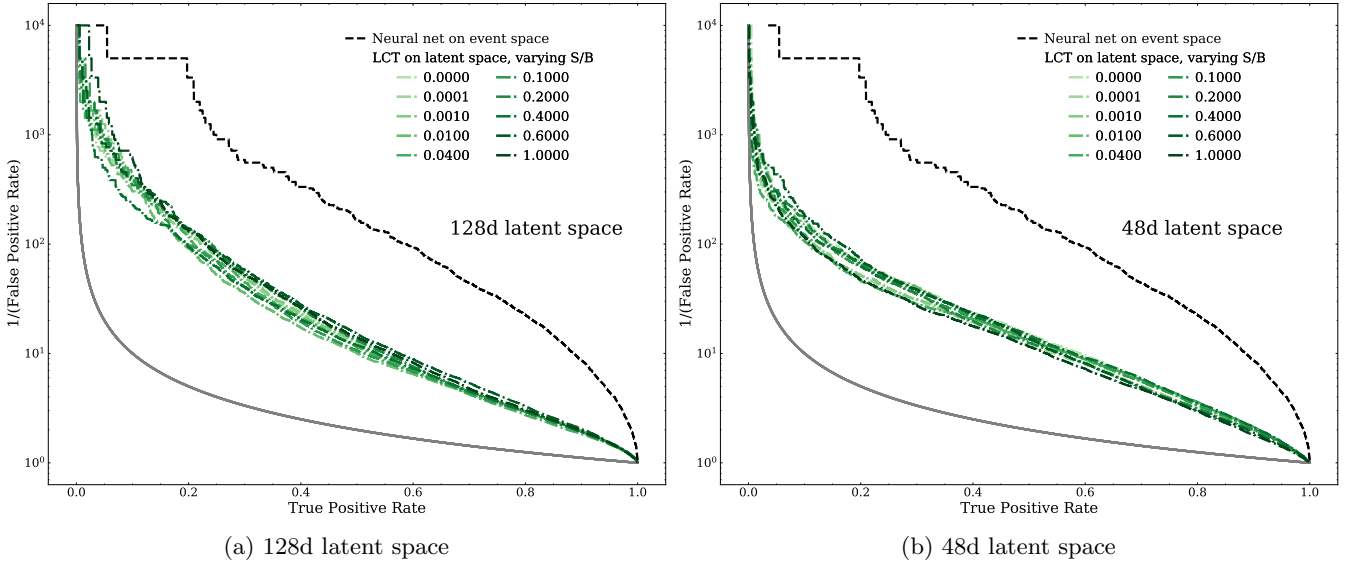


FIG. 10: Performance curves for a latent space linear classifier test trained on varying amounts of signal fraction. The classifier performance is again robust with respect to the signal fraction of the training data.

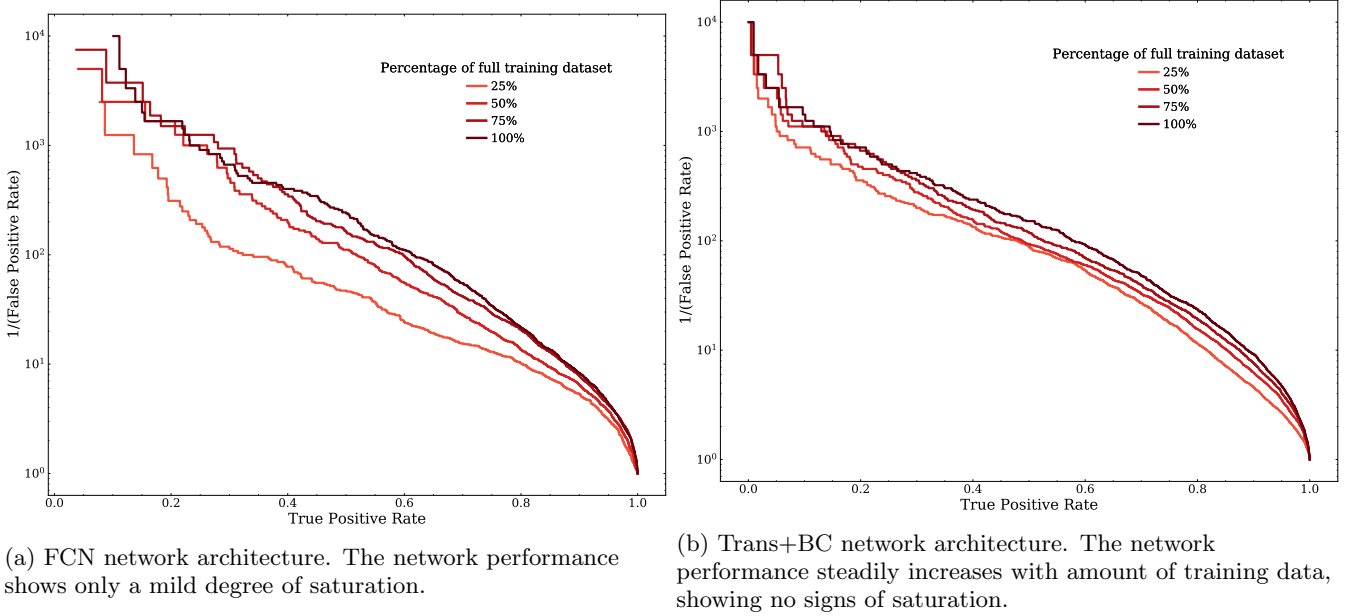


FIG. 11: Performance curves for a binary classifier trained on an increasing percentage of the full dijet dataset.

for a FCN. In both cases, the ROC AUCs of the trained binary classifiers do not appear to be saturated when trained on the full dataset.

- [1] Nathaniel Craig, Patrick Draper, Kyoungchul Kong, Yvonne Ng, and Daniel Whiteson, “The unexplored landscape of two-body resonances,” *Acta Phys. Polon. B* **50**, 837 (2019), [arXiv:1610.09392 \[hep-ph\]](#).
- [2] Jeong Han Kim, Kyoungchul Kong, Benjamin Nachman, and Daniel Whiteson, “The motivation and status of two-body resonance decays after the LHC Run 2 and beyond,” *JHEP* **04**, 030 (2020), [arXiv:1907.06659 \[hep-ph\]](#).
- [3] Gregor Kasieczka *et al.*, “The LHC Olympics 2020 a community challenge for anomaly detection in high energy physics,” *Rept. Prog. Phys.* **84**, 124201 (2021), [arXiv:2101.08320 \[hep-ph\]](#).
- [4] T. Aarrestad *et al.*, “The Dark Machines Anomaly Score Challenge: Benchmark Data and Model Independent Event Classification for the Large Hadron Collider,” (2021), [arXiv:2105.14027 \[hep-ph\]](#).
- [5] Georgia Karagiorgi, Gregor Kasieczka, Scott Kravitz, Benjamin Nachman, and David Shih, “Machine Learning in the Search for New Fundamental Physics,” (2021), [arXiv:2112.03769 \[hep-ph\]](#).
- [6] Matthew Feickert and Benjamin Nachman, “A Living Review of Machine Learning for Particle Physics,” (2021), [arXiv:2102.02770 \[hep-ph\]](#).
- [7] ATLAS Collaboration, “Dijet resonance search with weak supervision using 13 TeV pp collisions in the ATLAS detector,” (2020), [10.1103/PhysRevLett.125.131801](#), [arXiv:2005.02983 \[hep-ex\]](#).
- [8] Blaž Bortolato, Barry M. Dillon, Jernej F. Kamenik, and Aleks Smolkovič, “Bump Hunting in Latent Space,” (2021), [arXiv:2103.06595 \[hep-ph\]](#).
- [9] Md Abul Hayat, George Stein, Peter Harrington, Zarija Lukić, and Mustafa Mustafa, “Self-supervised representation learning for astronomical images,” (2020).
- [10] Barry M. Dillon, Gregor Kasieczka, Hans Olschlager, Tilman Plehn, Peter Sorrenson, and Lorenz Vogel, “Symmetries, Safety, and Self-Supervision,” (2021), [arXiv:2108.04253 \[hep-ph\]](#).
- [11] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun, “Deep residual learning for image recognition,” (2015).
- [12] Eric M. Metodiev, Benjamin Nachman, and Jesse Thaler, “Classification without labels: Learning from mixed samples in high energy physics,” *JHEP* **10**, 174 (2017), [arXiv:1708.02949 \[hep-ph\]](#).
- [13] Jack H. Collins, Kiel Howe, and Benjamin Nachman, “Anomaly Detection for Resonant New Physics with Machine Learning,” *Phys. Rev. Lett.* **121**, 241803 (2018), [arXiv:1805.02664 \[hep-ph\]](#).
- [14] Jack H. Collins, Kiel Howe, and Benjamin Nachman, “Extending the search for new resonances with machine learning,” *Phys. Rev. D* **99**, 014038 (2019), [arXiv:1902.02634 \[hep-ph\]](#).
- [15] Gregor Kasieczka, Benjamin Nachman, and David Shih, “Official Datasets for LHC Olympics 2020 Anomaly Detection Challenge (Version v6) [Data set].” (2019), <https://doi.org/10.5281/zenodo.4536624>.
- [16] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton, “A simple framework for contrastive learning of visual representations,” (2020), [arXiv:2002.05709 \[cs.LG\]](#).
- [17] Torbjorn Sjostrand, Stephen Mrenna, and Peter Z. Skands, “PYTHIA 6.4 Physics and Manual,” *JHEP* **05**, 026 (2006), [arXiv:hep-ph/0603175](#).
- [18] Torbjörn Sjöstrand, Stefan Ask, Jesper R. Christiansen, Richard Corke, Nishita Desai, Philip Ilten, Stephen Mrenna, Stefan Prestel, Christine O. Rassmusen, and Peter Z. Skands, “An introduction to PYTHIA 8.2,” *Comput. Phys. Commun.* **191**, 159–177 (2015), [arXiv:1410.3012 \[hep-ph\]](#).
- [19] J. de Favereau, C. Delaere, P. Demin, A. Giammanco, V. Lemaitre, A. Mertens, and M. Selvaggi (DELPHES 3), “DELPHES 3, A modular framework for fast simulation of a generic collider experiment,” *JHEP* **02**, 057 (2014), [arXiv:1307.6346 \[hep-ex\]](#).
- [20] Matteo Cacciari, Gavin P. Salam, and Gregory Soyez, “FastJet User Manual,” *Eur. Phys. J. C* **72**, 1896 (2012), [arXiv:1111.6097 \[hep-ph\]](#).
- [21] Matteo Cacciari and Gavin P. Salam, “Dispelling the N^3 myth for the k_t jet-finder,” *Phys. Lett. B* **641**, 57–61 (2006), [arXiv:hep-ph/0512210](#).
- [22] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala, “Pytorch: An imperative style, high-performance deep learning library,” in *Advances in Neural Information Processing Systems 32* (Curran Associates, Inc., 2019) pp. 8024–8035.
- [23] Diederik P. Kingma and Jimmy Ba, “Adam: A method for stochastic optimization,” (2014).
- [24] Raffaele Tito D’Agnolo and Andrea Wulzer, “Learning New Physics from a Machine,” *Phys. Rev. D* **99**, 015014 (2019), [arXiv:1806.02350 \[hep-ph\]](#).
- [25] Raffaele Tito D’Agnolo, Gaia Grosso, Maurizio Pierini, Andrea Wulzer, and Marco Zanetti, “Learning Multivariate New Physics,” (2019), [10.1140/epjc/s10052-021-08853-y](#), [arXiv:1912.12155 \[hep-ph\]](#).
- [26] Raffaele Tito d’Agnolo, Gaia Grosso, Maurizio Pierini, Andrea Wulzer, and Marco Zanetti, “Learning New Physics from an Imperfect Machine,” (2021), [arXiv:2111.13633 \[hep-ph\]](#).
- [27] Benjamin Nachman and David Shih, “Anomaly Detection with Density Estimation,” *Phys. Rev. D* **101**, 075042 (2020), [arXiv:2001.04990 \[hep-ph\]](#).
- [28] Anders Andreassen, Benjamin Nachman, and David Shih, “Simulation Assisted Likelihood-free Anomaly Detection,” *Phys. Rev. D* **101**, 095004 (2020), [arXiv:2001.05001 \[hep-ph\]](#).
- [29] George Stein, Uros Seljak, and Biwei Dai, “Unsupervised in-distribution anomaly detection of new physics through conditional density estimation,” in *34th Conference on Neural Information Processing Systems* (2020) [arXiv:2012.11638 \[cs.LG\]](#).
- [30] Kees Benkendorfer, Luc Le Pottier, and Benjamin Nachman, “Simulation-Assisted Decorrelation for Resonant Anomaly Detection,” (2020), [arXiv:2009.02205 \[hep-ph\]](#).
- [31] Anna Hallin, Joshua Isaacson, Gregor Kasieczka, Claudius Krause, Benjamin Nachman, Tobias Quadfasel, Matthias Schlaffer, David Shih, and Manuel Sommerhalder, “Classifying Anomalies Through Outer Density

- Estimation (CATHODE),” (2021), [arXiv:2109.00546 \[hep-ph\]](#).
- [32] John Andrew Raine, Samuel Klein, Debajyoti Sengupta, and Tobias Golling, “CURTAINS for your Sliding Window: Constructing Unobserved Regions by Transforming Adjacent Intervals,” (2022), [arXiv:2203.09470 \[hep-ph\]](#).
- [33] Fabian Pedregosa *et al.*, “Scikit-learn: Machine learning in python,” (2012), [10.48550/arXiv.1201.0490](#).