RECEIVED: December 10, 2021

REVISED: January 19, 2022

ACCEPTED: January 21, 2022

PUBLISHED: February 10, 2022

Autoencoders for semivisible jet detection

**Florencia Canelli,^a Annapaola de Cosa,^b Luc Le Pottier,^c Jeremi Niedziela,^b
Kevin Pedro^d and Maurizio Pierini^e**

^a*Department of Physics, University of Zurich,
Zurich, Switzerland*

^b*ETH Zurich, Institute for Particle Physics and Astrophysics,
Zurich, Switzerland*

^c*University of California,
Berkeley, U.S.A.*

^d*Fermi National Accelerator Laboratory,
Batavia, IL 60510, U.S.A.*

^e*Experimental Physics Department, European Organization for Nuclear Research (CERN),
Meyrin, Switzerland*

E-mail: florencia.canelli@physik.uzh.ch, adecosa@phys.ethz.ch,
luclepot@berkeley.edu, jeremi.niedziela@cern.ch, pedrok@fnal.gov,
maurizio.pierini@cern.ch

ABSTRACT: The production of dark matter particles from confining dark sectors may lead to many novel experimental signatures. Depending on the details of the theory, dark quark production in proton-proton collisions could result in semivisible jets of particles: collimated sprays of dark hadrons of which only some are detectable by particle collider experiments. The experimental signature is characterised by the presence of reconstructed missing momentum collinear with the visible components of the jets. This complex topology is sensitive to detector inefficiencies and mis-reconstruction that generate artificial missing momentum. With this work, we propose a signal-agnostic strategy to reject ordinary jets and identify semivisible jets via anomaly detection techniques. A deep neural autoencoder network with jet substructure variables as input proves highly useful for analyzing anomalous jets. The study focuses on the semivisible jet signature; however, the technique can apply to any new physics model that predicts signatures with anomalous jets from non-SM particles.

KEYWORDS: Jet substructure, Beyond Standard Model, Dark Matter, Hadron-Hadron Scattering, Jets

ARXIV EPRINT: [2112.02864](https://arxiv.org/abs/2112.02864)

Contents

1	Introduction	1
2	Data samples	2
2.1	Generation	2
2.2	Preselection	3
2.3	Feature selection	3
3	Models	5
3.1	Autoencoders	5
3.2	Alternative models	6
3.3	Model optimization	6
4	Results	7
4.1	Reconstructed distributions	7
4.2	Reconstruction loss	7
4.3	Discrimination power	7
4.4	Robustness	9
4.5	Sensitivity to semivisible jet models	10
5	Conclusion	11
A	m_{dark} variations	12

1 Introduction

Hidden Valley models of dark matter allow for a strong interaction in the dark sector [1]. Such a strong interaction would cause showering of the dark particles, producing both invisible dark matter and Standard Model (SM) hadrons. As a result, so-called semivisible (SV) jets would be formed, containing some fraction r_{inv} of invisible particles [2, 3]. Depending on the mass of the mediator $m_{Z'}$, SV jets could be produced at the Large Hadron Collider (LHC) energies and registered in detectors [4].

Events containing such jets would be characterised by missing transverse energy $\cancel{E}_T$ being aligned with a jet, which is a scenario difficult to distinguish from detector issues and therefore usually neglected in LHC data analyses. Moreover, the details of the Hidden Valley physics depend on a number of parameters. There are four parameters that most influence the kinematic behavior in the final state:

- α_{dark} — the coupling constant of the dark sector strong interaction equivalent,
- $m_{Z'}$ — the mass of the mediator,
- m_{dark} — the mass of the dark hadrons,
- r_{inv} — the fraction of stable, invisible dark hadrons.

These parameters could be a function of other parameters in a fully-specified Hidden Valley model, as well as parameters that impact the modeling of strong dynamics in the event generator (hadronization, fragmentation). This leads to an extremely large model space with a huge number of possible scenarios that can easily evade any constraints from e.g. cosmological measurements. It is clearly impractical to perform dedicated searches for all possible model variations.

In order to gain robustness against detector effects, as well as details of the model implementation, anomaly detection techniques such as unsupervised deep learning can be used. Those are designed to detect objects significantly different from the training sample, without prior knowledge of signal characteristics. The presented approach is suitable to search for a wide class of new physics models containing anomalous jets, including the SV jets stemming from Hidden Valley models, which are the main focus of this work.

Here, an autoencoder-based analysis for SV jets detection will be presented. Unlike previous attempts [5–10] that use jet images (four-momenta of jet constituents), we incorporate high-level jet features and substructure variables computed based on the four-momenta of the jet constituents: Energy Flow Polynomials (EFPs) [11], Energy Correlation Functions (ECFs) [12] and their ratios: C_2 and D_2 , as well as the jet p_T dispersion p_TD [13] and jet axes [14]. These prove to be highly useful in discriminating between the QCD background and the SV jet signal [15].

The trained model can then be used to tag anomalous jets, as the autoencoder has been trained to recognize QCD jets in the signal region, but not SV signal jets (or any other anomalous jets, such as emerging jets [16]).

It is worth noting that, while the performance is assessed using SV jet models as a benchmark, the training process is independent of the signal implementation, unlike in case of the supervised learning [17]. The aforementioned choice of substructure variables may have an impact on sensitivity to different anomalous jet models. However, selected features capture various properties of jet substructure, such as number of sub-jets, jet width or distribution of momenta among constituents, therefore making this approach suitable for a wide range of new physics models. Given that the landscape of possible Hidden Valley scenarios is so vast, using one classifier (even parametric [18]) per signal model becomes impractical.

2 Data samples

2.1 Generation

Both multi-jet QCD and dark sector samples are generated using PYTHIA8 [19] with model parameters identical as in [20] and reconstructed with DELPHES [21] providing a detector response approximating that of CMS. Proton-proton collisions at a centre-of-mass energy of 13 TeV are considered, with approximately 50 collisions per bunch crossing. The particle flow reconstruction algorithm distributed with DELPHES is used. In particular, jets are clustered from reconstructed particles and pileup reduction is applied.

To examine the robustness of autoencoder models to varying signal parameters, multiple SV jet samples were generated with values of the fraction of stable jet hadrons $r_{\text{inv}} \in \{0.3, 0.5, 0.7\}$ and the Z' boson mass $m_{Z'} \in \{1.5, 2.0, 2.5, 3.0, 3.5, 4.0\}$. The number

of events generated varies depending on the sample, accounting for the different selection efficiency (see section 2.2):

- QCD: $3.7 \cdot 10^7$ events,
- SV jets, $m_{Z'} > 1.5$ TeV: 50k events per sample,
- SV jets, $m_{Z'} = 1.5$ TeV, $r_{\text{inv}} \in \{0.3, 0.5\}$: 100k events per sample,
- SV jets, $m_{Z'} = 1.5$ TeV, $r_{\text{inv}} = 0.7$: 140k events per sample.

2.2 Preselection

A set of preselection criteria (based on those outlined in ref. [2]) is applied to both training and testing samples on an event-by-event basis. These requirements are aimed at isolating the SV jet signal, as well as reproducing the realistic triggering capabilities of LHC experiments. The signal region is defined to include events passing the following criteria:

- at least 2 jets with $|\eta| < 2.4$ and $p_T > 200$ GeV
- then, for the two leading jets:
 - $|\Delta\eta| < 1.5$,
 - $m_T > 1500$ GeV,
 - $\cancel{E}_T/m_T > 0.25$,

where η is the pseudorapidity, p_T is the transverse momentum for each jet, $\Delta\eta$ is the absolute dijet η difference, $\cancel{E}_T$ is the missing momentum in the transverse plane for the event, and m_T is the transverse mass of the leading dijet system and the $\cancel{E}_T$.

The selection efficiency was found to be at the level of 0.13% for the QCD events and between 0.2 and 15.3% for SV jets, depending on the signal model parameters. The number of QCD jets after applying selections is ≈ 100 k (split between training, validation, and testing) and varies between 500 and 15k for SV jets, depending on the model parameters.

2.3 Feature selection

Since the goal of this study is to find anomalies on the basis of tagging individual anomalous jets rather than anomalous events, a set of jet-level and jet substructure variables is determined for the training.

The coordinates of each jet (η and ϕ) are included in the training to allow the network to learn about problematic regions of the detector and avoid tagging noise or other detector failures as anomalous signals.

In order to avoid bias, the transverse momentum of the jet was not included in the training. Moreover, the jet p_T distribution was flattened (by applying appropriate weights) for the training.

A number of other jet-level and substructure variables have been considered for the training. A fraction of them were rejected due to poor signal-background discrimination.

Figure 1. Distributions of input variables for QCD background and selected signal models.

The correlations between remaining variables were tested, as autoencoders tend to perform better with uncorrelated inputs. It was found that for QCD and SV jets, the EFPs are fully correlated with each other, as well as with the girth [22], Les-Houches Angularity (LHA) [23], and normalized ECFs: e_2 and e_3 . Because of that, it was decided to only keep one of those variables for the training, namely EFP_1 [11] which exploits pairwise relations between jet constituents.

The jet constituent four-momenta have also been considered; however, they proved not to bring any significant gain in performance, while increasing the complexity (≈ 100 instead of ≈ 10 input features) and training time of the autoencoder. This can be understood since EFPs and ECFs are included in the training and provide a linear basis for characterizing jet constituent substructure, making up for the discrimination that is lost when using high-level jet features rather than jet constituents.

Ultimately, the following set of jet-related input variables is selected: η , ϕ , invariant mass m_j , jet fragmentation function $p_T D$, jet ellipse minor and major axes, EFP_1 , and ECF ratios: C_2 and D_2 . The distributions of those variables for the QCD background and selected signals is presented in figure 1.

The first three of these features are provided directly by the DELPHES reconstruction algorithm. The remaining ones are calculated using a list of jet constituents, which consists of all particles in the event which lie within a cone of radius $R = 0.8$ from the jet axis.

All of the selected high level features discussed here provide some meaningful discrimination power, with the exception of the angular coordinates, η and ϕ . Since those are

randomly distributed, they provide no meaningful information about QCD jet structure. However, given the ability of autoencoders to learn a compressed representation of input signals, it is feasible that they might also learn the locations and behaviors of defective detector components. For instance, they could learn location of energy spikes in certain detector locations as normal behavior that does not indicate an anomaly associated with a potential signal. Such detector malfunctions are not included in the **DELPHES** reconstruction; however, they may be present in the collision data, on which the autoencoder can be trained.

3 Models

All unsupervised models presented in the following sections are trained using the QCD data sample described in section 2 as background, with a 70:15:15 split between training, validation, and testing, respectively.

3.1 Autoencoders

An autoencoder is a neural network which approximates the identity transformation as the application of a function f from the input space to a latent space (encoding) and a function f^{-1} from the latent space back to the input space (decoding), which provides an approximation of the input as the output. The constraint forces the autoencoder to find hidden relationships between features in the input data set. In the case of a neural autoencoder, this constraint takes the form of a bottleneck in the network architecture: an input vector $\vec{x} \in \mathbb{R}^n$ is compressed into a latent space vector $\vec{l} \in \mathbb{R}^m$ with $m < n$, enforced by an m -node layer within the architecture of an autoencoder with input dimension n . This latent space vector is then reconstructed by a decoding network in an attempt to match the input values.

The autoencoder is trained to minimize the reconstruction error of each sample, defined for an n -dimensional input feature vector $\vec{x}$ with components $x_1, \dots, x_n$. A number of loss functions were considered and the mean absolute error (MAE) was selected, as it had the best performance on the considered benchmark data sets:

$$\text{MAE} = \frac{1}{N} \sum_i^N |x'_i - x_i| \quad (3.1)$$

where x_i is the original value and x'_i is the reconstructed value for the i^{th} feature of a jet. The default choice of the loss function for autoencoders is often the mean squared error, which in our case indicated around 7% decrease in performance compared to MAE. This can be understood as MAE tends to be more robust against outliers and most input variables have non-gaussian distributions with long tails (see figure 1).

The training data consists of the simulated QCD background events that fall within our signal region as per the cuts described in section 2.

Once a model is trained to recognize the background, one can evaluate it on unlabeled data and use the reconstruction errors as a measure of similarity of a given sample to the training data set. Typically, a model trained on QCD will return low reconstruction loss for samples similar to those it was trained on, and high loss for samples that are more complex [24, 25]. In case of a search for new physics, an autoencoder would be trained and evaluated, with the largest 10%, 1%, and 0.1% of its reconstruction losses isolated and analyzed.

Figure 2. Neural Network architecture.

3.2 Alternative models

Along with neural autoencoders — the simplest form of autoencoder — other types of autoencoders were considered:

- Variational Autoencoders [26] (VAE): provide probabilistic representation of the input data, which can give higher performance in some applications.
- Principal Component Analysis (PCA) [27]: can be viewed as a linear autoencoder with no hidden layers, which allows determining the weights analytically.

The results obtained with the VAE and PCA will be reported in this work. Sparse Autoencoders (SAE) [28] were investigated in the early stages of this study and found to give worse results than the dense autoencoder; therefore, they will not be included.

3.3 Model optimization

In order to determine the optimal architecture, normalization, and hyper-parameters for the SV jet search, a large number of models were trained and evaluated. The optimization was focused on maximizing the Area Under Receiver Operating Characteristic (ROC) Curve (AUC). For each set of hyper-parameters and each choice of architecture, 80% of best models were selected and the average AUC over all signal samples and all models was used to choose the highest performing one.

It was found that the best results are obtained by standardizing the data, i.e., by normalizing them such that the mean of the distribution is at zero and the standard deviation is unity. An extensive architecture scan was performed and the architecture shown in figure 2, with two hidden layers of size 10, followed by the bottleneck of size 6, was found to be optimal. The activation function of the hidden layers and the bottleneck is ‘elu’ [29] and the output layer uses linear activation.

The autoencoder models were implemented in Keras/TensorFlow [30, 31]. Each model was trained for 200 epochs (with early stopping enabled), with a batch size of 256. The learning rate was set to 10^{-6} and the optimizer found to yield the best performance was ‘Nadam’ [32].

4 Results

In this section, the results from an average autoencoder will be presented (unless stated otherwise). The average autoencoder was found by training 150 models with the optimal hyper-parameters and varying random seed. We select a model closest to the mean μ of AUC distribution (averaged out over all signal scenarios for given autoencoder model) and use models closest to $\mu \pm \sigma$, where σ is the standard deviation, to estimate the uncertainty. No signs of overfitting were observed, and the loss evolution was smooth and saturated after around 100 epochs.

4.1 Reconstructed distributions

The differences between input and reconstructed jet features are shown for both SVJ and QCD data in figure 3. One can observe that, in almost all cases, the difference is larger for the signals than for the background. What is more, the distributions for the signal do not change with varying mass; however there is a slight difference in shape when moving from $r_{\text{inv}} = 0.3$ to 0.7.

4.2 Reconstruction loss

An important figure of merit is the distribution of losses for the background and signal samples. As shown in figure 4, QCD tends to have lower loss values than SVJ, providing good discrimination.

In order to achieve the best possible performance, alternative approaches to autoencoder training and evaluation were studied. For those, reference distributions were built for each node of the network from the training data set, using the reconstructed variables, MAE values, or even values from the latent space nodes. The performance was then evaluated by comparing values from the tested jets to those reference distributions.

The comparison was done using the negative log-likelihood, as well as the Mahalanobis distance [33]. It was found that none of the six combinations performed better than the MAE loss function.

4.3 Discrimination power

In order to compare discrimination power of different architectures, one can examine the ROC curves and AUC values for each generated signal. As a reference for the best-case performance, we considered a Boosted Decision Trees (BDT), providing a strong baseline for supervised algorithms. The BDT was implemented using SKLearn’s AdaBoostClassifier, trained on a mixture of all signals, and had its hyper-parameters optimized for maximum AUC. The input features were the same as for the autoencoder. For the final comparison, 10 models were trained and the best one was selected (although the performance spread between different models was negligible).

Figure 3. Distributions of difference between input and reconstructed jet features for QCD and a few example SVJ signals.

Figure 4. Reconstruction loss for QCD and different SVJ signals.

Figure 5. Comparison of the ROC curves for the autoencoder (mean and best models), PCA, BDT and VAE (the best model out of 15 trained) for a few selected signals.

A comparison of ROC curves for the autoencoder, BDT, PCA and VAE is shown in figure 5. As can be seen, the BDT classifier provides the best performance regardless of the signal considered, although it is rather closely followed by the autoencoder. What is more, as will be shown in the further part of this section, the AE can become more powerful than the BDT if the latter was trained on a wrong signal hypothesis. The VAE has been evaluated on MAE reconstruction loss exclusively, neglecting the KL divergence [34]. The coefficient determining proportions of reconstruction and KL loss during the training has been scanned and optimized for maximum performance. It was found optimal to completely exclude KL loss during training, which leads to variance collapse to zero, showing that VAE tested on reconstruction loss only does not improve the results with respect to a regular AE. The PCA performance is significantly worse than those of autoencoders and BDT.

Figure 6 presents a comparison of the AUC values between the autoencoder and BDT for each of the signals. It is clear from this plot that signals with lower values of r_{inv} are much more efficiently identified, although all signals can be tagged at relatively high rates. Moreover, even though the autoencoder performs worse than the BDT in all signal cases, it is still quite good, with the added benefit of requiring no particular signal to train against.

4.4 Robustness

As an additional test, the autoencoder performance was compared to that of a BDT trained on the wrong signal hypothesis. This allows to study the generalization power of the autoencoder against that of a supervised model and assess their ability to discover “by accident” a signal other than the targeted one. As can be seen in the rightmost panel of figure 6, a clear decrease in the performance can be observed in two cases, up to 19% worse when trained on $m_{Z'} = 2 \text{ TeV}$, $r_{\text{inv}} = 0.7$, but tested on $m_{Z'} = 4 \text{ TeV}$, $r_{\text{inv}} = 0.3$. What is remarkable is that in this particular case, the BDT performs worse than the autoencoder, which was trained without any information about the signal characteristics.

Figure 6. Left and middle panels: comparison of AUC values of the autoencoder and a BDT trained on a mixture of all signal models with $m_{\text{dark}} = 20$ GeV. Right panel: AUC values for a BDT trained on a signal with parameters different from those it was tested on. E.g., the AUC value presented in the top left corner of the table comes from a model trained on the lower right corner sample.

Another robustness test was to evaluate both the BDT and the AE on signal samples with varying m_{dark} values, with $m_{Z'}$ fixed to 3 TeV. The results of this study are presented in figure 7. Here, one can see again that in certain cases the BDT trained on a wrong signal hypothesis can be outperformed by the AE. In this scenario, the AE becomes more powerful than a classifier for $m_{\text{dark}} = 100$ GeV, when the BDT training was performed on a mixture of 18 signal samples (with varying $m_{Z'}$ and r_{inv}) with $m_{\text{dark}} = 20$ GeV. More details on the signal samples with varying m_{dark} values can be found in appendix A.

We noticed that training the BDT on a specific signal point with low r_{inv} leads to higher efficiency for signals with higher r_{inv} with respect to training on a mixture of all signal points. This result comes from the fact that low r_{inv} jets are more similar to background jets. Such conclusions suggest that it might be a better strategy to train a classifier on low r_{inv} signals rather than on a mixture of low and high parameter values.

From this test, we conclude that it would be preferable for an LHC signal-specific search to have a signal-agnostic equivalent exploiting autoencoders to complement model-dependent searches, as opposed to relying on accidental generalization properties of the signal-specific selection algorithm. The synergy of the two approaches would enhance the chances of discovering a signal.

4.5 Sensitivity to semivisible jet models

As a last step, a threshold was optimized for the SVJ signal samples to tag jets with loss > 0.06 as anomalous. Doing so, one can mimic the impact of the AE algorithm on a real

Figure 7. Comparison of AUC values of the autoencoder and BDT with varying m_{dark} values. The BDT was trained on a mixture of all signals with $m_{\text{dark}} = 20$ GeV.

analysis, estimating expected upper limits. Events passing the preselection were grouped in categories with 0, 1 or 2 SV jets, and dijet transverse mass distributions for the signal and background in these categories were used to set approximate exclusion limits on the cross section, depending on the Z' mass (while m_{dark} was fixed to 20 GeV). The results, which take into account basic uncertainties (luminosity and trigger), are presented in figures 8a and 8b and compared with the example theoretical cross section. The results without applying an autoencoder-based tagger are also shown.

As expected, the limits are stronger for lower values of r_{inv} and the improvement from using the tagger is more pronounced in this region. The presented results show that one can expect an approach based on the autoencoder to be able to cover a large area of the phase space.

5 Conclusion

This study shows that the usage of a neural autoencoder based on physics-motivated high-level features, despite its simplicity, is extremely effective as an anomaly detection model. This is especially true in cases of a signal with widely varying parameters, such as semivisible jets, as the model is trained on the background sample only and is therefore highly model independent. The projections of the exclusion limits for this particular class of models show a 40–60% improvement in the search reach compared to a study without such an anomalous jet tagger. What is more, the autoencoder could be trained on the real data, accounting for any effects of an imperfect detector simulation.

This result demonstrates the usefulness of the autoencoder network and jet substructure variables such as Energy Flow Polynomials, Energy Correlation Functions, the jet p_T dispersion $p_T D$, and the jet ellipse axes for the detection of anomalous jets. It is also worthwhile to acknowledge certain aspects that should be further investigated in the future.

Figure 8. The cross section limits vs. $m_{Z'}$ for different values of r_{inv} and $m_{\text{dark}} = 20$ GeV, scaled to the luminosity of the full Run 2 of the LHC.

The assumption that the network should be insensitive to detector effects could be tested by simulating such issues and contaminating the training sample with affected events. Next, in order to train autoencoder on the collision data, one has to verify that a small admixture of signal events in the training sample does not influence the performance. Finally, one could try to further exploit the four-momenta of the jet constituents using convolutional or graph autoencoders, which should preserve information about the spatial distribution of particles within the jet.

Acknowledgments

M. Pierini is supported by the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation program (Grant Agreement No. 772369). A. de Cosa and J. Niedziela are supported by the Swiss National Science Foundation (SNFS) under the SNSF Eccellenza program (PCEFP2_186878). L. Le Pottier is supported by the Swiss State Secretariat for Education, Research and Innovation (SERI) under the ThinkSwiss program. K. Pedro is supported by the Fermi National Accelerator Laboratory, managed and operated by Fermi Research Alliance, LLC under Contract No. DE-AC02-07CH11359 with the U.S. Department of Energy.

A m_{dark} variations

Figure 9 displays the input variables used in training the tagger algorithms for signal samples with different m_{dark} values. In particular, the effect of a large m_{dark} value can be seen in the m_j distribution. Figure 10 displays the reconstruction error for these variables and signal samples after applying the AE. The signal distributions show only moderate differences over a large range of m_{dark} values. These figures can be compared, respectively, to figures 1 and 3, which show the variations in r_{inv} and $m_{Z'}$.

Figure 9. Distributions of input variables for QCD background and signal models with $m_{Z'} = 3 \text{ TeV}$ and varying m_{dark} and r_{inv} values.

Figure 10. Distributions of difference between input and reconstructed jet features for QCD and signal models at $m_{Z'} = 3 \text{ TeV}$ and varying m_{dark} and r_{inv} values.

Open Access. This article is distributed under the terms of the Creative Commons Attribution License ([CC-BY 4.0](https://creativecommons.org/licenses/by/4.0/)), which permits any use, distribution and reproduction in any medium, provided the original author(s) and source are credited.

References

- [1] M.J. Strassler and K.M. Zurek, *Echoes of a hidden valley at hadron colliders*, *Phys. Lett. B* **651** (2007) 374 [[hep-ph/0604261](#)] [[INSPIRE](#)].
- [2] T. Cohen, M. Lisanti and H.K. Lou, *Semivisible Jets: Dark Matter Undercover at the LHC*, *Phys. Rev. Lett.* **115** (2015) 171804 [[arXiv:1503.00009](#)] [[INSPIRE](#)].
- [3] T. Cohen, M. Lisanti, H.K. Lou and S. Mishra-Sharma, *LHC Searches for Dark Sector Showers*, *JHEP* **11** (2017) 196 [[arXiv:1707.05326](#)] [[INSPIRE](#)].
- [4] CMS collaboration, *Search for resonant production of strongly coupled dark matter in proton-proton collisions at 13 TeV*, [arXiv:2112.11125](#) [[INSPIRE](#)].
- [5] T. Heimel, G. Kasieczka, T. Plehn and J.M. Thompson, *QCD or What?*, *SciPost Phys.* **6** (2019) 030 [[arXiv:1808.08979](#)] [[INSPIRE](#)].
- [6] M. Farina, Y. Nakai and D. Shih, *Searching for New Physics with Deep Autoencoders*, *Phys. Rev. D* **101** (2020) 075021 [[arXiv:1808.08992](#)] [[INSPIRE](#)].
- [7] G. Kasieczka et al., *The LHC Olympics 2020 a community challenge for anomaly detection in high energy physics*, *Rept. Prog. Phys.* **84** (2021) 124201 [[arXiv:2101.08320](#)] [[INSPIRE](#)].
- [8] T. Aarrestad et al., *The Dark Machines Anomaly Score Challenge: Benchmark Data and Model Independent Event Classification for the Large Hadron Collider*, *SciPost Phys.* **12** (2022) 043 [[arXiv:2105.14027](#)] [[INSPIRE](#)].
- [9] T. Cheng, J.-F. Arguin, J. Leissner-Martin, J. Pilette and T. Golling, *Variational Autoencoders for Anomalous Jet Tagging*, [arXiv:2007.01850](#) [[INSPIRE](#)].
- [10] O. Cerri, T.Q. Nguyen, M. Pierini, M. Spiropulu and J.-R. Vlimant, *Variational Autoencoders for New Physics Mining at the Large Hadron Collider*, *JHEP* **05** (2019) 036 [[arXiv:1811.10276](#)] [[INSPIRE](#)].
- [11] P.T. Komiske, E.M. Metodiev and J. Thaler, *Energy flow polynomials: A complete linear basis for jet substructure*, *JHEP* **04** (2018) 013 [[arXiv:1712.07124](#)] [[INSPIRE](#)].
- [12] A.J. Larkoski, G.P. Salam and J. Thaler, *Energy Correlation Functions for Jet Substructure*, *JHEP* **06** (2013) 108 [[arXiv:1305.0007](#)] [[INSPIRE](#)].
- [13] CMS collaboration, *Performance of quark/gluon discrimination in 8 TeV pp data*, *CMS-PAS-JME-13-002* (2013).
- [14] R. Kogler et al., *Jet Substructure at the Large Hadron Collider: Experimental Review*, *Rev. Mod. Phys.* **91** (2019) 045003 [[arXiv:1803.06991](#)] [[INSPIRE](#)].
- [15] D. Kar and S. Sinha, *Exploring jet substructure in semi-visible jets*, *SciPost Phys.* **10** (2021) 084 [[arXiv:2007.11597](#)] [[INSPIRE](#)].
- [16] P. Schwaller, D. Stolarski and A. Weiler, *Emerging Jets*, *JHEP* **05** (2015) 059 [[arXiv:1502.05409](#)] [[INSPIRE](#)].
- [17] E. Bernreuther, T. Finke, F. Kahlhoefer, M. Krämer and A. Mück, *Casting a graph net to catch dark showers*, *SciPost Phys.* **10** (2021) 046 [[arXiv:2006.08639](#)] [[INSPIRE](#)].

- [18] P. Baldi, K. Cranmer, T. Faucett, P. Sadowski and D. Whiteson, *Parameterized neural networks for high-energy physics*, *Eur. Phys. J. C* **76** (2016) 235 [[arXiv:1601.07913](#)] [[INSPIRE](#)].
- [19] T. Sjöstrand et al., *An introduction to PYTHIA 8.2*, *Comput. Phys. Commun.* **191** (2015) 159 [[arXiv:1410.3012](#)] [[INSPIRE](#)].
- [20] CMS collaboration, *Search for resonant production of strongly coupled dark matter in proton-proton collisions at 13 TeV*, *HEPData* (2021).
- [21] DELPHES 3 collaboration, *DELPHES 3, A modular framework for fast simulation of a generic collider experiment*, *JHEP* **02** (2014) 057 [[arXiv:1307.6346](#)] [[INSPIRE](#)].
- [22] J. Gallicchio and M.D. Schwartz, *Quark and Gluon Jet Substructure*, *JHEP* **04** (2013) 090 [[arXiv:1211.7038](#)] [[INSPIRE](#)].
- [23] J.R. Andersen et al., *Les Houches 2015: Physics at TeV Colliders Standard Model Working Group Report*, in *9th Les Houches Workshop on Physics at TeV Colliders*, (2016) [[arXiv:1605.04692](#)] [[INSPIRE](#)].
- [24] T. Finke, M. Krämer, A. Morandini, A. Mück and I. Oleksiyuk, *Autoencoders for unsupervised anomaly detection in high energy physics*, *JHEP* **06** (2021) 161 [[arXiv:2104.09051](#)] [[INSPIRE](#)].
- [25] B.M. Dillon, T. Plehn, C. Sauer and P. Sorrenson, *Better Latent Spaces for Better Autoencoders*, *SciPost Phys.* **11** (2021) 061 [[arXiv:2104.08291](#)] [[INSPIRE](#)].
- [26] D.P. Kingma and M. Welling, *Auto-Encoding Variational Bayes*, [arXiv:1312.6114](#) [[INSPIRE](#)].
- [27] I. Jolliffe, *Principal Component Analysis*, Springer Berlin Heidelberg, Berlin, Heidelberg (2011), pp. 1094–1096.
- [28] A. Borghesi, A. Bartolini, M. Lombardi, M. Milano and L. Benini, *Anomaly Detection using Autoencoders in High Performance Computing Systems*, [arXiv:1811.05269](#).
- [29] D.-A. Clevert, T. Unterthiner and S. Hochreiter, *Fast and accurate deep network learning by exponential linear units (ELUs)*, in *4th International Conference on Learning Representations, ICLR 2016*, San Juan, Puerto Rico, May 2–4, 2016, Conference Track Proceedings (2016) [[arXiv:1511.07289](#)].
- [30] F. Chollet et al., *Keras*, <https://github.com/fchollet/keras> (2015).
- [31] M. Abadi et al., *TensorFlow: Large-scale machine learning on heterogeneous systems*, <https://www.tensorflow.org/> (2015).
- [32] T. Dozat, *Incorporating Nesterov momentum into Adam*, in *Workshop track — ICLR 2016*, <https://openreview.net/pdf/OM0jvwB8jIp57ZJjtNEZ.pdf>.
- [33] G. McLachlan, *Mahalanobis distance*, *Resonance* **4** (1999) 20.
- [34] J.M. Joyce, *Kullback-Leibler Divergence*, Springer Berlin Heidelberg, Berlin, Heidelberg (2011), pp. 720–722.