Available on CMS information server

**CMS NOTE -2023/013**

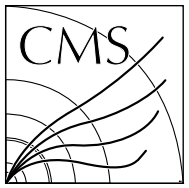

**The Compact Muon Solenoid Experiment**

# CMS Note

Mailing address: CMS CERN, CH-1211 GENEVA 23, Switzerland

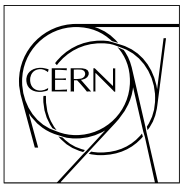

**10 November 2023 (v2)**

# Machine learning techniques for model-independent searches in dijet final states

Phil Harris, Patrick McCormack, Sang Eon Park, Tobias Quadfasel, Manuel Sommerhalder, Louis Moureaux, Gregor Kasieczka, Oz Amram, Petar Maksimovic, Benedikt Maier, Maurizio Pierini, Kinga Wozniak, Thea Aarrestad, Jennifer Ngadiuba, Irene Zoi, Samuel Bright-Thonney, David Shih, and Aritra Bal

**Abstract**

We present the performance of Machine Learning–based anomaly detection techniques for extracting potential new physics phenomena in a model-agnostic way with the CMS Experiment at the Large Hadron Collider. We introduce five distinct outlier detection or density estimation techniques, namely CWoLa, Tag N' Train, CATHODE, QUAK, and QR-VAE, tailored for the identification of anomalous jets originating from the decay of unknown heavy particles. We demonstrate the utility of these diverse approaches in enhancing the sensitivity to a wide variety of potential signals and assess their comparative performance in simulation.

# 1 Introduction

One of the core components of the research program of high-energy particle physics experiments at the Large Hadron Collider, like ATLAS [1] and CMS [2] is to search for physics beyond the Standard Model.

One of the most generic and signal-agnostic type of search to be performed at colliders is a search for new heavy resonances decaying into jets [3–10]. As many models of new physics predict such dijet resonances, this search can provide sensitivity to a wide range of signal models. However, quantum chromodynamics (QCD) processes also produce copious pairs of high energy jets, which represent a large background to these searches and limit their sensitivity. However, some signals may produce jets with properties differing from those of standard QCD jets. Therefore, a large class of searches has also been performed targeting specific signal models in which jet substructure is used to reject QCD background and improve upon the sensitivity of the inclusive search [11]. However, each of these searches are typically optimized based on Monte Carlo simulation of the targeted signal, and makes a selection targeted towards a specific jet substructure signal. Each search therefore only has sensitivity to a small subset of potential dijet resonance signals, and numerous possibilities remain unexplored.

To improve the analysis sensitivity beyond the inclusive dijet search, while preserving a high degree of signal independence, several methods for performing signal-agnostic searches in dijet topologies using Machine Learning (ML) have been proposed [12]. These searches use various ML techniques to design discriminating variables that differentiate QCD jets from signals producing jets with 'anomalous' substructure. Many of these methods are entirely data-driven; utilizing data events to train the ML model used to identify potential anomalies and estimate backgrounds. In this paper, we describe five ML-driven anomaly detection methods used to enhance sensitivity of dijet resonances with anomalous substructure.

Though these methods are agnostic about the specific substructure characteristics of the signal, several assumptions are made about the overall event topology. Signal events assumed to arise from the decay of a narrow generic resonance A with a mass of O( TeV) to two other generic particles B and C which can be either single quarks or gluons, or resonances of unknown mass decaying to hadrons. In the case in which B and C are massive resonances, we consider masses of the A for which B and C are produced at high Lorentz boost such that their decay products are merged into single, large-radius jets. The substructure of the jets can thus be exploited for an efficient signal-independent tagger able to suppress the background rate by several order of magnitude.

The anomaly detection algorithms we explore are complementary to one another both in choice of architecture and training paradigm; un-, weakly- and semi-supervised. Their performance and correlations in simulation is compared.

# 2 ML techniques

Five different ML-based anomaly detection techniques are used in order to suppress the overwhelming QCD background, while enhancing a potential New Physics signal in a model-independent way. A demonstration of this principle is shown in Figure 1, using the *Tag N' Train* (TNT) method described in the following. This figure shows the data (pseudodata) distribution after injection of a New Physics signal without a cut on an anomaly detection score (left) and after a cut on an anomaly detection score (center). The background is significantly suppressed, while the signal is persistent. Ensuring these methods do not change the shape of

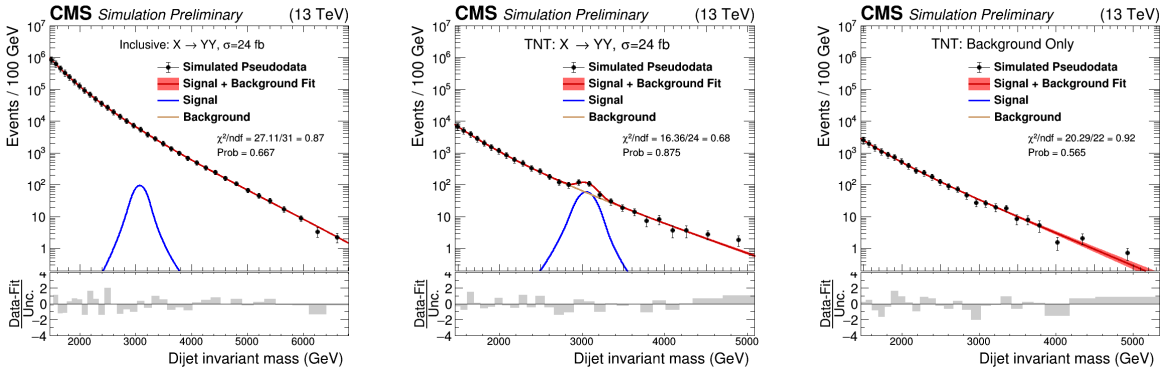


Figure 1: Dijet mass distribution of a simulated set of QCD background events injected with 24 fb of the X→YY signal before any cut on the anomaly score (left) and after cutting on the anomaly score of the TNT algorithm (middle). The distribution after cutting on the TNT anomaly score in a background only sample is shown on the right. In both cases the background distribution after the anomaly remains smooth and is well modeled with a parametric function. Cutting on the TNT score removes a significant amount of background events, resulting in a substantially enhanced signal peak in the middle plot.

the background distribution into something non-smooth is also crucial, as the final statistical analysis involves assuming that the background can be described by a smoothly falling function. This is shown in the plot on the right. A full explanation of the method used to produce this signal-sensitive data distribution, as well as four other similar methods, will be described in the following.

The anomaly detection methods we use are based on three different training paradigms for ML based anomaly detection: un-supervised, weakly-supervised and semi-supervised learning.

The un-supervised learning attempts to construct a model to identify anomalous jets without using any labeled examples. The method employed here consists of a Variational Autoencoder (VAE) trained on a data sample dominated by QCD jets and a quantile regression network (QR) used to decorrelate the anomaly score with the dijet mass. This method is referred to as VAE-QR. Autoencoders are a type of neural network which are trained to compress inputs into a smaller representation and decompress to recover the original inputs. The VAE employed here takes as input the 100 highest-$p_{\mathrm{T}}$ constituents of a jet, with the ordering obtained from a C/A reclustering of the components. Each particle is represented as a set of three features, which are the $x$, $y$, and $z$ component of its momentum $p$. The VAE is trained using jets from the signal-depleted control region. It therefore learns how to perform this compression and decompression on QCD background jets, but should not be able to perform this task as well on anomalous jets not present in the training sample. Therefore the difference between the original and reconstructed data can be used as an effective anomaly score, with higher values corresponding to more signal-like events. To decorrelate this anomaly score from the variable of interest (in this case the dijet invariant mass), a quantile regression [13] method is used. The quantile regression is trained to find the cut on the anomaly score as a function of $m_{\mathrm{jj}}$ which corresponds to a fixed data efficiency in the signal region. A cut on the decorrelated anomaly score is then applied to both jets in the signal region. A cut corresponding to the 10% most anomalous data is used.

Three methods based on weak supervision are employed: CWoLa Hunting [14], TNT [15] and CATHODE [16]. Weakly supervised training [17] is entirely data-driven, and allows one to train a signal versus background classifier by using labels for groups of data events rather than

labels for individual events. This training paradigm requires two mixed samples of data events. The samples are chosen such that one is a mixture of potential signal events and background events, and the other is nearly background pure. The classifier is trained to distinguish between events in these two samples. If there sufficient signal is present in the dataset, the classifier will learn to distinguish signal from background, provided that the background composition in the two samples is the same. These methods train directly on data events in the signal region of the analysis, and thus learn the specific characteristics of the signal if it is present in the dataset. All of the weakly supervised methods pursued in this analysis assume the signal is a narrow resonance, and use the invariant mass of the two jets in their definition of the mixed samples. These methods therefore assume a particular hypothesis for the mass of the new resonance in order to construct the samples for training. If a narrow resonance exists inside the hypothesized region in sufficient abundance, the weakly supervised training procedure will produce a classifier able to discriminate between it and the QCD background. In the absence of a signal, the two samples will be indistinguishable in the training procedure and the resulting classifier will produce random results.

For all 3 weakly supervised methods, the training procedure is repeated for multiple hypotheses of the signal mass in order to scan over the full dijet mass spectrum. The signal region is split into 8 orthogonal $m_{jj}$ windows to be used for the training procedure. The size of these windows is chosen to be significantly larger than the dijet mass resolution, such that we would expect a narrow resonance to be nearly fully contained in a single bin. An additional set of windows, shifted half a bin width with respect to the original set, are defined in order to ensure no signals close to a bin boundary are missed.

For CWoLa Hunting, the two mixed samples are obtained using windows in dijet invariant mass. A potentially signal-rich sample is defined as all events falling in a dijet invariant mass window, and the background-rich sample is defined as events in the neighboring sideband regions. The upper and lower sidebands are reweighted in the training to have the same total weight. Two per-jet classifiers, one for the heavier jet and one for the lighter jet in the event, are trained rather than a single classifier for the full dijet system as was done in Refs. [14, 18]. This allows an additional reweighting procedure during the training: jets in the background region are reweighted to match the $p_T$ distribution of jets in the signal region. This procedure mitigates correlation between the final anomaly score and $m_{jj}$.

The per-jet classifiers are fully connected feed forward neural networks. They take as input the softdrop mass of the jet $m_{SD}$ [19], the n-subjettiness variables $\tau_{21}$,$\tau_{32}$ and $\tau_{43}$ [20], the number of PF candidates inside the jet $n_{PF}$, the lepton subjet fraction (LSF) $LSF_3$ [21] and the maximum b-tagging from the DeepCSV algorithm [22] of to the two leading subjets of the large jet.

In TNT, the two mixed samples are defined in a similar way to CWoLa Hunting, with one addition: an additional unsupervised classifier, a jet-based autoencoder, is used to increase the purity of the signal-rich sample. In addition to the dijet mass, the anomaly score of the autoencoder, evaluated on one of the jets the event at time, is used to sort events into the signal-rich or background-rich categories. The autoencoder used is based on an image representation of the jet [23, 24] and uses a CNN architecture. A separate model is trained for each signal region based on using events in the corresponding sidebands.

To construct the signal-rich and background-rich samples, the two jets in each event are first randomly sorted into two groups. The dijet mass of the event and the autoencoder score of the jet from the first group are used to categorize jets of the second group into mixed samples. The signal-rich mixed sample is defined as events in which the jet from the first group is in the top 20% of the autoencoder scores and have $m_{jj}$ within the signal region. The background-rich

mixed sample is defined as jets coming from events in the dijet mass sidebands or have jets from the first group in the bottom 40% of the autoencoder scores. This process is then repeated: the autoencoder is evaluated on the second group of jets and is used to define mixed samples of jets in the first group. The samples of categorized jets from both groups are then merged together for a single weakly supervised training. The TNT classifier uses identical network architecture, reweighting schemes and input variables as used for the CWoLa Hunting method.

For CATHODE, a dijet mass window is also used to define the potentially signal-rich sample but a different approach is used for the background-rich sample. First, the conditional probability density of background events as function of invariant mass is learned using a normalizing flow architecture [25–33] trained on all events outside the signal window. This probability density is then interpolated into the signal window and used to generate a sample of synthetic background events which are used as the background-rich sample. This approach fully learns correlations between the input features and $m_{\mathrm{jj}}$, allowing the use of variables that are significantly correlated with $m_{\mathrm{jj}}$. The input variables used in CATHODE are the mass of the heavier jet $m_{j1}$, the mass difference between the two jets $\Delta m_{\mathrm{j_1 j_2}} = m_{\mathrm{j_1}} - m_{\mathrm{j_2}}$, and the n-subjettiness variable $\tau_{41}$ for each jet.

Finally, a semi-supervised algorithm referred to as QUAK [34] is used. This method seeks a middle ground between the fully model-agnostic approach of the previous methods and a standard dedicated search. In QUAK, density estimators are used to encode a 'prior' about what the signature of a new physics signal is likely to be based on labeled signal Monte Carlo samples. An additional density estimator is trained using simulated background events to help reject QCD events.

Each density estimator is trained per-event, utilizing the sub-structure information of both jets. The substructure variables used are the same as those used in the CWoLa Hunting method except a modified $N$-subjettiness metric, $\tau_s = \frac{\sqrt{\tau_{21}}}{\tau_1}$, and $\rho = \frac{M}{p_{\mathrm{T}}}$ are used instead of the LSF variable and softdrop mass variables respectively.

A two-dimensional 'QUAK space' is formed, in which each event's position is defined by its negative log probability in the background- and signal-trained flows. We will refer to this probability as a "loss" value; a lower probability indicates a more out-of-distribution event. The most anomalous events are expected to populate the region with high background loss and low signal loss. For a given mass hypothesis, $m_H$, an approximate template of background events is created by considering events in sidebands $(m_H - 900,\ m_H - 400)$ and $(m_H + 200,\ m_H + 700)$. The signal region is taken to be $(m_H - 400,\ m_H + 200)$. The template itself is a binned 2D histogram of the QUAK space. QUAK space bins with an excess of events in the signal region as opposed to the sidebands are chosen for the selection. Events from the full dijet mass spectrum which fall into into these selected QUAK bins are then utilized for final statistical tests for the presence of a potential signal. A general search is performed using QUAK by using a combination of all the benchmark signal samples as a signal prior.

## 3 Performance

The algorithms described above are tested on a simulated dataset of QCD and other minor backgrounds representing 27 $\mathrm{fb}^{-1}$ of luminosity. A basic pre-selection requiring two anti-$k_{\mathrm{T}}$ [35] R = 0.8 jets groomed with the PUPPI algorithm [36] with $p_{\mathrm{T}} > 300$ GeV is applied. A requirement of $|\Delta\eta_{jj}| < 1.3$ is applied to target s-channel resonances and the dijet mass is required to be larger than 1455 GeV to be above trigger thresholds. Two benchmark signals are used to test performance. The first consists of a heavy resonance (X) of mass 3 TeV, decaying to

two daughter resonances (Y and Y′) each of mass 170 GeV which each decay into two quarks. The second consists of a W′ boson of mass 3 TeV decaying to a top quark and a B′ of mass 400 GeV. The B′ subsequently decays into a bottom quark and a Z boson. The performance of these anomaly detection algorithms is verified and compared in several ways: First, it is verified that none of the methods distorts either the background or signal dijet mass shape, as shown in Fig. 1.

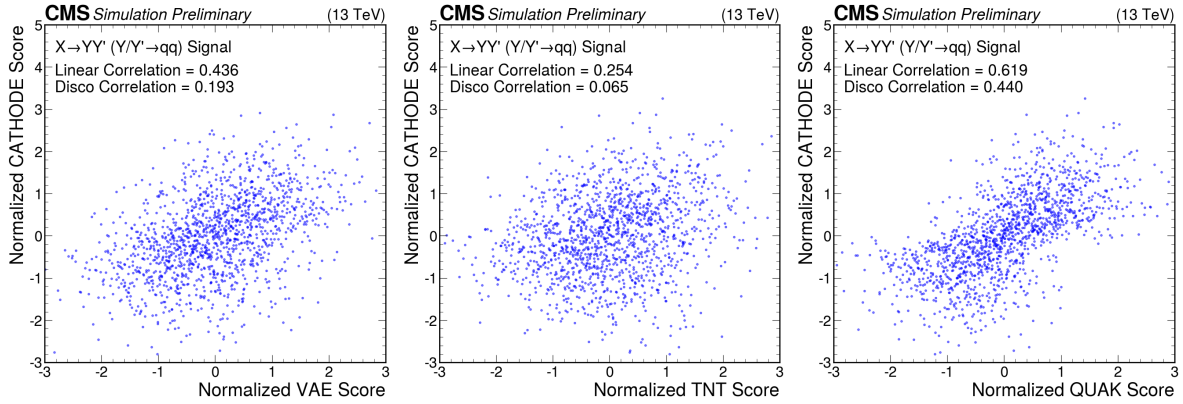


Figure 2: Scatter plots showing the correlation between the anomaly score of the CATHODE method as compared to that of VAE-QR (left), TNT (middle) and QUAK (right) on events from the X→YY′ signal. The Pearson linear correlation coefficient as well as the DISCO correlation [37] are computed in each case.

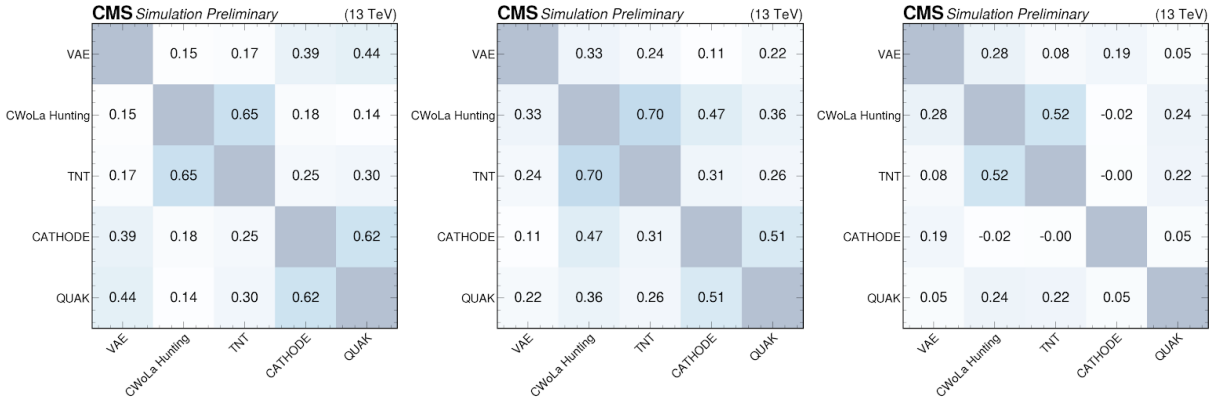


Figure 3: Summary plots showing the Pearson correlation coefficient for each pair of anomaly detection algorithms as evaluated on events from the X→YY′ signal (left), W′ → B′t signal (middle), and QCD background (right). In many cases, the correlations are weak, indicating there is complementarity information in the different approaches.

Secondly, the correlations of the anomaly scores of the different methods are compared in Fig. 2 and Fig. 3. In Fig. 2, we show the CATHODE anomaly score versus three of the other methods (VAE-QR, TNT and QUAK) and compute their Pearson and DISCO [37] correlation coefficients. It is found that the anomaly scores of the different approaches have relatively low correlations, indicating their complementary. The full Pearson correlation matrix between the different methods is shown in Fig. 3.

Finally, the ability of these anomaly scores to enhance search sensitivity to a unknown resonances is assessed and compared to traditional methods in Figure 4. The performance of the anomaly detection methods was verified in a simulated mock dataset. The mock dataset was constructed by selectively sampling events from simulations of different background processes

according to their cross section, instead the commonly used practice of applying event weights. This sampling procedure better captures the dataset size and statistical fluctuations that would present when being applied to data, both of which effect the performance achieved when training a neural networks. Versions of the mock dataset with different amounts of injected signal events were constructed, and the search procedure was repeated on each version. It was verified that no method produces artificial excesses in the absence of signal.

Mock datasets with injected signals were used to test the sensitivities of the anomaly detection methods based on the expected statistical significance of the signal as a function of the size of the injected signal. Their sensitivity was compared to that of an inclusive dijet search, that used only the basic selection criteria and did not feature a cut on an anomaly score, as well as several model-dependent event selections. The first (second) model-specific event selection was tailored to two-pronged (three-pronged) signals and required $\tau_{21} < 0.4$ ($\tau_{32} < 0.65$) and $m_{\mathrm{SD}} > 50\,\mathrm{GeV}$ for both jets in the event. All search strategies utilize the same basic selection criteria, fitting procedure and statistical analysis as employed by the anomaly detection methods. A comparison of the extracted p-value as a function of the signal cross section is shown for two candidate signals, $X \to YY'$ ($M_X = 3\,\mathrm{TeV}$, $M_Y = 170\,\mathrm{GeV}$, $M_{Y'} = 170\,\mathrm{GeV}$) and $W' \to tB'$ ($M_{W'} = 3\,\mathrm{TeV}$, $M_{B'} = 400\,\mathrm{GeV}$) , is shown in Fig. 4.

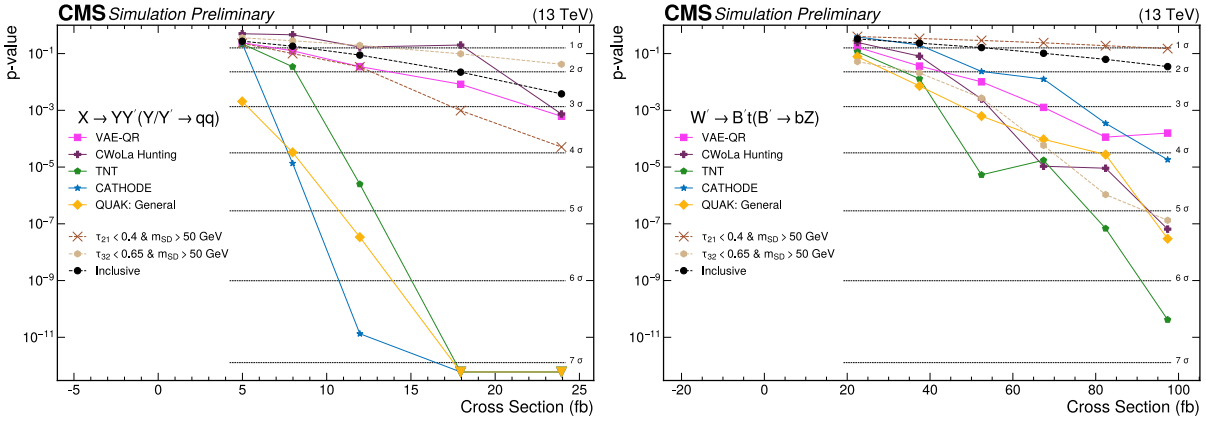


Figure 4: $p$-values as a function of injected signal cross sections for the different anomaly detection methods (solide lines) for two different signals: X→YY′ with $M_X$= 3 TeV, $M_Y$= 170 GeV and $M_{Y'}$= 170 GeV (left), and W′ →tB′ with $M_{W'}$= 3 TeV and $M_{B'}$= 400 GeV (right). The performance of the anomaly detectors are compared to several reference methods (dashed lines): an inclusive search (black), a traditional two-prong-targeted event selection (brown), and a traditional three-prong-targeted selection (beige)

The two-prong targeted and three-pronged targeted selections were found to achieve similar performance to the anomaly detection methods on the signal that matched the targeted number of prongs, but were found to be significantly worse than the inclusive selection on the opposite case. In contrast, all anomaly detection methods were able to demonstrate increased sensitivity above an inclusive search for both signals. For both signals, signal cross sections that resulted in $\leq 2\sigma$ significances for the inclusive search led to $\geq 5\sigma$ significances for multiple anomaly detection methods, illustrating the enhanced discovery potential of these approaches. The sensitivities of the weakly supervised methods depend non-linearly on the signal cross section because the amount of signal present in the data affects the training procedure and therefore the signal selection efficiency.

# 4 Summary

This note described the performance of several signal-independent methods to identify heavy resonances decaying to anomalous jets. Approaches based on weakly supervised, unsupervised, and semi-supervised paradigms were explored. When applying these methods to simulated samples, it was found that all methods were successfully able to identify anomalous jets as distinct from QCD backgrounds, and are therefore able to enhance the discovery potential of these signals in a model-agnostic fashion. It was found that the anomaly scores produced by these methods generally have low correlations, underscoring their complementary and the need for a robust set of tools for such an open-ended objective such as anomaly tagging. These methods lay the groundwork for future model-agnostic explorations of LHC data.

## References


[1] ATLAS Collaboration, "The atlas experiment at the cern large hadron collider", *Journal of Instrumentation* **3** (2008), no. 08, S08003.

[2] CMS Collaboration, "The cms experiment at the cern lhc", *Journal of Instrumentation* **3** (2008), no. 08, S08004.

[3] UA 1 Collaboration, "Two-jet mass distributions at the CERN proton-antiproton collider", *Physics Letters B* **209** (1988), no. 1, 127–134, doi:https://doi.org/10.1016/0370-2693(88)91843-6.

[4] UA 2 Collaboration, "A measurement of two-jet decays of the W and Z bosons at the CERN $p\bar{p}$ collider", *Zeitschrift für Physik C Particles and Fields* **49** (1991), no. 1, 17–28, doi:10.1007/BF01570793.

[5] CDF Collaboration, "Two-jet invariant-mass distribution at $\sqrt{s} = 1.8$ TeV", *Phys. Rev. D* **41** (Mar, 1990) 1722–1725, doi:10.1103/PhysRevD.41.1722.

[6] DØ Collaboration, "Search for new particles in the two-jet decay channel with the DØ detector", *Phys. Rev. D* **69** (Jun, 2004) 111101, doi:10.1103/PhysRevD.69.111101.

[7] ATLAS Collaboration, "Search for new particles in two-jet final states in 7 TeV proton-proton collisions with the ATLAS detector at the LHC", *Phys. Rev. Lett.* **105** (Oct, 2010) 161801, doi:10.1103/PhysRevLett.105.161801.

[8] CMS Collaboration, "Search for new phenomena in dijet mass and angular distributions from pp collisions at $\sqrt{s} = 13$ TeV with the atlas detector", *Physics Letters B* **754** (2016) 302–322, doi:https://doi.org/10.1016/j.physletb.2016.01.032.

[9] CMS Collaboration, "Search for dijet resonances in 7 TeV $pp$ collisions at cms", *Phys. Rev. Lett.* **105** (Nov, 2010) 211801, doi:10.1103/PhysRevLett.105.211801.

[10] CMS Collaboration, "Search for narrow resonances decaying to dijets in proton-proton collisions at $\sqrt{s} = 13$ TeV", *Phys. Rev. Lett.* **116** (Feb, 2016) 071801, doi:10.1103/PhysRevLett.116.071801.

[11] CMS Collaboration, "Search for new heavy resonances decaying to ww, wz, zz, wh, or zh boson pairs in the all-jets final state in proton-proton collisions at s=13tev", *Physics Letters B* **844** (2023) 137813.

[12] G. Kasieczka et al., "The LHC Olympics 2020 a community challenge for anomaly detection in high energy physics", *Rept. Prog. Phys.* **84** (2021), no. 12, 124201, doi:10.1088/1361-6633/ac36b9, arXiv:2101.08320.

[13] K. A. Woźniak et al., "New physics agnostic selections for new physics searches", *EPJ Web Conf.* **245** (2020) 06039, doi:10.1051/epjconf/202024506039.

[14] J. H. Collins, K. Howe, and B. Nachman, "Extending the search for new resonances with machine learning", *Phys. Rev. D* **99** (Jan, 2019) 014038, doi:10.1103/PhysRevD.99.014038.

[15] O. Amram and C. M. Suarez, "Tag N' Train: a technique to train improved classifiers on unlabeled data", *JHEP* **01** (2021) 153, doi:10.1007/JHEP01(2021)153, arXiv:2002.12376.

[16] A. Hallin et al., "Classifying Anomalies THrough Outer Density Estimation (CATHODE)", arXiv:2109.00546.

[17] E. M. Metodiev, B. Nachman, and J. Thaler, "Classification without labels: learning from mixed samples in high energy physics", *Journal of High Energy Physics* **2017** (2017), no. 10, 174, doi:10.1007/JHEP10(2017)174.

[18] ATLAS Collaboration, "Dijet resonance search with weak supervision using $\sqrt{s} = 13$ TeV $pp$ collisions in the ATLAS detector", *Phys. Rev. Lett.* **125** (2020), no. 13, 131801, doi:10.1103/PhysRevLett.125.131801, arXiv:2005.02983.

[19] A. J. Larkoski, S. Marzani, G. Soyez, and J. Thaler, "Soft drop", *Journal of High Energy Physics* **2014** (2014), no. 5, 146, doi:10.1007/JHEP05(2014)146.

[20] J. Thaler and K. Van Tilburg, "Identifying boosted objects with n-subjettiness", *Journal of High Energy Physics* **2011** (2011), no. 3, 15, doi:10.1007/JHEP03(2011)015.

[21] C. Brust et al., "Identifying boosted new physics with non-isolated leptons", *Journal of High Energy Physics* **2015** (2015), no. 4, 79, doi:10.1007/JHEP04(2015)079.

[22] CMS Collaboration, "Identification of heavy-flavour jets with the CMS detector in pp collisions at 13 TeV", *Journal of Instrumentation* **13** (may, 2018) P05011, doi:10.1088/1748-0221/13/05/P05011.

[23] M. Farina, Y. Nakai, and D. Shih, "Searching for New Physics with Deep Autoencoders", *Phys. Rev. D* **101** (2020), no. 7, 075021, doi:10.1103/PhysRevD.101.075021, arXiv:1808.08992.

[24] T. Heimel, G. Kasieczka, T. Plehn, and J. M. Thompson, "QCD or What?", *SciPost Phys.* **6** (2019), no. 3, 030, doi:10.21468/SciPostPhys.6.3.030, arXiv:1808.08979.

[25] D. J. Rezende and S. Mohamed, "Variational inference with normalizing flows", 2016. arXiv:1505.05770.

[26] G. Papamakarios, T. Pavlakou, and I. Murray, "Masked autoregressive flow for density estimation", 2018. arXiv:1705.07057.

[27] M. Germain, K. Gregor, I. Murray, and H. Larochelle, "Made: Masked autoencoder for distribution estimation", 2015. arXiv:1502.03509.

[28] C. Durkan, A. Bekasov, I. Murray, and G. Papamakarios, "Neural spline flows", 2019. arXiv:1906.04032.

[29] R. van den Berg, L. Hasenclever, J. M. Tomczak, and M. Welling, "Sylvester normalizing flows", 2019. arXiv:1803.05649.

[30] L. Dinh, J. Sohl-Dickstein, and S. Bengio, "Density estimation using real nvp", 2017. arXiv:1605.08803.

[31] D. P. Kingma and P. Dhariwal, "Glow: Generative flow", 2018. arXiv:1807.03039.

[32] L. Dinh, D. Krueger, and Y. Bengio, "Nice: Non-linear independent components estimation", 2015. arXiv:1410.8516.

[33] J. Brehmer and K. Cranmer, "Flows for simultaneous manifold learning and density estimation", 2020. arXiv:2003.13913.

[34] S. E. Park et al., "Quasi Anomalous Knowledge: Searching for new physics with embedded knowledge", *JHEP* **21** (2020) 030, doi:10.1007/JHEP06(2021)030, arXiv:2011.03550.

[35] M. Cacciari, G. P. Salam, and G. Soyez, "The anti-$k_t$ jet clustering algorithm", *JHEP* **04** (2008) 063, doi:10.1088/1126-6708/2008/04/063, arXiv:0802.1189.

[36] D. Bertolini, P. Harris, M. Low, and N. Tran, "Pileup Per Particle Identification", *JHEP* **10** (2014) 59, doi:10.1007/JHEP10(2014)059, arXiv:1407.6013.

[37] G. Kasieczka and D. Shih, "Robust Jet Classifiers through Distance Correlation", *Phys. Rev. Lett.* **125** (2020), no. 12, 122001, doi:10.1103/PhysRevLett.125.122001, arXiv:2001.05310.