

**Simulation assisted likelihood-free anomaly detection**Anders Andreassen<sup>1,\*</sup>, Benjamin Nachman<sup>2,†</sup> and David Shih<sup>2,3,4,‡</sup><sup>1</sup>Google, Mountain View, California 94043, USA<sup>2</sup>Physics Division, Lawrence Berkeley National Laboratory, Berkeley, California 94720, USA<sup>3</sup>NHETC, Department of Physics and Astronomy, Rutgers, Piscataway, New Jersey 08854, USA<sup>4</sup>Berkeley Center for Theoretical Physics, University of California, Berkeley, California 94720, USA

(Received 9 February 2020; accepted 13 April 2020; published 6 May 2020)

Given the lack of evidence for new particle discoveries at the Large Hadron Collider (LHC), it is critical to broaden the search program. A variety of model-independent searches have been proposed, adding sensitivity to unexpected signals. There are generally two types of such searches: those that rely heavily on simulations and those that are entirely based on (unlabeled) data. This paper introduces a hybrid method that makes the best of both approaches. For potential signals that are resonant in one known feature, this new method first learns a parametrized reweighting function to morph a given simulation to match the data in sidebands. This function is then interpolated into the signal region, and then the reweighted background-only simulation can be used for supervised learning as well as for background estimation. The background estimation from the reweighted simulation allows for nontrivial correlations between features used for classification and the resonant feature. A dijet search with jet substructure is used to illustrate the new method. Future applications of Simulation Assisted Likelihood-free Anomaly Detection (SALAD) include a variety of final states and potential combinations with other model-independent approaches.

DOI: [10.1103/PhysRevD.101.095004](https://doi.org/10.1103/PhysRevD.101.095004)**I. INTRODUCTION**

An immense search effort by the LHC collaborations has successfully probed many extreme regions of the Standard Model phase space [1–7]. Despite strong theoretical and noncollider experimental motivation, there is currently no convincing evidence for new particles or forces of nature from the LHC searches. However, many final states are uncovered [8,9] and the full hypervariate phase space accessible by modern detector technology is only starting to be probed holistically with deep learning methods [10–13]. There is a great need for new searches that can identify unexpected scenarios.

Until recently, nearly all model independent searches relied heavily on simulation. Generically, these searches operate by comparing data with background-only simulation in a large number of phase space regions. Such searches have been performed without machine learning at D0 [14–17], H1 [18,19], ALEPH [20], CDF [21–23], CMS [24,25], and ATLAS [26–28]. A recent phenomenological

study proposed extending this idea to deep learning classifiers [29,30]. While independent of signal models, these approaches are dependent on the fidelity of the background model simulation for both signal sensitivity and background accuracy. If the background simulation is inaccurate, then differences between simulation and (background-only) data will hide potential signals. Even if a biased simulation can find a signal, if the background is mismodeled, then the signal specificity will be poor.

A variety of approaches have been proposed to enhance signal sensitivity without simulations. Such proposals are based on clustering or nearest neighbor algorithms [31–33], autoencoders [34–39], probabilistic modeling [40], weak supervision [41,42], density estimation [43], and others [44]. These approaches must also be combined with a background estimation strategy. If simulation is used to estimate the background, then the specificity is the same as the model-dependent searches. Many of these approaches can be combined with a resonance search, as explicitly demonstrated in Refs. [41,42]. The background estimation strategy may impose additional constraints on the learning, such as the need for decorrelation between a resonant feature and other discriminative features [45–55]. A detailed overview of model independent approaches can be found in Ref. [43].

While it is desirable to be robust to background model inaccuracies, it is also useful to incorporate information from Standard Model simulations. Even though these

\*ajandreassen@google.com

†bpnachman@lbl.gov

‡shih@physics.rutgers.edu

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International license](https://creativecommons.org/licenses/by/4.0/). Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI. Funded by SCOAP<sup>3</sup>.

simulations are only an approximation to nature, they include an extensive set of fundamental and phenomenological physics models describing the highest energy reactions all the way to signal formation in the detector electronics. This paper describes a method that uses a background simulation in a way that depends as little on that simulation as possible. In particular, a model based on the *deep neural networks using classification for tuning and reweighting* (DCTR) procedure [56] is trained in a region of phase space that is expected to be devoid of signals. In a resonance search, there is one feature where the signal is known to be localized and the sideband can be used to train the DCTR model. This reweighting function learns to morph the simulation into the data and is parametrized in the feature(s) used to mask potential signals. Then, the model is interpolated to the signal-sensitive region and the reweighted background simulation can be used for both enhancing signal sensitivity and estimating the Standard Model background. As deep learning classifiers can naturally probe high dimensional spaces, this reweighting model can in principle exploit the full phase space for both enhancing signal sensitivity and specificity.

This paper is organized as follows. Section II introduces the *Simulation Assisted Likelihood-free Anomaly Detection* (SALAD) method. A dijet search at the LHC is emulated to illustrate the new method. The simulation and deep learning setup are introduced in Sec. III, and then the application of DCTR is shown in Sec. IV. The signal sensitivity and specificity are presented in Secs. V and VI, respectively. The paper ends with conclusions and outlook in Sec. VII.

## II. METHODS

Let  $m$  be a feature (or set of features) that can be used to localize a potential signal in a signal region (SR). Furthermore, let  $x$  be another set of features which are useful for isolating a potential signal. The prototypical example is a resonance search where  $m$  is the single resonant feature, such as the invariant mass of two jets, while  $x$  are other properties of the event, such as the substructure of the two jets. The SALAD method then proceeds as follows:

- (1) Train a classifier  $f$  to distinguish data and simulation for  $m \notin \text{SR}$ . This classifier is parametrized in  $m$  by simply augmenting  $x$  with  $m$ ,  $f = f(x, m)$  [57,58]. If  $f$  is trained using the binary cross entropy or the mean squared error loss, then asymptotically, a weight function  $w(x|m)$  is defined by

$$\begin{aligned} w(x|m) &\equiv \frac{f(x)}{1 - f(x)} \\ &= \frac{p(x|\text{data})}{p(x|\text{simulation})} \times \frac{p(\text{data})}{p(\text{simulation})}, \end{aligned} \quad (2.1)$$

where the last factor in Eq. (2.1) is an overall constant that is the ratio of the total amount of data to the total amount of simulation. This property of neural networks to learn likelihood ratios has been exploited for a variety of full phase space reweighting and parameter estimation proposals in high energy physics [56,57,59–62].

- (2) Simulated events in the SR are reweighted using  $w(x|m)$ . The function  $w(x|m)$  is interpolated automatically by the neural network. A second classifier  $g(x)$  is used to distinguish the reweighted simulation from the data. This can be achieved in the usual way with a weighted loss function such as the binary cross-entropy:

$$\begin{aligned} \text{loss}(g(x)) &= - \sum_{m_i \in \text{SR}_{\text{data}}} \log g(x_i) \\ &\quad - \sum_{m_i \in \text{SR}_{\text{simulation}}} w(x_i|m_i) \log(1 - g(x_i)). \end{aligned} \quad (2.2)$$

Events are then selected with large values of  $g(x)$ . Asymptotically,<sup>1</sup>  $g(x)$  will be monotonically related with the optimal classifier:

$$\frac{g(x)}{1 - g(x)} \propto \frac{p(x|\text{signal} + \text{background})}{p(x|\text{background})}. \quad (2.3)$$

It is important that the same data are not used for training and testing. The easiest way to achieve this is by using different partitions of the data for these two tasks. One can make use of more data with a cross-validation procedure [41,42].

- (3) One could combine the previous step with a standard data-driven background estimation technique such as a sideband fit or the ABCD method. However, one can also directly use the weighted simulation to predict the number of events that should pass a threshold requirement on  $g(x)$ :

$$N_{\text{predicted}}(c) = \sum_{m_i \in \text{SR}_{\text{simulation}}} w(x_i|m_i) \mathbb{I}[g(x_i) > c], \quad (2.4)$$

for some threshold value  $c$  and where  $\mathbb{I}[\cdot]$  is the indicator function that is one when its argument is true and zero otherwise. The advantage of Eq. (2.4) over other data-based methods is that  $g(x)$  could be correlated with  $m$ ; for sideband fits, threshold requirements on  $g$  cannot sculpt local features in the  $m$  spectrum.

<sup>1</sup>Sufficiently flexible neural network architecture, enough training data, and an effective optimization procedure.

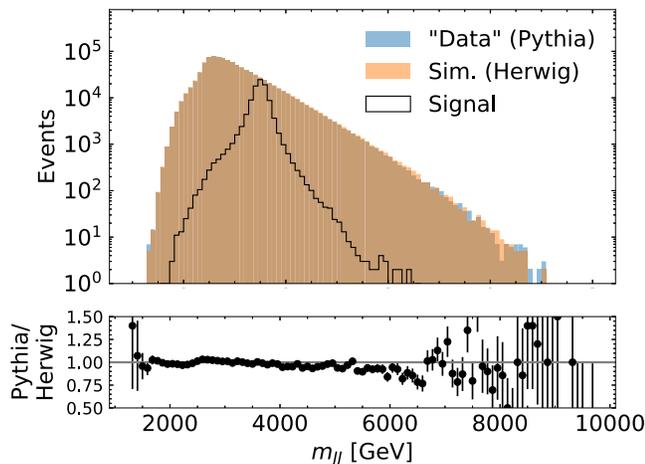


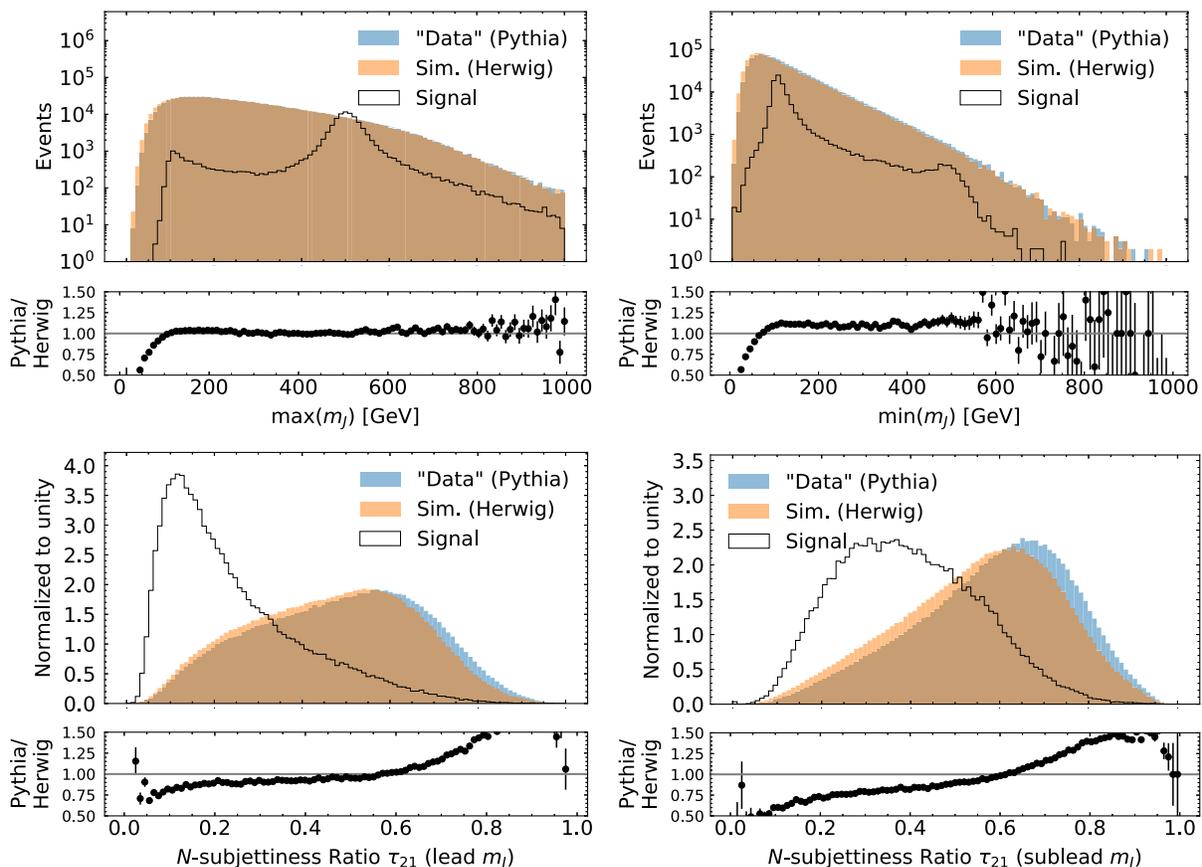
FIG. 1. The invariant mass of the leading two jets.

### III. SIMULATION

A large-radius dijet resonance search is used to illustrate the SALAD method. The simulations are from the LHC Olympics 2020 community challenge R&D dataset [63]. The background process is generic  $2 \rightarrow 2$  parton scattering

(labeled QCD for quantum chromodynamics) and the signal is a hypothetical  $W'$  boson that decays into an  $X$  boson and a  $Y$  boson. Each of the  $X$  and  $Y$  decays to quarks. The masses of the  $W'$ ,  $X$ , and  $Y$  particles are 3.5, 0.5, and 0.1 TeV, respectively. The mass hierarchy between the  $W'$  particle and its decay products means that the  $X$  and  $Y$  particles are Lorentz boosted in the lab frame, and therefore their two-prong decay products are captured inside a single large-radius jet. Particle-level simulations are produced with PYTHIA8 [64,65] or HERWIG++ [66] without pileup or multiple parton interactions and a detector simulation is performed with DELPHES3.4.1 [67–69]. Particle flow objects are clustered into jets using the FASTJET [70,71] implementation of the anti- $k_t$  algorithm [72] using  $R = 1.0$  as the jet radius. Events are selected by requiring at least one such jet with  $p_T > 1.3$  TeV. In the remaining studies, the PYTHIA dataset is treated as “data,” while the HERWIG dataset is treated as “simulation,” to mimic the scenario in practice where the simulation is different from data.

Figure 1 presents the invariant mass of the leading two jets. The  $p_T$  selection is evident from the peak around 3 TeV. The signal peaks around the  $W'$  mass, and aside from the kinematic feature from the jet selection, the

FIG. 2. The four features used for machine learning: jet mass (top) and the  $N$ -subjettiness ratios  $\tau_{21}$  (bottom) for the more massive jet (left) and the less massive jet (right).

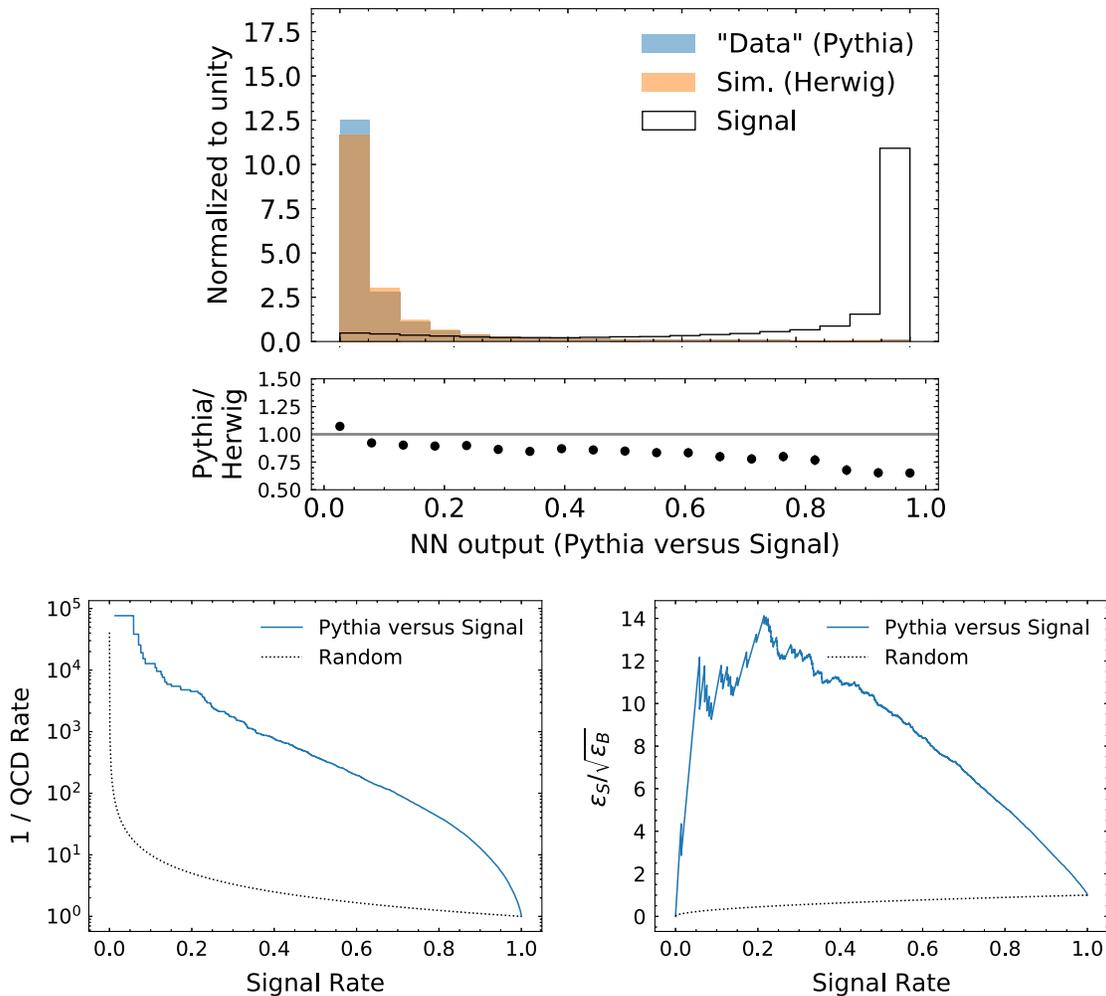


FIG. 3. A supervised classifier trained to distinguish signal from PYTHIA QCD. The top plot is a histogram of the neural network output, the left bottom plot is a receiver operating characteristic (ROC) curve, and the right bottom plot is a significance improvement (SIC) curve.  $\epsilon_S$  is the signal efficiency or true positive rate and  $\epsilon_B$  is the background efficiency or false positive rate.

background distribution is featureless. The spectra from PYTHIA and HERWIG are nearly identical, which may be expected since the invariant mass is mostly determined by hard-scatter matrix elements and not final state effects.

To demonstrate the SALAD approach, two features<sup>2</sup> about each of the leading jets are used for classification. The first feature is the jet mass, and the second feature is the  $N$ -subjettiness ratio [74,75]  $\tau_{21} = \tau_2/\tau_1$ . This second feature is the most widely used feature for differentiating jets that have two hard prongs (as in the signal) from jets that have only one hard prong (as for most of the background). The two jets are ordered by their mass and the four features used for machine learning are presented in Fig. 2. As expected, the signal mass distributions show peaks at the  $X$  and  $Y$  masses and the  $\tau_{21}$

<sup>2</sup>In principle, SALAD can readily accommodate the full phase space as used by the original DCTR method [56] based on particle flow networks [73]; this will be explored in future studies.

distributions are small, indicating two-prong substructure. PYTHIA and HERWIG differ mostly at low mass and across the entire  $\tau_{21}$  distribution.

The baseline performance for classifying signal versus the QCD background is presented in Fig. 3. As is the case for all neural networks presented in the following sections, three fully connected layers with 100 hidden nodes on each intermediate layer are implemented using KERAS [76] and TENSORFLOW [77] with the ADAM [78] optimization algorithm. Rectified linear units are the activation function for all intermediate layers and the sigmoid is used for the final output layer. Networks are trained with binary cross entropy for 50 epochs with early stopping (with patience 10). The supervised classifier presented in Fig. 3 effectively differentiates signal from background, with a maximum significance improvement of about 10. It is expected that the performance of any model independent approach will be bounded from above by the performance of this classifier.

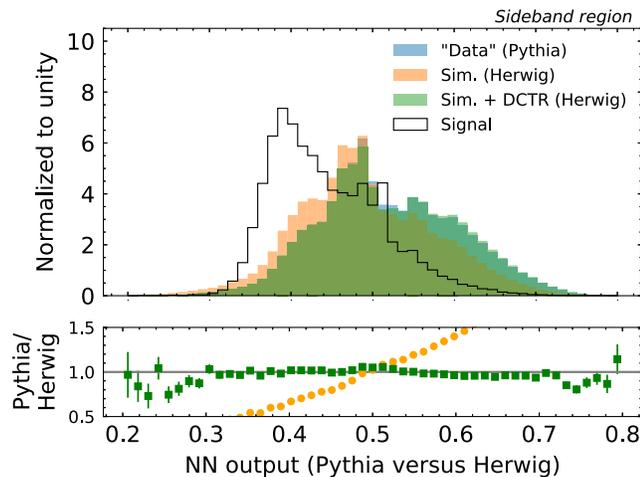


FIG. 4. A histogram of the classifier output for a neural network trained to distinguish data (PYTHIA) and simulation (HERWIG) in the sideband region. The ratio between the simulation (HERWIG) or simulation + DCTR and data (PYTHIA) is depicted by orange circles (green squares) in the lower panel.

#### IV. PARAMETRIZED REWEIGHTING WITH DCTR

The first step of the DCTR reweighting procedure is to train a classifier to distinguish the data (PYTHIA) from the simulation (HERWIG) in a sideband region. The output of such a classifier is shown in Fig. 4, where the signal region is defined as  $m_{jj} \notin [3250, 3750]$  GeV. There are about 850k events in the sideband region and 150k events in the signal region. Unlike the classifier in Fig. 3, the separation in Fig. 4 is not as dramatic because PYTHIA and HERWIG are much more similar than the signal is with QCD. While the background data and simulation are still significantly different, it is important that the simulation be relatively close to the data. One can always ensure that there are no regions of phase space where the data probability density is nonzero while the simulation probability density is zero by augmenting the simulation with e.g., uniform phase space. However, this will lead to very large or very small weights in the DCTR method and therefore a large effective statistical uncertainty in the background prediction and a suboptimal classifier.

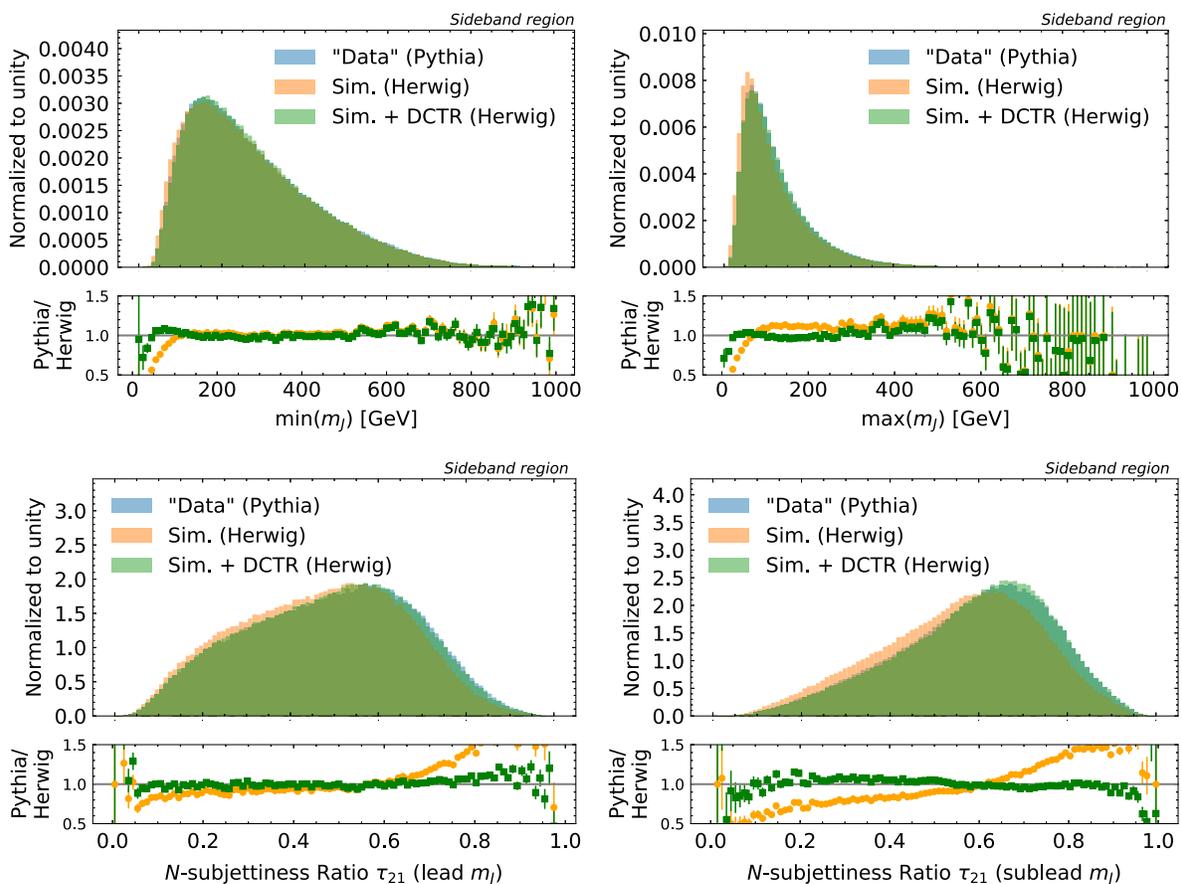


FIG. 5. The four features used for machine learning in the sideband region, before and after applying DCTR: jet mass (top) and the  $N$ -subjettiness ratios  $\tau_{21}$  (bottom) for the more massive jet (left) and the less massive jet (right). The ratio between the simulation (HERWIG) or simulation + DCTR and data (PYTHIA) is depicted by orange circles (green squares) in the lower panels.

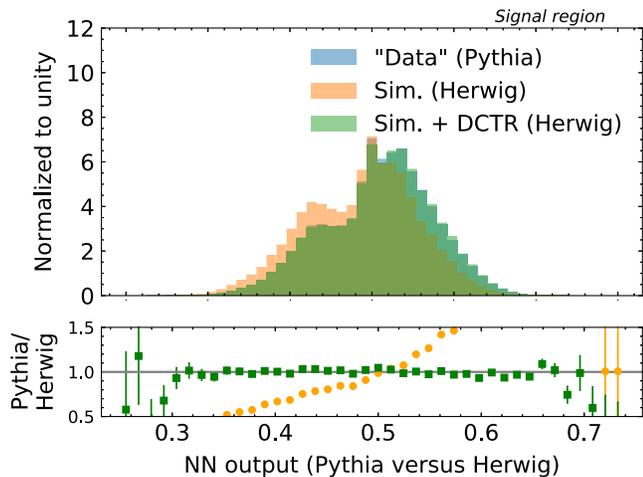


FIG. 6. A histogram of the classifier output for a neural network trained to distinguish data (PYTHIA) and simulation (HERWIG) in the signal region. The ratio between the simulation (HERWIG) or simulation + DCTR and data (PYTHIA) is depicted by orange circles (green squares) in the lower panel.

As expected, the network is a linear function of the likelihood ratio so the ratio plot in Fig. 4 is linear. Interestingly, the signal is more HERWIG-like than PYTHIA-like. The reweighting function is applied to the HERWIG in Fig. 4 to show that the reweighted simulation (Sim + DCTR) looks nearly identical to the data. All of the events used for Fig. 4 are independent from the ones used for training the network. Figure 5 shows that this reweighting works for all of the input distributions to the neural network as well.

The next step for SALAD is to interpolate the reweighting function. The neural network presented in Fig. 4 is trained to be conditional on  $m_{jj}$ , and so it can be evaluated in the SR for values of the invariant mass that were not available during the network training. Note that the signal region must be chosen large enough so that the signal contamination in the sideband does not bias the reweighting function. For this example, for 25% signal fraction in the signal region, the contribution in the sideband is about 1% and has no impact on the DCTR model. Figure 6 shows a classifier trained to distinguish data and simulation in the signal region before and after the application of the interpolated DCTR model. There is excellent closure, also for each of the input features to the classifier as shown in Fig. 7.

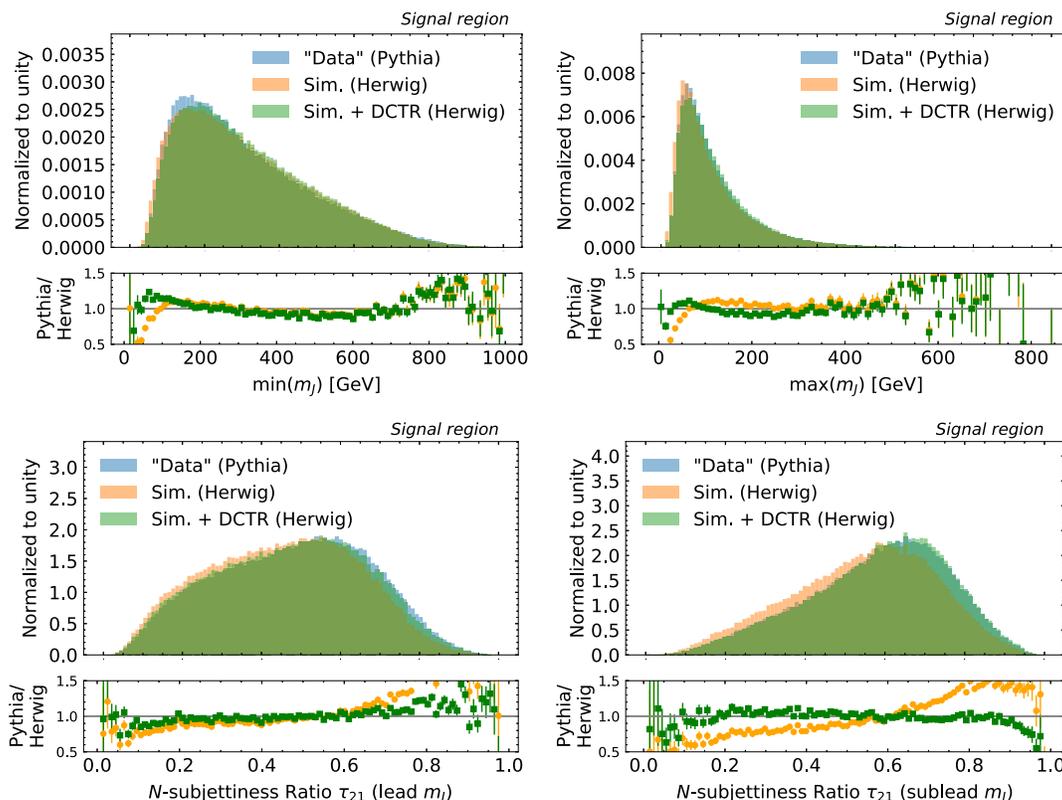


FIG. 7. The four features used for machine learning in the signal region, before and after applying DCTR: jet mass (top) and the  $N$ -subjettiness ratios  $\tau_{21}$  (bottom) for the more massive jet (left) and the less massive jet (right). The ratio between the simulation (HERWIG) or simulation + DCTR and data (PYTHIA) is depicted by orange circles (green squares) in the lower panels.

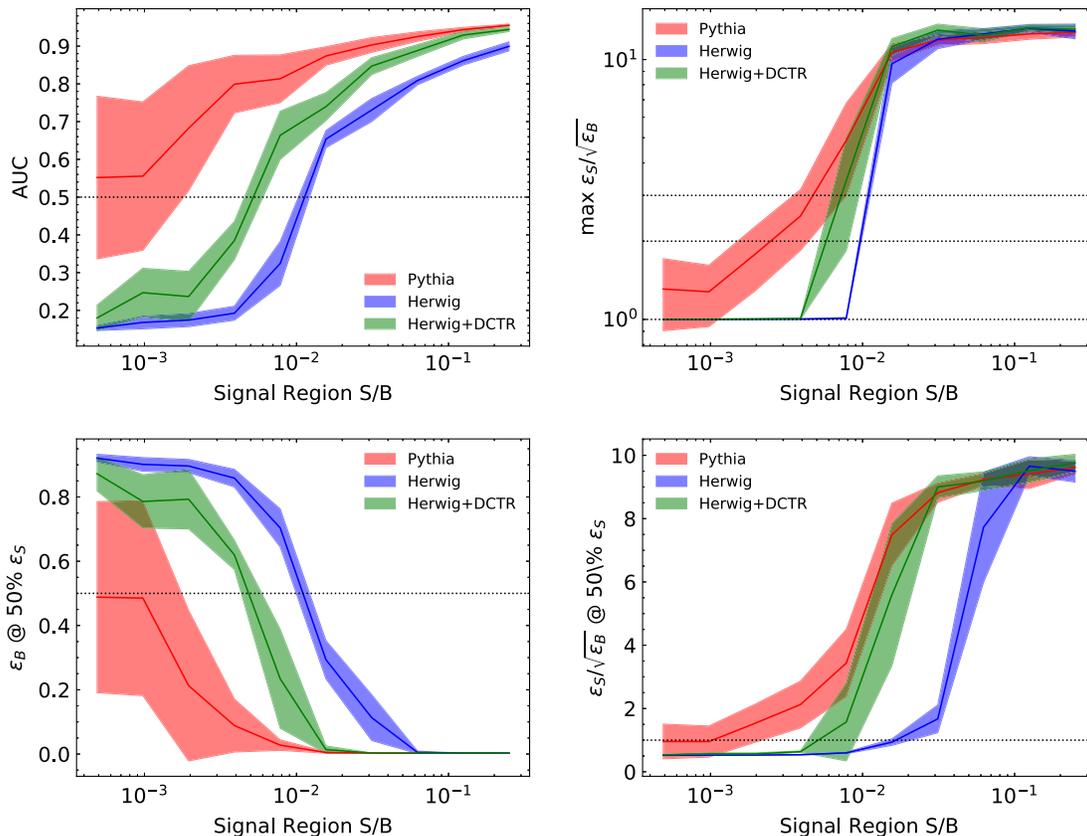


FIG. 8. Four metrics for the sensitivity of the SALAD classifier as a function of the signal-to-background ratio ( $S/B$ ) in the signal region: the AUC in the top left, the maximum significance improvement (top right), the false positive rate at a fixed 50% signal efficiency (bottom left), and the significance improvement at the same fixed 50% signal efficiency (bottom right). The evaluation of these metrics requires signal labels, even though the training of the classifiers themselves do not have signal labels. Error bars correspond to the standard deviation from training five different classifiers. Each classifier is itself the truncated mean over ten random initializations.

## V. SENSITIVITY

After reweighting the signal region to match the data, the next step of the search is to train a classifier to distinguish the reweighted simulation from the data in the signal region. If the reweighting works exactly, then this new classifier will asymptotically learn  $p(\text{signal} + \text{background})/p(\text{background})$ , which is the optimal classifier by the Neyman-Pearson lemma [79]. If the reweighting is sub-optimal, then some of the classifier capacity will be diverted to learning the residual difference between the simulation and background data. If the reweighted simulation is nothing like the data, then all of the capacity will go toward this task and it will not be able to identify the signal. There is therefore a tradeoff between how different the (reweighted) simulation is from the data and how different the signal is from the background. If the signal is much more different from the background than the simulation is from the background data, it is possible that a suboptimally reweighted simulation will still be able to identify the signal (see Sec. VI for problems with background estimation).

Figure 8 shows the sensitivity of the SALAD tagger to signal as a function of the signal-to-background ratio ( $S/B$ )

in the signal region. In all cases, the background is the QCD simulation using PYTHIA.<sup>3</sup> The PYTHIA lines correspond to the case where the simulation follows the same statistics as the data (= PYTHIA). The area under the curve (AUC) should be as close to one as possible, and a tagger that is operating uniformly at random will produce an AUC of 0.5. Antitagging (preferentially tagging events that are not signal-like) results in an AUC of less than 0.5. The maximum significance improvement is calculated as the largest value of  $\epsilon_S/\sqrt{\epsilon_B + 0.01\%}$ , where the 0.01% offset regulates statistical fluctuations at low efficiency.

When the  $S/B \sim \mathcal{O}(1)$ , then the performance in Fig. 8 is similar to the fully supervised classifier presented in Sec. III. As  $S/B \rightarrow 0$ , the PYTHIA curves approach the random classifier, with an AUC of 0.5 and a max

<sup>3</sup>Note that the full one million PYTHIA events are divided in two pieces, one that acts as the test set for all methods and one that is used for further study. The remaining half is further split in half to represent the data or the simulation for the lines marked “PYTHIA” in Fig. 8. For a fair comparison, the HERWIG statistics are comparable to 25% of the full PYTHIA dataset.

significance improvement of unity. The HERWIG curve has an AUC less than 0.5 as  $S/B \rightarrow 0$  because the signal is more HERWIG-like than PYTHIA-like (see Fig. 4), and thus a tagger that requires the features to be datalike (data = PYTHIA) will antitag the signal. Likewise, the efficiency of the tagger on the simulation is higher than 50% when placing a threshold on the  $NN$  that keeps 50% of the events in data. The maximum significance improvement quickly drops to unity for HERWIG when  $S/B \lesssim 1\%$ , indicating the network is spending more capacity on differentiating PYTHIA from HERWIG than finding a signal.

For all four metrics, SALAD significantly improves the performance of the HERWIG-only approach. In particular, the SALAD tagger is effective to about  $S/B \lesssim 0.5\%$ , whereas the HERWIG-only tagger is only able to provide useful discrimination power down to about  $S/B \sim 1\%$ . For the significance improvement and false positive rate at a fixed true positive rate, the SALAD tagger tracks the PYTHIA tagger almost exactly down to below 1%. The AUC about halfway between PYTHIA and HERWIG at high  $S/B$ , which is indicative of poor performance at low efficiency.

## VI. BACKGROUND ESTIMATION

The performance gains from Sec. V can be combined with a sideband background estimation strategy, as long as threshold requirements on the classifier do not sculpt bumps in the  $m_{jj}$  spectrum. However, there is also an opportunity to use SALAD to directly estimate the background from the interpolated simulation. Figure 9 illustrates the efficacy of the background estimation for a single classifier trained in the absence of a signal. Without the DCTR reweighting, the predicted background rate is too low by a factor of 2 or more below 10% data efficiency. With

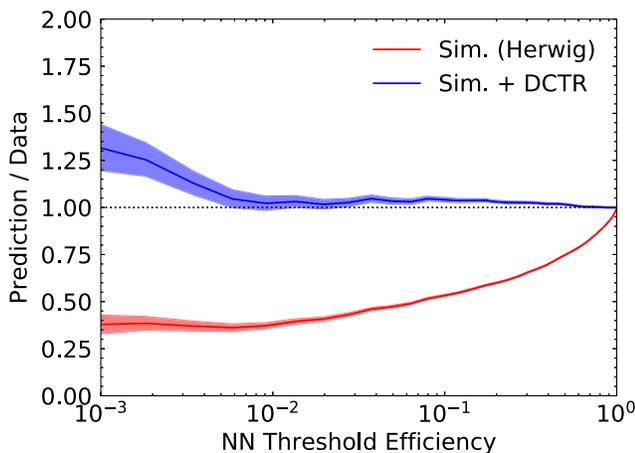


FIG. 9. The predicted efficiency normalized to the true data efficiency in the signal region for various threshold requirements on the  $NN$ . The  $x$  axis is the data efficiency from the threshold. The error bars are due to statistical uncertainties.

the interpolated reweighting function, the background prediction is accurate within a few percent down to about 1% data efficiency.

In practice, the difficulty in using SALAD to directly estimate the background is the estimation of the residual bias. One may be able to use validation regions between the signal region and sideband region, but it will never require as much interpolation as the signal region itself. One can rely on simulation variations and auxiliary measurements to estimate the systematic uncertainty from the direct SALAD background estimation, but estimating high-dimensional uncertainties is challenging [80,81]. With a low-dimensional reweighting or with a proper high-dimensional systematic uncertainty estimate, the parametrized reweighting used in SALAD should result in a lower uncertainty than directly estimating the uncertainty from simulation. In particular, any nuisance parameters that affect the sideband region and the signal region in the same way will cancel when reweighting and interpolating.

## VII. CONCLUSIONS

This paper has introduced SALAD, a new approach to search for resonant anomalies by using parametrized reweighting functions for classification and background estimation. The SALAD approach uses information from simulation in a way that is nearly background-model independent while remaining signal-model agnostic. The only requirement for the signal is that there is one feature where the signal is known to be localized. In the example presented in the paper, this feature was the invariant mass of two jets. The location of the resonance need not be known ahead of time and can be scanned using a series of signal and sideband regions. This scanning will result in a trials factor per nonoverlapping signal region. An additional look elsewhere effect is incurred by scanning the threshold on the neural network. In practice, one could use a small number of widely separated thresholds to be broadly sensitive. As long as the data used for training and testing are independent, there is no additional trials factor for the feature space used for classification. Strategies for maximally using the data for training can be found in Refs. [41,42].

While the numerical SALAD results presented here did not fully achieve the performance of a fully supervised classifier trained directly with inside knowledge about the data, there is room for improvement. In particular, a detailed hyperparameter scan could improve the quality of the reweighting. Additionally, calibration techniques could be used to further increase the accuracy [57]. Future work will investigate the potential of SALAD to analyze higher-dimensional feature spaces as well as classifier features that are strongly correlated with the resonant feature. It will also be interesting to compare SALAD with other recently proposed model independent methods. When the nominal background simulation is an excellent

model of nature, SALAD should perform similarly to the methods presented in Refs. [29,30] and provide a strong sensitivity to new particles. In other regimes where the background simulation is biased, SALAD should continue to provide a physics-informed but still mostly background/signal model-independent approach to extend the search program for new particles at the LHC and beyond.

### CODE AND DATA AVAILABILITY

The code can be found at <https://github.com/bnachman/DCTRHunting>, and the datasets are available on Zendo as part of the LHC Olympics [63].

### ACKNOWLEDGMENTS

B. N. thanks Jack Collins for countless discussions about anomaly detection, including ideas related to SALAD. Some of those discussions happened at the Aspen Center for Physics, which is supported by National Science Foundation Grant No. PHY-1607611. This work was supported by the U.S. Department of Energy, Office of Science under Contract No. DE-AC02-05CH11231. D. S. is supported by DOE Grant No. DOE-SC0010008. D. S. thanks LBNL, Berkeley Center for Theoretical Physics (BCTP), and Berkeley Center for Cosmological Physics (BCCP) for their generous support and hospitality during his sabbatical year. B. N. thanks NVIDIA for providing Volta GPUs for neural network training.

- 
- [1] ATLAS Collaboration, Exotic physics searches (2019), <https://twiki.cern.ch/twiki/bin/view/AtlasPublic/ExoticsPublicResults>.
- [2] ATLAS Collaboration, Supersymmetry searches (2019), <https://twiki.cern.ch/twiki/bin/view/AtlasPublic/SupersymmetryPublicResults>.
- [3] ATLAS Collaboration, Higgs and Diboson searches (2019), <https://twiki.cern.ch/twiki/bin/view/AtlasPublic/HDBSPublicResults>.
- [4] CMS Collaboration, CMS exotica public physics results (2019), <https://twiki.cern.ch/twiki/bin/view/CMSPublic/PhysicsResultsEXO>.
- [5] CMS Collaboration, CMS supersymmetry physics results (2019), <https://twiki.cern.ch/twiki/bin/view/CMSPublic/PhysicsResultsSUS>.
- [6] CMS Collaboration, CMS Beyond-two-Generations (B2G) public physics results (2019), <https://twiki.cern.ch/twiki/bin/view/CMSPublic/PhysicsResultsB2G>.
- [7] LHCb Collaboration, Publications of the QCD, electroweak and Exotica Working Group (2019), [http://lhcbproject.web.cern.ch/lhcbproject/Publications/LHCbProjectPublic/Summary\\_QEE.html](http://lhcbproject.web.cern.ch/lhcbproject/Publications/LHCbProjectPublic/Summary_QEE.html).
- [8] J. H. Kim, K. Kong, B. Nachman, and D. Whiteson, The motivation and status of two-body resonance decays after the LHC Run 2 and beyond, [arXiv:1907.06659](https://arxiv.org/abs/1907.06659).
- [9] N. Craig, P. Draper, K. Kong, Y. Ng, and D. Whiteson, The unexplored landscape of two-body resonances, *Acta Phys. Pol. B* **50**, 837 (2019).
- [10] A. J. Larkoski, I. Moutl, and B. Nachman, Jet substructure at the Large Hadron Collider: A review of recent advances in theory and machine learning, *Phys. Rep.* **841**, 1 (2020).
- [11] D. Guest, K. Cranmer, and D. Whiteson, Deep learning and its application to LHC physics, *Annu. Rev. Nucl. Part. Sci.* **68**, 161 (2018).
- [12] M. Abdughani, J. Ren, L. Wu, J. M. Yang, and J. Zhao, Supervised deep learning in high energy phenomenology: A mini review, *Commun. Theor. Phys.* **71**, 955 (2019).
- [13] A. Radovic, M. Williams, D. Rousseau, M. Kagan, D. Bonacorsi, A. Himmel, A. Aurisano, K. Terao, and T. Wongjirad, Machine learning at the energy and intensity frontiers of particle physics, *Nature (London)* **560**, 41 (2018).
- [14] B. Knuteson, Ph. D. thesis, University of California at Berkeley, 2000.
- [15] B. Abbott *et al.* (D0 Collaboration), Search for new physics in  $e\mu X$  data at  $D\bar{O}$  using Sherlock: A quasi model independent search strategy for new physics, *Phys. Rev. D* **62**, 092004 (2000).
- [16] V. M. Abazov *et al.* (D0 Collaboration), A quasi model independent search for new physics at large transverse momentum, *Phys. Rev. D* **64**, 012004 (2001).
- [17] B. Abbott *et al.* (D0 Collaboration), A Quasi-Model-Independent Search for New High  $p_T$  Physics at  $D\bar{O}$ , *Phys. Rev. Lett.* **86**, 3712 (2001).
- [18] F. D. Aaron *et al.* (H1 Collaboration), A general search for new Phenomena at HERA, *Phys. Lett. B* **674**, 257 (2009).
- [19] A. Aktas, V. Andreev, T. Anthonis, A. Asmone, A. Babaev, S. Backovic, J. Bähr, P. Baranov, E. Barrelet, and W. Bartel (H1 Collaboration), A general search for new phenomena in ep scattering at HERA, *Phys. Lett. B* **602**, 14 (2004).
- [20] K. S. Cranmer, Searching for new physics: Contributions to LEP and the LHC, Ph.D. Thesis, Wisconsin University, 2005.
- [21] T. Aaltonen *et al.* (CDF Collaboration), Model-independent and Quasi-model-independent search for new physics at CDF, *Phys. Rev. D* **78**, 012002 (2008).
- [22] T. Aaltonen *et al.* (CDF Collaboration), Model-independent Global search for new high-p(T) physics at CDF, [arXiv:0712.2534](https://arxiv.org/abs/0712.2534).
- [23] T. Aaltonen *et al.* (CDF Collaboration), Global search for new physics with  $2.0 \text{ fb}^{-1}$  at CDF, *Phys. Rev. D* **79**, 011101 (2009).
- [24] CMS Collaboration, MUSiC, a Model Unspecific Search for New Physics, in pp Collisions at  $\sqrt{s} = 8 \text{ TeV}$ , CERN Report No. CMS-PAS-EXO-14-016, 2017.
- [25] CMS Collaboration, Model unspecific search for new physics in pp collisions at  $\sqrt{s} = 7 \text{ TeV}$ , CERN Report No. CMS-PAS-EXO-10-021, 2011.

- [26] M. Aaboud *et al.* (ATLAS Collaboration), A strategy for a general search for new phenomena using data-derived signal regions and its application within the ATLAS experiment, *Eur. Phys. J. C* **79**, 120 (2019).
- [27] ATLAS Collaboration, A general search for new phenomena with the ATLAS detector in pp collisions at  $\sqrt{s} = 8$  TeV, CERN Report No. ATLAS-CONF-2014-006, 2014.
- [28] ATLAS Collaboration, A general search for new phenomena with the ATLAS detector in  $pp$  collisions at  $\sqrt{s} = 7$  TeV, CERN Report No. ATLAS-CONF-2012-107 (2012).
- [29] R. T. D’Agnolo and A. Wulzer, Learning new physics from a machine, *Phys. Rev. D* **99**, 015014 (2019).
- [30] R. T. D’Agnolo, G. Grosso, M. Pierini, A. Wulzer, and M. Zanetti, Learning multivariate new physics, [arXiv:1912.12155](https://arxiv.org/abs/1912.12155).
- [31] A. De Simone and T. Jacques, Guiding new physics searches with unsupervised learning, *Eur. Phys. J. C* **79**, 289 (2019).
- [32] A. Mullin, H. Pacey, M. Parker, M. White, and S. Williams, Does SUSY have friends? A new approach for LHC event analysis, [arXiv:1912.10625](https://arxiv.org/abs/1912.10625).
- [33] G. M. Alessandro Casa, Nonparametric semisupervised classification for signal detection in high energy physics, [arXiv:1809.02977](https://arxiv.org/abs/1809.02977).
- [34] M. Farina, Y. Nakai, and D. Shih, Searching for new physics with deep autoencoders, [arXiv:1808.08992](https://arxiv.org/abs/1808.08992).
- [35] T. Heimel, G. Kasieczka, T. Plehn, and J. M. Thompson, QCD or What?, *SciPost Phys.* **6**, 030 (2019).
- [36] T. S. Roy and A. H. Vijay, A robust anomaly finder based on autoencoder, [arXiv:1903.02032](https://arxiv.org/abs/1903.02032).
- [37] O. Cerri, T. Q. Nguyen, M. Pierini, M. Spiropulu, and J.-R. Vlimant, Variational autoencoders for new physics mining at the Large Hadron Collider, *J. High Energy Phys.* **05** (2019) 036.
- [38] A. Blance, M. Spannowsky, and P. Waite, Adversarially-trained autoencoders for robust unsupervised new physics searches, *J. High Energy Phys.* **10** (2019) 047.
- [39] J. Hajer, Y.-Y. Li, T. Liu, and H. Wang, Novelty detection meets collider physics, [arXiv:1807.10261](https://arxiv.org/abs/1807.10261).
- [40] B. M. Dillon, D. A. Faroughy, and J. F. Kamenik, Uncovering latent jet substructure, *Phys. Rev. D* **100**, 056002 (2019).
- [41] J. H. Collins, K. Howe, and B. Nachman, Anomaly Detection for Resonant New Physics with Machine Learning, *Phys. Rev. Lett.* **121**, 241803 (2018).
- [42] J. H. Collins, K. Howe, and B. Nachman, Extending the search for new resonances with machine learning, *Phys. Rev. D* **99**, 014038 (2019).
- [43] B. Nachman and D. Shih, Anomaly detection with density estimation, [arXiv:2001.04990](https://arxiv.org/abs/2001.04990).
- [44] J. A. Aguilar-Saavedra, J. H. Collins, and R. K. Mishra, A generic anti-QCD jet tagger, *J. High Energy Phys.* **11** (2017) 163.
- [45] G. Louppe, M. Kagan, and K. Cranmer, Learning to Pivot with adversarial networks, [arXiv:1611.01046](https://arxiv.org/abs/1611.01046).
- [46] J. Dolen, P. Harris, S. Marzani, S. Rappoccio, and N. Tran, Thinking outside the ROCs: Designing Decorrelated Taggers (DDT) for jet substructure, *J. High Energy Phys.* **05** (2016) 156.
- [47] I. Moutl, B. Nachman, and D. Neill, Convolved substructure: Analytically decorrelating jet substructure observables, *J. High Energy Phys.* **05** (2018) 002.
- [48] J. Stevens and M. Williams, uBoost: A boosting method for producing uniform selection efficiencies from multivariate classifiers, *J. Instrum.* **8**, P12013 (2013).
- [49] C. Shimmin, P. Sadowski, P. Baldi, E. Weik, D. Whiteson, E. Goul, and A. Sogaard, Decorrelated jet substructure tagging using adversarial neural networks, *Phys. Rev. D* **96**, 074034 (2017).
- [50] L. Bradshaw, R. K. Mishra, A. Mitridate, and B. Ostdiek, Mass agnostic jet taggers, *SciPost Phys.* **8**, 011 (2020).
- [51] ATLAS Collaboration, Performance of mass-decorrelated jet substructure observables for hadronic two-body decay tagging in ATLAS, CERN Report No. ATL-PHYS-PUB-2018-014, 2018.
- [52] L.-G. Xia, QBDT, A new boosting decision tree method with systematical uncertainties into training for high energy physics, *Nucl. Instrum. Methods Phys. Res., Sect. A* **930**, 15 (2019).
- [53] C. Englert, P. Galler, P. Harris, and M. Spannowsky, Machine learning uncertainties with adversarial neural networks, *Eur. Phys. J. C* **79**, 4 (2019).
- [54] S. Wunsch, S. Jörger, R. Wolf, and G. Quast, Reducing the dependence of the neural network function to systematic uncertainties in the input space, *Comput. Software Big Sci.* **4**, 5 (2020).
- [55] G. Kasieczka and D. Shih, DisCo fever: Robust networks through distance correlation, [arXiv:2001.05310](https://arxiv.org/abs/2001.05310).
- [56] A. Andreassen and B. Nachman, Neural networks for full phase-space reweighting and parameter tuning, [arXiv:1907.08209](https://arxiv.org/abs/1907.08209).
- [57] K. Cranmer, J. Pavez, and G. Louppe, Approximating likelihood ratios with calibrated discriminative classifiers, [arXiv:1506.02169](https://arxiv.org/abs/1506.02169).
- [58] P. Baldi, K. Cranmer, T. Faucett, P. Sadowski, and D. Whiteson, Parameterized neural networks for high-energy physics, *Eur. Phys. J. C* **76**, 235 (2016).
- [59] J. Brehmer, G. Louppe, J. Pavez, and K. Cranmer, Mining gold from implicit models to improve likelihood-free inference, *Proc. Natl. Acad. Sci. U.S.A.* **117**, 5242 (2020).
- [60] J. Brehmer, K. Cranmer, G. Louppe, and J. Pavez, A guide to constraining effective field theories with machine learning, *Phys. Rev. D* **98**, 052004 (2018).
- [61] J. Brehmer, K. Cranmer, G. Louppe, and J. Pavez, Constraining Effective Field Theories with Machine Learning, *Phys. Rev. Lett.* **121**, 111801 (2018).
- [62] A. Andreassen, P. T. Komiske, E. M. Metodiev, B. Nachman, and J. Thaler, OmniFold: A method to simultaneously unfold all observables, [arXiv:1911.09107](https://arxiv.org/abs/1911.09107).
- [63] G. Kasieczka, B. Nachman, and D. Shih, R&D dataset for LHC Olympics 2020 anomaly detection challenge, <https://doi.org/10.5281/zenodo.2629073>.
- [64] T. Sjöstrand, S. Mrenna, and P. Z. Skands, PYTHIA 6.4 Physics and manual, *J. High Energy Phys.* **05** (2006) 026.
- [65] T. Sjöstrand, S. Mrenna, and P. Z. Skands, A brief introduction to PYTHIA 8.1, *Comput. Phys. Commun.* **178**, 852 (2008).

- [66] M. Bahr *et al.*, Herwig++ physics and manual, *Eur. Phys. J. C* **58**, 639 (2008).
- [67] J. de Favereau, C. Delaere, P. Demin, A. Giammanco, V. Lemaitre, A. Mertens, and M. Selvaggi (DELPHES3 Collaboration), DELPHES3, A modular framework for fast simulation of a generic collider experiment, *J. High Energy Phys.* **02** (2014) 057.
- [68] A. Mertens, New features in DELPHES3, *J. Phys. Conf. Ser.* **608**, 012045 (2015).
- [69] M. Selvaggi, DELPHES3: A modular framework for fast-simulation of generic collider experiments, *J. Phys. Conf. Ser.* **523**, 012033 (2014).
- [70] M. Cacciari, G. P. Salam, and G. Soyez, FastJet user manual, *Eur. Phys. J. C* **72**, 1896 (2012).
- [71] M. Cacciari and G. P. Salam, Dispelling the  $N^3$  myth for the  $k_t$  jet-finder, *Phys. Lett. B* **641**, 57 (2006).
- [72] M. Cacciari, G. P. Salam, and G. Soyez, The anti- $k_t$  jet clustering algorithm, *J. High Energy Phys.* **04** (2008) 063.
- [73] P. T. Komiske, E. M. Metodiev, and J. Thaler, Energy flow networks: Deep sets for particle jets, *J. High Energy Phys.* **01** (2019) 121.
- [74] J. Thaler and K. Van Tilburg, Maximizing boosted top identification by minimizing N-subjettiness, *J. High Energy Phys.* **02** (2012) 093.
- [75] J. Thaler and K. Van Tilburg, Identifying boosted objects with N-subjettiness, *J. High Energy Phys.* **03** (2011) 015.
- [76] F. Chollet, Keras, <https://github.com/fchollet/keras> (2017).
- [77] M. Abadi, P. Barham, J. Chen, Z. Chen, A. Davis, J. Dean, M. Devin, S. Ghemawat, G. Irving, M. Isard *et al.*, TensorFlow: A system for large-scale machine learning, in *OSDI* (USENIX, Savannah, 2016), Vol. 16, p. 265.
- [78] D. Kingma and J. Ba, Adam: A method for stochastic optimization, [arXiv:1412.6980](https://arxiv.org/abs/1412.6980).
- [79] J. Neyman and E. S. Pearson, On the problem of the most efficient tests of statistical hypotheses, *Phil. Trans. R. Soc. A* **231**, 289 (1933).
- [80] B. Nachman, A guide for deploying deep learning in LHC searches: How to achieve optimality and account for uncertainty, [arXiv:1909.03081](https://arxiv.org/abs/1909.03081).
- [81] B. Nachman and C. Shimmin, AI safety for high energy physics, [arXiv:1910.08606](https://arxiv.org/abs/1910.08606).