

Autoencoders for unsupervised anomaly detection in high energy physics

Thorben Finke, Michael Krämer, Alessandro Morandini, Alexander Mück and Ivan Oleksiyuk

Institute for Theoretical Particle Physics and Cosmology (TTK), RWTH Aachen University, D-52056 Aachen, Germany

E-mail: finke@physik.rwth-aachen.de, mkraemer@physik.rwth-aachen.de, morandini@physik.rwth-aachen.de, mueck@physik.rwth-aachen.de, ivan.oleksiyuk@rwth-aachen.de

ABSTRACT: Autoencoders are widely used in machine learning applications, in particular for anomaly detection. Hence, they have been introduced in high energy physics as a promising tool for model-independent new physics searches. We scrutinize the usage of autoencoders for unsupervised anomaly detection based on reconstruction loss to show their capabilities, but also their limitations. As a particle physics benchmark scenario, we study the tagging of top jet images in a background of QCD jet images. Although we reproduce the positive results from the literature, we show that the standard autoencoder setup cannot be considered as a model-independent anomaly tagger by inverting the task: due to the sparsity and the specific structure of the jet images, the autoencoder fails to tag QCD jets if it is trained on top jets even in a semi-supervised setup. Since the same autoencoder architecture can be a good tagger for a specific example of an anomaly and a bad tagger for a different example, we suggest improved performance measures for the task of model-independent anomaly detection. We also improve the capability of the autoencoder to learn non-trivial features of the jet images, such that it is able to achieve both top jet tagging and the inverse task of QCD jet tagging with the same setup. However, we want to stress that a truly model-independent and powerful autoencoder-based unsupervised jet tagger still needs to be developed.

KEYWORDS: Jets, QCD Phenomenology

ARXIV EPRINT: [2104.09051](https://arxiv.org/abs/2104.09051)

Contents

1	Introduction	1
2	Autoencoder limitations	4
2.1	Jet data and autoencoder architecture	4
2.2	Limited reconstruction	5
2.3	Complexity bias	8
2.4	Tagging performance	10
3	Improving the autoencoder performance	11
3.1	Intensity remapping	11
3.2	Kernel MSE	11
3.3	Autoencoder performance	12
3.4	Tagging performance	17
4	Conclusion	19
A	Jet simulation and neural network architecture	21
B	Intensity distributions after remapping	22
C	Further results	23

1 Introduction

Having discovered the Higgs boson in 2012, the experimental foundation of the Standard Model of particle physics has been completed at the Large Hadron Collider (LHC) [1, 2]. However, working at the energy frontier, new particles and interactions, not covered by the Standard Model, could be measured by accumulating more and more statistics, or may already hide in the huge amount of data taken by the LHC experiments. From a machine learning perspective, the search for new physics is a quest for finding anomalous data, usually called signal in the physics context, in the vast background of collider events described by Standard Model interactions. New physics would contribute the out-of-distribution data with respect to the data expected in the Standard Model. Hence, it is not surprising that the recent advances in machine learning have also had a huge impact in the LHC context [3–8]. In addition to improving classical search strategies, machine learning may even open up completely new ways to look for anomalous data, i.e. new physics, in a model-independent way.

Single collider events, i.e. data instances, usually cannot simply be labeled as signal and background due to the intrinsically probabilistic nature of quantum mechanics. However,

identifying so-called reducible backgrounds, which share some but not all features of the signal events, amounts to a straightforward classification task, where one can try to either employ supervised or unsupervised machine learning techniques to separate the two classes.

Supervised classifiers are extremely powerful tools but have a limited applicability since labeled data has to be available. In the collider-physics context, they cannot be directly applied to experimental data since collider data is not labeled. However, labeled data is available from the sophisticated simulation of collider events. If the transition from simulation to measured data is understood well enough, supervised machine learning can be successfully applied in particle physics. Still, traditional supervised algorithms are always model-dependent and can only efficiently detect the signals simulated for training. Interesting ideas to overcome those shortcomings have been recently discussed in the literature [9–15].

In unsupervised machine learning, an algorithm does not learn from labeled examples but should understand the structure of the data in some other way in order to sort it into a fixed or variable number of classes. Using unsupervised methods, which can be directly applied to LHC data, is a much more difficult but also an even more rewarding task. Designing an unsupervised tagger which is able to detect anomalies independently of the new physics model is certainly the ultimate vision.

For anomaly or outlier detection (see e.g. [16, 17] for reviews), it is usually assumed that only a few anomalous data instances are to be found in the data distribution. In semi-supervised machine learning, which is slightly less ambitious, a data sample consisting only of the background class is provided during training. The task is to tag signal data when testing on a mixed sample. This might be also an option in particle physics applications since signal free data from control regions might be available. Many different unsupervised and semi-supervised algorithms have been applied in the physics context (see e.g. [18, 19] for reviews).

One particularly promising unsupervised method for anomaly detection is based on the autoencoder (AE) architecture [20]. An AE consists of two parts: an encoder and a decoder. The encoder is a neural network that compresses the input data into a few latent space variables which are also often called the bottleneck. The decoder is a neural network which reconstructs the initial data from the latent space variables. By choosing a suitable loss function both parts are trained together as one neural network to reconstruct the input data as well as possible.

The main idea is to train an AE on a dataset consisting purely (for semi-supervised) or mostly (for unsupervised anomaly detection) of the background class. By having much fewer variables in the latent space than in the input and output layers, the AE is not able to learn the identity transformation. Instead, to minimize the loss, it is forced to extract in its bottleneck the variables that correspond to the most prominent features of the background class. Ideally it extracts correlations in the input data which allow for an efficient data compression. The AE learns a representation that uses the structure of the training data and is therefore specific for this set. Hence, if an AE encounters data that has features somewhat different from the background class, it should not be able to effectively encode and decode these features. The loss of the reconstruction is expected

to be larger. Therefore, a trained AE can be used as an anomaly tagger with its loss function as the anomaly score. This simple idea and its variants have been successfully explored in different machine learning applications (see e.g. [16, 21, 22] for reviews). Also in particle physics autoencoders have been successfully used for anomaly detection [23–29], in particular for top tagging [30–32] which will be the benchmark application in this work.

An unsupervised or semi-supervised machine learning algorithm for anomaly detection has an advantage compared to supervised methods if it is as model-independent as possible. In particular it should not be tailored for a specific kind of anomaly, i.e. a specific new physics model, but it should work for any, or at least a wide variety of possible signals. Ideally, it should also be able to detect an unexpected new physics signature.

Specifically, the setup should also be a working tagger if the background data and the anomalous signal interchange their role. While such an inverse tagger is not of direct physical importance, it offers a simple test for model-independence. If such an inverse tagger performs significantly worse (or better) than the original tagger, there is a bias in the setup favoring a particular kind of anomaly. Such a bias can be beneficial if it improves the tagging for a class of models of particular physics interest. However, the bias will limit the model-independence of the tagging capabilities and, hence, the usefulness of a given unsupervised machine learning method. It implies the existence of potential anomalies for which the tagging performance is diminished or the tagger may even fail completely. In machine learning applications outside of particle physics this kind of bias has recently been investigated mostly in the context of anomaly tagging using deep generative models [33–38]. Once an algorithm is known to work as an anomaly tagger in specific examples, it is a natural next step to study and understand potential biases in order to evaluate the performance and to improve the method.

In this work, we explicitly perform this investigation for an autoencoder which is used as a semi-supervised binary classifier to find an unknown signal within a Standard Model background. As a well-known benchmark example [39–44], we distinguish boosted top jets from QCD jets employing a simple convolutional autoencoder working on jet images. Although top-tagging is not a true use case for an anomaly tagger, it has been often employed as a test case and we confirm findings from the literature [30, 31] that an autoencoder can indeed find top jets as anomalies without having seen them during training. However, the ability to detect this particular anomaly does not imply that the autoencoder works in general. Indeed, in section 2 we will show that the very same autoencoder fails on the inverse task of finding QCD jets as anomalies when being exposed only to top jets during training; it performs worse than picking anomalies randomly. Our anomaly tagging setup employing the autoencoder is strongly biased to label top jets as anomalous no matter what the AE has seen during training. To understand this behavior we investigate what the autoencoder actually learns and why it fails under certain circumstances. It will become clear why a functional autoencoder can be a bad anomaly tagger and vice versa, in particular in the particle physics application at hand. Using those insights, we propose several improvements for our particular autoencoder architecture, such that it works as an anomaly tagger in both ways to distinguish QCD and top jets.

This work is structured as follows: after defining our setup in section 2.1, we scrutinize the autoencoder performance in sections 2.2 and 2.3, and discuss its tagging performance in section 2.4. In sections 3.1 and 3.2 we introduce improved data preprocessings and loss functions. We discuss the impact of these improvements on the autoencoder performance in general and on the tagging performance in particular in sections 3.3 and 3.4, respectively. We conclude in section 4 and present additional material in the appendices.

2 Autoencoder limitations

In this section, we introduce our AE architecture and investigate its training and performance. We either train the AE on a pure sample of QCD jets and call it a direct tagger, or we train the AE on a pure sample of top jets and call it an inverse tagger. While the former setup is designed to perform the well-known task of tagging top jets as anomalies, the latter setup is designed to perform the inverse task, i.e. tagging QCD jets as anomalies in a background sample of top jets. Hence, we always use a semi-supervised setup, i.e. we assume that a dataset consisting only of background data is available.

Before we discuss the tagging performance in section 2.4, we scrutinize how and to what extent the AE actually learns to reconstruct jet images during training and how these AE capabilities affect its performance as a model-independent tagger. In particular, we are able to explain the success of the direct tagger and the failure of the inverse tagger by the interplay between an insufficient AE performance and the different complexity in the images of the two jet classes.

2.1 Jet data and autoencoder architecture

Throughout this work, we use top tagging as an example, which is a well established benchmark in collider physics: simulated QCD jets (initiated by light-flavor quarks or gluons) have to be distinguished from hadronically decaying boosted top-quarks which form single jets. We use the benchmark dataset from ref. [45], which is publicly available at [46]. A subset of 100k and 40k jets of one of the two classes is used for training and validation, respectively, and two 40k subsets of each type for testing.

In the following, we describe our default setup which will be modified and improved in section 3. We employ the preprocessing introduced in ref. [42]. The four momenta of the jet constituents are converted into two-dimensional jet images with 40×40 pixels. The details of the jet generation and the preprocessing are discussed in appendix A. Figure 1 shows the average images of both classes. For QCD jets, most of the intensity is concentrated in the central pixels. For top jets, there is a clearly visible three-prong structure (as expected for top-quark decays after preprocessing). Of course, individual jets are harder to distinguish than their average images may indicate.

As our anomaly detection algorithm, we use a convolutional autoencoder with an architecture similar to the one in ref. [30]. We implement our AE with TENSORFLOW 2.4.1 [47], relying on the built in version of KERAS [48]. Several convolution layers with 4×4 kernel and average pooling layers with 2×2 kernel are applied before the image is flattened and a fully connected network reduces the input further into the bottleneck latent

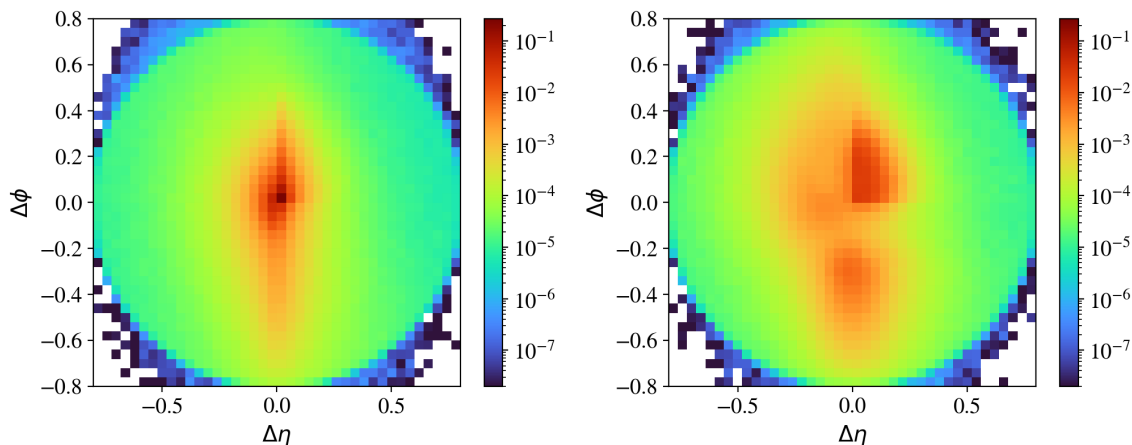


Figure 1. Average of 40k QCD (left) and top jet images (right) in the test dataset, according to the standard preprocessing as described in appendix A.

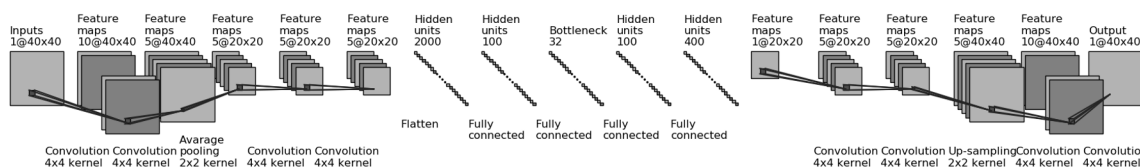


Figure 2. Architecture of our autoencoder, see also ref. [30].

space with 32 nodes. The Parametric ReLU activation function is used in all layers. The described encoder structure is inverted to form the corresponding decoder which is used to reconstruct the original image from its latent space description. Our architecture is defined in figure 2; the hyperparameter settings are described in more detail in appendix A.

Following ref. [30], to evaluate the reconstruction of the input picture we use the mean squared error (MSE), i.e. the average of the squared error of each reconstructed pixel with respect to its input value, as a loss function. During testing the value of the loss function is also used as the discriminator between signal and background. An event is tagged as signal/anomaly if the value of the loss function is larger than a given threshold. Changing the threshold value, one obtains the usual receiver operating characteristic (ROC) curve.

2.2 Limited reconstruction

We first investigate what the AE is actually learning, as it is trained for the reconstruction of the input and not as an anomaly tagger. Figure 3 shows the learning history of an exemplary top jet in terms of its reconstruction after training on top jets for a given number of epochs. To guide the eye, we also show the evolution of the squared error per pixel and highlight the reconstruction of the intensity of the brightest pixels in detail. Moreover, the MSE of the reconstruction, i.e. the loss function of the AE, is given as a measure for the overall reconstruction improvement.

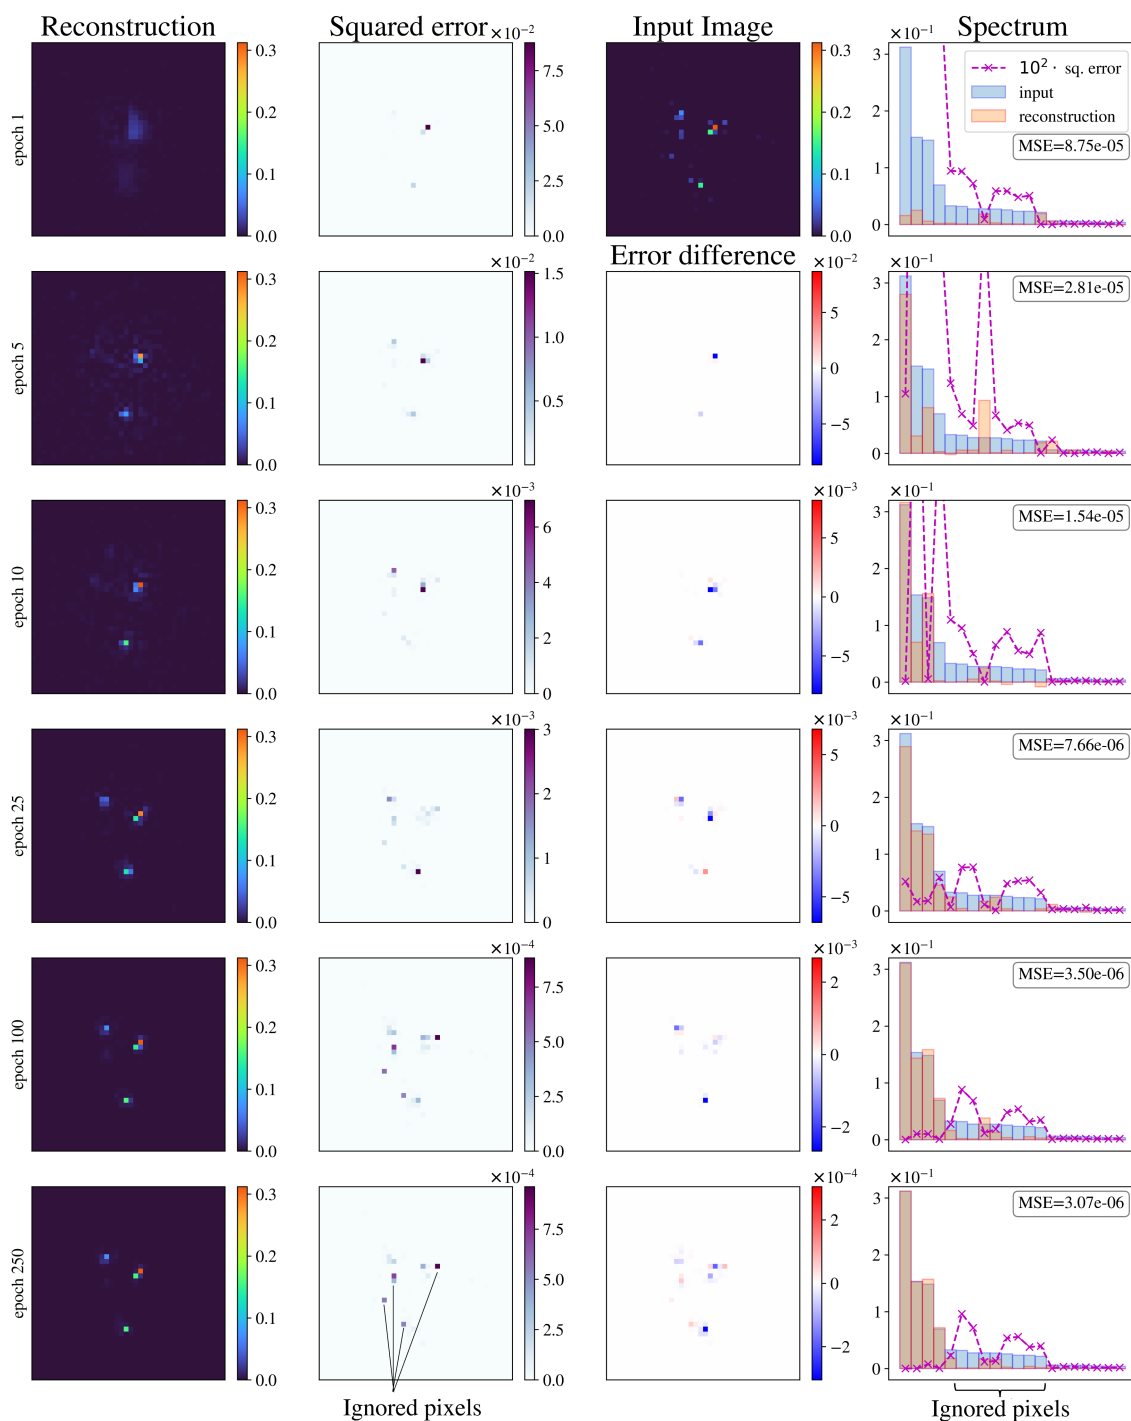


Figure 3. Reconstruction of an exemplary image (1st column) after 1, 5, 10, 25, 100, 250 (top to bottom) epochs of training. We also show the squared error per pixel between input and reconstructed image (2nd column) and its difference w.r.t. the previous row (3rd column). The 4th column shows the intensity of the 20 brightest input pixels (blue) together with the reconstructed intensity (orange) and the corresponding squared error (purple crosses).

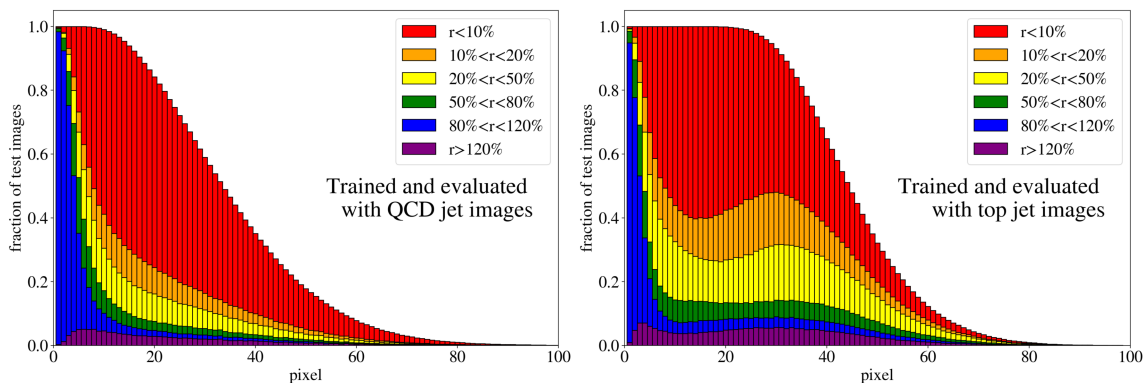


Figure 4. Stacked histogram for different categories of the ratio r of the reconstructed and the input intensity of the non-zero pixels. Pixels are ordered by intensity from left to right for each of the 40k test jets. We show results for both QCD jets (left) and top jets (right), where the AE has been trained on the corresponding training set.

During the first epoch, the AE learns to reconstruct the average top jet image, as can be seen by comparing the reconstructed image to the average image in figure 1. (Note that the images in figure 1 are shown on a logarithmic scale.) In the following epochs, the squared error is dominated by the brightest pixels. The AE improves the loss by improving their reconstruction, so its trainable weights are updated accordingly. After 10 epochs, the AE recovers a smeared reconstruction of the brightest pixels. After roughly 25 epochs, the weights are learned to reconstruct the brightest pixels so well that the corresponding error becomes small compared to the error of the remaining pixels. However, the AE further improves the reconstruction of the brightest pixels (and the surrounding zero-intensity pixels) instead of changing its focus to the dimmer pixels. The latter would probably harm the previously learned reconstruction of the brightest pixels and therefore increase the loss. After around 100 epochs the four brightest pixels are very well reconstructed, while most of the dimmer pixels are still ignored. The dimmer pixels dominate the total error, as can be seen in the right column of figure 3. The AE is apparently trapped in a local minimum of the loss function, and training longer does not change the picture significantly.

The previous discussion applies to both training on QCD and top images. To show that the example presented in figure 3 is rather generic, the reconstruction capabilities of the AE for the whole test dataset are summarized in figure 4, where we quantify the quality of the reconstruction for individual pixels of a jet image. Specifically, for each of the 40k jets of the test data sample, we determine the ratio of the reconstructed and input intensities, r , for each pixel, and show the fraction of test jet images where the brightest, next-to-brightest etc. pixel (from left to right on the horizontal axis) is reconstructed with a certain quality r . In figure 4 (left) this fraction is shown for the AE trained and tested on QCD jets. For the majority of the jet images, the brightest pixels are reconstructed well (blue histogram, corresponding to a ratio $80\% \leq r \leq 120\%$). On the other hand, the dimmer a pixel the more likely it will be reconstructed insufficiently, e.g. with an intensity of less than 10% of the input intensity (red histogram). Note that the overall number of

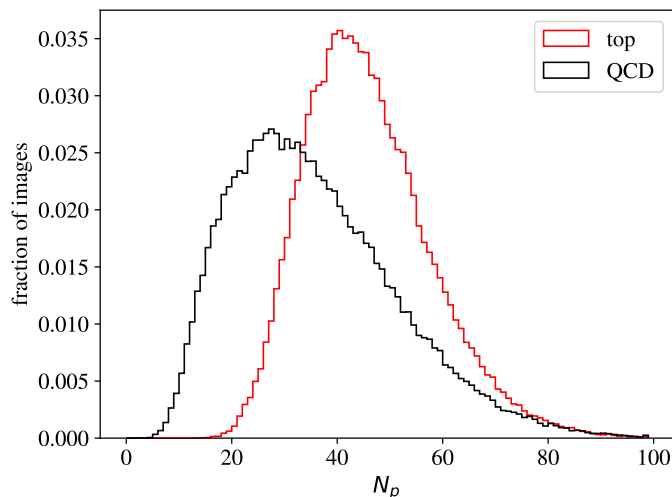


Figure 5. Distribution of the QCD and top jet images in the number of non-zero pixels N_p .

jets in the training data, and thus the fraction of jet images, decreases as one requires an increasing number of non-zero pixels. A qualitatively similar picture emerges for the AE trained and tested on top jets (see figure 4, right). In the next section, we will explore why such a limited AE can nevertheless tag top jets as anomalies in a QCD background.

2.3 Complexity bias

We can gain additional insight into the AE performance by investigating the reconstruction loss as a function of the number of non-zero pixels, N_p , of the jet images. We show the N_p distribution for QCD and top images in figure 5. When training on QCD images, the reconstruction loss of QCD images on average increases strongly with N_p , as shown in figure 6, left. QCD jets with small N_p are simply easier to reconstruct. When training on top images, top jets are also harder to reconstruct for larger N_p , see figure 6, right. However, in this case the correlation is less pronounced. Since, on average, top images have more non-zero pixels than QCD images, see figure 5, the top-trained AE sees less training data with small N_p . Given this correlation between the reconstruction loss and the number of non-zero pixels N_p , we conclude that N_p at least partially describes the complexity of the images.

However, complexity is not only related to N_p . For fixed N_p , QCD jets are on average also easier to reconstruct than top jets. This can be understood from the underlying physics. A top jet is naturally composed of three sub-jets initiated by the top-quark decay products. Hence, this intrinsic three-prong structure leads to more complex structures in the jet images.

These results show that the AE has a strong complexity bias: images which would intuitively be labeled as simpler, are reconstructed better. It has been discussed in the literature, especially for natural images, that simpler images may be harder to identify as anomalous [33–38, 49, 50]. In the context of natural images, it was noted that the

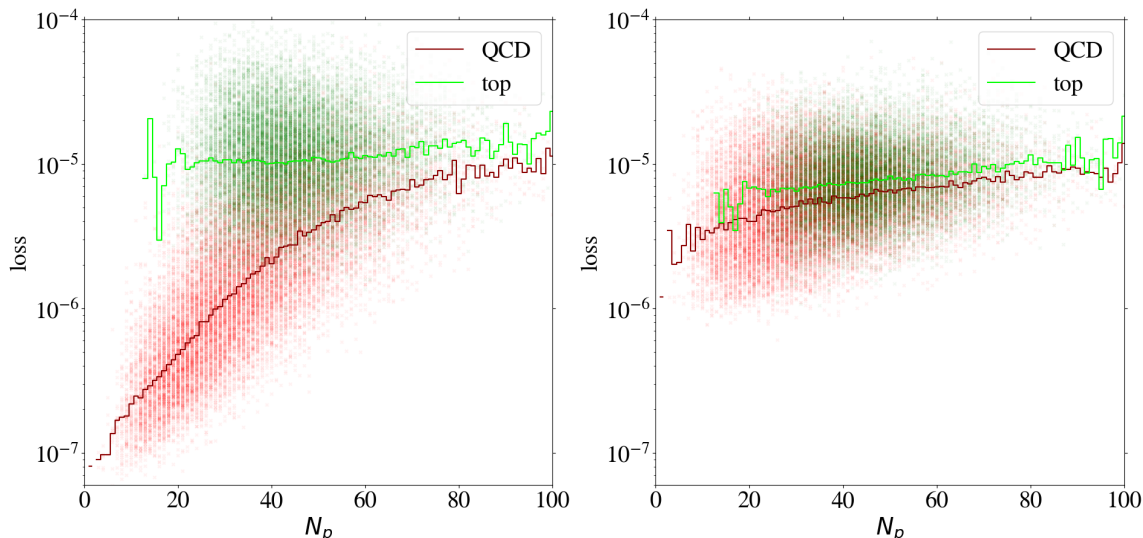


Figure 6. Reconstruction loss of individual test jets (only half of the test set is shown as points in the scatter plot) as a function of the number of non-zero pixels N_p , for an AE trained on QCD jets (left) and top jets (right). QCD (top) jets are shown in red (green). The solid lines show the median loss for a given number of non-zero pixels.

algorithm tends to learn features that are not representative of the specific training set, like local pixel correlations.

The strong complexity bias in our application might be mainly due to the limited reconstruction performance of the AE discussed in section 2.2. It is not surprising that the AE is not able to reconstruct the complex structure of top images when ignoring all but a few pixels. Limitations to the reconstruction of structures in high energy physics have also been noted recently in [51, 52].

The correlation between the complexity of an image and the reconstruction loss is dangerous if the AE is to be used as a model-independent anomaly tagger. One should not assume that a new physics signal looks more complex than the SM background. An anomaly tagger should only reconstruct well those images it has been trained on. On the other hand, our AE has the tendency to better reconstruct simple images. For a functional model-independent anomaly tagger the first effect should be much more pronounced than the second. For the case at hand, these two effects go in the same direction for the direct tagger, since it is trained on simple QCD images, but go in opposite directions for the inverse tagger, which is trained on more complex top images. This can be seen already in figure 6, where we also show the loss of top jets in the test set for the AE trained on QCD jets and vice versa. Indeed, when training on top jets, even QCD images with the same number of non-zero pixels are better reconstructed than top jets, due to the complexity issue. Although the inverse tagger has improved the top reconstruction through its training on top jet images, it cannot overcome the complexity bias. Given the limited reconstruction capabilities, after learning to reproduce top images, the inverse tagger has also learned to interpolate QCD images. On the other hand, the direct tagger is never exposed to the

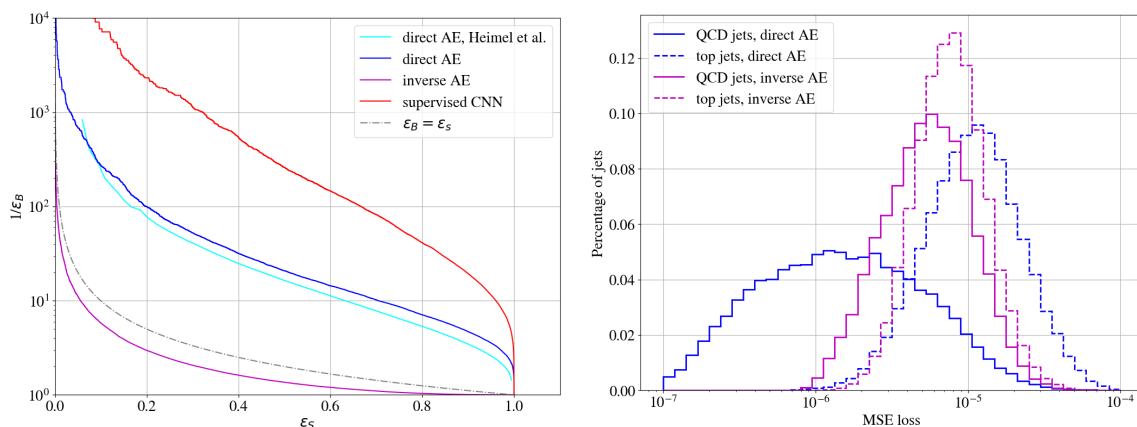


Figure 7. Left: ROC curve of our direct (dark blue) and inverse (purple) tagger. For comparison we also show the performance of a supervised CNN tagger (red), the results of ref. [30] for the direct tagger, and a random tagger (grey dashed). Right: loss distribution for QCD (solid) and top (dashed) images for the direct (blue) and inverse (purple) taggers.

complex features characterizing top images during training and thus fails to extrapolate to the corresponding top jet structures.

2.4 Tagging performance

Having investigated the limitations in the AE reconstruction power and the complexity bias, we show the performance of the direct and inverse taggers in figure 7. The ROC curves are constructed by varying the loss threshold for tagging a jet as anomalous. We also show the ROC curve for a supervised CNN tagger inspired by ref. [43] (see appendix A). Our direct AE tagger reproduces the results for the direct unsupervised AE tagger in ref. [30], which has used a similar setup. As expected, it performs worse than a supervised tagger, but much better than random guessing. A similar performance for the unsupervised AE tagger has also been obtained in ref. [31] in a slightly different setup. Hence, our direct tagger works as expected.

As can be seen from figure 7, the inverse tagger performs worse than randomly tagging jets as anomalous. Our inverse autoencoder fails to tag an anomaly (QCD jets in this case) that is simpler than the background. Through its training on top jets, the inverse tagger learns to reconstruct top jets better than the direct tagger (see right plot in figure 7), and its performance for the reconstruction of QCD jets is diminished. However, this is not enough to overcome the complexity bias.

To summarize, even with a limited reconstruction capability, an autoencoder can be a good anomaly tagger if there is a bias to reconstruct the background better. However, if the bias works against anomaly detection, the learning capabilities of the AE may not be sufficient to overcome the bias. Only a powerful AE with a background specific data compression in the latent space could potentially be able to overcome such a bias. Improvements of our setup to partly achieve these goals are discussed in section 3.

It should be noted that a perfect AE, which is always able to perfectly reconstruct the input via the identity mapping regardless of the input data, would be useless as a tagger as it would always interpolate perfectly from the learned data to the anomalies.

3 Improving the autoencoder performance

Given the limited performance of the AE setup described in section 2.1, we investigate possible improvements. One approach would be to change the AE architecture. There are unlimited possibilities which are worth investigating. However, here we want to point out some generic improvements concerning the complexity bias and the limited learning capabilities which might be helpful for any AE architecture. These improvements are introduced in sections 3.1 and 3.2, and their impact on the AE performance and tagging capabilities is quantified in sections 3.3 and 3.4, respectively.

3.1 Intensity remapping

As shown in section 2.2, the AE focuses on reconstructing only the brightest pixels of an image. The contrast between bright and dim pixels is apparently so large that the dim pixels are not recognized during the early learning epochs. To tackle this problem, we propose several remappings of the pixel intensities, such that the intensities of the dim pixels are enhanced compared to those of the bright pixels. Such a remapping should encourage the AE to ignore less pixels and learn more of the intrinsic structure and features of the given training images. Hence, in particular the complexity of the simplest pictures (having only one or two very bright pixels) is increased, and we expect to also reduce the complexity bias. The intensity remapping can be considered as part of an alternative preprocessing.

Due to the normalization of the input features in the standard preprocessing described in appendix A, all pixel intensities lie in the interval between 0 and 1. To highlight dim pixels by remapping this interval onto itself, we use the functions

$$R_1(x) = \sqrt{x}, \quad R_2(x) = \sqrt[4]{x}, \quad R_3(x) = \frac{\log(\alpha x + 1)}{\log(\alpha + 1)} \quad \text{and} \quad R_4(x) = \Theta(x), \quad (3.1)$$

which are displayed in figure 17 in appendix B. For the mapping R_3 we use $\alpha = 1000$; the Θ -function maps each non-zero pixel to one. The original images correspond to the identity remapping $R_0(x) = x$ and are labeled accordingly in the following. After the remapping the image is normalized again to have a pixel sum equal to one. While remappings R_1 to R_3 are invertible, information is lost through R_4 , which should be thought of as a limiting case. The effects of the remappings are exemplified and quantified further in appendix B.

3.2 Kernel MSE

As discussed in section 2, it is difficult for the AE to learn the structure of an image. Instead, it only learns the exact position of the few brightest pixels. This is actually not surprising, since the AE is not rewarded by the loss function to reconstruct the image structure. It rather has to reproduce the exact pixel position, which seems to be an unnecessarily tough task. For example, shifting the input image by one pixel would be considered by a human as a valid reconstruction, but is strongly penalized for the standard MSE loss.

In fact, the similarity of two images is only poorly represented by the MSE loss. To compare two distributions, i.e. the intensity distribution in our case, the Wasserstein distance, also called energy or earth mover’s distance [53], would be a much better choice. However, its full calculation includes an optimization problem that would need to be solved for each image during training. Approximate estimates of this distance, e.g. the sliced Wasserstein distance [54], are also computationally expensive. We thus propose a simple loss function which nevertheless provides a better notion of similarity between two pictures than the MSE loss.

To introduce a notion of neighborhood for each pixel in the loss function we convolve the whole image with a smearing kernel. Hence, neighboring pixels have a partly overlapping smeared distribution which can be recognized by the loss function. The optimal reconstruction, i.e. vanishing loss, is still achieved by reconstructing the original (unsmeared) input image.

On the discrete set of pixels in each image the convolution is defined as

$$S_{ij} = \sum_{k,l} K_{ij,kl} I_{kl},$$

where I_{kl} denotes the intensity of the pixel with coordinates (k, l) of the original image, and S_{ij} refers to the smeared pixels. The kernel is defined as

$$K_{ij,kl} = \frac{L+1}{L} \left(\frac{1}{1 + \sqrt{(i-k)^2 + (j-l)^2}} - \frac{1}{1+L} \right)$$

with the additional constraint that it is set to zero for negative values. Hence, it is only non-zero for an Euclidean distance to the center pixel which is smaller than L pixels and has an approximately circular shape. We choose $L = 8$. The normalization is chosen such that the intensity of the central pixel is unchanged. To avoid boundary effects, the original picture is padded with zeros on each side and becomes 54×54 pixels. There is an infinite amount of possible choices for this kernel. However, we do not expect the following results to depend on its details. The reconstruction loss is defined as the MSE of the smeared input and the reconstructed image. We refer to it as kernel MSE or KMSE in the following. Since the smearing is a standard matrix convolution its computational cost is negligible for training and testing.

3.3 Autoencoder performance

To evaluate the effect of the intensity remapping on the autoencoder performance, we again investigate the distribution of the ratio r of the reconstructed intensity and the input intensity for the leading pixels. The corresponding histogram, i.e. the equivalent of figure 4, is shown for the R_2 -remapping in figure 8. When training on QCD jets (left figure), the fraction of well-reconstructed pixels is increased by the remapping. This suggests that the AE indeed learns more of the jet features. When training on top jets (right figure), the fraction of well-reconstructed jets is not increased by much, but it extends to more of the brighter pixels, also signaling that more features are potentially accessible. The decrease of performance beyond the 20th pixel (and in particular the disappearance of the peak)

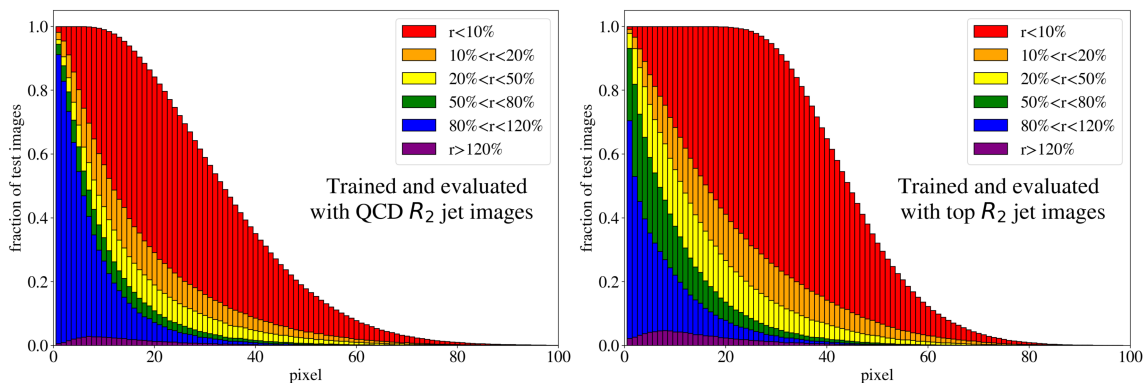


Figure 8. Same as figure 4 but using the R_2 intensity remapping.

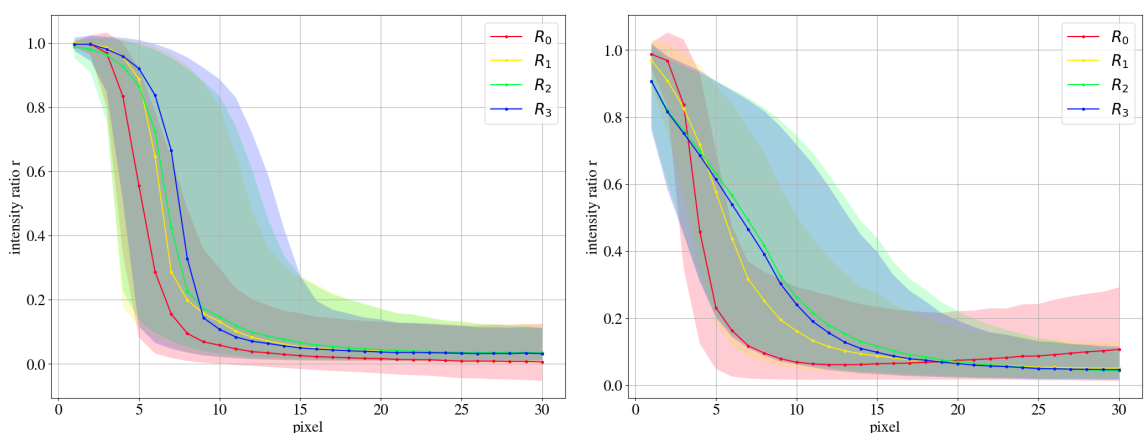


Figure 9. Median and bands for the 25%/75% quantiles for the distribution of r for the leading pixels, where r is the ratio of the reconstructed and the input intensity of a given pixel in a given jet. Results are shown for QCD jets (left) and top jets (right). The AE is trained on the respective training set with a given remapping.

compared to the performance without remapping, figure 4 (right), is not discouraging, since these pixels correspond to very soft particles.¹

In figure 9, we compare the reconstruction quality for the different remappings. The distribution of the intensity ratio r (displayed in figures 4 and 8 by stacked histograms) is represented by its median and the 25%/75% quantiles. As long as the distribution is dominated by values of r close to one, the leading pixels are reconstructed well. For both the QCD and the top case, the median curves shift to the right for the remappings, i.e. more of the leading pixels are on average taken into account by the AE. Hence, the AE ignores less pixels and thus learns more features of the jet images. For the AE trained and tested on top jets, right figure, this is only achieved by reducing the reconstruction precision of the brightest pixels.

¹They are only accidentally well reconstructed for R_0 by a tiny non-vanishing intensity which is typically found in the reconstructed image in the vicinity of bright pixels. Since the R_2 -remapping highlights their intensity, this mechanism does not work any more.

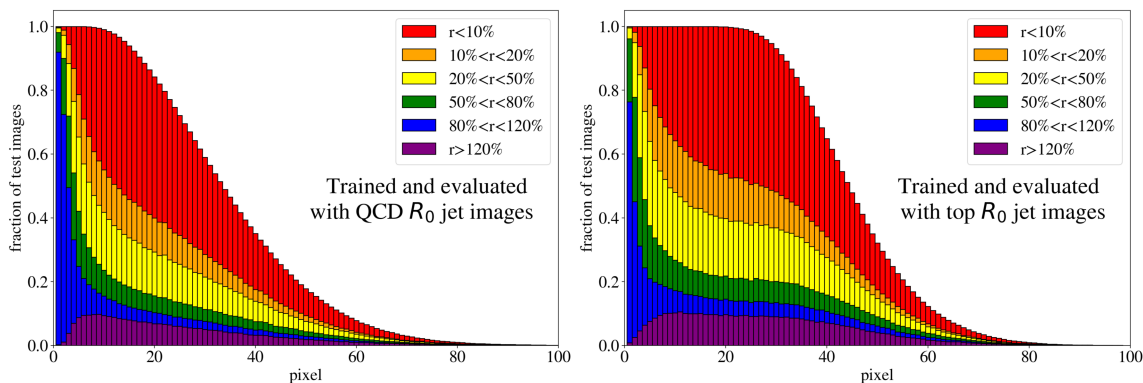


Figure 10. Same as figure 4 but using KMSE as the loss function during training. No smearing kernel is applied for the data shown during testing.

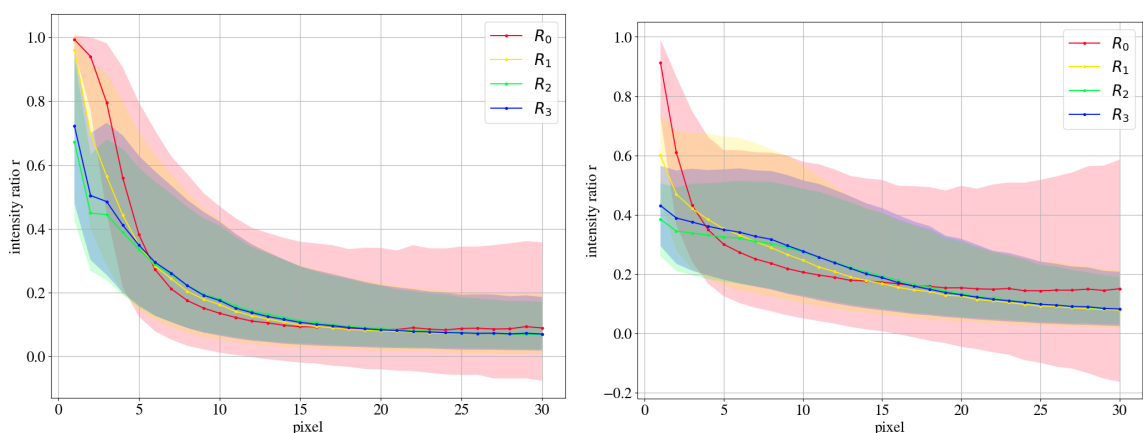


Figure 11. Same as figure 9 but for training with KMSE loss. Results are shown for QCD jets (left) and top jets (right).

In figure 10, we show the reconstruction performance for training on R_0 images using the KMSE loss introduced in section 3.2. Note that the KMSE loss function is only used during training. To evaluate the pixel reconstruction in the plots no smearing kernel is applied. As one expects, the exact intensity reconstruction of the leading pixels is traded for paying more attention to dimmer pixels. The strong focus to reduce the squared error of the brightest pixels as much as possible is removed. The same can be seen in figure 11 where we again show the distributions of the intensity ratio r in terms of the median and quantiles for the different remappings. Although the brightest pixels are not very precisely reconstructed, dim pixels often have a significant part of their intensity reconstructed. Also when training on top jets using the KMSE loss functions, the remappings lead to more pixels being taken into account.

The stacked histograms for all remappings with and without KMSE loss function, which contain additional information on the r -distribution, are displayed in appendix C.

The reconstruction of top jets (when training on top jets) with the R_0 and R_2 remappings using the standard MSE and the KMSE loss functions is illustrated in figure 12.

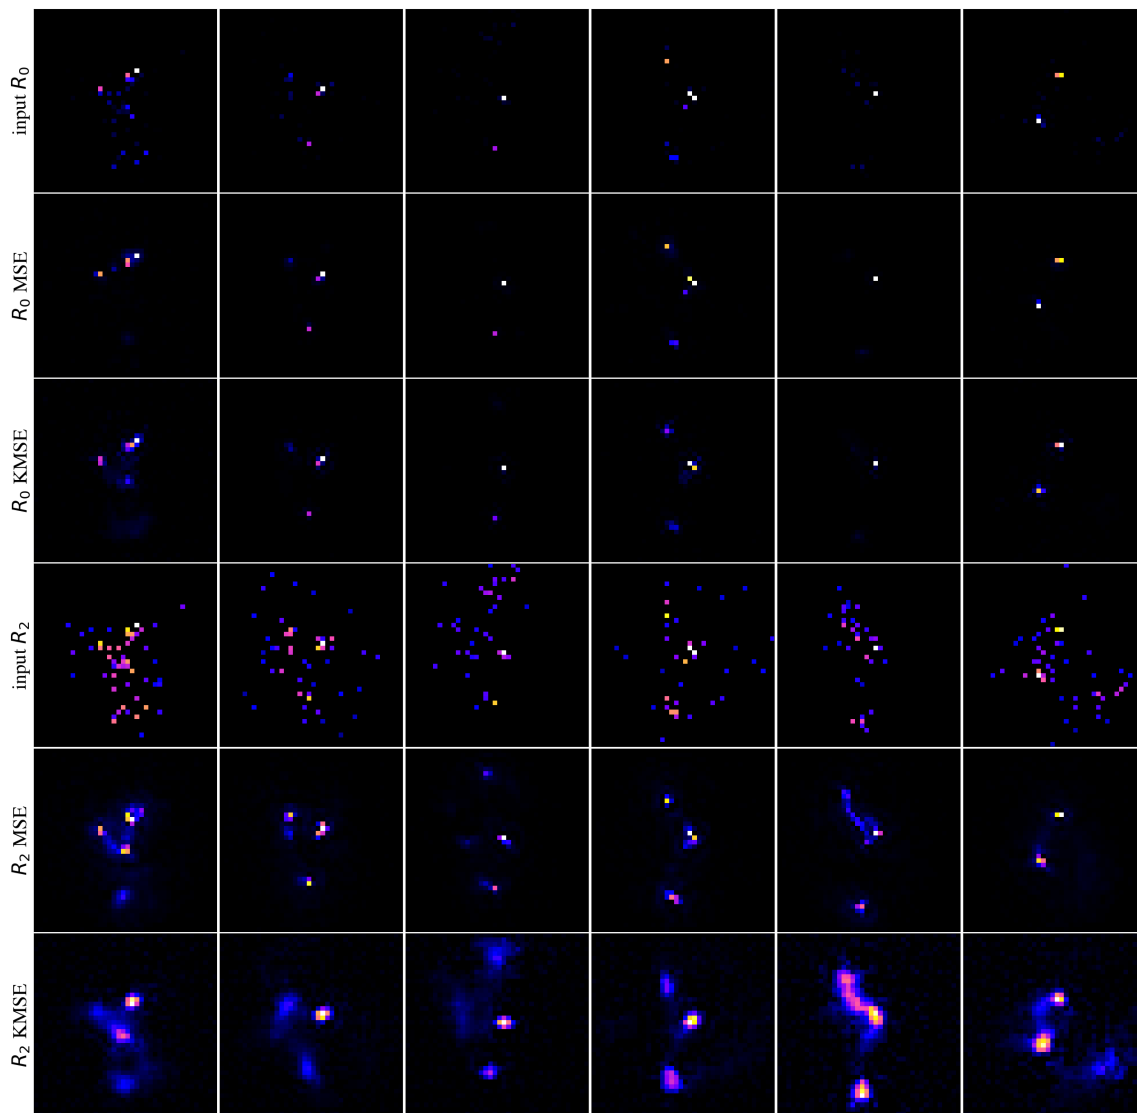


Figure 12. Six top jet input R_0 images (first row) and the corresponding R_2 images (fourth row) together with the AE reconstructions using the MSE (second and fifth row, respectively) or KMSE loss (third and sixth row, respectively). Each image is individually normalized to have the same maximum pixel intensity. Small negative pixel intensities in the reconstructions are set to zero.

Using the KMSE loss, the AE focuses more on the reconstruction of a continuous distribution of the intensity than on the reconstruction of individual pixels. Even relatively dim regions far from the center receive some attention. Thus, as also shown in figure 11, the dim pixels are partially reconstructed as part of this continuous distribution. Therefore, we expect the KMSE autoencoder to extract more useful features of the overall jet structure compared to the standard AE that focuses almost exclusively on the brightest pixels.

To understand the implications of the intensity remapping with respect to the complexity bias, we show the loss of R_2 images as a function of the number of non-zero pixels N_p in figure 13 for the AE trained on QCD (left figure) or top images (right figure) using

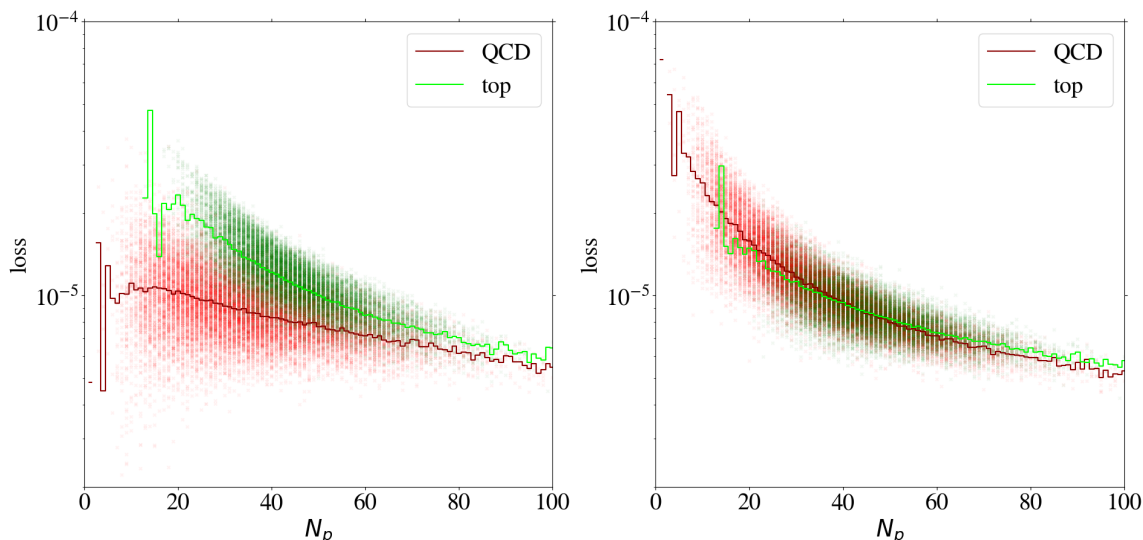


Figure 13. Reconstruction loss of individual test jets (only half of the test set is shown as points in the scatter plot) as a function of non-zero pixels for an AE trained on QCD jets (left) and top jets (right) using the R_2 intensity remapping. QCD (top) jets are shown in red (green). The solid lines show the median loss for a given number of non-zero pixels.

the MSE loss function. Comparing to figure 6, the bias for the QCD-trained AE is reduced a lot and even reversed in direction. For the AE trained on top jets, the bias is also reversed, but additionally increased. This can be understood since the inverse AE sees only few of the training images with small N_p , the complexity of which has been increased by the intensity remapping.

Training on QCD images, the AE still learns to reconstruct images with low pixel number rather efficiently. Looking at the reconstruction of top images by the direct tagger in the same plot, it still cannot extrapolate well to this unseen data, so that the direct tagger will perform well (see next section). Training on top images with R_2 remapping, the AE has lost its ability to simply interpolate the QCD images with small N_p , because their complexity has increased w.r.t. the R_0 case. Hence, this AE is a working inverse tagger. However, comparing the average loss for top and QCD jets in the right panel of figure 13, only in a region below 40 active pixels, top images are on average reconstructed better than QCD images. Moreover, the difference is small compared to the width of the distribution and the slope of the bias. Most of the inverse tagging performance is based on the QCD images with small N_p , due to the reversed bias and still not due to a well-performing AE.

Using the KMSE loss, figure 14 shows that the loss as a function of N_p for the direct tagger remains very similar to what we have seen in the MSE case. However, we observe that QCD jets with low N_p are not automatically learned to be reconstructed well when training the inverse tagger on top images. Moreover, for $N_p < 35$ top images are again reconstructed on average better than QCD images. Using both a remapping and the KMSE loss function for training, the behavior gets more pronounced. The corresponding distributions are shown in figure 19 in appendix C for all combinations of remappings and loss functions.

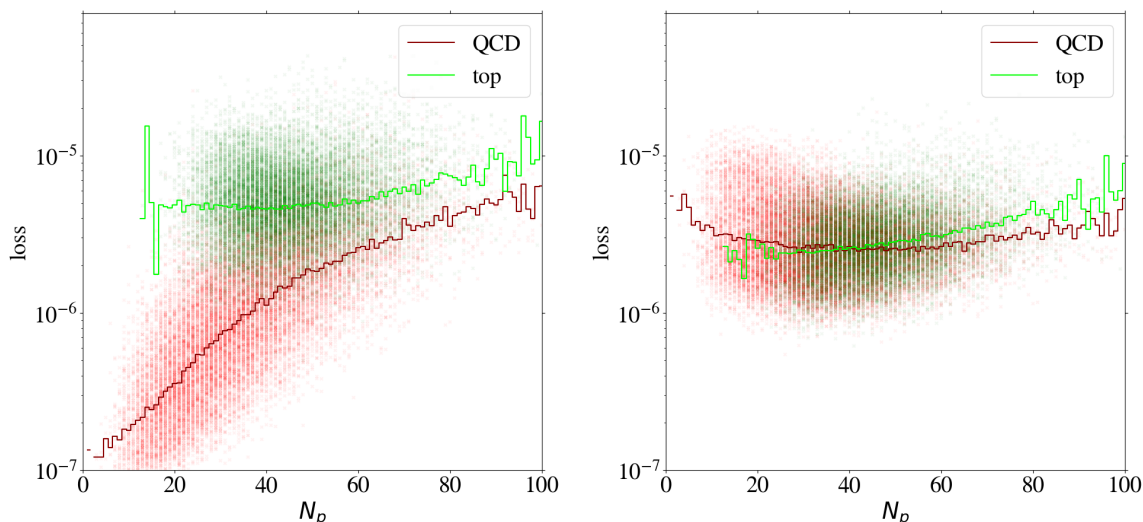


Figure 14. Same as figure 6 but for KMSE instead of MSE loss (the R_0 -remapping is used).

3.4 Tagging performance

The performance for the direct and inverse taggers using the different intensity remappings introduced in section 3.1 are summarized in terms of their ROC curves in figure 15. Derived performance measures are shown in table 1 in appendix C. For the direct tagger, QCD jets are background ($\epsilon_B = \epsilon_{\text{QCD}}$) and top jets are signal ($\epsilon_S = \epsilon_{\text{top}}$), so that the area under the curve plotted in figure 15 directly corresponds to the AUC. For the inverse tagger, signal and background interchange their role ($\epsilon_S = \epsilon_{\text{QCD}}$ and $\epsilon_B = \epsilon_{\text{top}}$) and the area under the plotted curves should be as small as possible, since it corresponds to 1-AUC.

Plotting the ROC curves in this way, the area between the ROC curves of the direct and the inverse tagger, called ABC in the following, can be interpreted as the background specific learning achieved by the AE. Furthermore, the difference between the two AUC values for the direct and the inverse tagger, called ΔAUC in the following, is a measure for how biased the AE is. For an AE with small ABC or large ΔAUC one cannot expect a model-independent tagging capability. On the other hand, a large ABC and a small ΔAUC do not guarantee model-independence. It is only an indication that a step towards model-independence has been made.

Using the standard preprocessing without intensity remapping (R_0), we find a small $\text{ABC} = 0.28$ and a large $\Delta\text{AUC} = 0.55$. As discussed in section 2.4, the direct tagger only works due to a large complexity bias, and the inverse tagger fails.

For R_1 -remapping we see a minor decrease in the direct tagging performance, but a significant improvement in the inverse tagging performance with $\text{ABC} = 0.48$. This improvement is strong enough to enable inverse tagging although $\Delta\text{AUC} = 0.35$ is still large. In contrast, R_2 -remapping shows little bias ($\Delta\text{AUC} = 0.06$) and an $\text{ABC} = 0.43$. Both direct and inverse tagging have a modest AUC close to 0.7. R_3 has a little more bias ($\Delta\text{AUC} = 0.18$) but the best $\text{ABC} = 0.50$. R_4 -remapping shows that too much highlighting of dim pixels is also counterproductive. It leads to $\Delta\text{AUC} = -0.3$ and a very poor $\text{ABC} = 0.11$.

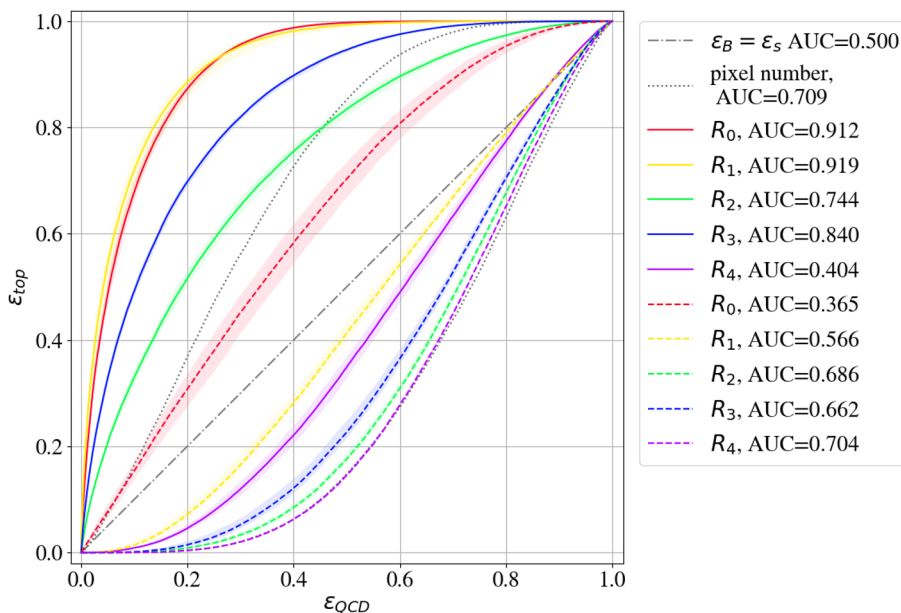


Figure 15. ROC curves of the direct (solid) and inverse (dashed) taggers using the different intensity mappings R_0 to R_4 . The dotted lines represent the ROC curves of the two taggers based on the number of the non-zero pixels in the image.

Although these results are a promising first step towards a more model-independent anomaly tagging, we want to recall that the inverse tagging is mainly caused by the correlation of the AE loss with the number of non-zero pixels N_p . The AE is still not able to reconstruct a top jet on which it has been trained better than a QCD jet if both have the same N_p . Moreover, in particular the tagging performance of the inverse taggers is still very poor. To illustrate this, we also show the two taggers which only use the distribution of N_p for both classes, as shown in figure 5, to tag anomalies with either a large or a small N_p . None of our inverse taggers can beat the performance of the trivial tagger looking for few non-zero pixels. This shows once more that the inverse tagging for the intensity remappings mostly relies on the correlation of the loss with the number of pixels.

Figure 16 shows the ROC curves for direct and inverse taggers using the KMSE loss and the different intensity remappings (see also table 2 in appendix C for additional performance measures). The KMSE loss is also used as the anomaly score. On the one hand, the direct R_0 -tagger using the KMSE loss performs similarly to the MSE one. On the other hand, also the KMSE autoencoder is not able to provide a working inverse tagger, although the performance of the inverse tagger is significantly better than the AE with MSE loss. With $ABC = 0.41$ and $\Delta AUC = 0.44$ it is less biased towards reconstructing QCD better.

For the intensity remappings, we find $ABC = 0.52$ and $\Delta AUC = 0.27$ for R_1 , hence the KMSE loss improves the overall performance. The remapping R_2 results in $ABC = 0.27$ and $\Delta AUC = -0.09$, i.e. the bias is still small but the performance is reduced. The remapping R_3 shows almost no bias ($\Delta AUC = 0.02$) but also reduced $ABC = 0.32$. The remapping R_4 still fails completely. In contrast to the MSE loss, using KMSE the inverse taggers

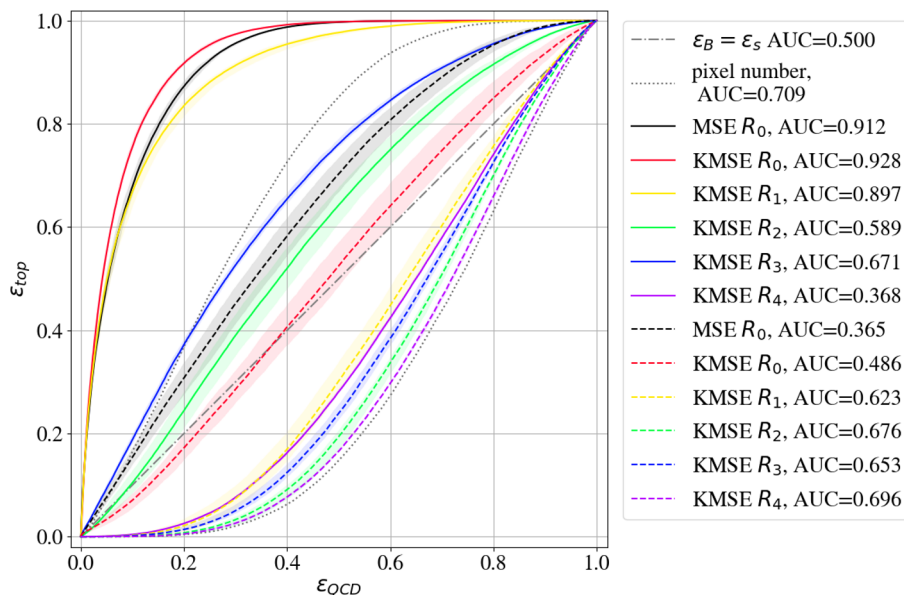


Figure 16. ROC curves of the direct (solid) and inverse (dashed) taggers using the different intensity mappings R_0 to R_4 for the KMSE and R_0 for the MSE loss function. The dotted lines represent the ROC curves of the two taggers based on the number of the non-zero pixels in the image.

based on R_1 to R_4 all show a performance very close to the pixel number tagger. While we have investigated the ability of the KMSE-based AE to reconstruct more jet features in section 3.3, the tagging performance shows that these features are not necessarily correlated with only one of the two classes.

4 Conclusion

In this work we have investigated a specific autoencoder architecture, based on a convolutional neural network, for tagging top jets in a background of QCD jets — and vice versa — using the reconstruction loss as the anomaly score. We observe a rather limited performance of the autoencoder with respect to its image reconstruction capabilities, see section 2.2. Nevertheless, we confirm findings from the literature [30, 31] that such an autoencoder is a powerful top tagger. These apparently contradicting observations are a consequence of a strong complexity bias of the autoencoder-based tagger as discussed in section 2.3. The sparsity of the jet images in combination with the underlying physics of QCD and top jets allow for a more successful reconstruction of the QCD jet images with respect to the mean squared error of all pixels, no matter which training data (i.e. QCD jets or top jets) are used.

Having made this observation in a very specific benchmark scenario, it is nevertheless quite obvious that it might generalize to true anomalies, other autoencoder architectures and most probably also to some applications with completely different data. Similar observations have been recently made for anomaly detection in natural images [33–38]. A

poor autoencoder can be a good anomaly tagger if there is a strong bias which favours the reconstruction of the background data for a given anomaly example. Hence, a good tagging performance for a specific example does not imply a functional model-independent unsupervised tagger. On the other hand, a perfect autoencoder is useless as a tagger if it can also interpolate to reconstruct anomalies which have not (in the semi-supervised case) or rarely (in the unsupervised case) been seen during training.

As the top-tagging example shows, a biased autoencoder is not necessarily a bad thing. If, for example, one is only interested in anomalies with more complex jet structures, the bias helps tagging anomalies, and the setup is certainly also robust in an unsupervised approach working with a training sample including signal contamination [30, 31]. However, model-independence will always be lost to a certain extent. For example, dark-matter jets from a strongly interacting sector are impossible to tag as anomalies in this case, as discussed in [55, 56].

How model-independent an autoencoder based tagger is, is difficult to answer. After all, the autoencoder tagger should exactly be designed to find the sort of anomaly which is not a priori thought of. One way to approach the problem is to test the tagger on as many physics cases as there are available. However, this might be a tedious task, if many physics cases are at hand or if they would have to be designed for this purpose. We propose to at least evaluate the autoencoder by inverting a given test task, i.e. interchanging the role of background and signal and investigating the corresponding performance. Possible performance measures are the difference of the corresponding AUCs or the area between the ROC curves in a plot like figure 15 in section 3.4.

Having understood at least some of the limitations of the original setup, we have investigated two modifications to improve the autoencoder. An intensity remapping for the jet images, as introduced in section 3.1, and the kernel MSE loss function, as introduced in section 3.2, are designed to help the autoencoder learn more relevant features within the jet images. Our modifications are shown to be promising first steps, both with respect to the autoencoder performance (section 3.3) and with respect to the model-independence of the tagging performance (section 3.4). However, the true progress is hard to quantify and better performance measures are needed in the future. Moreover, the intensity remapping is also affecting the bias of the autoencoder for reconstructing one of the two classes better than the other, irrespective of the training data. In particular, the inverse tagging, i.e. finding QCD jets in a top jet background, is to some extent due to a reversed bias, which is helpful for this particular test case but cannot be claimed as a model-independent advantage.

Future directions include the investigation of improved autoencoder architectures which we have not touched upon at all in this work. It would also be interesting to explore how anomaly detection based on the latent space of (variational) autoencoders or even completely different architectures is impacted by the complexity bias or other related biases when trained on the sparse jet images.

To summarize, this work provides valuable insights into the interplay of the autoencoder performance and anomaly tagging as well as first steps for improvements. However, we want to stress that a powerful and truly model-independent autoencoder for unsupervised anomaly tagging on jet pictures still needs to be developed.

Note added. The submission of this paper has been coordinated with [57, 58] and with [59], which address the challenges of using autoencoders for anomaly detection in complementary ways.

Acknowledgments

We are grateful to Anja Butter, Tilman Plehn and the participants of the Machine Learning mini workshops organized within the Collaborative Research Center TRR 257 “Particle Physics Phenomenology after the Higgs Discovery” for useful discussions. This work has been funded by the Deutsche Forschungsgemeinschaft through the CRC TRR 257 under Grant 396021762 — TRR 257 and the Research Training Group GRK 2497 “The physics of the heaviest particles at the Large Hadron Collider” under grant 400140256 — GRK 2497. Simulations were performed with computing resources granted by RWTH Aachen University.

A Jet simulation and neural network architecture

The benchmark dataset from ref. [45] is publicly available [46]. The jets are obtained for a center-of-mass energy of 14 TeV using PYTHIA8 [60] and fast detector simulation with the default ATLAS detector card of DELPHES [61]. Multiple interactions and pile-up are ignored. The jets are clustered with FASTJET [62] using the anti- k_T algorithm [63] with a jet radius of $R = 0.8$. Moreover, the jets fulfil $p_T \in [550 \text{ GeV}, 650 \text{ GeV}]$ and $|\eta| < 2$. For top jets, a parton-level top and its decay partons are required within $\Delta R = 0.8$ of the jet axis. The 200 jet constituents leading in p_T are stored. Jets with less constituents are padded with zeros.

According to the preprocessing in ref. [42], the jets are first centered in the η - ϕ -plane. Afterwards, they are rotated such that the principle axis is vertical and finally mirrored along both axes to obtain the hardest component in the first quadrant. After these steps, the jets are converted into two-dimensional jet images with 40×40 pixels corresponding to a two-dimensional transverse-momentum histogram of the jet constituents in the rotated $\Delta\eta$ - $\Delta\phi$ -plane. Finally, each jet image is divided by its total transverse-momentum, i.e. the pixel intensities sum to one.

Here, we explain in more detail the AE architecture given in figure 2. Each convolutional layer has a stride of 1 in both directions and a 4×4 kernel. Padding is performed in all convolutional layers such that the dimension of the output is the same as of the input. First, we use two convolutional layers with 10 and 5 filters for feature extraction. We reduce the dimension using average pooling with stride 2 and a 2×2 kernel, and add two more convolutions with 5 filters each. The feature map of the last convolution is flattened into a vector of 2000 nodes that are fully connected to 100 nodes in the next layer, which in turn are fully connected to the bottleneck layer with 32 nodes. These layers together form the encoder part of the AE. For the decoder we first add two fully connected layers with 100 and 400 nodes, respectively. The output of the latter is reshaped into a 20×20 feature

map. Afterwards, we perform two convolutions (with 5 filters each) followed by a 2×2 up-sampling layer to restore the dimensions of the image. We complete the decoder by adding three more convolutional layers with 5, 10 and 1 filter, respectively, resulting in a 40×40 dimensional output, matching the dimensions of the input. We apply the parameterized ReLU activation function in all hidden layers (convolutional and fully connected) with the corresponding α parameter initialized from a uniform distribution in the interval $[-1, 1]$. For the final layer, i.e. the last convolutional layer resulting in the AE reconstruction, no activation function is applied.

The training is done on the designated training set consisting of 100k images using the ADAM optimizer [64] and a batch size of 500. As we observe a saturation of the reconstruction loss for this amount of training data, we conclude that it is sufficient. After each epoch of training, validation is performed on 40k images. Early stopping is employed if no improvement of the validation loss is achieved in 10 consecutive epochs. Early stopping always terminates the training before the maximum of 1200 epochs is reached, and we use the weights of the epoch with the lowest validation loss for testing. To improve the training with the KMSE loss function, we reduce the learning rate to 10^{-4} and the batch size to 64.

To establish an upper limit for the unsupervised learning and to test our framework, we first use a supervised approach, i.e. we apply a convolutional neural network to this classification problem. The architecture of the CNN is based on the one used in [43]. The corresponding ROC curve in figure 7 is obtained by training for 100 epochs on 100k QCD and 100k top jet images, validating on 40k QCD and 40k top images and using the model with the lowest validation loss. The performance of our CNN is comparable to the results from the literature [43, 65].

B Intensity distributions after remapping

In this appendix we exemplify and quantify the effects of the intensity remapping as introduced in section 3.1. The remapping functions,

$$R_1(x) = \sqrt{x}, \quad R_2(x) = \sqrt[4]{x}, \quad R_3(x) = \frac{\log(\alpha x + 1)}{\log(\alpha + 1)} \quad \text{and} \quad R_4(x) = \Theta(x), \quad (\text{B.1})$$

are displayed in figure 17. The identity mapping, corresponding to the original images, is denoted as $R_0(x) = x$. The strength of the highlighting of the dimmer pixels is different for the different functions. Moreover, the derivative of the logarithmic remapping, R_3 , does not diverge at the origin, in contrast to the root functions.

In the left column of figure 18, we show the effect of the remapping for an exemplary top jet. The dim region in the bottom right of the R_0 image (upper panel) may be the third sub-jet in the top jet structure. Without highlighting it is barely visible and will most likely be ignored by the AE. By highlighting these dim pixels, the AE can more easily learn the information and features stored in them. Learning more distinctive features of the training set, the AE is also expected to perform better at distinguishing anomalies.

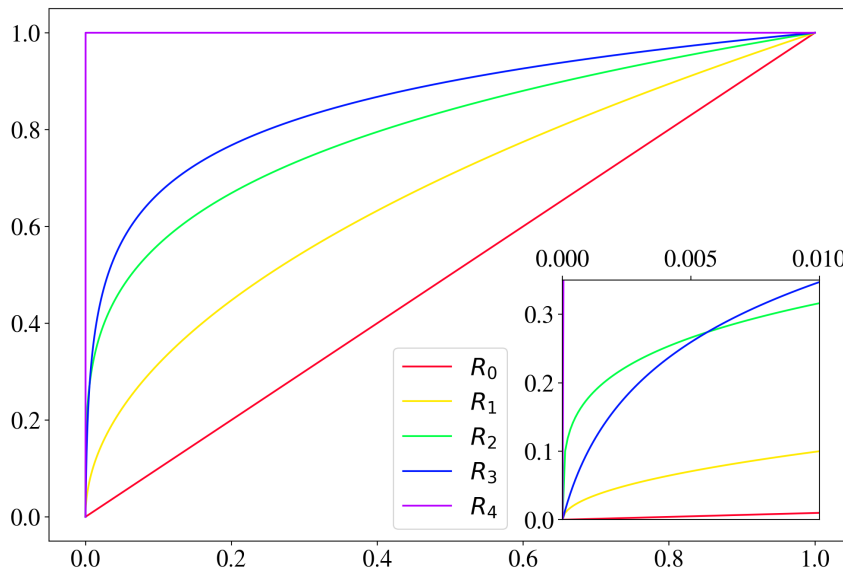


Figure 17. Remapping functions for highlighting dim pixels, see eq. (B.1).

To quantify the effect of the highlighting on the whole set of images, we show the intensity distribution of the non-zero pixels for the 100k top jet images of the training set in the right column of figure 18. In the original images (upper panel) the distribution of the intensities of the non-zero pixels is rather broad, with a median that is more than an order of magnitude below the average value. The intensity remapping compresses the distribution on a logarithmic scale. The distributions after the R_1 and R_3 remapping span across around 1.5 orders of magnitude, and the distribution after the R_2 remapping lies nearly fully within one order of magnitude. Due to the normalization to sum to one, the distribution of pixel intensities for the R_4 -images reflects the distribution of the overall number of non-zero pixels. Note that the remappings do not change the number of non-zero pixels. Because of the final normalization, the average intensity of the non-zero pixels is also unchanged.

C Further results

In this appendix, we complete the display of the results discussed in the main part of this work.

We first collect the performance measures for the direct and inverse taggers (see section 3, figures 15 and 16) in tables 1 and 2. The measures E10 and E100 are defined as $\epsilon_S(1/\epsilon_B = 10)$ and $\epsilon_S(1/\epsilon_B = 100)$, respectively.

In figures 19 and 20 we finally show the reconstruction loss for individual jet images as a function of the number of non-zero pixels and the distribution of the ratio of reconstructed and input intensity for the leading pixels, respectively, for the different remappings and for the MSE and KMSE loss functions.

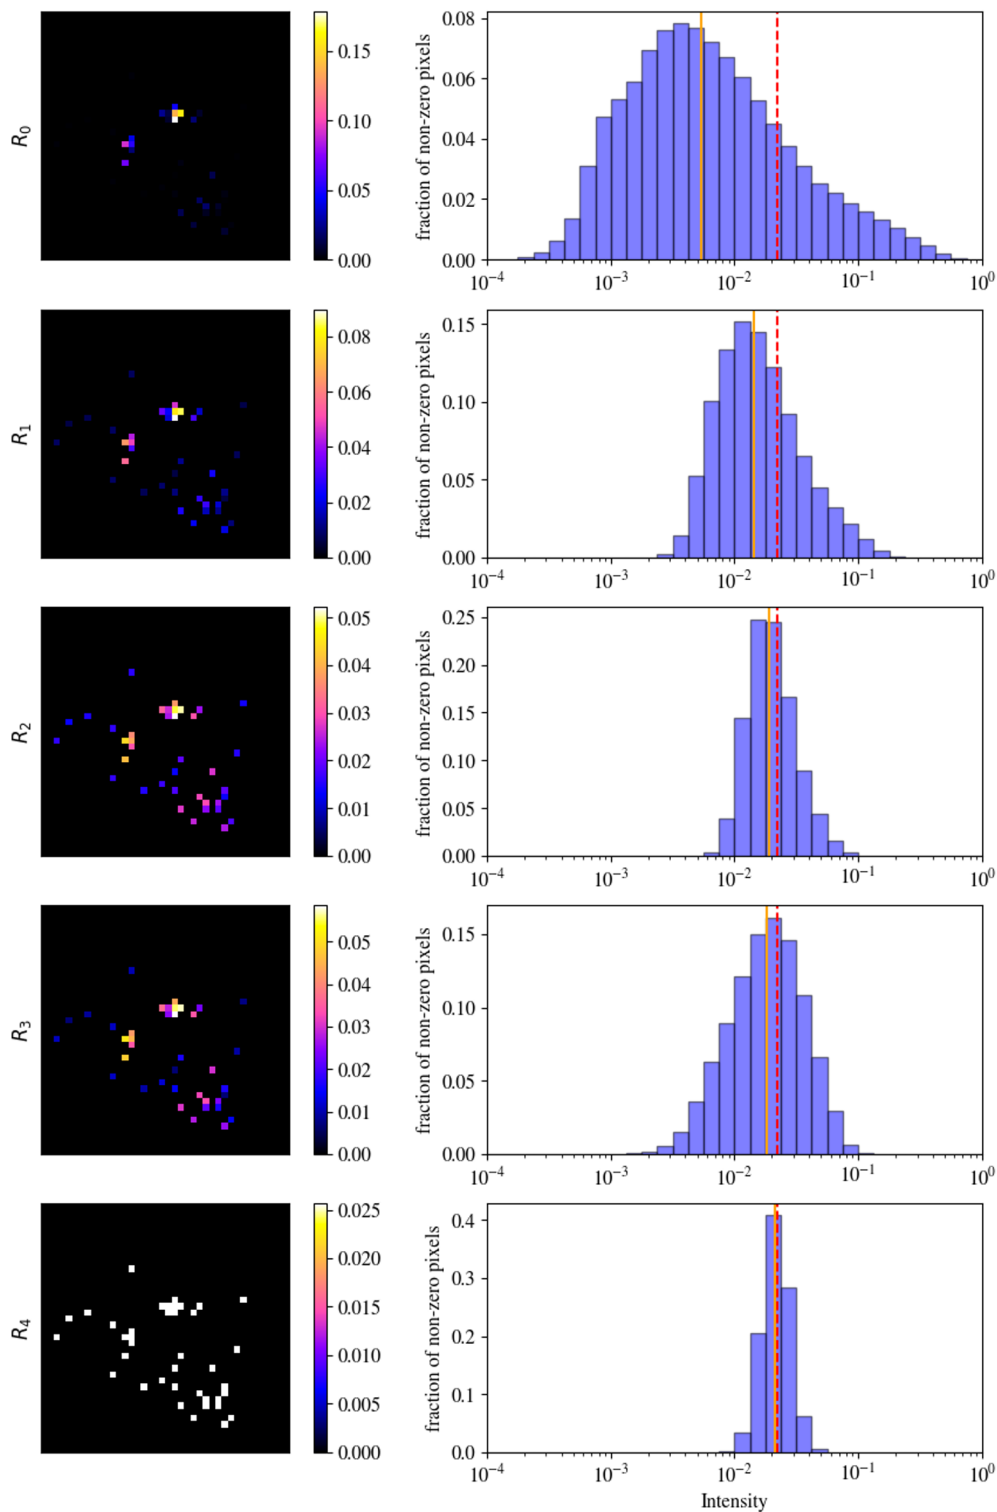


Figure 18. Example of a top jet image (left), and the distribution of intensities of non-zero pixels for 100k top jet images (right) for different intensity remappings (from top to bottom R_0 , R_1 , R_2 , R_3 and R_4). The orange and red-dashed vertical lines are the median and the average of each distribution, respectively.

background	signal	R	AUC [%]	E10 [%]	E100 [%]	$1/\epsilon_B$ ($\epsilon_S = 0.3$)
QCD	top	R0	$91.2^{+0.4}_{-0.4}$	$67.8^{+2.1}_{-1.9}$	$17.5^{+1.9}_{-1.4}$	46^{+3}_{-3}
		R1	$91.9^{+0.6}_{-1.4}$	$72.0^{+2.6}_{-4.9}$	$22.7^{+3.0}_{-5.2}$	64^{+12}_{-19}
		R2	$74.4^{+0.8}_{-1.0}$	$33.0^{+1.4}_{-1.5}$	$6.0^{+0.5}_{-0.6}$	12^{+1}_{-1}
		R3	$84.0^{+0.3}_{-0.5}$	$49.2^{+0.9}_{-1.0}$	$10.4^{+0.9}_{-0.8}$	23^{+1}_{-1}
		R4	$40.4^{+1.1}_{-1.1}$	$0.9^{+0.1}_{-0.1}$	$0.0^{+0.0}_{-0.0}$	$2.2^{+0.1}_{-0.1}$
top	QCD	R0	$36.5^{+2.4}_{-2.0}$	$6.6^{+0.9}_{-1.0}$	$0.6^{+0.2}_{-0.2}$	$2.2^{+0.2}_{-0.2}$
		R1	$56.6^{+0.7}_{-1.1}$	$23.4^{+0.9}_{-1.6}$	$8.6^{+1.2}_{-1.2}$	$6.1^{+0.3}_{-0.6}$
		R2	$68.6^{+0.4}_{-0.6}$	$42.1^{+1.0}_{-1.0}$	$20.9^{+1.3}_{-1.3}$	31^{+5}_{-4}
		R3	$66.2^{+0.6}_{-0.9}$	$37.3^{+1.4}_{-2.2}$	$18.0^{+1.2}_{-2.4}$	19^{+3}_{-4}
		R4	$70.4^{+0.2}_{-0.2}$	$45.2^{+0.5}_{-0.5}$	$25.2^{+0.3}_{-0.4}$	50^{+3}_{-2}

Table 1. Performance measures for AE-based taggers for different remappings of the jet images trained using MSE. Each value is an average of values for 4 autoencoders trained using the same hyperparameters but different initialization. The mentioned uncertainties denote the envelope of each value.

background	signal	R	AUC [%]	E10 [%]	E100 [%]	$1/\epsilon_B$ ($\epsilon_S = 0.3$)
QCD	top	R0	$92.8^{+0.1}_{-0.1}$	$75.5^{+1.1}_{-0.9}$	$18.5^{+3.0}_{-1.4}$	54^{+9}_{-4}
		R1	$89.7^{+0.4}_{-1.0}$	$66.8^{+1.5}_{-3.1}$	$18.0^{+1.5}_{-2.2}$	48^{+6}_{-7}
		R2	$58.9^{+2.3}_{-2.4}$	$10.4^{+2.3}_{-2.1}$	$0.7^{+0.3}_{-0.2}$	$4.2^{+0.5}_{-0.4}$
		R3	$67.1^{+0.7}_{-0.5}$	$18.8^{+1.0}_{-0.9}$	$2.0^{+0.2}_{-0.2}$	$6.3^{+0.2}_{-0.2}$
		R4	$36.8^{+0.3}_{-0.2}$	$0.4^{+0.1}_{-0.0}$	$0.0^{+0.0}_{-0.0}$	$1.9^{+0.1}_{-0.1}$
top	QCD	R0	$48.6^{+2.5}_{-3.0}$	$13.3^{+2.0}_{-1.5}$	$1.1^{+0.7}_{-0.4}$	$3.5^{+0.3}_{-0.4}$
		R1	$62.3^{+1.3}_{-2.1}$	$33.2^{+2.4}_{-3.0}$	$16.3^{+3.9}_{-3.3}$	14^{+4}_{-4}
		R2	$67.6^{+1.1}_{-0.5}$	$41.1^{+2.5}_{-1.1}$	$21.8^{+3.8}_{-1.7}$	32^{+14}_{-6}
		R3	$65.3^{+0.6}_{-0.9}$	$37.2^{+1.1}_{-1.8}$	$18.0^{+1.0}_{-2.6}$	20^{+2}_{-4}
		R4	$69.6^{+0.3}_{-0.3}$	$43.1^{+0.4}_{-0.4}$	$23.3^{+0.7}_{-0.6}$	38^{+3}_{-3}

Table 2. Performance measures for AE-based taggers for different remappings of the jet images trained using KMSE. Each value is an average of values for 4 autoencoders trained using the same hyperparameters but different initialization. The mentioned uncertainties denote the envelope of each value.

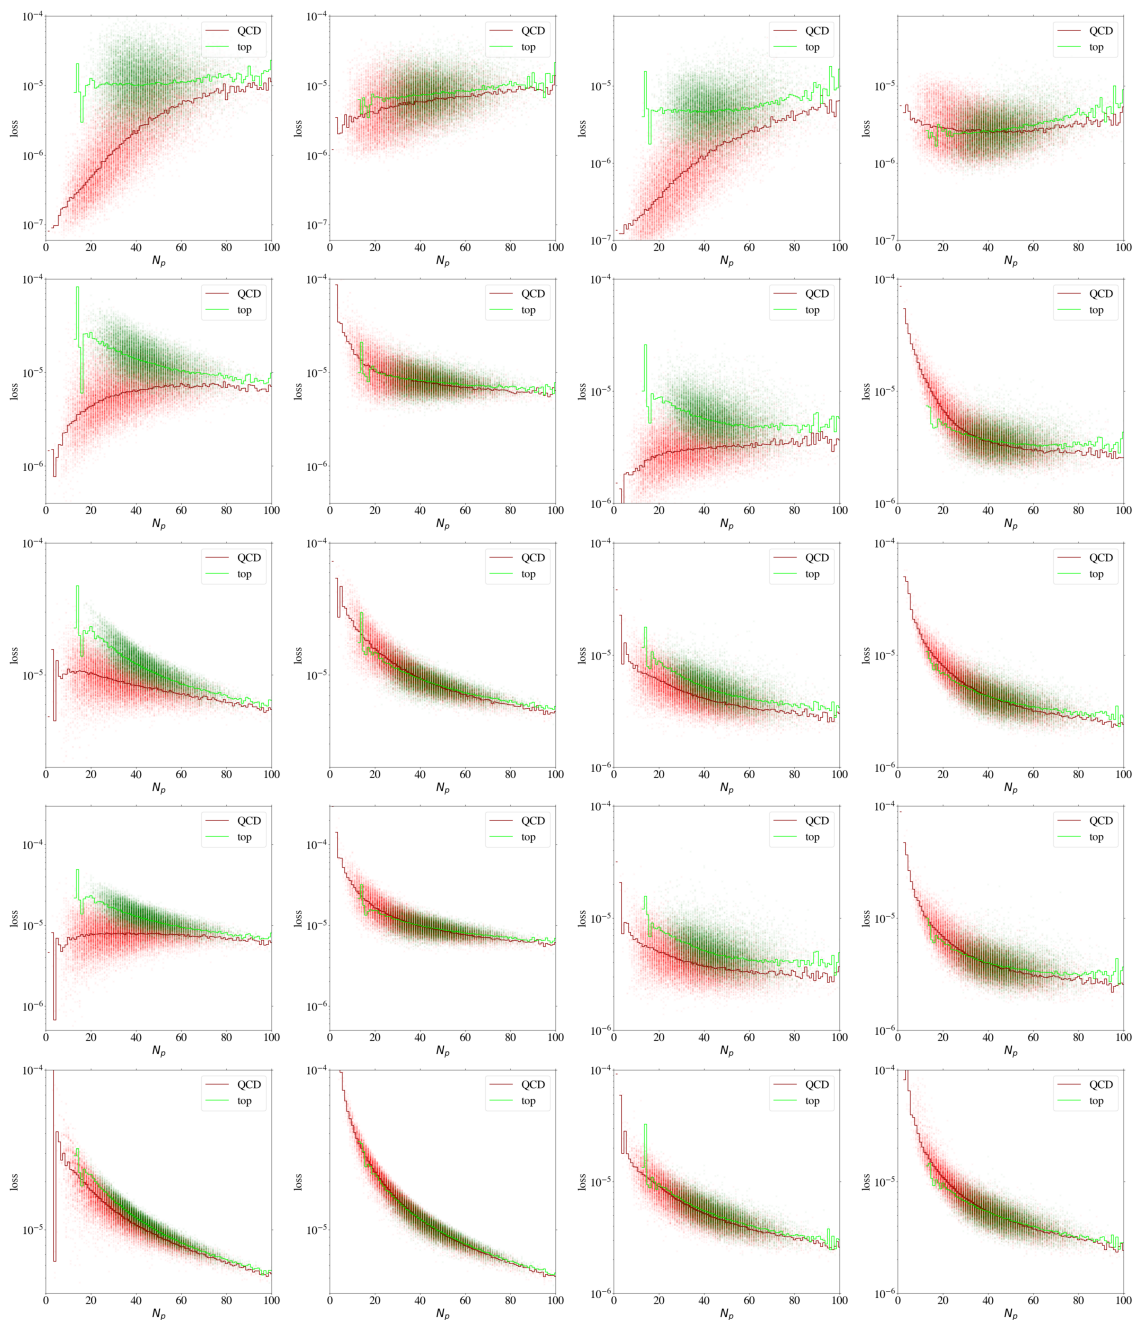


Figure 19. Reconstruction loss of individual test jets as a function of the number of non-zero pixels N_p , cf. figure 6. The first and second columns represent the results for an AE trained on QCD and top jet images, respectively, using MSE loss, while the third and fourth columns show the corresponding results using the KMSE loss function. The rows correspond to the different image remappings R_0 , R_1 , R_2 , R_3 and R_4 from top to bottom.

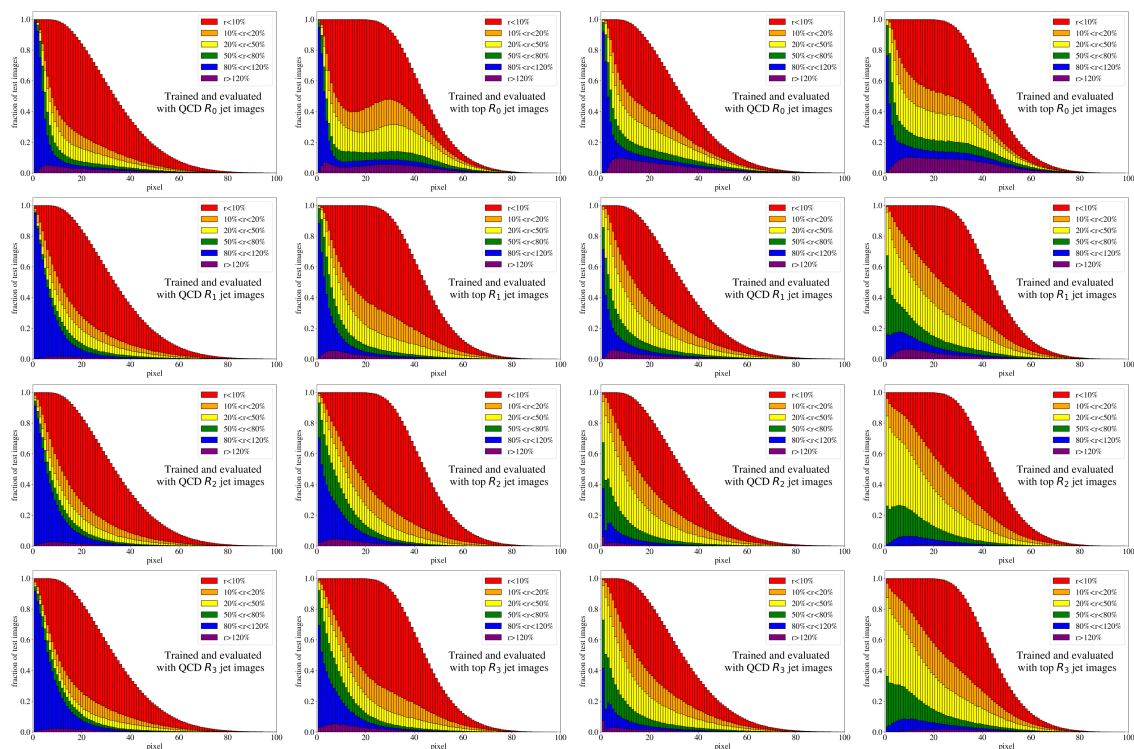


Figure 20. Stacked histogram for different categories of the ratio r of the reconstructed and the input intensity of the non-zero pixels. Pixels are ordered by intensity from left to right for each of the 40k test jets, cf. figure 4. The first and second columns represent the results for an AE trained on QCD and top jet images, respectively, using MSE loss, while the third and fourth columns show the corresponding results using the KMSE loss function. The rows correspond to the different image remappings R_0 , R_1 , R_2 , and R_3 from top to bottom.

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] ATLAS collaboration, *Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC*, *Phys. Lett. B* **716** (2012) 1 [[arXiv:1207.7214](https://arxiv.org/abs/1207.7214)] [[INSPIRE](#)].
- [2] CMS collaboration, *Observation of a New Boson at a Mass of 125 GeV with the CMS Experiment at the LHC*, *Phys. Lett. B* **716** (2012) 30 [[arXiv:1207.7235](https://arxiv.org/abs/1207.7235)] [[INSPIRE](#)].
- [3] M. Feickert and B. Nachman, *A Living Review of Machine Learning for Particle Physics*, [arXiv:2102.02770](https://arxiv.org/abs/2102.02770) [[INSPIRE](#)].
- [4] M.D. Schwartz, *Modern Machine Learning and Particle Physics*, [arXiv:2103.12226](https://arxiv.org/abs/2103.12226) [[INSPIRE](#)].
- [5] D. Bourilkov, *Machine and Deep Learning Applications in Particle Physics*, *Int. J. Mod. Phys. A* **34** (2020) 1930019 [[arXiv:1912.08245](https://arxiv.org/abs/1912.08245)] [[INSPIRE](#)].
- [6] D. Guest, K. Cranmer and D. Whiteson, *Deep Learning and its Application to LHC Physics*, *Ann. Rev. Nucl. Part. Sci.* **68** (2018) 161 [[arXiv:1806.11484](https://arxiv.org/abs/1806.11484)] [[INSPIRE](#)].
- [7] K. Albertsson et al., *Machine Learning in High Energy Physics Community White Paper*, *J. Phys. Conf. Ser.* **1085** (2018) 022008 [[arXiv:1807.02876](https://arxiv.org/abs/1807.02876)] [[INSPIRE](#)].
- [8] A.J. Larkoski, I. Moult and B. Nachman, *Jet Substructure at the Large Hadron Collider: A Review of Recent Advances in Theory and Machine Learning*, *Phys. Rept.* **841** (2020) 1 [[arXiv:1709.04464](https://arxiv.org/abs/1709.04464)] [[INSPIRE](#)].
- [9] L.M. Dery, B. Nachman, F. Rubbo and A. Schwartzman, *Weakly Supervised Classification in High Energy Physics*, *JHEP* **05** (2017) 145 [[arXiv:1702.00414](https://arxiv.org/abs/1702.00414)] [[INSPIRE](#)].
- [10] T. Cohen, M. Freytsis and B. Ostdiek, *(Machine) Learning to Do More with Less*, *JHEP* **02** (2018) 034 [[arXiv:1706.09451](https://arxiv.org/abs/1706.09451)] [[INSPIRE](#)].
- [11] E.M. Metodiev, B. Nachman and J. Thaler, *Classification without labels: Learning from mixed samples in high energy physics*, *JHEP* **10** (2017) 174 [[arXiv:1708.02949](https://arxiv.org/abs/1708.02949)] [[INSPIRE](#)].
- [12] P.T. Komiske, E.M. Metodiev, B. Nachman and M.D. Schwartz, *Learning to classify from impure samples with high-dimensional data*, *Phys. Rev. D* **98** (2018) 011502 [[arXiv:1801.10158](https://arxiv.org/abs/1801.10158)] [[INSPIRE](#)].
- [13] M. Borisyak and N. Kazeev, *Machine Learning on data with sPlot background subtraction*, *2019 JINST* **14** P08020 [[arXiv:1905.11719](https://arxiv.org/abs/1905.11719)] [[INSPIRE](#)].
- [14] O. Amram and C.M. Suarez, *Tag N' Train: a technique to train improved classifiers on unlabeled data*, *JHEP* **01** (2021) 153 [[arXiv:2002.12376](https://arxiv.org/abs/2002.12376)] [[INSPIRE](#)].
- [15] J.S.H. Lee, S.M. Lee, Y. Lee, I. Park, I.J. Watson and S. Yang, *Quark Gluon Jet Discrimination with Weakly Supervised Learning*, *J. Korean Phys. Soc.* **75** (2019) 652 [[arXiv:2012.02540](https://arxiv.org/abs/2012.02540)] [[INSPIRE](#)].
- [16] L. Ruff et al., *A Unifying Review of Deep and Shallow Anomaly Detection*, *Proc. IEEE* **109** (2021) 756 [[arXiv:2009.11732](https://arxiv.org/abs/2009.11732)].

- [17] R. Chalapathy and S. Chawla, *Deep Learning for Anomaly Detection: A Survey*, [arXiv:1901.03407](#).
- [18] B. Nachman, *Anomaly Detection for Physics Analysis and Less than Supervised Learning*, [arXiv:2010.14554](#) [INSPIRE].
- [19] G. Kasieczka et al., *The LHC Olympics 2020: A Community Challenge for Anomaly Detection in High Energy Physics*, [arXiv:2101.08320](#) [INSPIRE].
- [20] P. Baldi and K. Hornik, *Neural networks and principal component analysis: Learning from examples without local minima*, *Neural Networks* **2** (1989) 53.
- [21] Y. Bengio, A. Courville and P. Vincent, *Representation learning: A review and new perspectives*, *IEEE Trans. Pattern Anal. Machine Intell.* **35** (2013) 1798 [[arXiv:1206.5538](#)].
- [22] G. Pang, C. Shen, L. Cao and A.V.D. Hengel, *Deep Learning for Anomaly Detection*, *ACM Computing Surveys* **54** (2021) 1 [[arXiv:2007.02500](#)].
- [23] J. Hajer, Y.-Y. Li, T. Liu and H. Wang, *Novelty Detection Meets Collider Physics*, *Phys. Rev. D* **101** (2020) 076015 [[arXiv:1807.10261](#)] [INSPIRE].
- [24] M. Crispim Romão, N.F. Castro and R. Pedro, *Finding New Physics without learning about it: Anomaly Detection as a tool for Searches at Colliders*, *Eur. Phys. J. C* **81** (2021) 27 [[arXiv:2006.05432](#)] [INSPIRE].
- [25] S. Alexander et al., *Decoding Dark Matter Substructure without Supervision*, [arXiv:2008.12731](#) [INSPIRE].
- [26] A. Blance, M. Spannowsky and P. Waite, *Adversarially-trained autoencoders for robust unsupervised new physics searches*, *JHEP* **10** (2019) 047 [[arXiv:1905.10384](#)] [INSPIRE].
- [27] 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].
- [28] T. Cheng, J.-F. Arguin, J. Leissner-Martin, J. Pilette and T. Golling, *Variational Autoencoders for Anomalous Jet Tagging*, [arXiv:2007.01850](#) [INSPIRE].
- [29] B. Bortolato, B.M. Dillon, J.F. Kamenik and A. Smolkovič, *Bump Hunting in Latent Space*, [arXiv:2103.06595](#) [INSPIRE].
- [30] T. Heimel, G. Kasieczka, T. Plehn and J.M. Thompson, *QCD or What?*, *SciPost Phys.* **6** (2019) 030 [[arXiv:1808.08979](#)] [INSPIRE].
- [31] M. Farina, Y. Nakai and D. Shih, *Searching for New Physics with Deep Autoencoders*, *Phys. Rev. D* **101** (2020) 075021 [[arXiv:1808.08992](#)] [INSPIRE].
- [32] T.S. Roy and A.H. Vijay, *A robust anomaly finder based on autoencoders*, [arXiv:1903.02032](#) [INSPIRE].
- [33] E. Nalisnick, A. Matsukawa, Y.W. Teh, D. Gorur and B. Lakshminarayanan, *Do Deep Generative Models Know What They Don't Know?*, [arXiv:1810.09136](#).
- [34] R.T. Schirmer, Y. Zhou, T. Ball and D. Zhang, *Understanding Anomaly Detection with Deep Invertible Networks through Hierarchies of Distributions and Features*, [arXiv:2006.10848](#).
- [35] P. Kirichenko, P. Izmailov and A.G. Wilson, *Why Normalizing Flows Fail to Detect Out-of-Distribution Data*, [arXiv:2006.08545](#).

- [36] J. Ren et al., *Likelihood Ratios for Out-of-Distribution Detection*, [arXiv:1906.02845](#).
- [37] J. Serrà et al., *Input complexity and out-of-distribution detection with likelihood-based generative models*, [arXiv:1909.11480](#).
- [38] A. Tong, G. Wolf and S. Krishnaswamy, *Fixing Bias in Reconstruction-based Anomaly Detection with Lipschitz Discriminators*, [arXiv:1905.10710](#).
- [39] L.G. Almeida, M. Backović, M. Cliche, S.J. Lee and M. Perelstein, *Playing Tag with ANN: Boosted Top Identification with Pattern Recognition*, *JHEP* **07** (2015) 086 [[arXiv:1501.05968](#)] [[INSPIRE](#)].
- [40] G. Kasieczka, T. Plehn, M. Russell and T. Schell, *Deep-learning Top Taggers or The End of QCD?*, *JHEP* **05** (2017) 006 [[arXiv:1701.08784](#)] [[INSPIRE](#)].
- [41] J. Pearkes, W. Fedorko, A. Lister and C. Gay, *Jet Constituents for Deep Neural Network Based Top Quark Tagging*, [arXiv:1704.02124](#) [[INSPIRE](#)].
- [42] S. Macaluso and D. Shih, *Pulling Out All the Tops with Computer Vision and Deep Learning*, *JHEP* **10** (2018) 121 [[arXiv:1803.00107](#)] [[INSPIRE](#)].
- [43] A. Butter et al., *The Machine Learning landscape of top taggers*, *SciPost Phys.* **7** (2019) 014 [[arXiv:1902.09914](#)] [[INSPIRE](#)].
- [44] J.Y. Araz and M. Spannowsky, *Combine and Conquer: Event Reconstruction with Bayesian Ensemble Neural Networks*, *JHEP* **04** (2021) 296 [[arXiv:2102.01078](#)] [[INSPIRE](#)].
- [45] A. Butter, G. Kasieczka, T. Plehn and M. Russell, *Deep-learned Top Tagging with a Lorentz Layer*, *SciPost Phys.* **5** (2018) 028 [[arXiv:1707.08966](#)] [[INSPIRE](#)].
- [46] G. Kasieczka, T. Plehn, J. Thompson and M. Russel, *Top Quark Tagging Reference Dataset*, [10.5281/zenodo.2603256](#) (2019).
- [47] M. Abadi et al., *TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems*, <https://www.tensorflow.org/> (2015).
- [48] F. Chollet et al., *Keras*, <https://github.com/fchollet/keras> (2015).
- [49] B. Zong et al., *Deep Autoencoding Gaussian Mixture Model for Unsupervised Anomaly Detection*, in *International Conference on Learning Representations*, Vancouver Convention Center, Vancouver, BC, Canada, April 30–May 3, 2018 [<https://openreview.net/forum?id=BJJLHbb0->].
- [50] D. Gong et al., *Memorizing Normality to Detect Anomaly: Memory-augmented Deep Autoencoder for Unsupervised Anomaly Detection*, [[arXiv:1904.02639](#)].
- [51] J. Batson, C.G. Haaf, Y. Kahn and D.A. Roberts, *Topological Obstructions to Autoencoding*, *JHEP* **04** (2021) 280 [[arXiv:2102.08380](#)] [[INSPIRE](#)].
- [52] J.H. Collins, P. Martín-Ramiro, B. Nachman and D. Shih, *Comparing Weak- and Unsupervised Methods for Resonant Anomaly Detection*, [arXiv:2104.02092](#) [[INSPIRE](#)].
- [53] Y. Rubner, C. Tomasi and L.J. Guibas, *The Earth Mover’s Distance as a Metric for Image Retrieval*, *Int. J. Comput. Vision* **40** (2000) 99.
- [54] N. Bonneel, J. Rabin, G. Peyré and H. Pfister, *Sliced and Radon Wasserstein Barycenters of Measures*, *J. Math. Imag. Vis.* **51** (2015) 22.
- [55] T. Finke, *Deep Learning for New Physics Searches at the LHC*, Master Thesis, RWTH Aachen University (2020).

- [56] I. Oleksiyuk, *Unsupervised learning for tagging anomalous jets at the LHC*, Bachelor Thesis, RWTH Aachen University (2021).
- [57] B.M. Dillon, T. Plehn, C. Sauer and P. Sorrenson, *Better Latent Spaces for Better Autoencoders*, [arXiv:2104.08291](#) [INSPIRE].
- [58] B.M. Dillon, *Learning the latent structure of collider events*, in *Anomaly Detection Mini-Workshop — LHC Summer Olympics*, (2020) [<https://indico.desy.de/event/25341/contributions/56828/>].
- [59] Y. Gershtein, D. Jaroslowski, K. Nasha, D. Shih and M. Tran, *Anomaly detection with convolutional autoencoders and latent space analysis*, in *Anomaly Detection Mini-Workshop — LHC Summer Olympics*, (2020) and publication in preparation [<https://indico.desy.de/event/25341/contributions/56829/>].
- [60] T. Sjöstrand et al., *An introduction to PYTHIA 8.2*, *Comput. Phys. Commun.* **191** (2015) 159 [[arXiv:1410.3012](#)] [INSPIRE].
- [61] DELPHES 3 collaboration, *DELPHES 3, A modular framework for fast simulation of a generic collider experiment*, *JHEP* **02** (2014) 057 [[arXiv:1307.6346](#)] [INSPIRE].
- [62] M. Cacciari, G.P. Salam and G. Soyez, *FastJet User Manual*, *Eur. Phys. J. C* **72** (2012) 1896 [[arXiv:1111.6097](#)] [INSPIRE].
- [63] M. Cacciari, G.P. Salam and G. Soyez, *The anti- k_t jet clustering algorithm*, *JHEP* **04** (2008) 063 [[arXiv:0802.1189](#)] [INSPIRE].
- [64] D.P. Kingma and J. Ba, *Adam: A Method for Stochastic Optimization*, [arXiv:1412.6980](#) [INSPIRE].
- [65] 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].