

Learning to see R -parity violating scalar top decays

Gerrit Bickendorf^{✉*} and Manuel Drees^{✉†}

Bethe Center for Theoretical Physics and Physikalisches Institut, University of Bonn, Bonn, Germany



(Received 16 June 2024; accepted 8 August 2024; published 4 September 2024)

With this article, we introduce recent, improved machine learning methods from computer vision to the problem of event classification in particle physics. Supersymmetric scalar top decays to top quarks and weak-scale bino-like neutralinos, where the neutralinos decay via the UDD operator to three quarks, are difficult to search for and therefore weakly constrained. The jet substructure of the boosted decay products can be used to differentiate signal from background events. We apply the transformer-based computer vision models CoATNet and MaxViT to images built from jet constituents and compare the classification performance to a more classical convolutional neural network (CNN). We find that results from computer vision translate well to physics applications, and both transformer-based models perform better than the CNN. By replacing the CNN with MaxViT, we find an improvement of S/\sqrt{B} by a factor of almost 2 for some neutralino masses. We show that combining this classifier with additional features results in a strong separation of background and signal. We also find that replacing a CNN with a MaxViT model in a simple mock analysis can push the 95% C.L. exclusion limit of stop masses by about 100 and 60 GeV for neutralino masses of 100 and 500 GeV.

DOI: [10.1103/PhysRevD.110.056006](https://doi.org/10.1103/PhysRevD.110.056006)

I. INTRODUCTION

The minimal supersymmetric extension of the standard model (MSSM) is a promising candidate for physics beyond the SM [1–6] that might solve the hierarchy problem. Despite many experimental searches, most notably by the ATLAS and CMS collaborations at the CERN LHC in recent years, no conclusive evidence of its realization in nature has been found, pushing the parameter space to ever higher masses (see Ref. [7] for an overview). In models with conserved R parity (RPC), many searches leverage the large p_T^{miss} due to the stable lightest supersymmetric particle (LSP) leaving the experiment undetected [8]. Once the RPC assumption is dropped, these strategies often become insensitive. In the context of the R -parity-violating (RPV) MSSM, new terms are added to the superpotential that break lepton- or baryon-number conservation [9]. These additional terms imply that drastically different search strategies are needed [10], especially when prompt decays of supersymmetric particles become drowned by the hadronic activity inside the detector.

At a hadron collider like the LHC, the production of strongly interacting superparticles has the largest cross section for a given mass. Among these, the stops—the superpartners of the top quark—are often assumed to be the lightest. On the one hand, for equal squark masses at some very high (e.g., grand unified or Planckian) energy scale, renormalization group effects reduce the masses of the stops; mixing between the masses of the $SU(2)$ doublet and singlet stops will reduce the mass of the lightest eigenstate even more [6]. On the other hand, simple naturalness arguments [11–13] prefer not too heavy stop squarks, but allow much heavier first- and second-generation squarks. This motivates the analysis of scenarios where the mass of the lighter stop squark lies well below those of the other strongly interacting superparticles.

The same naturalness arguments also prefer rather small supersymmetric contributions to the masses of the Higgs bosons. In most (though not all [14]) versions of the MSSM, this implies rather light Higgsinos, typically below the stop. Since the mass splitting between the three Higgsino-like mass eigenstates is small, they all behave similarly if the LSP is Higgsino-like. In particular, in the kind of RPV scenario we consider, all three states would lead to very similar “fat jets” when produced in stop decays; the recognition of such jets by exploiting recent developments in computer vision is one of the central points of our paper, which would apply equally to all three Higgsino states. However, about half of all stop decays would then produce a bottom, rather than a top, together with a Higgsino, thereby complicating the analysis of the

*Contact author: bickendorf@th.physik.uni-bonn.de

†Contact author: drees@th.physik.uni-bonn.de

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI. Funded by SCOAP³.

remainder of the final state. Moreover, Higgsinos, being $SU(2)$ doublets, have a sizable direct production rate. Their nonobservation therefore leads to significant constraints on parameter space, especially (but not only) if the bino has mass comparable to or smaller than the Higgsinos [15–20].

In order to avoid such complications, we consider the pair production of scalar top quarks, which decay to top quarks plus two neutralinos with unit branching ratio. The neutralinos in turn decay promptly via the UDD R -parity-breaking term, which is fairly difficult to constrain [21]. Since each neutralino may form a (fat) jet, one can use the substructure to differentiate it from background processes [22]. One may represent the jets as images made of calorimeter cell hits that can be used by computer vision techniques. Using a convolutional neural net (CNN) has already been shown to work well on these images [23–34].

In recent years, computer vision techniques have improved drastically with novel approaches such as the vision transformer [35]. In standardized computer vision tasks, these models have been shown to outperform CNN-based models for large data sets. Fortunately, generating large sets of simulated events is relatively cheap in particle physics, which motivates the use of these new techniques. Transformers have already been applied to classification in particle physics scenarios [36–44], although these focus on representing the jet as a set of particles, instead of as an image.

In this article, we for the first time apply two modern transformer-based computer vision techniques to find neutralinos from scalar top quark decays and compare the results to a classical CNN to see if the gain in performance translates to detector images. Using gradient-boosted decision trees (GBDTs), we combine the data from both neutralinos tagged in this way and add further high-level features to construct our final event classifier.

The remainder of this article is structured as follows. In Sec. II we describe the specifics of the signal model we use. In Sec. III we show how we generate the data sets to which we apply the preprocessing outlined in Sec. IV. The novel computer vision architectures we wish to adapt are described in Sec. V. Section VI outlines the generation of data sets from neutralino decay and background events, which are used to train the neutralino taggers in Sec. VII. The performance of these taggers is discussed in Sec. VIII. In Sec. IX we combine information from both tagged neutralinos into even more powerful classifiers. We also demonstrate the power of combining the neutralino taggers with other high-level features in Sec. X. The improved stop mass reach is shown in Sec. XI, while Sec. XII contains a brief summary and some conclusions.

II. SIGNAL MODEL

The MSSM contains two scalar top quarks, which are mixtures of the $SU(2)$ doublet \tilde{t}_L and singlet \tilde{t}_R weak gauge eigenstates. We work with breaking parameters such that

the lighter mass eigenstate \tilde{t}_1 contains mainly the right-handed top squark, which decays promptly into a top quark and the bino-like neutralino $\tilde{\chi}_1^0 \equiv \tilde{\chi}$. We consider all other scalar quarks to be decoupled. In order to avoid the constraints from missing E_T -based searches, we add

$$W_{\not{p}} = \frac{1}{2} \lambda''_{ijk} U_i^c D_j^c D_k^c \quad (1)$$

to the superpotential, where U_i^c and D_i^c are the up- and down-type $SU(2)$ singlet chiral quark superfields and i, j, k are generation indices. Clearly, this term violates baryon number conservation. In Eq. (1), antisymmetrization over color (i.e., contraction with the totally antisymmetric tensor in color space) is implied, and hence the coupling has to be antisymmetric in the last two indices. Therefore, there are in general nine independent coupling constants λ''_{ijk} . When $i = 3$, this would allow the stop to decay directly to two lighter quarks, which has already been extensively studied [45–48]. A coupling with $i \neq 3$ allows even a light neutralino to decay into three quark jets via an off-shell squark. The process we are interested in is shown in Fig. 1.

We also note that a mostly \tilde{t}_R eigenstate decaying into a bino-like neutralino produces a predominantly right-handed top quark. The same is true for a \tilde{t}_L decaying into a neutral Higgsino. In contrast, a \tilde{t}_L decaying into a bino or a \tilde{t}_R decaying into a neutral Higgsino would produce a left-handed top quark. Since we do not try to reconstruct the polarization of the top (anti)quark in the final state, all four reactions would have very similar signatures and could be treated with the methods developed in this paper. However, as already noted in the Introduction, a light neutral Higgsino implies the existence of a nearly mass-degenerate charged Higgsino (and a second neutral Higgsino), thereby reducing the branching ratio for $\tilde{t} \rightarrow t + \tilde{\chi}$ decays. Moreover, by $SU(2)$ invariance, a mostly \tilde{t}_L stop eigenstate would be close in mass to \tilde{b}_L , leading to additional signals from \tilde{b}_L pair production. By focusing on a mostly \tilde{t}_R lighter stop and a bino-like LSP, we avoid these complications.

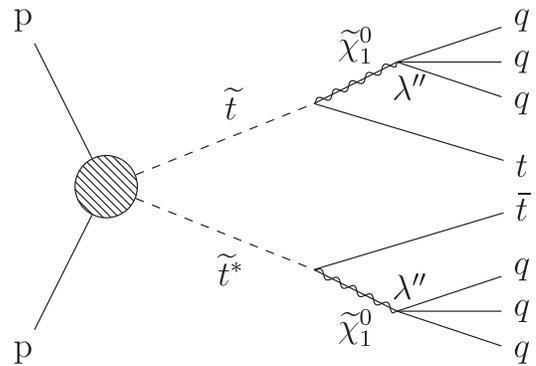


FIG. 1. Stop pair production with each stop decaying to a top quark and a neutralino. The neutralinos decay via the RPV UDD operator with nonzero λ'' .

III. DATA GENERATION AND PRESELECTION

For baseline selections, we follow roughly the CMS search for this signal process [49]. We impose the following preselection cuts:

- (1) One muon with $p_T > 30$ GeV or electron with $p_T > 37$ GeV and $|\eta| < 2.4$.
- (2) The lepton must be isolated within a cone radius depending on the p_T of the lepton like

$$R = \begin{cases} 0.2, & p_T < 50 \text{ GeV}, \\ 10 \text{ GeV}/p_T, & 50 \text{ GeV} < p_T < 200 \text{ GeV}, \\ 0.05, & p_T > 200 \text{ GeV}. \end{cases}$$

Together with the first cut, this isolation requirement implies that in almost all events the lepton originates from the semileptonic decay of one of the top (anti)quarks in the final state. These two cuts satisfy the requirements of the single lepton trigger. Note that the events must contain exactly one such isolated lepton; this largely removes $Z + \text{jets}$ backgrounds.

- (3) We define “AK04 jets” via the anti- k_T (AK) jet clustering algorithm with distance parameter $R = 0.4$, requiring $p_T > 30$ GeV and $|\eta| < 2.4$ for each jet. We demand that the event contains at least seven such AK04 jets, at least one of which is b tagged. We note that our signal events contain at least two b (anti)quarks from top decay. Moreover, even if both t and \bar{t} decay semileptonically, signal events contain eight energetic quarks even in the absence of QCD radiation. They should therefore pass this cut with high efficiency, except for very light neutralinos where several of their decay products might end up in the same (quite narrow) AK04 jet. On the other hand, SM $t\bar{t}$ events with one top decaying semileptonically contain only four hard quarks. Hence, at least three additional jets would have to be produced by QCD radiation, significantly reducing the $t\bar{t}$ background and reducing the $W + \text{jets}$ background even more.
- (4) $H_T > 300$ GeV, where H_T is the scalar sum of the transverse momenta of all AK04 jets. This cut is mostly effective against $W, Z + \text{jets}$ backgrounds.
- (5) At least one combination of b -tagged jet and isolated lepton must have an invariant mass between 50 and 250 GeV. Most events where the lepton and the b quark originate from the decay of the same t quark pass this cut, which helps to further reduce the $W + \text{jets}$ background.
- (6) At least one AK08 jet (defined with distance parameter $R = 0.8$), with $p_T > 100$ GeV. We will later try to tag these “fat jets” as coming from neutralino decay. However, a boosted, hadronically decaying top (anti)quark can also produce such a jet.

We will also consider even fatter jets. Since (nearly) all particles inside an AK08 jet will end up inside the same jet if $R > 0.8$ is used in the jet clustering, while these fatter jets will contain additional “nearby” particles, they will automatically also have $p_T > 100$ GeV.

After these cuts, the remaining background is almost exclusively due to top-quark pair production, as can be seen in the original CMS publication [49]. In our simulation, we therefore only consider this background process.

For the signal model, we set the masses of squarks (except that of the stop), gluinos, and wino- and Higgsino-like neutralinos to 5 TeV. We only set one RPV coupling nonzero, $\lambda''_{223} = -\lambda''_{232} = 0.75$; this leads to prompt neutralino $\tilde{\chi} \rightarrow csb$ decay even if the exchanged squark has a mass of 5 TeV, $\tau_{\tilde{\chi}} \sim 10^{-18} \text{ s} \cdot [m_{\tilde{\chi}}/(100 \text{ GeV})]^{-5}$. We scan over the stop mass from $m_{\tilde{t}} = 700$ GeV to $m_{\tilde{t}} = 1200$ GeV in steps of 25 GeV. We also scan over the neutralino mass from $m_{\tilde{\chi}} = 100$ GeV to $m_{\tilde{\chi}} = 500$ GeV in 10 GeV steps.

Background and signal events are simulated using MadGraph5_aMC@NLO 3.2.0 [50]. The $t\bar{t}$ background is generated with between zero and three additional matrix element partons, while the signal events contain up to two additional partons. The NNPDF3.1 PDF set [51] is used. We use Pythia 8.306 [52] for parton showering and hadronization; background events are showered with the CP5 tune, while signal events are showered with the CP2 tune [53]. Events with different matrix element-level final-state parton multiplicities are merged with the MLM prescription [54], in order to avoid double counting events where the parton shower produces additional jets. Finally, detector effects are simulated with the CMS card of Delphes 3.5.0 [55,56].

IV. PREPROCESSING

The main novelty of this paper is the adoption of very recent computer vision techniques to tag the hadronically decaying neutralinos. To that end, we first have to translate the simulated detector data to images.

The objects we are interested in are jets clustered with the anti- k_T jet algorithm as implemented by the FastJet package [57]. Choosing the optimal distance parameter R for a given purpose can be somewhat nontrivial. A small value of R means that most particles inside a sufficiently hard jet originated from the same parton, but some of the energy of that parton might not be counted in this jet due to final-state showering. On the other hand, a large R likely leads to jets that capture all daughter particles while also muddying the waters by including unrelated objects, e.g., from initial-state showering. One can use the fact that the decay products of a resonance with a fixed mass m and transverse momentum p_T spread roughly like $\Delta R = \sqrt{\Delta\phi^2 + \Delta\eta^2} \propto m/p_T$ and the typical energy scale of the process to arrive at a *best guess* for an optimal R

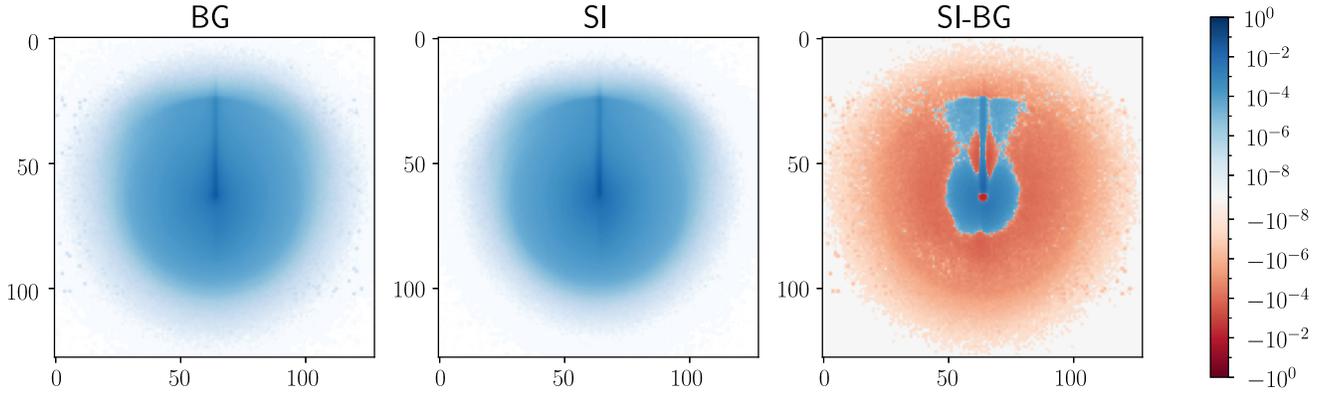


FIG. 2. Signal and background AK14 jet image averaged over the entire training data set. All three channels are aggregated by summation. The rightmost plot shows the difference between the average signal and the average background jet image. Signal events are more concentrated at the origin, while background jets are more spread out.

parameter. This can be aided by the use of jet clustering algorithms with variable R (e.g., [58,59]). In the case at hand, this optimal value of R would depend on both the stop and the neutralino mass. We therefore do not work with a single fixed value of R but instead cluster each event using several values of R , and ensemble the resulting jet images to get a better per-event classification. Because we consider rather large neutralino masses, $m_{\tilde{\chi}} \geq 100$ GeV, we consider AK08 ($R = 0.8$), AK10 ($R = 1.0$), and AK14 ($R = 1.4$) jets. This also allows us to keep the technique general, i.e., to use the same algorithm over the entire parameter space. Recall that the resulting fat jet has to satisfy $p_T > 100$ GeV and $|\eta| < 2.4$.

In order to get images out of the jets, we now consider the calorimeter towers and tracks as jet constituents in the (η, ϕ) plane. As in the construction of top taggers [27], we do not use the energy E of the calorimeter towers directly but rather opt for the transverse energy $E_T = E/\cosh\eta$. The relevant features are more readily learned by the classifier if we normalize the coordinates. First, we calculate the E_T weighted center of the calorimeter towers via

$$\bar{\eta} = \frac{\sum_i E_{Ti} \eta_i}{\sum_i E_{Ti}}, \quad (2)$$

$$\bar{\phi} = \frac{\sum_i E_{Ti} \phi_i}{\sum_i E_{Ti}}, \quad (3)$$

where the sums run over all constituents of a given fat jet. We then shift the coordinates $\eta_i \rightarrow \eta_i - \bar{\eta}$ and $\phi_i \rightarrow \phi_i - \bar{\phi}$ so that the image is centered on the origin. Next, we rotate the coordinate system around the origin so that the calorimeter tower with the highest E_T points vertically from the origin. We use the last degree of freedom to make sure that the calorimeter tower with the second highest E_T lies in the right half of the coordinate by flipping along the vertical axis if necessary.

Next, we pixelate the coordinates to a 0.04×0.04 grid. The brightness/intensity of each pixel is given as the measured E_T . We use three channels, corresponding to E_T in the electromagnetic calorimeter (ECAL), hadron calorimeter (HCAL), and p_T of the tracks, analogous to three color channels in classical images.¹ We divide each pixel by the maximal value found in this image so that each intensity is between 0 and 1. This makes learning more efficient. It also removes information about the p_T and mass of the jet, which are powerful discriminators. We partly remedy this by giving the classifier the mass of the fat jet as another input; this will be explained in more detail in a following section. We also note that we will later introduce additional high-level features to our final classifier, which will reintroduce information about the overall E_T scale of the event.

In the last preprocessing step we crop the image to a square centered around the origin with side lengths chosen as 64 pixels for AK08 and AK10 jets and 128 pixels for AK14 jets. This size is chosen to contain most of the constituents, while also being a power of 2 which aids in the application of the computer vision techniques. The resulting images after all preprocessing steps, averaged over the entire event sample, are shown in Fig. 2. These average images look quite similar for signal and background, at least to the human eye; however, taking the difference between the average images does reveal some differences.

Moreover, there is more information available to the computer vision techniques than can be displayed in the figure. For instance, the number of nonzero pixel values is useful for classification. On average the signal images

¹In principle, the tracks have a much higher resolution compared to the calorimeter towers and could thus be pixelated into a finer grid. However, we do not expect these very fine details to improve the discrimination between signal and background. We therefore use the same grid spacing for all three channels.

contain more nonzero pixels than the background images do. Just cutting on this quantity allows, for one set of parameters, to reach an accuracy of 74% when applied to a sample containing an equal number of signal and background events. Of course, our final classifier should perform much better than this.²

V. ARCHITECTURES

In this section, we describe how we process a single fat jet, the goal being to distinguish jets due to the three-body decay of neutralinos from SM background (“LSP tagging”). At the core is one of three architectures adapted from computer vision and described in more detail in the following three subsections. In all three cases, the output of this architecture is concatenated with the measured jet mass and fed into the same multilayer perceptron classification network. It is built from a dense layer with 256 neurons followed by another dense layer with 128 neurons, which connect to two output neurons. Between all three layers, the ReLU activation function is used. The two neurons of the last layer are passed into the softmax activation function, such that the output can be interpreted as the predicted probability of the image belonging to either the signal or the background; since these probabilities should add to 1, the values of the two output neurons are almost completely correlated. In the following, we denote this as *MLP Head*.

All three architectures use convolution layers. Such a layer convolutes c_o sets of learnable weights, the kernels, with dimensions $c_i \times k_h \times k_w$ across the height and width of the image, where k_h and k_w are the height and width of the weights and c_i is the number of input channels. Recall that we start with three input channels, akin to the three colors commonly used in vision: E_T measured in the ECAL, E_T measured in the HCAL, and p_T of tracks. The information of the three “colors” is therefore merged already in the first convolution layer of a given architecture. The results of each convolution are stacked such that the output has c_o channels. All architectures are built and trained within the PyTorch [60] deep learning library.

A. CNN

The first architecture is a (comparatively) simple CNN. We follow loosely an existing model used for top tagging [23]. The first layers of the CNN are two blocks, each containing a convolutional layer with 128 kernels of

²We note in passing that this multiplicity does contain some information on the hardness of the event, since it correlates positively with both the mass and the transverse momentum of the fat jet. Hence, the normalization step described above does not completely remove the information on these quantities. However, these dependencies are only logarithmic and subject to large event-by-event fluctuations. Explicitly adding the jet mass as an input variable can therefore still be expected to aid in the classification task.

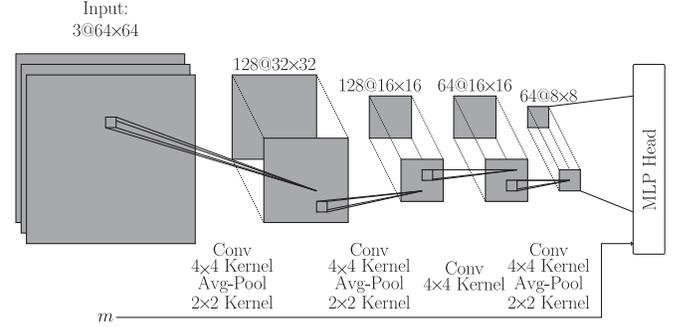


FIG. 3. Architecture of the CNN for AK08 and AK10 jets.

size 4×4 , stride 1, zero padding to keep the image dimensions, the ReLU activation function, and average pooling with kernel size 2 and stride 2. This halves the spatial image dimensions. Next, we apply the same block with only 64 kernels and without pooling. The last convolution block contains again 64 kernels but this time with the pooling operation. For AK08 and AK10 jet images, this network produces outputs of shape $64 \times 8 \times 8$; here, 8×8 refers to the size of one image after three convolutions, and we use 64 different “filters” (i.e., sets of weights in the convolution) for each image. In order to put AK14 images of size 128×128 on the same footing, we repeat the last block one more time. The output of shape $64 \times 8 \times 8$ is then flattened into 4096 features and fed into the MLP head. In total, there are 1 350 850 learnable parameters for AK08 and AK10 jets, while the model for AK14 jets contains 1 416 450 learnable parameters.

The full architecture for AK08 and AK10 jets is shown in Fig. 3. As already noted, this kind of architecture is already being used for similar tasks; it serves as our baseline, against which we compare the more advanced architectures described in the following two subsections.

B. CoAtNet

Since their inception in the context of natural language processing [61], transformer models have been shown to also be applicable for computer vision tasks. Because the attention mechanism inside the transformer is computationally expensive it is often impractical to apply it directly to the entire input image. In order to counteract this, the vision transformer (ViT) model [35] splits the input image into manageable patches. They are then flattened, linearly projected, equipped with a position embedding, and fed into the transformer encoder. Since both translation equivariance and locality are explicitly broken in this approach one cannot expect to outperform CNNs that are known to leverage both of these features and therefore generalize well onto the unseen test set. In our tests, the vanilla ViT indeed performs poorly, so we do not pursue it further.

However, it is possible to construct a model that combines the global receptive field of view from self-attention with the aforementioned advantages of CNNs.

This is the aim of CoAtNet [62] (the name derives from the combination of *depth-wise convolution* and *self-attention*). Let us now briefly describe how CoAtNet is built; we refer the interested reader to the original publication for a more detailed description.

The model is constructed in five stages. The first stage consists of three convolution layers which 3×3 kernels, where the first has a stride of 2. This halves the spatial resolution of the input image. This is followed by two stages of three MBConv blocks [63], which are computationally cheaper while maintaining most of the performance of full convolutional layers. In both stages, the first layers perform downsampling again with a stride size of 2. By now the width and height of the input image are shrunk by a factor of 2^3 so global attention is feasible even for the large AK14 jet. Thus, the last two stages consist of five and two transformer blocks, respectively. In each transformer, 2D relative attention is used, which adds a learnable weight that depends only on the relative 2D position. These steps were performed using the publicly available code.³ Finally, we average pool the outputs and feed the 768 features into the classification head. This architecture contains the most parameters. The version for AK08 and AK10 jets are built of 17 356 666 learnable parameters, while the model for AK14 jets contains 17 364 346 learnable parameters.

C. Multi-axis vision transformer

The third architecture we will use is the multi-axis vision transformer (MaxViT) [64]. This approach circumvents the shortcomings of the vanilla ViT differently. Instead of splitting the image in patches and only applying attention to these flattened local patches, the goal is to apply it locally and globally consecutively by decomposing the image with two strategies into smaller bits. First, the image is separated into equal-sized, nonoverlapping blocks along the spatial dimensions. Self-attention is applied within each block using only the local information. In order to leverage global information, the next strategy groups every pixel that is reached by a step of fixed size into the object to which self-attention is applied. (For a nice visualization, we refer to Fig. 3 of [64].) This approach essentially uses high-resolution local information before using low-resolution global information. Both attention mechanisms are applied after a MBConv block forming the so-called MaxViT block.

For our application, we chose the publicly available implementation that is shipped with PyTorch,⁴ with the only change being the replacement of the MLP head that takes the 256 output features. Concretely, the first stage consists of two convolutional layers with $64 \ 3 \times 3$ kernels each, where the first has stride 2, reducing the spatial dimensions by half. This is followed by three stages with two MaxViT

blocks each. The convolution is strided with size 2 for the first block of each stage. The partitions are of size 4×4 each. The MaxViT stages have 64 128 and 256 channels, respectively. Self-attention uses 32-dimensional heads. This model contains 6 014 810 learnable parameters for all jet sizes.

VI. DATA SET CREATION

In this section, we describe how the data set used to train the LSP taggers is defined. In most events, there is more than one fat jet that passes the selection criteria. Since we investigate pair production, not all of the information useful for event classification can be expected to be contained in the hardest jet. It is therefore expected to be useful to combine information from more than one jet in the analysis. Wide jets produced from the $\tilde{\chi}$ decays are expected to be hard because of the large stop mass. Therefore, the two jets with the largest p_T are expected to be signal enriched. The preselection requirements imply that one top quark from stop decay will generally decay semileptonically. However, the second top quark might decay fully hadronically, resulting in a third wide jet with large p_T . Since the top quarks and the LSPs have very similar p_T distributions, the third largest- p_T jet may also well be from an LSP.

In order to design taggers that perform well on all three leading jets, and hence for a wide range of p_T , we therefore include samples of all three leading fat jets in our training data set in the ratios that the respective number of jets are present in the full events. To this end, we add up to three fat jets present in an event as images to the data set. Of course, an event may also contain only one or two such jets; in fact, this is generally the case for background events. We generate as many events as required to reach the desired size of the training set for each jet size.

VII. TRAINING THE LSP TAGGERS

We start by verifying that the different Pythia tunes that we adopt do not significantly influence our results. To this end, we train the CNN model described in Sec. VA to differentiate not between signal and background samples, but between background events generated with the CP2 tune and background samples generated with the CP5 tune. We combine 1 000 000 jet images generated with each tune into a data set and split it equally between training and test sets for each jet size. The initial learning rate η_L is chosen as 5×10^{-4} . This value worked best in tests of the LSP taggers. At the end of each epoch, the learning rate is lowered by a factor of 0.7 and the entire training data set is shuffled. The batch size is 64. We minimize the averaged binary cross-entropy loss

$$l = -\frac{1}{N} \sum_i y_i \ln x_i + (1 - y_i) \ln(1 - x_i), \quad (4)$$

³<https://github.com/chinhsuanwu/coatnet-pytorch>.

⁴<https://github.com/pytorch/vision/blob/main/torchvision/models/maxvit.py>.

where the index i runs through all $N = 64$ images in the batch, y_i is the true label, and x_i is the predicted label. Adam [65] is chosen as the optimizer. All taggers are trained for a total of 15 epochs.

The minimum validation losses for AK08, AK10, and AK14 jets are found to be 0.6917, 0.6918, and 0.6920, respectively. When the classifier is tasked to assign the label 0 to the first class (e.g., CP2 tune) and the label 1 to the second class (CP5 tune) and the classifier is perfectly confused (i.e., unable to distinguish between the classes), it will assign labels close to 0.5 regardless of the true class. The binary cross-entropy per image is then $\ln(2) = 0.6931$. Evidently, our observed losses are only very slightly below the value expected for a classifier that learns nothing. We therefore conclude that the difference in Pythia tunes can be neglected in the following.

We now turn to the actual training of the taggers to select LSP-like fat jets. Signal samples are generated for $\tilde{\chi}$ masses between 100 and 500 GeV in 10 GeV steps, and for stop masses between 700 and 1200 GeV in 25 GeV steps. For each combination of stop and neutralino mass, we take 4750 sample images from $\tilde{\tau}_1 \tilde{\tau}_1^*$ signal events. In order to generate an almost pure signal sample for training, we only include images of jets that are within $\Delta R < 0.5$ of a parton level $\tilde{\chi}$. Since we want the LSP tagger to work for all combinations of $m_{\tilde{\tau}_1}$ and $m_{\tilde{\chi}}$, we combine all $41 \times 21 \times 4750 = 4\,089\,750$ images into a single training set. We take the same number of background images, 4 089 750, from $t\bar{t}$ + jets events.

Finally, we split the 8 179 500 images into 5 725 650 images for training and 2 453 850 images for validation. After training the model state at the epoch with the lowest validation loss is selected to define the tagger.

VIII. RESULTS FOR NEUTRALINO TAGGERS

In order to compare the performance of our classifiers, we neglect any systematic uncertainties and define the signal significance Z as

$$Z = \frac{S}{\sqrt{B}} = \frac{\epsilon_S}{\sqrt{\epsilon_B}} \cdot \frac{\sigma_S}{\sqrt{\sigma_B}} \sqrt{\mathcal{L}_{\text{int}}}, \quad (5)$$

where S and B are the number of signal and background samples passing a cut (e.g., on the value of an output neuron of the MLP), $\epsilon_{S/B}$ is the selection efficiency of this cut, $\sigma_{S/B}$ is the fiducial production cross section, and \mathcal{L}_{int} is the integrated luminosity of the data set considered. Instead of comparing signal significances directly, we compare the significance improvement $\epsilon_S/\sqrt{\epsilon_B}$, which captures the gain due to the sophisticated event classifiers and is independent of the assumed luminosity.

Figure 4 shows the performance of all neutralino taggers on the entire test data set, i.e., with all signal masses present and with the three leading jets mixed, as mentioned in

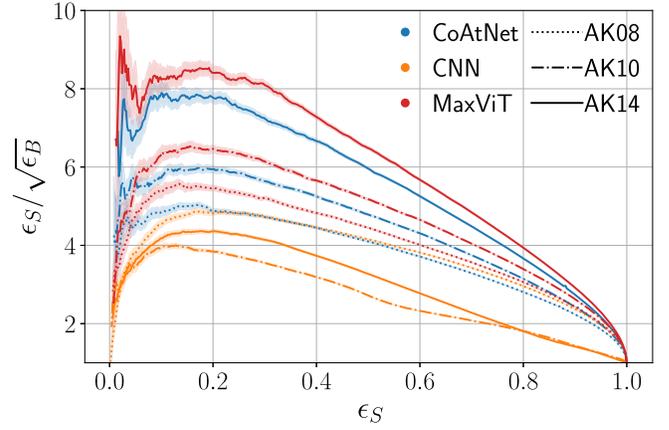


FIG. 4. Significance improvement curves for all three neutralino taggers for all single jet samples in the test data set. The shaded regions show one bootstrapped standard deviation.

Sec. VII. As a working point for the following analysis, we choose the cut on the MLP output neuron such that $\epsilon_S = 0.3$. Even lower values of ϵ_S can still increase ϵ_S/ϵ_B , but the significance improvement is already close to the maximum at the chosen point. Moreover, for smaller ϵ_S the background efficiency ϵ_B becomes so small that the statistical uncertainty on the accepted background becomes sizable, in spite of the large number of generated background events.

Both CoAtNet and MaxViT show superior performance in classical image classification tasks compared to CNN-based models, as was reported in the respective original publications. We expect this to carry over to jet classification. Indeed, this is the case here and both models outperform the classical CNN by up to a factor of 2 for AK14 jets. The most performant classifiers are the transformer-based models trained on the large-radius jets. These large jets still contain the entire narrow jets from small LSP masses, while the small jets might miss important features for larger neutralino masses. We also observe that MaxViT performs slightly better than CoAtNet, as is the case in the original MaxViT publication. Evidently, improvements in modern computer vision translate well to the classification of jet images. Even the worst transformer-based model (i.e., AK08 CoAtNet) matches the best CNN. Interestingly, despite the transformer models showing a clear hierarchy (the larger a jet is, the better), this is not the case for the CNN, which performs best for AK08 jets.

So far we have considered classification of singlet jets. In the next section, we will show how this can be used for event classification.

IX. BOOSTED CLASSIFIERS

As previously mentioned, our signal model always produces two neutralinos that subsequently decay into three quarks (plus possible gluons from final-state

radiation). It is therefore instructive to combine multiple jet images into our predictions. To this end, we apply one of our LSP taggers described above to the three leading fat jets in an event; from now on, we drop the merging requirement since it is not meaningful anymore. The three resulting MLP outputs are used as inputs for a GBDT classifier. If an event contains less than three fat jets with $p_T > 100$ GeV we assign the label -1 for the missing jets. The GBDT is implemented using the XGBoost [66] package. We use 120 trees with a learning rate of 0.1 with other hyperparameters left unchanged from the default values. In order to train the GBDT and calculate its results, we use 3000 and 2500 events, respectively, for each combination of stop and LSP masses. This corresponds to a total of 2 583 000 signal events for training and 2 152 500 signal events for calculating results. We again generate an equal number of $t\bar{t}$ + jets background events.

Figure 5 shows the significance improvement after a cut on the signal probability given by the GBDT. The difference in performance between the two transformer-based models has shrunk significantly for all jet sizes, especially for AK10 and AK08 jets. Comparing this with Fig. 4, the gain by combining the three jets is not very large. One has to keep in mind that the merging requirement is now dropped. If we calculate the significance improvement for only the jet with the highest p_T without requiring it to be close to a (truth-level) LSP, MaxViT reaches 6.79 ± 0.05 at $\epsilon_S = 0.3$. Comparing this with 9.92 ± 0.12 for the same base model after combining the LSP tagger output for the three hardest jets shows an improvement of almost 50%, equivalent to doubling the integrated luminosity in Eq. (5). The CNN now also works best with AK14 jets, even though the AK08 version is still better than the AK10 version, contrary to the hierarchy of the other models.

Overall, the level of improvement between the results of Fig. 5, which use information from up to three jets per

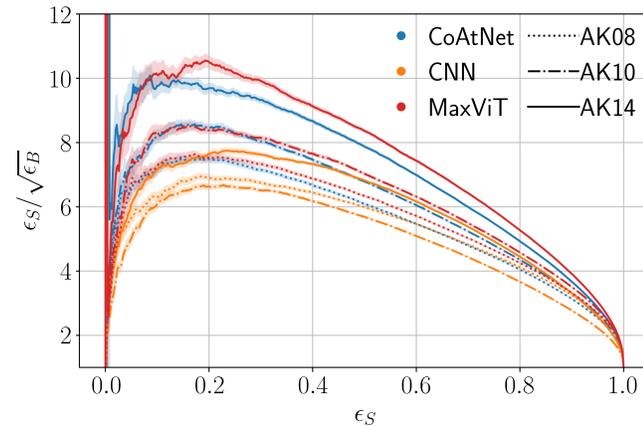


FIG. 5. Significance improvement curves for all GBDT classifiers built to combine the LSP tagger outputs for the three highest- p_T jets. The shaded regions are one bootstrapped standard deviation.

event, and Fig. 4 for single jets, might seem somewhat disappointing. After all, in the absence of QCD radiation a $t\bar{t}$ signal event contains two signal jets plus one fat background jet from the hadronically decaying top quark, whereas a generic $t\bar{t}$ event with one top quark decaying semileptonically contains only a single background fat jet. In such a situation simply requiring at least one fat jet to be tagged as signal would increase the signal efficiency (for $\epsilon_S \gg \epsilon_B$) from ϵ_S to $1 - (1 - \epsilon_S)^2$, while the background efficiency remains unchanged. However, recall that we require each event to contain at least seven AK04 jets. This greatly reduces the $t\bar{t}$ background since at least three additional partons need to be emitted for the event to pass this cut; on the other hand, it also means that background events frequently contain several fat jets, in which case a simple single tag requirement would not increase the significance. In any case, as noted above, there is a significant improvement in performance when information of the three leading fat jets is combined using a GBDT; of course, the GBDT output is not equivalent to simply demanding a fixed number of jets in a given event being tagged as LSP-like.

In Fig. 6 we show how the performance of the GBDT depends on the LSP mass. For small masses, the two transformer-based models perform comparably for all three jet sizes. Here the decay products are usually contained even in the AK08 jet so all three jet sizes contain the necessary information for our task. Evidently, the transformer networks are able to filter out the noise from particles not related to LSP decay that are present in the AK10 and AK14 jets, while the simpler CNN cannot; hence, the GBDT using the CNN applied to AK10 or AK14 jets performs relatively poorly for a small LSP mass. On the other hand, for an LSP mass above 200 GeV the GBDT performs significantly worse when used on the smaller jets,

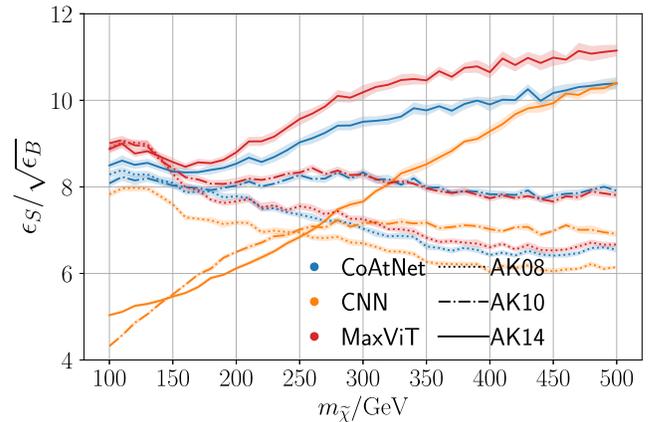


FIG. 6. Significance improvements depending on the LSP mass for all GBDT classifiers built to combine the LSP tagger outputs for the three highest- p_T jets. The cut on the GBDT output has been set such that $\epsilon_S = 0.3$ for each given LSP mass. The shaded regions are one bootstrapped standard deviation.

which no longer contain all particles originating from LSP decay.

We also note that using the CNN applied to AK14 jets performs far worse than the other models for small LSP mass, but matches the performance of the CoAtNet-based model for $m_{\tilde{\chi}}$ between 450 and 500 GeV. This curve also shows the strongest LSP mass dependence. We will revisit this point later in this section.

Finally, while the MaxViT architecture with AK10 and AK14 jets again shows the best overall performance, the resulting $\epsilon_S/\sqrt{\epsilon_B}$ shows a shallow minimum at $m_{\tilde{\chi}} \simeq m_t$. For a given p_T , fat jets originating from LSP and top decay will then have similar overall features, and the additional information about the jet mass will not help at all; moreover, recall that in our scenario the LSP decay products contain exactly one b quark, just like nearly all jets from top decay. Nevertheless, the model performs quite well even in this difficult mass region. Presumably, it exploits the fact that top decays into three quarks proceed via two two-body decays with a color-singlet on-shell W boson in the intermediate state, whereas the LSP decays via the exchange of a (far) off-shell squark.

At this point, we still have nine predictions for each event (the output of three architectures applied to AK08, AK10, and AK14 jets). Of course, these nine numbers are highly correlated. Nevertheless, a further improvement of the performance might be possible by either combining results from different jet definitions within a given architecture or vice versa. Comparing these results might also allow us to infer in which aspect a single model has room for improvements that might be gained by another architecture.

We start by combining LSP tagger outputs for different jet sizes. We show the results in Table I and compare the performance to that of the best single jet definition, which is achieved for AK14 jets, as we saw in Fig. 5. Evidently, the improvement is barely statistically significant for the two transformer-based models. These models extract most of the useful information from the images of the large AK14 jets, even when there is a lot of clutter present. The improvement is larger for the CNN-based classifier, which, however, still performs somewhat worse than the other models. It seems to benefit from the multiple jet definitions intended to extract high-level features such as the mass in

TABLE I. Significance improvement, $\epsilon_S/\sqrt{\epsilon_B}$, for $\epsilon_S = 0.3$ when only using AK14 jets (second column) and when combining the LSP tagger outputs on AK08, AK10, and AK14 jets using a larger GBDT (third column). The uncertainties are bootstrapped standard deviations.

Model	AK14	Combined jets
CoAtNet	9.32 ± 0.10	9.63 ± 0.10
MaxViT	9.91 ± 0.11	10.09 ± 0.12
CNN	7.62 ± 0.06	9.16 ± 0.09

TABLE II. Significance improvement with $\epsilon_S = 0.3$ when feeding the outputs of all three LSP taggers simultaneously to the GBDT, keeping the jet definition fixed. For comparison, the third column shows the significance improvement for the MaxViT-based model, which performs best for all three jet sizes. The uncertainties are bootstrapped standard deviations.

Jet	Combined	Best single model
AK08	7.53 ± 0.06	7.34 ± 0.06 (MaxViT)
AK10	8.94 ± 0.09	8.15 ± 0.07 (MaxViT)
AK14	10.51 ± 0.11	9.91 ± 0.11 (MaxViT)

classical applications. In particular, the combination allows to compensate the degraded performance when using the large jets for an LSP mass below 250 GeV by information from the AK08 jets, which is more useful in this parameter region, as we saw in Fig. 6.

Next, we combine the outputs of different LSP taggers into a single GBDT, for fixed jet definition. The results are shown in Table II. This time the combination leads to a slight but significant improvement over the best single model (the one based on MaxViT). This shows that, even though the transformer-based models perform almost equally well for the AK14 jets while the CNN is noticeably weaker, each model misses complementary information that the GBDT can combine into a stronger classifier.

In Fig. 7 we show how the performance of various strategies to combine LSP taggers varies with the neutralino

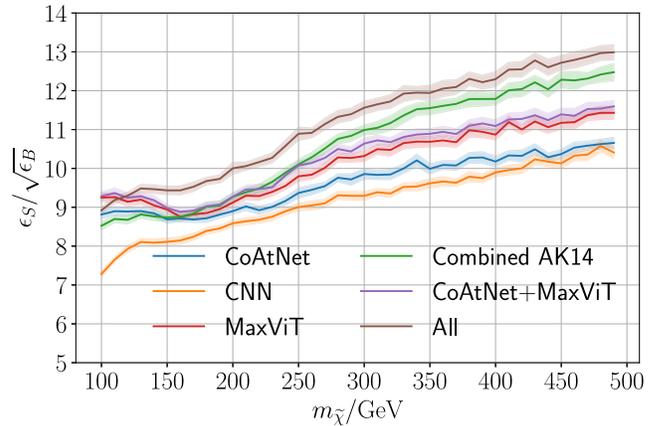


FIG. 7. Significance improvement as a function of the LSP mass for GBDT classifiers built by combining the output of different LSP taggers, with $\epsilon_S = 0.3$ in each case. The shaded regions are one bootstrapped standard deviation. The curves labeled CoAtNet (blue), CNN (orange), and MaxViT (red) show the performance of GBDTs built from combining the jet sizes for the given model, as in the third column of Table I. The green curve is for the GBDT that uses the outputs of all LSP taggers, but only for the AK14 jets, as in the third row of Table II. The purple line results from combining both transformer-based LSP taggers for all jet sizes, while the brown line is for a GBDT that combines all LSP taggers and all jet sizes.

mass. Combining all transformer-based predictions into a single GBDT does not show any significant improvement over the performance of the MaxViT -based tagger. This indicates that these models use the same features of the jet images and do not find complementary information. The combination of all CNN predictions is comparable to the weaker transformer-based model, CoAtNet , for an LSP mass above 200 GeV, while MaxViT is still more sensitive for all LSP masses.

Because our LSP taggers generally perform best on AK14 jets, we also show the combination of all three architectures using only AK14 jets, as in the last row of Table II. Comparing to Fig. 6, we see that for an LSP mass below ~ 160 GeV this combination does not further improve on the MaxViT -based model applied to AK14 jets. Between ~ 160 GeV and ~ 300 GeV, the performance closely follows that of the two combined transformer models shown in purple. Since we already showed that one does not gain much by combining the CoAtNet and MaxViT models, this shows that the CNN does not yield useful information in this region of parameter space either.

However, as we saw in Fig. 6, the CNN-based model applied to AK14 jets improves more with increasing LSP mass than the transformer-based models do, even matching CoAtNet at 500 GeV. The combination profits from this fact and outperforms above 300 GeV the GBDTs using only input from the transformer-based LSP taggers. This shows that the CNN learns something about the sample that the other models miss.

Finally, we show the result of a GBDT that is trained on the LSP tagger outputs of all three models and all three jet sizes, and thus has 27 inputs in total for each event. Compared to the AK14-only case, this does benefit from the inclusion of smaller jets, in particular at smaller LSP masses where the AK08 and AK10 jets already capture most LSP decay products. For larger LSP masses, the performance is only slightly better than that of the AK14-only case.

These various comparisons show that for the given signal process, the largest improvement in significance $\epsilon_S/\sqrt{\epsilon_B}$ is achieved by the transformer-based models applied to AK14 jets. Both models capture details of the jet images that the CNN misses. Nevertheless, also feeding the output of the CNN-based LSP tagger into a larger GBDT leads to a further slight improvement of the performance. This indicates that one might be able to find new architectures that perform even better than MaxViT .

X. ADDING HIGH-LEVEL FEATURES

The cuts discussed in Sec. III are only preselections. They ensure that the event passes the single lepton trigger and contains at least one fat jet to which the LSP tagger can be applied. They also reduce the background, but even after including information from the LSP tagger these cuts are not likely to yield the optimal distinction

between signal and background. A full event has additional features that allow to define additional, potentially useful cuts, even if they may show some correlation with the output of the LSP tagger.

In particular, so far the only dimensionful quantities we used in the construction of our classifier are the masses of the hardest three fat jets, which we use as input of the LSP tagger. Therefore, we now introduce as additional input variables for the final GBDT the sum of the masses of all AK14 jets [67],

$$M_J = \sum_{\text{AK14}} m, \quad (6)$$

and the total missing transverse momentum p_T^{miss} . In addition, we use the total number N_j of all AK04 jets as well as the scalar sum H_T of their transverse momenta.

Moreover, information about the angular separation of the jets might be helpful. Inspired by Ref. [49], we capture this information via the Fox-Wolfram moments [68] H_l , defined by

$$H_l = \sum_{i,j=1} \frac{p_{Ti} p_{Tj}}{(\sum_k p_{Tk})^2} P_l(\cos \Omega_{ij}), \quad (7)$$

where i, j, k run over all AK04 jets in the event, p_{Ti} is the p_T of the i th jet, P_l is the Legendre polynomial, and

$$\cos \Omega_{ij} = \cos \theta_i \cos \theta_j + \sin \theta_i \sin \theta_j \cos(\phi_i - \phi_j) \quad (8)$$

is the cosine of the opening angle between the jets i and j .

We combine these features into two sets. The first is the small set $\text{DS1} = [p_T^{\text{miss}}, H_T, M_J, N_j]$, which includes the most commonly used features for new physics searches in hadronic final states. In addition, we also consider a slightly larger set DS2 , which also includes the second to sixth Fox-Wolfram moments. We combine these features with the output of the LSP tagger based on MaxViT applied to AK14 jets (i.e., the most performant single model) and derive predictions with a similar GBDT as before.

Results are shown in Fig. 8. We see that even GBDT classifiers that only use the kinematic information of sets DS1 or DS2 are quite capable of separating signal from background, especially for larger LSP masses; this reconfirms the usefulness of these variables for new physics searches at the LHC. In fact, for an LSP mass above 300 GeV these classifiers even outperform the GBDT that only uses information from the MaxViT -based LSP tagger. On the other hand, except for $m_{\tilde{\chi}} = 100$ GeV, adding kinematic information to the output of the LSP tagger clearly improves the performance of the event classifier; the Fox-Wolfram moments prove useful for an LSP mass above 250 GeV or so.

Conversely, adding information from the LSP tagger to the purely kinematic variables raises the significance improvement by an amount that is nearly independent of

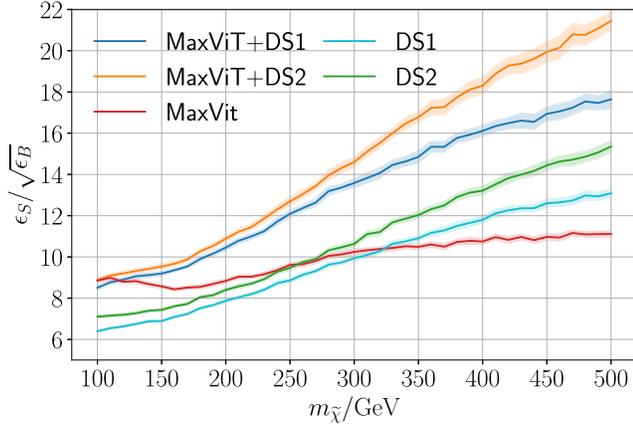


FIG. 8. Significance improvements as functions of the LSP mass for various GBDT classifiers. In all cases, the cut on the GBDT output has been set such that $\epsilon_S = 0.3$ for each given LSP mass. The upper two curves show results from classifiers that combine the output of the MaxViT-based LSP tagger applied to AK14 jets with additional kinematical features. The feature set DS1 contains $[p_T^{\text{miss}}, H_T, M_J, N_j]$, while DS2 contains, in addition, the second to sixth Fox-Wolfram moments. For comparison, the red curve results when using only LSP tagger information, as in Fig. 6, while the lower blue and green curves are for GBDTs that only use kinematical information. The shaded regions are one bootstrapped standard deviation.

the LSP mass. We expect the gain of performance to be even larger when compared to a classical selection based purely on kinematical cuts.

XI. APPLICATION AT 137 fb⁻¹

We are now ready to discuss how the different classifiers fare in terms of the reach in stop mass for exclusion or discovery. Here we set the integrated luminosity to $\mathcal{L}_{\text{int}} = 137 \text{ fb}^{-1}$, as in the original CMS publication [49]. For simplicity, we ignore the systematic uncertainty on the signal as well as the uncertainty from the finite size of our Monte Carlo samples. The former is much less important than the systematic error on the background estimate, and the latter should be much smaller than the statistical uncertainty due to the finite integrated luminosity. The $t\bar{t}$ background is normalized to the next-to-leading-order (NLO) production cross section [50]. The simulated stop pair samples are normalized to NLO + next-to-leading-logarithm accuracy [69]. This corresponds to 273 084 background events and a stop mass-dependent number of signal events. We calculate exclusion limits from the expected exclusion significance [70]:

$$Z_{\text{excl}} = \left[2S - 2B \ln\left(\frac{B+S+x}{2B}\right) - \frac{2B^2}{\Delta_B^2} \ln\left(\frac{B-S+x}{2B}\right) - (B+S-x) \frac{B+\Delta_B^2}{\Delta_B^2} \right]^{1/2}, \quad (9)$$

where B and S are the expected number of background and signal events, Δ_B is the absolute systematic uncertainty on the background, and

$$x = \sqrt{(S+B)^2 - \frac{4SB\Delta_B^2}{B+\Delta_B^2}}. \quad (10)$$

We chose $\Delta_B = 0.06B$, as described below. Z_{excl} is the expected number of standard deviations with which the predicted signal S can be excluded if the background-only hypothesis, described by the background B , is correct; note that $Z_{\text{excl}} \rightarrow S/\sqrt{B+\Delta_B^2}$ if $B \gg S$. This quantity is computed for every combination of stop and LSP masses introduced in Sec. III, using four different event classifiers.

The results are shown in Fig. 9. We again define signal-like events through a cut on the GBDT output corresponding to $\epsilon_S = 0.3$. The top-left frame is for a GBDT that uses only kinematic information about the AK04 jets, as in the green curve of Fig. 8. Associating the contour along $Z_{\text{excl}} = 1.645$ with the 95% confidence level exclusion bounds of this “traditional” analysis, we find an expected exclusion reach in $m_{\tilde{t}_1}$ of about 740 GeV for an LSP mass of 100 GeV. This is rather close to the expected reach of about 710 GeV for the same LSP mass achieved in the CMS search,⁵ which is based on a neural network “trained to recognize differences in the spatial distribution of jets and decay kinematic distributions” [49]. Unfortunately, they do not show results for other LSP masses. This agreement is not accidental; we chose the systematic background uncertainty, $\Delta_B = 0.06B$, accordingly. Presumably, even closer agreement would have been possible for somewhat larger Δ_B . However, it would then be significantly larger than the actual systematic error on the background estimate found by CMS, which is below 5%. We note that for $\Delta_B^2 \gg B$ the significance scales $\propto 1/B$, rather than $\propto 1/\sqrt{B}$, if Δ_B is a fixed percentage of B . A larger Δ_B therefore increases the relative improvement in reach achieved by including information from one of our LSP taggers; recall that this leads to a significant improvement of $\epsilon_S/\sqrt{\epsilon_B}$, and hence to an even bigger improvement in ϵ_S/ϵ_B .

The other three frames show results for GBDTs that also use the outputs of an LSP tagger as input variable; we apply this tagger to the three leading AK14 jets. We see that the simpler CNN-based tagger (top right) increases the reach in stop mass only by less than 10 GeV. Recall from Fig. 6 that the CNN tagger applied on AK14 jets does not perform well for a small LSP mass. For a larger LSP mass, and hence larger angular spread of the LSP decay products, the kinematic information on the AK04

⁵We note in passing that the actual CMS limit on the stop mass is only 670 GeV for this LSP mass, due to a small (not statistically significant) excess of events.

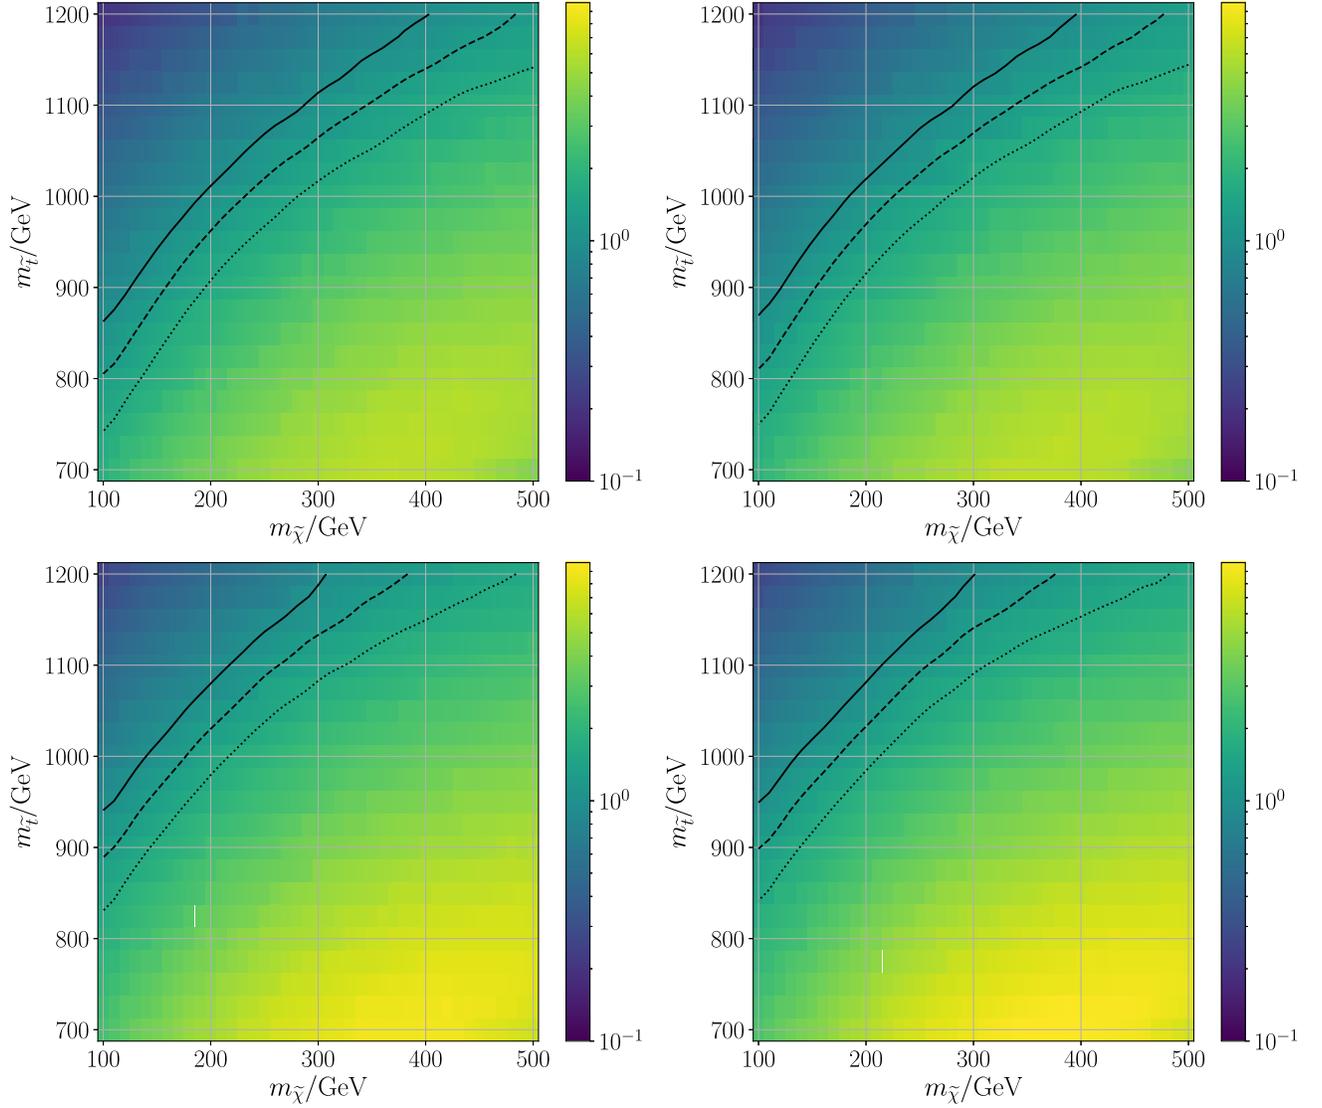


FIG. 9. Exclusion significance Z_{excl} defined in Eq. (9) for an integrated luminosity of 137 fb^{-1} . In all cases, the cut on the GBDT output has been chosen such that the signal efficiency $\epsilon_S = 0.3$ and $\Delta_B = 0.06B$. The top-left frame is for a GBDT using kinematical information only, corresponding to the green curve in Fig. 8. The other three frames are for GBDTs that also use the output of LSP taggers applied to AK14 jets, based on the CNN (top right), CoAtNet (bottom left), and MaxViT (bottom right). Solid, dashed, and dotted lines denote contour lines corresponding to a signal significance of 1, 1.281, and 1.645, respectively. These are smoothed by a Gaussian filter with standard deviations 10 and 25 GeV on the neutralino-mass and stop-mass axis, respectively, applied to the logarithm of the signal significances.

jets, many of which are components of AK14 jets, already seems to capture much of the physics found by the CNN. Recall that the kinematic GBDT includes information on the angular separation of these jets via the Fox-Wolfram moments of Eq. (7).

In contrast, using the transformer-based LSP taggers does improve the reach considerably. As before, MaxViT (bottom right) performs slightly better than CoAtNet (bottom left); the reach in stop mass increases by 100 GeV for $m_{\tilde{\chi}} = 100 \text{ GeV}$ and by about 60 GeV for $m_{\tilde{\chi}} = 500 \text{ GeV}$. This again indicates that the kinematic information on the AK04 jets allows some effective LSP tagging for large LSP masses.

For stop masses in the interesting range, the $\tilde{t}_1\tilde{t}_1^*$ production cross section of [69] can be roughly parametrized as

$$\sigma(pp \rightarrow \tilde{t}_1\tilde{t}_1^*) \simeq 0.08 \text{ pb} \cdot \left(\frac{m_{\tilde{t}_1}}{700 \text{ GeV}}\right)^{-7.8}. \quad (11)$$

Increasing the reach from 740 to 840 GeV (for $m_{\tilde{\chi}} = 100 \text{ GeV}$) thus corresponds to reducing the bound on the stop pair-production cross section by a factor of ~ 2.7 . Note that the limit-setting procedure is quite nonlinear because the background falls by nearly 2 orders of

magnitude when $m_{\tilde{\tau}_1}$ is increased from 700 to 1200 GeV, while keeping $\epsilon_S = 0.3$ fixed.

XII. CONCLUSION

The large hadronic activity in pp collisions makes the search for physics beyond the SM in purely hadronic processes at the LHC especially challenging. This problem can be mitigated by the use of sophisticated analysis methods. In particular, jet substructure has proved a powerful discriminator between various production processes.

In this article, we studied the feasibility of applying modern computer vision techniques in detecting RPV stop decays. As a benchmark, we used $\tilde{\tau}_1$ pair production, where each stop decays to a top and a neutralino LSP, which subsequently decays via the UDD operator to three quarks. For not too small mass splitting between the stop and the LSP, the decay products of the latter tend to reside in a single fat (e.g., AK14) jet. One can build images from the constituents of such jets by using the angle ϕ and pseudorapidity η as spatial positions and deposited energy into the detector as pixel intensity. One can then use computer vision techniques on this representation to build classifiers (“LSP taggers”) that aid in amplifying the signal process.

In recent years, transformer-based architectures have been shown to trump the performance of more classical convolutional neural network–based structures in standard classification tasks. We studied how well these novel architectures work on jet images by training LSP taggers based on MaxViT, CoAtNet, and a CNN architecture. The training was done on single jet images. The combined training time of all nine models was less than 1 month on a consumer-grade RTX 3060 Ti GPU. Therefore, the results can be obtained without specialized computing resources. We then combined the output of the LSP tagger applied to the three jets with the highest p_T using a gradient-boosted decision tree into a more robust classification score. We found that the CNN-based tagger improves the statistical significance of the signal by a factor between 5 and 10 for fixed signal efficiency $\epsilon_S = 0.3$, with the exact factor depending on the neutralino mass and the definition of the fat jets. In contrast, the transformer-based models lead to an improvement factor between 8 and 11, outperforming the CNN over the entire parameter space.

In addition, we showed a method to quantify the systematic uncertainty implied by the different Pythia tunes between background and signal events. For this, we trained a CNN to tell background events generated with the CP2 tune apart from background events with the CP5 tune. We found no significant difference between both tunes and concluded that the systematic uncertainty can be neglected.

Although the modern models are highly complex and therefore hard to interpret, one might speculate that the reason for the large improvement lies in the properties of the jet images. Due to the large neutralino masses, the

decay products are spread far apart. An efficient classifier needs to leverage the correlation of all decay products, which is therefore a global feature. This is seen by the transformer-based models earlier due to the global receptive field of view.

We also combined the predictions of all architectures for each jet size separately and found a modest improvement, hinting that even the transformer-based models do not use the entire information present in the images; hence, an investigation of further improvements of the architecture might be worthwhile.

Since the kinematic preselection cuts are not optimized for sensitivity, we also used high-level features such as Fox-Wolfram-moments, p_T^{miss} , H_T , M_J , and N_j as inputs to a GBDT, in combination with the output of one of our LSP taggers. This leads to a total gain of sensitivity by a factor of 20 for 500 GeV LSPs, on top of the effect due to the acceptance cuts.

Finally, we estimated the reach in stop and LSP mass that could be expected from the full Run-2 data set. We chose the systematic uncertainty on the background such that a GBDT that only uses kinematic information on AK04 jets leads to a reach (for an LSP mass of 100 GeV) similar to that found by CMS [49]. Using in addition the output of the relatively simple CNN-based LSP tagger then leads to almost no further improvement of the reach. By instead using the MaxViT-based tagger, one can improve the reach by 100 GeV (60 GeV) for neutralino masses of 100 GeV (500 GeV), under the assumption that the relative size of the systematic uncertainty remains the same. This corresponds to a reduction of the bound on the stop pair-production cross section by up to a factor of 2.7.

We conclude that LSP taggers built on modern transformer-based neural networks hold great promise in searches for supersymmetry with a neutralino LSP where R parity is broken by the UDD operator. This result can presumably be generalized to models with different LSPs, e.g., a gluino decaying via the same operator or a slepton decaying into a lepton and three jets via the exchange of a virtual neutralino.

In fact, it seems likely that these advanced techniques can also be used to build improved taggers for boosted, hadronically decaying top quarks or weak gauge or Higgs bosons. We did not attempt to construct such taggers ourselves since this field is already quite mature. Convincing progress would therefore have to be based on fully realistic detector-level simulations, for which we lack the computational resources. Moreover, a careful treatment of systematic uncertainties would be required, which ideally uses real data. However, we see no reason why the improvement relative to CNN-based taggers that we saw in our relatively simple simulations should not carry over to fully realistic ones.

- [1] J. Wess and B. Zumino, Supergauge transformations in four dimensions, *Nucl. Phys.* **B70**, 39 (1974).
- [2] H. P. Nilles, Supersymmetry, supergravity and particle physics, *Phys. Rep.* **110**, 1 (1984).
- [3] H. E. Haber and G. L. Kane, The search for supersymmetry: Probing physics beyond the standard model, *Phys. Rep.* **117**, 75 (1985).
- [4] P. Fayet and S. Ferrara, Supersymmetry, *Phys. Rep.* **32**, 249 (1977).
- [5] Manuel Drees, Rohini Godbole, and Probir Roy, *Theory and Phenomenology of Particles* (World Scientific, Singapore, 2005).
- [6] Stephen P. Martin, A supersymmetry primer, *Adv. Ser. Dir. High Energy Phys.* **18**, 1 (1998).
- [7] R. L. Workman *et al.*, Review of particle physics, *Prog. Theor. Exp. Phys.* **2022**, 083C01 (2022).
- [8] ATLAS Collaboration, The quest to discover supersymmetry at the ATLAS experiment, [arXiv:2403.02455](https://arxiv.org/abs/2403.02455).
- [9] R. Barbier, C. B erat, M. Besan on *et al.*, R —Parity-violating supersymmetry, *Phys. Rep.* **420**, 1 (2005).
- [10] Andreas Redelbach, Searches for prompt R —parity-violating supersymmetry at the LHC, *Adv. High Energy Phys.* **2015**, 982167 (2015).
- [11] Jonathan L. Feng, Konstantin T. Matchev, and Takeo Moroi, Multi—TeV scalars are natural in minimal supergravity, *Phys. Rev. Lett.* **84**, 2322 (2000).
- [12] Ryuichiro Kitano and Yasunori Nomura, Supersymmetry, naturalness, and signatures at the LHC, *Phys. Rev. D* **73**, 095004 (2006).
- [13] Michele Papucci, Joshua T. Ruderman, and Andreas Weiler, Natural SUSY endures, *J. High Energy Phys.* **09** (2012) 035.
- [14] Graham G. Ross, Kai Schmidt-Hoberg, and Florian Staub, Revisiting fine-tuning in the MSSM, *J. High Energy Phys.* **03** (2017) 021.
- [15] Howard Baer, Vernon Barger, Shadman Salam, Dibyashree Sengupta, and Xerxes Tata, The LHC higgsino discovery plane for present and future SUSY searches, *Phys. Lett. B* **810**, 135777 (2020).
- [16] Georges Aad *et al.*, Search for direct production of winos and higgsinos in events with two same-charge leptons or three leptons in pp collision data at $\sqrt{s} = 13$ TeV with the ATLAS detector, *J. High Energy Phys.* **11** (2023) 150.
- [17] Georges Aad *et al.*, Search for nearly mass-degenerate higgsinos using low-momentum mildly-displaced tracks in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector, *Phys. Rev. Lett.* **132**, 221801 (2024).
- [18] Georges Aad *et al.*, Search for pair production of higgsinos in events with two Higgs bosons and missing transverse momentum in $\sqrt{s} = 13$ TeV pp collisions at the ATLAS experiment, *Phys. Rev. D* **109**, 112011 (2024).
- [19] Armen Tumasyan *et al.*, Search for higgsinos decaying to two Higgs bosons and missing transverse momentum in proton-proton collisions at $\sqrt{s} = 13$ TeV, *J. High Energy Phys.* **05** (2022) 014.
- [20] Aram Hayrapetyan *et al.*, Combined search for electroweak production of winos, binos, higgsinos, and sleptons in proton-proton collisions at $\sqrt{s} = 13$ TeV, *Phys. Rev. D* **109**, 112001 (2024).
- [21] Jared A. Evans and Yevgeny Kats, LHC coverage of RPV MSSM with light stops, *J. High Energy Phys.* **04** (2013) 028.
- [22] Jonathan M. Butterworth, John R. Ellis, Are R. Raklev, and Gavin P. Salam, Discovering baryon-number violating neutralino decays at the LHC, *Phys. Rev. Lett.* **103**, 241803 (2009).
- [23] Sebastian Macaluso and David Shih, Pulling out all the tops with computer vision and deep learning, *J. High Energy Phys.* **10** (2018) 121.
- [24] Quark versus gluon jet tagging using jet images with the ATLAS detector, Technical Report, CERN, Geneva, 2017. All figures including auxiliary figures are available at <https://atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/PUBNOTES/ATL-PHYS-PUB-2017-017>.
- [25] Tilman Plehn, Anja Butter, Barry Dillon, and Claudius Krause, Modern machine learning for LHC physicists, [arXiv:2211.01421](https://arxiv.org/abs/2211.01421).
- [26] Patrick T. Komiske, Eric M. Metodiev, and Matthew D. Schwartz, Deep learning in color: Towards automated quark/gluon jet discrimination, *J. High Energy Phys.* **01** (2017) 110.
- [27] Gregor Kasieczka, Tilman Plehn, Michael Russell, and Torben Schell, Deep-learning top taggers or the end of QCD?, *J. High Energy Phys.* **05** (2017) 006.
- [28] Electron identification with a convolutional neural network in the ATLAS experiment, Technical Report, CERN, Geneva, 2023. All figures including auxiliary figures are available at <https://atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/PUBNOTES/ATL-PHYS-PUB-2023-001>.
- [29] A. M. Sirunyan, A. Tumasyan, W. Adam *et al.*, Identification of heavy, energetic, hadronically decaying particles using machine-learning techniques, *J. Instrum.* **15**, P06005 (2020).
- [30] Huifang Lv, Daohan Wang, and Lei Wu, Deep learning jet images as a probe of light Higgsino dark matter at the LHC, *Phys. Rev. D* **106**, 055008 (2022).
- [31] Jun Guo, Jinmian Li, Tianjun Li, Fangzhou Xu, and Wenxing Zhang, Deep learning for R —parity violating supersymmetry searches at the LHC, *Phys. Rev. D* **98**, 076017 (2018).
- [32] Jason Sang Hun Lee, Inkyu Park, Ian James Watson, and Seungjin Yang, Quark-gluon jet discrimination using convolutional neural networks, *J. Korean Phys. Soc.* **74**, 219 (2019).
- [33] Jakub Filipek, Shih-Chieh Hsu, John Kruper, Kirtimaan Mohan, and Benjamin Nachman, Identifying the quantum properties of hadronic resonances using machine learning, [arXiv:2105.04582](https://arxiv.org/abs/2105.04582).
- [34] Tao Han, Ian M. Lewis, Hongkai Liu, Zhen Liu, and Xing Wang, A guide to diagnosing colored resonances at hadron colliders, *J. High Energy Phys.* **08** (2023) 173.
- [35] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov *et al.*, An image is worth 16×16 words: Transformers for image recognition at scale, [arXiv:2010.11929](https://arxiv.org/abs/2010.11929).
- [36] Vinicius Mikuni and Florencia Canelli, Point cloud transformers applied to collider physics, *Mach. Learn. Sci. Tech.* **2**, 035027 (2021).
- [37] Francesco Armando Di Bello, Etienne Dreyer, Sanmay Ganguly *et al.*, Reconstructing particles in jets using set

- transformer and hypergraph prediction networks, *Eur. Phys. J. C* **83**, 596 (2023).
- [38] Luc Builtjes, Sascha Caron, Polina Moskvitina *et al.*, Attention to the strengths of physical interactions: Transformer and graph-based event classification for particle physics experiments, [arXiv:2211.05143](https://arxiv.org/abs/2211.05143).
- [39] A. Hammad, S. Moretti, and M. Nojiri, Multi-scale cross-attention transformer encoder for event classification, *J. High Energy Phys.* **03** (2024) 144.
- [40] Huilin Qu, Congqiao Li, and Sitian Qian, Particle transformer for jet tagging, [arXiv:2202.03772](https://arxiv.org/abs/2202.03772).
- [41] Thorben Finke, Michael Krämer, Alexander Mück, and Jan Tönshoff, Learning the language of QCD jets with transformers, *J. High Energy Phys.* **06** (2023) 184.
- [42] Minxuan He and Daohan Wang, Quark/gluon discrimination and top tagging with dual attention transformer, *Eur. Phys. J. C* **83**, 1116 (2023).
- [43] A. Hammad and Mihoko M. Nojiri, Streamlined jet tagging network assisted by jet prong structure, *J. High Energy Phys.* **06** (2024) 176.
- [44] A. Hammad, P. Ko, Chih-Ting Lu, and Myeonghun Park, Exploring exotic decays of the Higgs boson to multi-photons at the LHC via multimodal learning approaches, [arXiv:2405.18834](https://arxiv.org/abs/2405.18834).
- [45] V. Khachatryan, A. M. Sirunyan, A. Tumasyan *et al.*, Search for pair-produced resonances decaying to jet pairs in proton-proton collisions at $s = 8$ TeV, *Phys. Lett. B* **747**, 98 (2015).
- [46] G. Aad, B. Abbott, J. Abdallah *et al.*, A search for top squarks with R-parity-violating decays to all-hadronic final states with the ATLAS detector in $\sqrt{s} = 8$ TeV proton-proton collisions, *J. High Energy Phys.* **06** (2016) 067.
- [47] M. Aaboud, G. Aad, B. Abbott *et al.*, A search for pair-produced resonances in four-jet final states at $\sqrt{s} = 13$ TeV with the ATLAS detector, *Eur. Phys. J. C* **78**, 250 (2018).
- [48] A. M. Sirunyan, A. Tumasyan, W. Adam *et al.*, Search for pair-produced resonances decaying to quark pairs in proton-proton collisions at $\sqrt{s} = 13$ TeV, *Phys. Rev. D* **98**, 112014 (2018).
- [49] A. M. Sirunyan, A. Tumasyan *et al.*, Search for top squarks in final states with two top quarks and several light-flavor jets in proton-proton collisions at $\sqrt{s} = 13$ TeV, *Phys. Rev. D* **104**, 032006 (2021).
- [50] J. Alwall, R. Frederix, S. Frixione, V. Hirschi, F. Maltoni, O. Mattelaer, H.-S. Shao, T. Stelzer, P. Torrielli, and M. Zaro, The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations, *J. High Energy Phys.* **07** (2014) 079.
- [51] Richard D. Ball, Valerio Bertone, Stefano Carrazza *et al.*, Parton distributions from high-precision collider data, *Eur. Phys. J. C* **77**, 663 (2017).
- [52] Christian Bierlich, Smitya Chakraborty, Nishita Desai *et al.*, A comprehensive guide to the physics and usage of Pythia 8.3, *SciPost Phys. Codebases* **2022**, 8 (2022).
- [53] A. M. Sirunyan, A. Tumasyan, W. Adam *et al.*, Extraction and validation of a new set of CMS Pythia 8 tunes from underlying-event measurements, *Eur. Phys. J. C* **80**, 4 (2020).
- [54] Michelangelo L Mangano, Mauro Moretti, Fulvio Piccinini, and Michele Treccani, Matching matrix elements and shower evolution for top-pair production in hadronic collisions, *J. High Energy Phys.* **01** (2007) 013.
- [55] J. de Favereau, C. Delaere, P. Demin, A. Giammanco, V. Lemaître, A. Mertens, and M. Selvaggi, DELPHES 3: A modular framework for fast simulation of a generic collider experiment, *J. High Energy Phys.* **02** (2014) 057.
- [56] Alexandre Mertens, New features in DELPHES 3, *J. Phys. Conf. Ser.* **608**, 012045 (2015).
- [57] Matteo Cacciari, Gavin P. Salam, and Gregory Soyez, FastJet user manual, *Eur. Phys. J. C* **72**, 1896 (2012).
- [58] A. Chakraborty, S. Dasmahapatra, H. A. Day-Hall, B. G. Ford, S. Jain, S. Moretti, E. Olaiya, and C. H. Shepherd-Themistocleous, Revisiting jet clustering algorithms for new Higgs Boson searches in hadronic final states, *Eur. Phys. J. C* **82**, 346 (2022).
- [59] David Krohn, Jesse Thaler, and Lian-Tao Wang, Jets with variable R , *J. High Energy Phys.* **06** (2009) 059.
- [60] Adam Paszke, Sam Gross, Francisco Massa *et al.*, PyTorch: An imperative style, high-performance deep learning library, [arXiv:1912.01703](https://arxiv.org/abs/1912.01703).
- [61] Ashish Vaswani, Noam Shazeer, Niki Parmar *et al.*, Attention is all you need, [arXiv:1706.03762](https://arxiv.org/abs/1706.03762).
- [62] Zihang Dai, Hanxiao Liu, Quoc V. Le, and Mingxing Tan, CoAtNet: Marrying convolution and attention for all data sizes, [arXiv:2106.04803](https://arxiv.org/abs/2106.04803).
- [63] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen, MobileNetV2: Inverted residuals and linear bottlenecks, in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2019), [arXiv:1801.04381](https://arxiv.org/abs/1801.04381).
- [64] Zhengzhong Tu, Hossein Talebi, Han Zhang *et al.*, MaxViT: Multi-axis vision transformer, [arXiv:2204.01697](https://arxiv.org/abs/2204.01697).
- [65] Diederik P. Kingma and Jimmy Ba, Adam: A method for stochastic optimization, [arXiv:1412.6980](https://arxiv.org/abs/1412.6980).
- [66] Tianqi Chen and Carlos Guestrin, XGBoost: A scalable tree boosting system, in *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '16* (ACM, 2016), [arXiv:1603.02754](https://arxiv.org/abs/1603.02754).
- [67] Anson Hook, Eder Izaguirre, Mariangela Lisanti, and Jay G. Wacker, High multiplicity searches at the LHC using jet masses, *Phys. Rev. D* **85**, 055029 (2012).
- [68] Catherine Bernaciak, Malte Seán Andreas Buschmann, Anja Butter, and Tilman Plehn, Fox-Wolfgram moments in Higgs physics, *Phys. Rev. D* **87**, 073014 (2013).
- [69] Christoph Borschensky, Michael Krämer, Anna Kulesza, Michelangelo Mangano, Sanjay Padhi, Tilman Plehn, and Xavier Portell, Squark and gluino production cross sections in pp collisions at $\sqrt{s} = 13, 14, 33$ and 100 TeV, *Eur. Phys. J. C* **74**, 3174 (2014).
- [70] Nilanjana Kumar and Stephen P. Martin, Vectorlike leptons at the Large Hadron Collider, *Phys. Rev. D* **92**, 115018 (2015).