

Deep-learned Top Tagging with a Lorentz Layer

Anja Butter¹, Gregor Kasieczka², Tilman Plehn^{1*} and Michael Russell^{1,3}

¹ Institut für Theoretische Physik, Universität Heidelberg, Germany

² Institute for Particle Physics, ETH Zürich, Switzerland

³ School of Physics and Astronomy, University of Glasgow, Scotland

* plehn@uni-heidelberg.de

Abstract

We introduce a new and highly efficient tagger for hadronically decaying top quarks, based on a deep neural network working with Lorentz vectors and the Minkowski metric. With its novel machine learning setup and architecture it allows us to identify boosted top quarks not only from calorimeter towers, but also including tracking information. We show how the performance of our tagger compares with QCD-inspired and image-recognition approaches and find that it significantly increases the performance for strongly boosted top quarks.



Copyright A. Butter *et al.*

This work is licensed under the Creative Commons

[Attribution 4.0 International License](https://creativecommons.org/licenses/by/4.0/).

Published by the SciPost Foundation.

Received 17-01-2018

Accepted 19-09-2018

Published 26-09-2018

doi:[10.21468/SciPostPhys.5.3.028](https://doi.org/10.21468/SciPostPhys.5.3.028)



Check for
updates

Content

1	Introduction	1
2	Tagger	2
2.1	Combination layer	2
2.2	Lorentz layer	3
3	Performance	4
3.1	Calorimeter	5
3.2	Learning the Minkowski metric	6
3.3	Calorimeter and tracking	7
4	Conclusions	8
	References	8

1 Introduction

The classification of hadronic objects has become the main driving force behind machine learning techniques in LHC physics. The task is to identify the partonic nature of large-area jets or

fat jets. Such jets occur for instance in boosted hadronic decays of Higgs bosons [1], weak gauge bosons [2–4], or top quarks [5–22].

A widely debated, central question is how we can analyze these jet substructure patterns using a range of machine learning techniques. An early example were wavelets, describing patterns of hadronic weak boson decays [23, 24]. The most frequently used approach is image recognition applied to calorimeter entries in the azimuthal angle vs rapidity plane, so-called jet images. They can be used to search for hadronic decays of weak bosons [25–29] or top quarks [30, 31], or to distinguish quark-like from gluon-like jets [32]. Another approach is inspired by natural language recognition, applied to decays of weak bosons [33].

Top taggers inspired by image recognition rely on convolutional networks (CNN) [31, 34], which work well for numbers of pixels small enough to be analyzed by the network. We have shown that they can outperform multi-variate QCD-based taggers, but also that the CNN learns all the appropriate sub-jet patterns [31]. A major problem arises when we include tracking information with its much better experimental resolution, leading to too many, too sparsely distributed pixels [32].

We propose a new approach to jet substructure using machine learning: rather than relying on analogies to image or natural language recognition we analyze the constituents of the fat jet directly, only using the Lorentz group and Minkowski space-time. For our DEEPTOPLOLA tagger we introduce a Combination layer (CoLa) together with a Lorentz layer (LoLa) and two fully connected layers forming a novel deep neural network (DNN) architecture. In the standard setup the input 4-momenta correspond to calorimeter towers [35]. However, unlike other approaches the DEEPTOPLOLA tagger can be extended to include tracking information and particle flow objects with their full experimental resolution in a technically trivial way.

This flexible setup allows us to study how much performance gain tracking information actually gives. Moreover, it means that DEEPTOPLOLA can be immediately included in state-of-the-art ATLAS and CMS analyses and can be combined with b -tagging.

In this letter we first introduce our new machine learning setup. Using standard fat jets from hadronic top decays we compare its performance to multivariate QCD-inspired tagging and an image-based convolutional network [31]. We then extend the tagger to include particle flow information and estimate the performance gain compared to calorimeter information for mildly boosted and strongly boosted top quarks.

2 Tagger

The basic constituents entering any subjet analysis are a set of N measured 4-vectors sorted by p_T , for example organized as the matrix

$$(k_{\mu,i}) = \begin{pmatrix} k_{0,1} & k_{0,2} & \cdots & k_{0,N} \\ k_{1,1} & k_{1,2} & \cdots & k_{1,N} \\ k_{2,1} & k_{2,2} & \cdots & k_{2,N} \\ k_{3,1} & k_{3,2} & \cdots & k_{3,N} \end{pmatrix}. \quad (1)$$

We show a typical jet image for a hadronic top decay in Fig. 1, indicating that the calorimeters entries of a typical top decay form a sparsely filled image.

2.1 Combination layer

Our tagger consists of two physics-inspired modules. As a first step, we multiply the 4-vectors from Eq.(1) with a matrix C_{ij} . Inspired by the treatment of jet clustering in the Qjets ap-




Figure 1: Jet image illustrating a signal event, showing 20 4-vectors $k_{\mu,i}$ with an energy threshold $k_0 > 1$ GeV on the calorimeter level.

proach [36] this defines our Combination layer

$$k_{\mu,i} \xrightarrow{\text{CoLa}} \tilde{k}_{\mu,j} = k_{\mu,i} C_{ij} . \quad (2)$$

It returns M 4-vectors $\tilde{k}_j$, so $i = 1 \dots N$ and $j = 1 \dots M$. From many top tagging tests we known that an efficient tagger needs to find the mass drops associated with the top decay and the W decay [13, 14, 31]. For illustration purposes, we look at the two corresponding on-shell conditions in our framework,

$$\begin{aligned} \tilde{k}_{\mu,1}^2 &= (k_{\mu,1} + k_{\mu,2} + k_{\mu,3})^2 = m_t^2 \\ \tilde{k}_{\mu,2}^2 &= (k_{\mu,1} + k_{\mu,2})^2 = m_W^2 . \end{aligned} \quad (3)$$

They correspond to non-zero entries

$$C_{11} = C_{21} = C_{31} \quad \text{and} \quad C_{12} = C_{22} . \quad (4)$$

In general, the CoLa matrix in our neural network has the trainable form

$$C = \begin{pmatrix} 1 & 0 & \cdots & 0 & C_{1,N+2} & \cdots & C_{1,M} \\ 0 & 1 & & \vdots & C_{2,N+2} & \cdots & C_{2,M} \\ \vdots & \vdots & \ddots & 0 & \vdots & & \vdots \\ 0 & 0 & \cdots & 1 & C_{N,N+2} & \cdots & C_{N,M} \end{pmatrix} . \quad (5)$$

It guarantees that the set of M 4-momenta $\tilde{k}_j$ includes

1. each original momentum k_i ;
2. a trainable set of $M - N$ linear combinations.

These $\tilde{k}_j$ will be analyzed by a DNN. While one could use advanced pre-processing beyond some kind of ordering of the input 4-momenta, our earlier study [31] suggests that this is not necessary. For our numerical study we vary N according to physics scenario. We use 15 trainable combinations, or $M = 15 + N$, where we have checked that changing M has no effect.

2.2 Lorentz layer

From fundamental theory we know that the relevant distance measure between two substructure objects is the Minkowski metric. We use it to construct a weight function which makes it

easier for the DNN to learn the underlying features.* The Lorentz layer as the second part of the DNN first transforms the M 4-vectors $\tilde{k}_j$ into the same number of measurement-motivated objects $\hat{k}_j$,

$$\tilde{k}_j \xrightarrow{\text{LoLa}} \hat{k}_j = \begin{pmatrix} m^2(\tilde{k}_j) \\ p_T(\tilde{k}_j) \\ w_{jm}^{(E)} E(\tilde{k}_m) \\ w_{jm}^{(d)} d_{jm}^2 \end{pmatrix}. \quad (6)$$

The four entries illustrate different structures we can include in this Lorentz layer. The first two $\hat{k}_j$ map individual $\tilde{k}_j$ onto their invariant mass and transverse momentum. The invariant mass entry corresponds to the illustration in Eq.(4). The third entry constructs a linear combination of all energies, with a trainable vector of weights $w_{jm}^{(E)}$ with $m = 1 \dots M$. Different values of j give us a set of M copies of this linear combination. The fourth entry combines all $\tilde{k}_m$ with a fixed $\tilde{k}_j$, including a trainable vector of weights $w_{jm}^{(d)}$. We can either sum over or minimize over the internal index m , always keeping the external index j fixed. For the third entry with the trainable weights $w_{jm}^{(E)}$ we choose the sum over the internal index. For the last entry with the weights $w_{jm}^{(d)}$ we improve the performance of the network by including four copies with independently trainable weights. Two of these copies sum over the internal index and two of them minimize over it.

We have checked that neither the exact composition of the $\hat{k}_j$ nor the number of entries in Eq.(6) have an effect on the performance of our tagger. What is important is that we combine the invariant mass with an energy or transverse momentum and include the trainable weights. The first and last entries in Eq.(6) explicitly use the Minkowski distance,

$$d_{jm}^2 = (\tilde{k}_j - \tilde{k}_m)_\mu g^{\mu\nu} (\tilde{k}_j - \tilde{k}_m)_\nu. \quad (7)$$

The LoLa objects $\hat{k}_j$ are the input of the DNN. One can think of them as a rotation in the observable space, making the relevant information more accessible to the neural network, so the LoLa should be loss-less, provided the truncation in the number of input 4-vectors and the selection in Eq.(6) is carefully tested. Finally, the combined set of trainable weights in Eq.(5) and in Eq.(6) is large and can most likely be reduced for a given application. To maintain the general structure of our approach we decide to not apply this optimization.

3 Performance

For any proposed new analysis tool, a realistic and convincing comparison with the state-of-the-art tools is crucial. For our DEEPTOPLOLA tagger we compare its performance with a QCD-inspired top tagger and with an image-based top tagger, both working on calorimeter entries.

For our comparison we simulate a hadronic $t\bar{t}$ sample and a QCD di-jet sample with PYTHIA8.2.15 [37] for the 14 TeV LHC [38]. We ignore multi-parton interactions and in particular pile-up, which can eventually be removed [39–41]. Moreover, we assume that our top tagger can be trained on a pure sample of lepton-hadron top pair events with an identified leptonic top decay.

All events are passed through the fast detector simulation DELPHES3.3.2 [42], with calorimeter towers of size $\Delta\eta \times \Delta\phi = 0.1 \times 5^\circ$ and an energy threshold of 1 GeV. We cluster these

*We are grateful to Johann Brehmer for pointing out that this approach limits us to fat jets far from black holes.




Figure 2: Number of constituents (left) and mean of the transverse momentum (right) of the ranked constituents available as 4-vectors in Eq.(1). We show 4-vectors for the top signal from calorimeter cells or jet images (dashed) and from calorimeter and tracker information combined through particle flow (solid).

towers with FASTJET3.1.3 [43, 44] to anti- k_T [45] jets with $R = 1.5$. This defines a smooth outer shape and a jet area of the fat jet. The fat jets have to fulfill $|\eta_{\text{fat}}| < 1.0$, to guarantee that they are entirely in the central part of the detector and to justify our calorimeter tower size. For signal events, we require that the fat jet can be associated with a true top quark within $\Delta R < 1.2$. Unlike in our earlier study we do not re-cluster the anti- k_T jet constituents, because we eventually include tracking information and do not focus on a comparison with QCD-inspired taggers [31].

3.1 Calorimeter

We consider the two standard ranges, moderately boosted tops available in Standard Model processes and highly boosted tops in resonance searches,

$$\begin{aligned} p_{T,\text{fat}} &= 350 \dots 450 \text{ GeV} \\ p_{T,\text{fat}} &= 1300 \dots 1400 \text{ GeV}. \end{aligned} \quad (8)$$

In Fig. 2 we show the number of calorimeter-based 4-vectors $k_{\mu,i}$ as well as their ordered mean transverse momentum for the soft and hard fat jet selections of Eq.(8). For the soft and hard selections we have tested values $N = 10 \dots 60$ and find the using the leading $N = 40$ calorimeter constituents completely saturates the tagging performance. The remaining entries will typically be much softer than the top decay products and hence carry little signal or background information from the hard process.

For the softer fat jets we use 180,000 signal and 180,000 background events to train the network, 60,000 events each for tests during training, and 60,000 events each to estimate the performance. For technical reason the harder fat jets rely on a 10% smaller sample.

The network includes the CoLa, the LoLa, and two fully connected hidden layers, one with 100 and one with 50 nodes. It is trained using KERAS [46] with the THEANO [47] back-end, the ADAM optimizer, and a learning rate of 0.001. Training terminates either after 200 epochs or when the performance on the test sample does not improve for five epochs, typically after several tens of epochs. [†] We independently train five copies of the network, and compare their

[†]Using this setup, the training for the softer fat jets takes less than 15 minutes in total on a Tesla K80 using a p2.xlarge computing instance on Amazon Web Services.




Figure 3: ROC curve for the new DEEPTOPLoLA tagger, compared to the QCD-inspired MOTHEROFTAGGERS and the image-based DEEPTOP tagger [31]. In all cases we only use calorimeter information for soft fat jets, $p_{T,\text{fat}} = 350 \dots 450$ GeV.

performances on the independent validation sample.

Because of a long history of tests and applications on data, top taggers are especially useful to establish the performance of machine learning tools. In Fig. 3 we compare our DEEPTOPLoLA tagger to earlier benchmarks for the softer of the two selections in Eq.(8): a BDT of a large number of QCD-inspired observables and the image-based DEEPTOP tagger [31]. The QCD-inspired MOTHEROFTAGGERS consists of a boosted decision tree which includes a large, relatively well-understood set of observables, which can be linked to a systematic approach to including sub-jet correlations [48]. It includes the HEPTOPTAGGER mass drop algorithm [14] with an optimal choice of jet size [17], different jet masses including SoftDrop [49], as well as N-subjettiness [50]. As long as we only include calorimeter information we cannot expect the new method to significantly improve over these two approaches. On the other hand, the number of weights (inputs) of the LoLa-based DNN are lower by a factor of three to eight (ten to twenty) than what is used by the reference convolutional network. The proposed architecture is simpler, more flexible and physics-motivated but easily matches the convolutional network approach.

3.2 Learning the Minkowski metric

A technical challenge related to the Minkowski metric is that it combines two different features: two subjets are Minkowski-close if they are collinear or when one of them is soft ($k_{i,0} \rightarrow 0$). Because these two scenarios correspond to different, but possibly overlapping phase space regions, they are hard to learn for a DNN.

To see how our DEEPTOPLoLA tagger deals with this problem and to test what kind of structures drive the network output, we turn the problem around and ask the question if the Minkowski metric is really the feature distinguishing top decays and QCD jets. To this end, we define the invariant mass $m(\tilde{k}_j)$ and the distance d_{jm}^2 in Eq.(6) with a trainable diagonal metric. After applying a global normalization we find

$$g = \text{diag}(0.99 \pm 0.02 - 1.01 \pm 0.01, -1.01 \pm 0.02, -0.99 \pm 0.02), \quad (9)$$

where the errors are given by five independently trained copies. It is crucial for our physics understanding [48] that the distinguishing power of the DEEPTOPLoLA tagger is indeed the




Figure 4: ROC curve for the new DEEPTOPLoLA tagger operating on particle flow objects, compared to the its performance operating on calorimeter objects.

same mass drop [1] that drives many QCD-based top taggers [13, 14] and the image-based top tagger, as shown in detail in Ref. [31].

3.3 Calorimeter and tracking

A standard criticism of the jet image approach is that the pixelled image removes information from the original jet. For the calorimeter information alone this is not the case, because the image pixels are given by the calorimeter resolution. However, this identification is not possible for tracking information, because the tracking resolution of ATLAS and CMS is much finer than a jet image can realistically resolve [32]. This makes it hard to in general extend jet images to particle flow objects and to reliably determine how much performance can be gained through tracking information.

In contrast, for our LoLa-based approach this extension to particle flow constituents is straightforward: instead of defining one constituent or 4-vector per calorimeter cell we use all objects defined by the DELPHES3 particle flow algorithm in the same $p_{T,\text{fat}}$ range as in Eq.(8). The fat jet constituents at the particle flow level are different from the calorimeter case, which implies that for the same $p_{T,\text{fat}}$ range the underlying top quarks are around 5% softer for fat jets based on particle flow objects. Nevertheless, defining the signal and background events using Eq.(8) still is the best choice.

In Fig. 2 we show the number of constituents for the calorimeter-level and the particle flow approaches. The latter offers not only considerably more objects, the corresponding 4-momenta are also measured more precisely. We also show the mean transverse momentum for each of these constituents, indicating that the larger number of particle flow objects at least in part arises from splitting harder calorimeter entries into several objects at higher resolution. For our DEEPTOPLoLA tagger Fig. 2 implies that we could include more particle flow objects than calorimeter objects in Eq.(1). Again, we use $N = 40$ and confirm that an increase to $N = 60$ has no measurable effect on the performance.

Searching for possible improvements to our tagger, we first check that indeed the top quark kinematics are more precisely measured by the particle flow objects. However, the observed 5% improvement, for example in the resolution of the top transverse momentum, is unlikely to significantly improve our analysis.

In Fig. 4 we confirm that using the same neural network for calorimeter and particle flow

objects gives hardly any improvement for moderately boosted tops with $p_{T,\text{fat}} = 350 \dots 450$ GeV. The situation changes when we train and test our tagger at larger transverse momenta, $p_{T,\text{fat}} = 1300 \dots 1400$ GeV. Here the calorimeter resolution is no longer sufficient to separate the substructures [51]. For a fixed signal efficiency the background rejection including particle flow increases by a factor of two to three.

4 Conclusions

Based on a deep neural network working on Lorentz vectors of jet constituents we have built the new, simple, and flexible DEEPTOPLOLA tagger. It includes a Combination layer mimicking QCD-inspired jet recombination, a Lorentz layer translating the 4-vectors into appropriate kinematic observables, and two fully connected layers. The 4-vector input is not limited to a single detector output but allows us to add more information about a subject object in a straightforward manner.

We have compared the tagging performance to QCD-inspired taggers and to image-based convolutional network taggers using only calorimeter information for moderately boosted top quarks [31]. Figure 3 shows that the new tagger is competitive with either of these alternative approaches. Because we consider it crucial to control what machine learning methods actually exploit [48] we not only compared the DEEPTOPLOLA performance to an established QCD-inspired tagger [31], but also confirmed that the Minkowski metric related to a mass drop condition indeed drives the signal and background distinction.

Finally, we have used our tagger on particle flow objects, combining calorimeter and tracker information at their respective full experimental resolution. We have found that while for moderately boosted top quarks the performance gain from the tracker is negligible, it makes a big difference for strongly boosted top quarks.

The coverage of the full transverse momentum range and the possibility to include b -tagging through the tracking information should make the DEEPTOPLOLA tagger an excellent starting point to employ machine learning as the standard in ATLAS and CMS subjects analyses. It also opens a wide range of applications based on 4-vectors describing structures like for example matrix elements or phase space.

Acknowledgements

First, we would like to thank Anke Biekötter for her help on tracking simulation. A.B acknowledges support from the *Heidelberg Graduate School for Fundamental Physics* and the DFG research training group *Particle Physics Beyond the Standard Model*. M.R. was supported by the European Union Marie Curie Research Training Network MCnetITN, under contract PITN-GA-2012-315877. Our work was supported by a grant from the Swiss National Supercomputing Centre (CSCS) under project D61.

References

- [1] J. M. Butterworth, A. R. Davison, M. Rubin and G. P. Salam, *Jet substructure as a new Higgs-search channel at the Large Hadron Collider*, Phys. Rev. Lett. **100**, 242001 (2008), doi:[10.1103/PhysRevLett.100.242001](https://doi.org/10.1103/PhysRevLett.100.242001).
- [2] M. H. Seymour, *Searches for new particles using cone and cluster jet algorithms: a comparative study*, Z. Phys. C **62**, 127 (1994), doi:[10.1007/BF01559532](https://doi.org/10.1007/BF01559532).

- [3] J. M. Butterworth, B. E. Cox and J. R. Forshaw, *WW scattering at the CERN LHC*, Phys. Rev. D **65**, 096014 (2002), doi:[10.1103/PhysRevD.65.096014](https://doi.org/10.1103/PhysRevD.65.096014).
- [4] Y. Cui, Z. Han and M. D. Schwartz, *W-jet tagging: Optimizing the identification of boosted hadronically-decaying W bosons*, Phys. Rev. D **83**, 074023 (2011), doi:[10.1103/PhysRevD.83.074023](https://doi.org/10.1103/PhysRevD.83.074023).
- [5] W. Skiba and D. Tucker-Smith, *Using jet mass to discover vector quarks at the CERN LHC*, Phys. Rev. D **75**, 115010 (2007), doi:[10.1103/PhysRevD.75.115010](https://doi.org/10.1103/PhysRevD.75.115010).
- [6] B. Holdom, *t' at the LHC: The physics of discovery*, J. High Ener. Phys. **03**, 063 (2007) doi:[10.1088/1126-6708/2007/03/063](https://doi.org/10.1088/1126-6708/2007/03/063) [[arXiv:hep-ph/0702037](https://arxiv.org/abs/hep-ph/0702037)];
- [7] M. Gerbush, T. Jian Khoo, D. J. Phalen, A. Pierce and D. Tucker-Smith, *Color-octet scalars at the CERN LHC*, Phys. Rev. D **77**, 095003 (2008), doi:[10.1103/PhysRevD.77.095003](https://doi.org/10.1103/PhysRevD.77.095003).
- [8] D. E. Kaplan, K. Rehermann, M. D. Schwartz and B. Tweedie, *Top tagging: A method for identifying boosted hadronically decaying top quarks*, Phys. Rev. Lett. **101**, 142001 (2008), doi:[10.1103/PhysRevLett.101.142001](https://doi.org/10.1103/PhysRevLett.101.142001).
- [9] L. G. Almeida, S. J. Lee, G. Perez, I. Sung and J. Virzi, *Top quark jets at the LHC*, Phys. Rev. D **79**, 074012 (2009), doi:[10.1103/PhysRevD.79.074012](https://doi.org/10.1103/PhysRevD.79.074012).
- [10] L. G. Almeida, S. J. Lee, G. Perez, G. Sterman, I. Sung and J. Virzi, *Substructure of high-p_T jets at the LHC*, Phys. Rev. D **79**, 074017 (2009), doi:[10.1103/PhysRevD.79.074017](https://doi.org/10.1103/PhysRevD.79.074017).
- [11] L. G. Almeida, S. J. Lee, G. Perez, G. Sterman and I. Sung, *Template overlap method for massive jets*, Phys. Rev. D **82**, 054034 (2010), doi:[10.1103/PhysRevD.82.054034](https://doi.org/10.1103/PhysRevD.82.054034).
- [12] M. Backović and J. Juknevich, *Template Tagger v1.0.0: A template matching tool for jet substructure*, Comput. Phys. Commun. **185**, 1322 (2014), doi:[10.1016/j.cpc.2013.12.018](https://doi.org/10.1016/j.cpc.2013.12.018).
- [13] T. Plehn, G. P. Salam and M. Spannowsky, *Fat jets for a light Higgs boson*, Phys. Rev. Lett. **104**, 111801 (2010), doi:[10.1103/PhysRevLett.104.111801](https://doi.org/10.1103/PhysRevLett.104.111801).
- [14] T. Plehn, M. Spannowsky, M. Takeuchi and D. Zerwas, *Stop reconstruction with tagged tops*, J. High Energ. Phys. **10**, 078 (2010), doi:[10.1007/JHEP10\(2010\)078](https://doi.org/10.1007/JHEP10(2010)078).
- [15] T. Plehn, M. Spannowsky and M. Takeuchi, *How to improve top-quark tagging*, Phys. Rev. D **85**, 034029 (2012), doi:[10.1103/PhysRevD.85.034029](https://doi.org/10.1103/PhysRevD.85.034029).
- [16] C. Anders, C. Bernaciak, G. Kasieczka, T. Plehn and T. Schell, *Benchmarking an even better top tagger algorithm*, Phys. Rev. D **89**, 074047 (2014), doi:[10.1103/PhysRevD.89.074047](https://doi.org/10.1103/PhysRevD.89.074047).
- [17] G. Kasieczka, T. Plehn, T. Schell, T. Strebler and G. P. Salam, *Resonance searches with an updated top tagger*, J. High Energ. Phys. **06**, 203 (2015), doi:[10.1007/JHEP06\(2015\)203](https://doi.org/10.1007/JHEP06(2015)203).
- [18] D. E. Soper and M. Spannowsky, *Finding top quarks with shower deconstruction*, Phys. Rev. D **87**, 054012 (2013), doi:[10.1103/PhysRevD.87.054012](https://doi.org/10.1103/PhysRevD.87.054012).
- [19] A. Abdesselam et al., *Boosted objects: a probe of beyond the standard model physics*, Eur. Phys. J. C **71**, 1661 (2011), doi:[10.1140/epjc/s10052-011-1661-y](https://doi.org/10.1140/epjc/s10052-011-1661-y).
- [20] T. Plehn and M. Spannowsky, *Top tagging*, J. Phys. G: Nucl. Part. Phys. **39**, 083001 (2012), doi:[10.1088/0954-3899/39/8/083001](https://doi.org/10.1088/0954-3899/39/8/083001).

- [21] A. Altheimer *et al.*, *Boosted objects and jet substructure at the LHC. Report of BOOST2012, held at IFIC Valencia, 23rd–27th of July 2012*, Eur. Phys. J. C **74**, 2792 (2014), doi:[10.1140/epjc/s10052-014-2792-8](https://doi.org/10.1140/epjc/s10052-014-2792-8).
- [22] S. Schätzel, *Boosted top quarks and jet structure*, Eur. Phys. J. C **75**, 415 (2015), doi:[10.1140/epjc/s10052-015-3636-x](https://doi.org/10.1140/epjc/s10052-015-3636-x).
- [23] V. Rentala, W. Shepherd and T. M. P. Tait, *Tagging boosted W s with wavelets*, J. High Energ. Phys. **08**, 042 (2014), doi:[10.1007/JHEP08\(2014\)042](https://doi.org/10.1007/JHEP08(2014)042).
- [24] J. W. Monk, *Wavelet analysis: Event de-noising, shower evolution and jet substructure without jets* (2014), [arXiv:1405.5008 \[hep-ph\]](https://arxiv.org/abs/1405.5008).
- [25] J. Cogan, M. Kagan, E. Strauss and A. Schwartzman, *Jet-images: Computer vision inspired techniques for jet tagging*, J. High Energ. Phys. **02**, 118 (2015), doi:[10.1007/JHEP02\(2015\)118](https://doi.org/10.1007/JHEP02(2015)118).
- [26] L. de Oliveira, M. Kagan, L. Mackey, B. Nachman and A. Schwartzman, *Jet-images — deep learning edition*, J. High Energ. Phys. **07**, 069 (2016), doi:[10.1007/JHEP07\(2016\)069](https://doi.org/10.1007/JHEP07(2016)069).
- [27] P. Baldi, K. Bauer, C. Eng, P. Sadowski and D. Whiteson, *Jet substructure classification in high-energy physics with deep neural networks*, Phys. Rev. D **93**, 094034 (2016), doi:[10.1103/PhysRevD.93.094034](https://doi.org/10.1103/PhysRevD.93.094034).
- [28] J. Barnard, E. Noel Dawe, M. J. Dolan and N. Rajcic, *Parton shower uncertainties in jet substructure analyses with deep neural networks*, Phys. Rev. D **95**, 014018 (2017), doi:[10.1103/PhysRevD.95.014018](https://doi.org/10.1103/PhysRevD.95.014018).
- [29] L. de Oliveira, M. Paganini and B. Nachman, *Learning particle physics by example: Location-aware generative adversarial networks for physics synthesis*, Comput. Softw. Big Sci. **1**, 4 (2017), doi:[10.1007/s41781-017-0004-6](https://doi.org/10.1007/s41781-017-0004-6).
- [30] L. G. Almeida, M. Backović, M. Cliche, S. J. Lee and M. Perelstein, *Playing tag with ANN: Boosted top identification with pattern recognition*, J. High Energ. Phys. **07**, 086 (2015), doi:[10.1007/JHEP07\(2015\)086](https://doi.org/10.1007/JHEP07(2015)086).
- [31] G. Kasieczka, T. Plehn, M. Russell and T. Schell, *Deep-learning top taggers or the end of QCD?*, J. High Energ. Phys. **05**, 006 (2017), doi:[10.1007/JHEP05\(2017\)006](https://doi.org/10.1007/JHEP05(2017)006).
- [32] P. T. Komiske, E. M. Metodiev and M. D. Schwartz, *Deep learning in color: Towards automated quark/gluon jet discrimination*, J. High Energ. Phys. **01**, 110 (2017), doi:[10.1007/JHEP01\(2017\)110](https://doi.org/10.1007/JHEP01(2017)110).
- [33] G. Louppe, K. Cho, C. Becot and K. Cranmer, *QCD-aware recursive neural networks for jet physics* (2017), [arXiv:1702.00748 \[hep-ph\]](https://arxiv.org/abs/1702.00748).
- [34] Y. LeCun, Y. Bengio, and G. Hinton, *Deep learning*, Nature **521**, 436 (2015), doi:[10.1038/nature14539](https://doi.org/10.1038/nature14539).
- [35] J. Pearkes, W. Fedorko, A. Lister and C. Gay, *Jet constituents for deep neural network based top quark tagging* (2017), [arXiv:1704.02124 \[hep-ex\]](https://arxiv.org/abs/1704.02124).
- [36] S. D. Ellis, A. Hornig, T. S. Roy, D. Krohn and M. D. Schwartz, *Nondeterministic approach to tree-based jet substructure*, Phys. Rev. Lett. **108**, 182003 (2012), doi:[10.1103/PhysRevLett.108.182003](https://doi.org/10.1103/PhysRevLett.108.182003).

- [37] T. Sjöstrand, S. Ask, J. R. Christiansen, R. Corke, N. Desai, P. Ilten, S. Mrenna S. Prestel, C. O. Rasmussen and P. Z. Skands, *An introduction to PYTHIA 8.2*, Comput. Phys. Commun. **191**, 159 (2015), doi:[10.1016/j.cpc.2015.01.024](https://doi.org/10.1016/j.cpc.2015.01.024).
- [38] Our event samples are available for performance studies upon request.
- [39] D. Bertolini, P. Harris, M. Low and N. Tran, *Pileup per particle identification*, J. High Energ. Phys. **10**, 059 (2014), doi:[10.1007/JHEP10\(2014\)059](https://doi.org/10.1007/JHEP10(2014)059).
- [40] M. Cacciari, G. P. Salam and G. Soyez, *SoftKiller, a particle-level pileup removal method*, Eur. Phys. J. C **75**, 59 (2015), doi:[10.1140/epjc/s10052-015-3267-2](https://doi.org/10.1140/epjc/s10052-015-3267-2).
- [41] P. T. Komiske, E. M. Metodiev, B. Nachman and M. D. Schwartz, *Pileup mitigation with machine learning (PUMML)*, J. High Energ. Phys. **12**, 051 (2018), doi:[10.1007/JHEP12\(2017\)051](https://doi.org/10.1007/JHEP12(2017)051) [[arXiv:1707.08600](https://arxiv.org/abs/1707.08600) [hep-ph]].
- [42] J. de Favereau *et al.* [DELPHES 3 Collaboration], *DELPHES 3: A modular framework for fast simulation of a generic collider experiment*, J. High Energ. Phys. **02**, 057 (2014), doi:[10.1007/JHEP02\(2014\)057](https://doi.org/10.1007/JHEP02(2014)057).
- [43] M. Cacciari and G. P. Salam, *Dispelling the N^3 myth for the k_t jet-finder*, Phys. Lett. B **641**, 57 (2006), doi:[10.1016/j.physletb.2006.08.037](https://doi.org/10.1016/j.physletb.2006.08.037).
- [44] M. Cacciari, G. P. Salam and G. Soyez, *FastJet user manual*, Eur. Phys. J. C **72**, 1896 (2012), doi:[10.1140/epjc/s10052-012-1896-2](https://doi.org/10.1140/epjc/s10052-012-1896-2).
- [45] M. Cacciari, G. P. Salam and G. Soyez, *The anti- k_t jet clustering algorithm*, J. High Energ. Phys. **04**, 063 (2008), doi:[10.1088/1126-6708/2008/04/063](https://doi.org/10.1088/1126-6708/2008/04/063).
- [46] F. Chollet, <https://github.com/fchollet/keras> (2015).
- [47] [Theano Development Team], *Theano: A Python framework for fast computation of mathematical expressions* (2016), [arXiv:1605.02688](https://arxiv.org/abs/1605.02688) [cs.SC].
- [48] K. Datta and A. Larkoski, *How much information is in a jet?*, J. High Energ. Phys. **06**, 073 (2017), doi:[10.1007/JHEP06\(2017\)073](https://doi.org/10.1007/JHEP06(2017)073).
- [49] A. J. Larkoski, S. Marzani, G. Soyez and J. Thaler, *Soft drop*, J. High Energ. Phys. **05**, 146 (2014), doi:[10.1007/JHEP05\(2014\)146](https://doi.org/10.1007/JHEP05(2014)146).
- [50] J. Thaler and K. Van Tilburg, *Maximizing boosted top identification by minimizing N -subjettiness*, J. High Energ. Phys. **02**, 093 (2012), doi:[10.1007/JHEP02\(2012\)093](https://doi.org/10.1007/JHEP02(2012)093).
- [51] S. Schätzel and M. Spannowsky, *Tagging highly boosted top quarks*, Phys. Rev. D **89**, 014007 (2014), doi:[10.1103/PhysRevD.89.014007](https://doi.org/10.1103/PhysRevD.89.014007).