



CHORUS

This is the accepted manuscript made available via CHORUS. The article has been published as:

Parton shower uncertainties in jet substructure analyses with deep neural networks

James Barnard, Edmund Noel Dawe, Matthew J. Dolan, and Nina Rajcic

Phys. Rev. D **95**, 014018 — Published 18 January 2017

DOI: [10.1103/PhysRevD.95.014018](https://doi.org/10.1103/PhysRevD.95.014018)

Parton Shower Uncertainties in Jet Substructure Analyses with Deep Neural Networks

James Barnard,¹ Edmund Noel Dawe,¹ Matthew J. Dolan,¹ and Nina Rajcic¹

¹*ARC Centre of Excellence for Particle Physics at the Terascale,
School of Physics, University of Melbourne, 3010, Australia*

Machine learning methods incorporating deep neural networks have been the subject of recent proposals for new hadronic resonance taggers. These methods require training on a dataset produced by an event generator where the true class labels are known. However, this may bias the network towards learning features associated with the approximations to QCD used in that generator which are not present in real data. We therefore investigate the effects of variations in the modelling of the parton shower on the performance of deep neural network taggers using jet images from hadronic W-bosons at the LHC, including detector-related effects. By investigating network performance on samples from the Pythia, Herwig and Sherpa generators, we find differences of up to fifty percent in background rejection for fixed signal efficiency. We also introduce and study a method, which we dub zooming, for implementing scale-invariance in neural network-based taggers. We find that this leads to an improvement in performance across a wide range of jet transverse momenta. Our results emphasise the importance of gaining a detailed understanding of what aspects of jet physics these methods are exploiting.

I. INTRODUCTION

The past decade has seen an explosion of interest in understanding and exploiting the distribution of energy (substructure) within hadronic jets and boosted resonances at the Large Hadron Collider (LHC) [1–4]. The study of jet substructure and the ability to identify (‘tag’) the hadronic decay products of a wide variety of such resonances — the Higgs, W and Z bosons, top quarks, supersymmetric particles and other Beyond-the-Standard-Model (BSM) states — is crucial in the analysis of both Standard Model processes and in searches for BSM physics, which will only become more important now that the LHC is running at high energy and with future colliders on the horizon.

Since the foundational work in Ref. [5] on studying jet substructure in Higgs boson associated production, a multitude of taggers and variables related to substructure have been proposed [5–14] (further discussion of which can be found in the BOOST proceedings [1–4]). These generally exploit our knowledge of QCD to construct functions which effectively discriminate between signal and background. Some of these techniques have already been applied to the problem of identifying boosted massive vector bosons and top quarks by the ATLAS and CMS collaborations in Run 1 of the LHC [15–25].

Another approach currently under development involves the application of machine learning (ML) techniques to hadronic resonance tagging and searches for new physics. The ML community has made large strides in problems related to image recognition and computer learning, which may now also be applied to particle physics. Signals produced by the LHC detectors may be processed into pixelated jet images [26], and ML algorithms can be adapted to discriminate between a signal (such as $h \rightarrow b\bar{b}$ or boosted hadronic top decays [27]) and background. These algorithms have also been proposed

as classifiers in neutrino experiments [28, 29].

The use of ML, and neural networks in particular, has a long history in particle physics and the idea of using neural networks for quark-gluon discrimination [30–32], Higgs tagging [33] and track identification [34] goes back over twenty-five years. However, the development of efficient deep neural networks (DNNs) and the computing power associated with graphics processing units (GPUs) means that image recognition technology has become extremely powerful, driving the resurgence of interest in these techniques.

Recent work has seen the application of neural networks with two hidden layers to hadronic top-quark tagging [27], and deep convolutional neural networks (known to have excellent performance in image classification) to the problem of identification of hadronic W decays [35]. These initial papers focussed on demonstrating and understanding the network performance, and used truth-level Monte Carlo (MC). The effects of pile-up and detector resolution were explored in Ref. [36], which showed that despite the loss of resolution when these are taken into account the neural network is still somewhat superior to traditional techniques. There has also been work on extending these methods to jet flavour classification [37].

The theoretical study of these techniques and their utility in high-energy particle physics is still in its infancy, and there are a number of issues still to be clarified in how deep learning methods may be applied at the LHC. Some of these are related to the robustness of these techniques: how can we guarantee that a network is learning about the physics differences between signal and background, and not details particular to a specific MC event generator? How robust are taggers based on these networks against detector effects such as smearing and how do they degrade in the presence of pile-up? A particular concern is that the network achieves a substantial fraction of its

discriminatory power from soft features in the spectrum which are modelled phenomenologically rather than via perturbative QCD. This paper provides a study of some of these issues.

We study the behaviour of neural networks over a number of different event generators, and hence different parton showers and models for hadronisation. For simplicity, we will usually just refer to these collectively as the parton shower.

We find that varying the parton shower leads to changes in the background rejection efficiency of up to 50%, depending on the shower model and selected signal efficiency. We consider this to be large, and perhaps more than would be expected from perturbative uncertainties from the parton shower. We believe that caution is therefore required before these methods are applied on data, and our results emphasise the necessity of understanding what features of the jet images the neural networks are relying on to achieve their discriminatory power. We also find changes in the factorisation and renormalisation scales lead to negligible differences, while the addition of pile-up leads to an overall degradation in network performance (in agreement with Ref. [36]) but not to a change in our conclusions.

There has also been interest recently in the development of scale-invariant jet and substructure taggers [38–41], and we discuss how similar ideas may be implemented in DNN-based taggers by applying a p_T -dependent ‘zooming’ factor on the jet images. The addition of zooming leads to a slight improvement of around 10–20% in the network performance over a wide range of jet transverse momenta. While in this article we focus on discriminating between hadronically decaying W bosons and QCD jets as a ‘standard candle’, these methods should be applicable to a wide variety of tagging and substructure issues.

In Section II we outline the architecture and training of the neural networks and in Section III discuss how we construct jet images, and present an idea of how to implement a scale-invariant tagger. In Section IV we show the variability in the DNN performance across multiple event generators and parton shower models.

II. NETWORK ARCHITECTURE, TRAINING, AND PERFORMANCE EVALUATION

We follow Ref. [35] in our choice of network architecture, who have already investigated the performance of a variety of different neural networks. While we have investigated convolutional networks, all results we present here have been produced using the MaxOut [42] architecture. The network input consists of 625 units, equal to the number of pixels (25×25) present in each jet image. The input layer is followed by two dense MaxOut layers consisting of 256 and 128 units each. The next two layers are fully connected with 64 and 25 units and use a ReLU activation function [43]. The output layer consists of two

nodes and a sigmoid activation. Further discussion of network choices can be found in Ref. [35].

We used the Keras Deep Learning library [44] and the Adam algorithm [45] to train our networks on four NVIDIA Tesla K80 GPUs. After selecting jet images within a window on the jet mass, $50 < m < 110$ GeV, and transverse momentum, $200 < p_T < 500$ GeV, networks were trained with approximately 3M signal and 3M background images where the signal and background images have been weighted to produce flat p_T distributions. A portion (10%) of the training images were set aside to evaluate a cross-entropy loss function after each epoch and the network training terminated after 100 epochs or after 10 epochs without an improvement in the loss function. The Adam algorithm learning rate parameter was initially set to 0.001 and then reduced by 2% after each epoch. We obtained reasonable performance with a batch size of 100. We also implemented and tested a cross-validated Bayesian optimisation procedure to determine optimal parameter values but did not observe performance that was significantly better and so we have left such investigations for future work. Further optimising the DNN should anyway not affect our conclusions here as we probe the general variability of a DNN with reasonable performance over different parton shower models.

Finally, we evaluate the performance of a network by computing the inverse background efficiency as a function of signal efficiency across a binned likelihood ratio of the signal to background output of the network. This variant of the standard receiver operating characteristic (ROC) curve better displays differences in background rejection at low signal efficiency and can be constructed from arbitrary jet observables or combinations of observables through a (possibly multidimensional) binned likelihood ratio.

III. CONSTRUCTING A JET IMAGE

This section provides a complete description of how we construct jet images, from event generation to image output, along with the reasoning behind many of our choices. Our process closely follows the one described in Ref. [35] with the addition that we have also tested an optional zooming step to reduce p_T -dependence.

We have developed a new framework for jet image construction, network training and performance evaluation. Low-level Cython [46] wrappers have been developed for PYTHIA [47], Delphes [48], and FastJet [49] that allow these tools to be connected in the Python programming language, where particles, calorimeter towers, jets, and jet images are stored as structured NumPy [50] arrays and optionally written to files on disk in the HDF5 [51] format. This design provides the potential to train networks on jet images generated ‘on-the-fly’. For the studies presented here, however, we have created large HDF5 data sets of jet images once that are then split and used for network training and performance studies.

Aside from the direct interface with PYTHIA, the framework is able to study output from other event generators by reading intermediate HepMC [52] files.

Following event generation, particles are given to the Delphes detector simulator configured with ATLAS-like settings where calorimeter towers extend to a maximum absolute pseudorapidity of 4.9. Jets are then reconstructed from the calorimeter towers (referred to as jet constituents below) using the anti- k_t algorithm [53] as implemented by FastJet 3.1.3. We have selected a jet clustering size of $R = 1.0$ for all studies presented here. For boosted W bosons with two-body decays the characteristic maximal separation of the subjets scales according to

$$\Delta R = \frac{2m_W}{p_T^{\min}} \quad (\text{III.1})$$

where p_T^{\min} is the minimum transverse momentum of the jets to be considered in the analysis. We have studied jets with transverse momenta above 200 GeV making $R = 1.0$ a reasonable choice.

The highest p_T jet is selected and subjets are formed in a jet trimming [7] stage, that also serves to lessen contributions from soft radiation in the underlying event. Using the k_t algorithm we recluster the jet constituents into subjets with a fixed size of $r = 0.3$ and then discard all subjets with less than 5% of the original jet momentum to form a trimmed jet. All jet observables are computed with the trimmed jet.

The next stages are designed to remove spatial symmetries. First, all constituents of the trimmed jet are translated in $\eta - \phi$ space to place the leading subjet at the origin. We then define a grid of pixels with a resolution of 0.1×0.1 in $\eta - \phi$ space and a jet image is formed by taking the total transverse energy measured within each pixel

$$E_{T,i} = \sum_j \frac{E_j}{\cosh \eta_j} \quad (\text{III.2})$$

for all constituents j in pixel i , with energy E_j and original pseudorapidity η_j . This image is rotated, either to put the subleading subjet directly below the leading subjet or to align the principle component of the jet image along the vertical axis if only one subjet is present. It is then reflected, either to put the third-leading subjet on the right-hand side of the image or to ensure that the total image intensity is highest on the right-hand side if there are only two subjets.

After the reflection stage above we are left with an image in which the leading subjet is centred and the subleading subjet (if present) is directly below. The separation between the two subjets is not constant, but varies linearly with $2m/p_T$. By standardising this separation we can potentially improve the DNN performance over a wide range of jet p_T . The aim is to pick a scaling factor that enhances and standardises features in signal images, i.e. those from boosted W decays, without artificially

creating similar features in background, QCD images. This optional step is in addition to those detailed in Ref. [35].

Denoting the physical separation between the two leading subjets as ΔR_{act} , enlarging all jet images by a factor $R/\Delta R_{\text{act}}$ for some fixed R gives a standardised jet image in which the separation (in pixels) between the two leading subjets is fixed for all images. The downside of this approach is that this is true for both signal and background images, so an improvement in isolation of the subleading subjet is tempered by an enhancement of signal-like features in background images. For this reason we use the characteristic size assuming the W mass, $s = 2m_W/p_T$, and enlarge all jet images by a scaling factor $\max(R/s, 1)$ where R is the original jet clustering size. Images are enlarged by performing bicubic interpolation [54] at a higher resolution. For signal images $\Delta R_{\text{act}} \approx s$ so this rescaling is very similar to, if a little less effective than, using the actual separation of subjets to define the scaling factor. For background images this scaling is not strongly correlated with the subjet separation so the subleading subjet tends to be smeared out.

Jet images are then cropped at 25×25 pixels (whether having been zoomed or not) and are normalised such that the sum of the squared pixel intensities is one. As discussed in Ref. [35] this does not preserve the jet mass that can be calculated from the original jet image but our zooming step destroys this information anyway.

In summary, the full jet image construction and preprocessing steps are as follows:

- **Jet Clustering and Trimming:** Reconstruct jets from all calorimeter towers using the anti- k_t algorithm with a jet size $R = 1.0$ and select the leading jet. Trim the jet using the k_t algorithm with a subjet size $r = 0.3$.
- **Translation:** All jet constituents are translated in $\eta - \phi$ space to put the leading subjet at the origin.
- **Pixelisation:** Pixelise the transverse energy of the jet using pixels of size $(0.1, 0.1)$ in $\eta - \phi$ space. This produces a jet image.
- **Rotation:** Rotate the jet image to put the subleading subjet directly below the leading subjet. If no subjets are present rotate to align the principle component of the jet image along the vertical axis.
- **Reflection:** Reflect the jet image horizontally to put the third-leading subjet on the right-hand side. If there are only two subjets, reflect to ensure that the summed image intensity is highest on the right-hand side.
- **Zooming:** Optionally zoom the jet image by a factor that reduces dependence on the jet momentum.
- **Cropping and Normalisation:** Crop the jet image at 25×25 pixels and normalise pixel intensities to make the sum of their squares equal to one.

In Fig. 1 we show the average jet images for boosted W and QCD jets in the range $200 < p_T < 500$ GeV for the default PYTHIA shower using the standard preprocessing

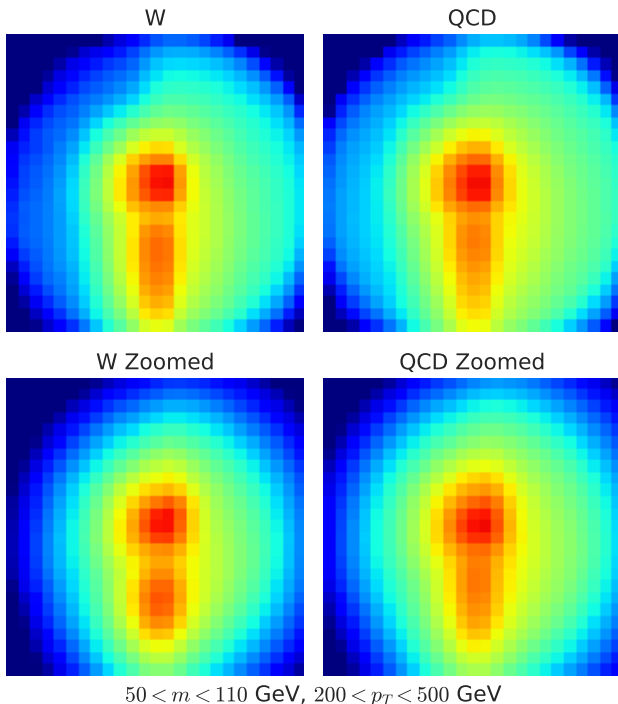


Figure 1: We show the average jet images obtained for hadronic W bosons and QCD as modelled by the PYTHIA default shower. The images on the top have been pre-processed in the standard way, while those on the bottom have also undergone the zooming procedure outlined in Section III. The axes are left unlabelled since they do not correspond to the physical η and ϕ dimensions following image rotations, reflections, and zooming. Pixels are coloured according to higher (red) and lower (blue) average normalised pixel intensities.

in the top panels, and using the zooming procedure in the bottom panels. For the W -jets we note the zooming procedure results in a more regular and compact average shower shape, and that the second (lower p_T) subject becomes better spatially defined as expected. While the average image of the QCD jets becomes more compact, the subjects remain somewhat smeared compared with the W -jets. Since the subjects do not originate in the decay of a heavy resonance and hence are not associated with a specific mass scale this is not a surprise.

An obvious conceptual advantage of using the zooming technique is that it makes the construction of scale-invariant taggers easier. Scale invariant searches [38–41] which are able to interpolate between the boosted and resolved parts of phase space have the advantage of being applicable over a broad range of masses and kinematics, allowing a single search or analysis to be effective where previously more than one may have been necessary.

We show in Fig. 2 the ROC curves for two different neural networks: the first (the solid blue line) was trained without zooming, while the second (the green dashed line) used zooming. Both networks were trained and tested on samples of jet images in the mass window $50 < m < 110 \text{ GeV}$ and a large p_T range, $200 < p_T <$




Figure 2: The ROC curves for the zoomed (solid blue) and unzoomed (dashed green) jet images for the PYTHIA default shower. The lower panel shows the ratio of the zoomed to unzoomed efficiencies, also showing the efficiency sliced in bins from $200 < p_T < 300 \text{ GeV}$ (dot-dashed red) and $300 < p_T < 500 \text{ GeV}$ (short dashed blue).

500 GeV. As predicted, the zoomed network outperforms the unzoomed one, particularly at low signal efficiency, where the background rejection rises by around 20%. We obtain similar results when we do not restrict the sample of jet images within a mass window. We find that the zooming has the greatest effect at high p_T . For less boosted W decays the enhancement in background rejection is around 10% which rises to just over 20% for $300 < p_T < 500 \text{ GeV}$.

IV. EVENT GENERATOR DEPENDENCY

The networks require supervised training prior to being applied on unlabelled data. Since it is difficult to isolate regions of very high signal purity, training on simulated data is necessary before application to real LHC events. However, all MC event generators and parton showers are only approximations to the full Standard Model. Understanding what features of QCD a DNN is learning about, and whether it is learning event generator-dependent approximations is thus an important question. Furthermore, there are features of real-world QCD such as colour reconnection which, while modelled in the parton shower, are in reality poorly understood. We will not attempt to quantify those effects in this work.

To gain an understanding of the systematic uncertainties in using networks trained on simulated data, we study the behaviour of networks across a variety of

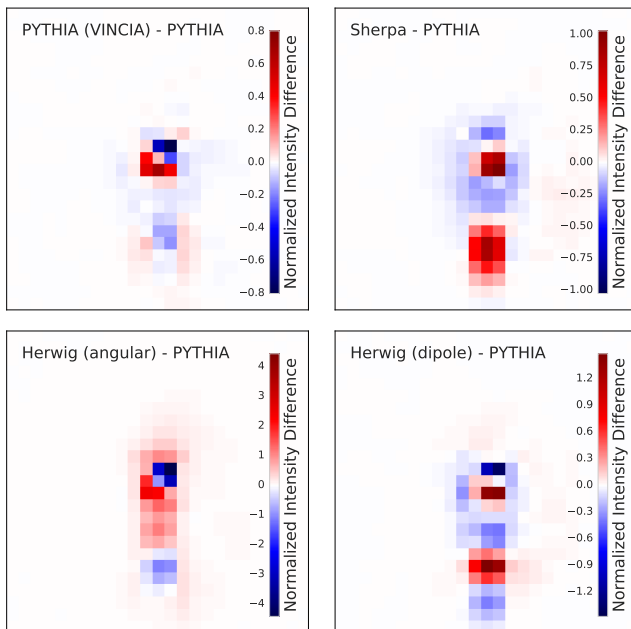


Figure 3: This figure shows the W-jet image differences between the default PYTHIA shower and the alternate VINCIA shower in PYTHIA (top left), the default SHERPA shower (top right), the default HERWIG angular shower (bottom left) and the HERWIG dipole shower (bottom right). The plots have been individually normalised.

different generators and parton showers which all provide an adequate description of current LHC data. We assume that given a number of different ROC curves derived from different generators and parton showers, the envelope of these curves provides an approximate uncertainty band associated with training the network on simulated, rather than real, data.

Recently, Ref. [55] has studied parton shower uncertainties in HERWIG 7. They divide the uncertainties into a number of classes: numerical, parametric, algorithmic, perturbative and phenomenological. Numerical uncertainties can be decreased by increasing the number of events, while parametric uncertainties are those external to the MC generator: masses, couplings, PDFs and so forth. The focus of our work in this section is on algorithmic uncertainties, those due to different choices of parton shower algorithm. The authors of Ref. [55] focus on perturbative and phenomenological uncertainties, which are from truncation of expansion series and parameters deriving from non-perturbative models. Our work is more in the spirit of the ‘Towards parton shower variations’ contribution to the 2015 SM Les Houches Proceedings [56]. Previous studies also exist within the HERWIG framework on the implications of MC uncertainties on jet substructure in the context of Higgs searches [57].

We generate background and signal events with three of the most widely used MC generators: PYTHIA 8.219 [47], SHERPA 2.0 [58, 59] and HERWIG 7.0 [60, 61]. For PYTHIA 8 we study both the default shower

and the VINCIA shower [62, 63], and for HERWIG we include both the default (angular ordered) and dipole showers [64, 65], giving us five different parton shower models to study.

The default HERWIG shower (known as QTilde) is based on $1 \rightarrow 2$ splittings using the DGLAP equations, with an angular ordering criterion [66]. The SHERPA shower is based on a Catani-Seymour dipole formalism [67]. In this case one particle of the dipole is the emitter which undergoes the splitting, while the other is a spectator which compensates for the recoil from the splitting and ensures that all particles remain on their mass-shells throughout the shower, leading to easier integration with matching and merging techniques. The default shower in PYTHIA 8 is also a dipole style shower [68], ordered in transverse momentum.

While parton showers have traditionally been based upon partonic DGLAP splitting functions, another possibility is to consider colour-connected parton pairs which undergo $2 \rightarrow 3$ branchings (note that this is distinct from Catani-Seymour dipoles used in SHERPA, where one parton is still an emitter, and the other recoils). In these so-called antenna showers, the 2-parton antenna is described with a single radiation kernel. This has the advantage, for instance, of explicitly including both the soft and collinear limits. We use the recently released VINCIA [62, 63] plug-in for PYTHIA 8 as a representative antenna shower.

These event generators also provide different treatments of the soft radiation from the underlying event which accompanies each hard partonic scattering. They also possess different implementations of the parton-to-hadron fragmentation process being based either around cluster fragmentation ideas (HERWIG and SHERPA) or the Lund string model (PYTHIA), giving us a wide range of QCD-related effects to probe. To incorporate detector effects such as smearing we pass all events through the Delphes 3 detector simulator [48]. In the studies presented here, our baseline shower is PYTHIA 8 with its default settings.

We construct average jet images for all five different generators and showers under investigation, and then subtract the default PYTHIA average jet image in order to see the differences in the average radiation patterns. The results are shown in Fig. 3 for the W-jet signal. We have normalised the intensity differences of the pixels so that red indicates a region of excess and blue a deficit relative to the PYTHIA default. While the VINCIA is roughly similar to the PYTHIA default, the SHERPA and HERWIG dipole showers exhibit more intense radiation in the resolved subjects and a substantial deficit in the region between the subjects. The HERWIG angular shower shows the opposite, with less radiation in the subject cores and more diffuse radiation. QCD radiation exhibits similar features.

Next we show ROC curves for the different showers in Fig. 4. We used the same network discussed in Section III trained on the default PYTHIA shower (without zooming),

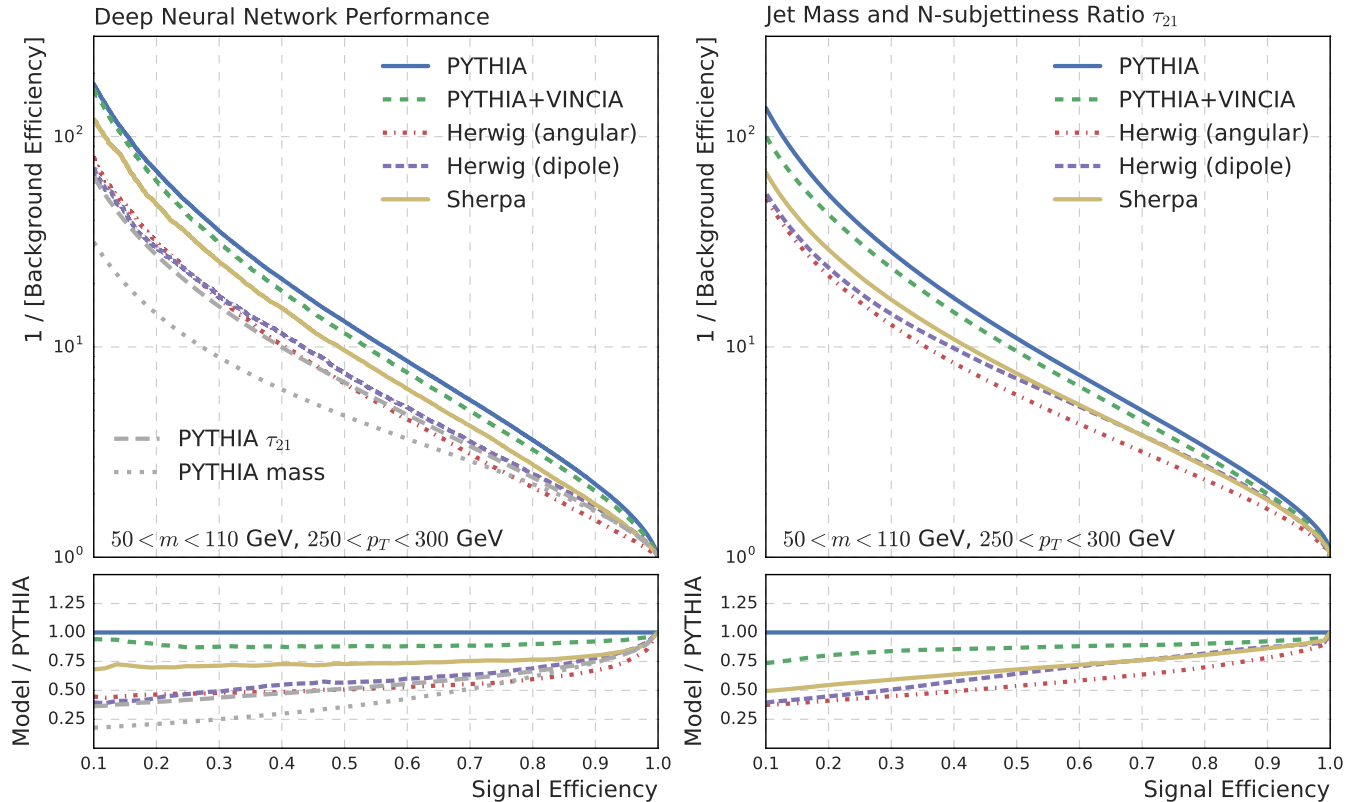


Figure 4: This figure shows the ROC curves of the **PYTHIA** (solid blue), **VINCIA** (dashed green), **HERWIG** angular (red dash-dot) and dipole (dashed purple), and **SHERPA** (solid gold) showers for the DNN output (left) and the combination of the jet mass and n -subjettiness ratio τ_{21} through a two-dimensional binned likelihood ratio (right). The lower panels show the ratio of the ROCs with the default **PYTHIA** shower. All ROC curves are computed using jet images within a window on the jet mass, $50 < m < 110$ GeV, and transverse momentum, $250 < p_T < 300$ GeV.

and then used events from the other generators and parton showers as input, *e.g.* we ask a neural network trained on the **PYTHIA** shower to discriminate between QCD and W-jets from **SHERPA**.

We do not extend the ROC curves down to zero signal efficiency since they are more statistically limited there. The **PYTHIA** ROC is higher than all other shower efficiency curves. While both the **SHERPA** and **HERWIG** dipole images exhibit superficial similarities in Fig. 3, the network is better at discriminating the **SHERPA** events. At a fixed low signal efficiency the **HERWIG** angular and dipole showers have the lowest background rejection, smaller than that obtained using the **PYTHIA** default by a factor of two. The **VINCIA** and **SHERPA** showers have a slightly lower rejection rate than the **PYTHIA** one. For signal efficiency of 50% the uncertainty from changing the event generator is around 40%.

For large background rejection rate we note that the network trained on the **PYTHIA** events has a lower efficiency for selecting signal events generated from the other showers, *i.e.* it is maximally efficient for the shower it was trained on. This may be due to the network learning some features associated specifically with the

PYTHIA shower and thus performing well on **PYTHIA**-like events.

We also show in Fig. 4 the ROC curves we obtain for the trimmed jet mass and the n -subjettiness ratio $\tau_{21} \equiv \tau_2/\tau_1$ [11] which is often used as a discriminating variable in studies of jet substructure [69]. We see that the neural network consistently outperforms these variables (in agreement with the conclusions already reached in Ref. [35]). This result stands independent of the uncertainty induced by the choice of event generator, although the results for the **HERWIG** showers are close to being degenerate with it.

In the right panel of Fig. 4 we show the ROC curves we obtain from the combined jet mass and τ_{21} observables for the different parton showers. We see that the parton shower uncertainties in this case are very similar to those obtained from the jet images. The uncertainties related to varying the parton shower for the jet images are thus of similar size to those associated other more common variables, such as those found in theoretical studies of the D2 tagger [70, 71] and those used by the ATLAS Collaboration in searches for boosted W-bosons [18].

Another possible independent source of uncertainty is

the dependence of the shower profile on the common renormalisation and factorisation scales, μ_R and μ_F respectively which for our purposes are set to be $\mu = p_T$. As is standard, we vary the scale μ upwards and downwards by a factor of two from its default value of $\mu = p_{T,W}$. We find that the changes in ROC curves due to this were negligible.

V. CONCLUSIONS

The use of deep neural networks to construct classifiers for hadronic substructure using jet images is an exciting proposal. However, it is important to quantify the dependence on the training dataset, and whether the network is learning the approximations inherent in the MC generator. We trained a network on the default parton shower from the PYTHIA generator and studied its performance on events from HERWIG and SHERPA. We found that the network performed better on test events also from the default PYTHIA shower, indicating that the network may be learning some PYTHIA-specific features. The change in performance through using different parton showers could be up to a factor of two in background rejection. Our results thus indicate that care is required to avoid over-interpretation of small changes in ROC curves, given the parton shower uncertainties. These uncertainties are relatively large, and further study is required to ascertain whether the network performance is truly being driven by features in the parton shower (which are under control in perturbative QCD) or by softer physics such as hadronisation modelling (which is not). Either way, our results demonstrate that caution is required in the application of machine learning techniques on simulated data. We intend to return to this issue in the near future.

There are many opportunities for further work in this area. One way to achieve event generator independency is by avoiding the use of training data through using data-driven unsupervised learning algorithms. Will these prove as powerful as the supervised techniques using

DNNs proposed thus far? It would also be desirable to incorporate more than just calorimeter information into jet images, in a similar vein to recent work on heavy flavour tagging [37]. We note that tracking information has been proposed in the context of substructure as being particularly important at high energies [72]. Since jet images can be both large and sparse, new algorithms may be required to render this feasible [73]. In an ideal world, it would be possible to use information from the whole detector to classify events into signal or background, as in event deconstruction [74].

The main outcome of our study is to emphasise the importance of what the neural networks are learning and how they use it to discriminate between signal and background. This lesson is true for any application of machine learning in particle physics, from widely used techniques such as boosted decision trees to newer methods like the deep neural network we have studied. However, provided cautious and detailed studies of the uncertainties involved lead to methods to constrain them in the analysis of real data, these methods may prove to be powerful and reliable analysis tools for future searches and measurements.

Acknowledgments

We thank Peter Skands, Michael Spannowsky and Phillip Urquijo for helpful discussions and comments. This work was supported in part by the Australian Research Council. We are grateful for the computational support provided by our colleagues in CoEPP and Research Compute Services at the University of Melbourne for generously providing access to the GPUs. We thank the Aspen Center for Physics, which is supported by National Science Foundation grant PHY-1066293, the Munich Institute for Astro and Particle Physics and the Mainz Institute of Theoretical Physics, where part of this work was completed, for their hospitality and support.

-
- [1] A. Abdesselam et al., *Boosted objects: A Probe of beyond the Standard Model physics*, *Eur. Phys. J.* **C71** (2011) 1661, [1012.5412].
 - [2] A. Altheimer et al., *Jet Substructure at the Tevatron and LHC: New results, new tools, new benchmarks*, *J. Phys.* **G39** (2012) 063001, [1201.0008].
 - [3] 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.* **C74** (2014) 2792, [1311.2708].
 - [4] D. Adams et al., *Towards an Understanding of the Correlations in Jet Substructure*, *Eur. Phys. J.* **C75** (2015) 409, [1504.00679].
 - [5] J. M. Butterworth, A. R. Davison, M. Rubin and G. P. Salam, *Jet substructure as a new Higgs search channel at the LHC*, *Phys. Rev. Lett.* **100** (2008) 242001, [0802.2470].
 - [6] T. Plehn, G. P. Salam and M. Spannowsky, *Fat Jets for a Light Higgs*, *Phys. Rev. Lett.* **104** (2010) 111801, [0910.5472].
 - [7] D. Krohn, J. Thaler and L.-T. Wang, *Jet Trimming*, *JHEP* **02** (2010) 084, [0912.1342].
 - [8] S. D. Ellis, C. K. Vermilion and J. R. Walsh, *Recombination Algorithms and Jet Substructure: Pruning as a Tool for Heavy Particle Searches*, *Phys. Rev.* **D81** (2010) 094023, [0912.0033].
 - [9] J. Gallicchio, J. Huth, M. Kagan, M. D. Schwartz, K. Black and B. Tweedie, *Multivariate discrimination and the Higgs + W/Z search*, *JHEP* **04** (2011) 069, [1010.3698].
 - [10] J. Gallicchio and M. D. Schwartz, *Seeing in Color: Jet Superstructure*, *Phys. Rev. Lett.* **105** (2010) 022001,

- [1001.5027].
- [11] J. Thaler and K. Van Tilburg, *Identifying Boosted Objects with N -subjettiness*, *JHEP* **03** (2011) 015, [1011.2268].
- [12] A. Hook, M. Jankowiak and J. G. Wacker, *Jet Dipolarity: Top Tagging with Color Flow*, *JHEP* **04** (2012) 007, [1102.1012].
- [13] A. J. Larkoski, G. P. Salam and J. Thaler, *Energy Correlation Functions for Jet Substructure*, *JHEP* **06** (2013) 108, [1305.0007].
- [14] A. J. Larkoski, S. Marzani, G. Soyez and J. Thaler, *Soft Drop*, *JHEP* **05** (2014) 146, [1402.2657].
- [15] ATLAS collaboration, G. Aad et al., *Search for dark matter in events with a hadronically decaying W or Z boson and missing transverse momentum in pp collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector*, *Phys. Rev. Lett.* **112** (2014) 041802, [1309.4017].
- [16] ATLAS collaboration, G. Aad et al., *Performance of jet substructure techniques for large- R jets in proton-proton collisions at $\sqrt{s} = 7$ TeV using the ATLAS detector*, *JHEP* **09** (2013) 076, [1306.4945].
- [17] ATLAS collaboration, G. Aad et al., *Search for high-mass diboson resonances with boson-tagged jets in proton-proton collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector*, *JHEP* **12** (2015) 055, [1506.00962].
- [18] ATLAS collaboration, G. Aad et al., *Identification of boosted, hadronically decaying W bosons and comparisons with ATLAS data taken at $\sqrt{s} = 8$ TeV*, *Eur. Phys. J.* **C76** (2016) 154, [1510.05821].
- [19] ATLAS collaboration, G. Aad et al., *Measurement of the charged-particle multiplicity inside jets from $\sqrt{s} = 8$ TeV pp collisions with the ATLAS detector*, *Eur. Phys. J.* **C76** (2016) 322, [1602.00988].
- [20] CMS collaboration, S. Chatrchyan et al., *Search for Anomalous $t\bar{t}$ Production in the Highly-Boosted All-Hadronic Final State*, *JHEP* **09** (2012) 029, [1204.2488].
- [21] CMS collaboration, *Boosted Top Jet Tagging at CMS*, Tech. Rep. CMS-PAS-JME-13-007, CERN, Geneva, 2014.
- [22] CMS collaboration, V. Khachatryan et al., *Identification techniques for highly boosted W bosons that decay into hadrons*, *JHEP* **12** (2014) 017, [1410.4227].
- [23] CMS collaboration, V. Khachatryan et al., *Search for vector-like charge $2/3$ T quarks in proton-proton collisions at $\sqrt{s} = 8$ TeV*, *Phys. Rev.* **D93** (2016) 012003, [1509.04177].
- [24] CMS COLLABORATION collaboration, *Search for massive resonances decaying into pairs of boosted W and Z bosons at $\sqrt{s} = 13$ TeV*, Tech. Rep. CMS-PAS-EXO-15-002, CERN, Geneva, 2015.
- [25] CMS collaboration, V. Khachatryan et al., *Jet energy scale and resolution in the CMS experiment in pp collisions at 8 TeV*, Submitted to: *JINST* (2016), [1607.03663].
- [26] J. Cogan, M. Kagan, E. Strauss and A. Schwartzman, *Jet-Images: Computer Vision Inspired Techniques for Jet Tagging*, *JHEP* **02** (2015) 118, [1407.5675].
- [27] L. G. Almeida, M. Backovi, M. Cliche, S. J. Lee and M. Perelstein, *Playing Tag with ANN: Boosted Top Identification with Pattern Recognition*, *JHEP* **07** (2015) 086, [1501.05968].
- [28] E. Racah, S. Ko, P. Sadowski, W. Bhimji, C. Tull, S.-Y. Oh et al., *Revealing Fundamental Physics from the Daya Bay Neutrino Experiment using Deep Neural Networks*, 1601.07621.
- [29] A. Aurisano, A. Radovic, D. Rocco, A. Himmel, M. D. Messier, E. Niner et al., *A Convolutional Neural Network Neutrino Event Classifier*, Submitted to: *JINST* (2016), [1604.01444].
- [30] L. Lonnblad, C. Peterson and T. Rognvaldsson, *Finding Gluon Jets With a Neural Trigger*, *Phys. Rev. Lett.* **65** (1990) 1321–1324.
- [31] L. Lonnblad, C. Peterson and T. Rognvaldsson, *Using neural networks to identify jets*, *Nucl. Phys.* **B349** (1991) 675–702.
- [32] C. Peterson, T. Rognvaldsson and L. Lonnblad, *JETNET 3.0: A Versatile artificial neural network package*, *Comput. Phys. Commun.* **81** (1994) 185–220.
- [33] P. Chiappetta, P. Colangelo, P. De Felice, G. Nardulli and G. Pasquariello, *Higgs search by neural networks at LHC*, *Phys. Lett.* **B322** (1994) 219–223, [hep-ph/9401343].
- [34] B. H. Denby, *Neural Networks and Cellular Automata in Experimental High-energy Physics*, *Comput. Phys. Commun.* **49** (1988) 429–448.
- [35] L. de Oliveira, M. Kagan, L. Mackey, B. Nachman and A. Schwartzman, *Jet-Images – Deep Learning Edition*, 1511.05190.
- [36] P. Baldi, K. Bauer, C. Eng, P. Sadowski and D. Whiteson, *Jet Substructure Classification in High-Energy Physics with Deep Neural Networks*, *Phys. Rev.* **D93** (2016) 094034, [1603.09349].
- [37] D. Guest, J. Collado, P. Baldi, S.-C. Hsu, G. Urban and D. Whiteson, *Jet Flavor Classification in High-Energy Physics with Deep Neural Networks*, 1607.08633.
- [38] A. Azatov, M. Salvarezza, M. Son and M. Spannowsky, *Boosting Top Partner Searches in Composite Higgs Models*, *Phys. Rev.* **D89** (2014) 075001, [1308.6601].
- [39] M. Gouzevitch, A. Oliveira, J. Rojo, R. Rosenfeld, G. P. Salam and V. Sanz, *Scale-invariant resonance tagging in multijet events and new physics in Higgs pair production*, *JHEP* **07** (2013) 148, [1303.6636].
- [40] J. Dolen, P. Harris, S. Marzani, S. Rappoccio and N. Tran, *Thinking outside the ROCs: Designing Decorrelated Taggers (DDT) for jet substructure*, 1603.00027.
- [41] M. Schlaffer, M. Spannowsky and A. Weiler, *Searching for supersymmetry scalelessly*, 1603.01638.
- [42] I. J. Goodfellow, D. Warde-Farley, M. Mirza, A. Courville and Y. Bengio, *Maxout Networks*, *ArXiv e-prints* (Feb., 2013), [1302.4389].
- [43] Y. B. Xavier Glorot, Antoine Bordes, *Deep sparse rectifier neural networks*, *JMLR* **15** (2011).
- [44] F. Chollet, “Keras.” <https://github.com/fchollet/keras>, 2015.
- [45] D. P. Kingma and J. Ba, *Adam: A method for stochastic optimization*, *CoRR abs/1412.6980* (2014).
- [46] S. Behnel, R. Bradshaw, C. Citro, L. Dalcin, D. Seljebotn and K. Smith, *Cython: The best of both worlds*, *Computing in Science Engineering* **13** (March–April, 2011) 31–39.
- [47] T. Sjostrand, S. Ask, J. R. Christiansen, R. Corke, N. Desai, P. Ilten et al., *An Introduction to PYTHIA 8.2*, *Comput. Phys. Commun.* **191** (2015) 159–177, [1410.3012].
- [48] DELPHES 3 collaboration, J. de Favereau, C. Delaere,

- P. Demin, A. Giammanco, V. Lematre, A. Mertens et al., *DELPHES 3, A modular framework for fast simulation of a generic collider experiment*, *JHEP* **02** (2014) 057, [1307.6346].
- [49] M. Cacciari, G. P. Salam and G. Soyez, *FastJet User Manual*, *Eur. Phys. J.* **C72** (2012) 1896, [1111.6097].
- [50] S. v. d. Walt, S. C. Colbert and G. Varoquaux, *The numpy array: A structure for efficient numerical computation*, *Computing in Science and Engg.* **13** (Mar., 2011) 22–30.
- [51] The HDF Group, *Hierarchical Data Format, version 5*, 1997–2016.
- [52] M. Dobbs and J. B. Hansen, *The HepMC C++ Monte Carlo event record for High Energy Physics*, *Comput. Phys. Commun.* **134** (2001) 41–46.
- [53] M. Cacciari, G. P. Salam and G. Soyez, *The Anti- $k(t)$ jet clustering algorithm*, *JHEP* **04** (2008) 063, [0802.1189].
- [54] W. H. Press, S. A. Teukolsky, W. T. Vetterling and B. P. Flannery, *Numerical Recipes in C (2nd Ed.): The Art of Scientific Computing*. Cambridge University Press, New York, NY, USA, 1992.
- [55] J. Bellm, G. Nail, S. Platzer, P. Schichtel and A. Sidmok, *Parton Shower Uncertainties with Herwig 7: Benchmarks at Leading Order*, 1605.01338.
- [56] J. R. Andersen et al., *Les Houches 2015: Physics at TeV Colliders Standard Model Working Group Report*, in *9th Les Houches Workshop on Physics at TeV Colliders (PhysTeV 2015) Les Houches, France, June 1-19, 2015*, 2016. 1605.04692.
- [57] P. Richardson and D. Winn, *Investigation of Monte Carlo Uncertainties on Higgs Boson searches using Jet Substructure*, *Eur. Phys. J.* **C72** (2012) 2178, [1207.0380].
- [58] T. Gleisberg and S. Hoeche, *Comix, a new matrix element generator*, *JHEP* **12** (2008) 039, [0808.3674].
- [59] T. Gleisberg, S. Hoeche, F. Krauss, M. Schonherr, S. Schumann, F. Siegert et al., *Event generation with SHERPA 1.1*, *JHEP* **02** (2009) 007, [0811.4622].
- [60] M. Bahr et al., *Herwig++ Physics and Manual*, *Eur. Phys. J.* **C58** (2008) 639–707, [0803.0883].
- [61] J. Bellm et al., *Herwig 7.0 / Herwig++ 3.0 Release Note*, 1512.01178.
- [62] N. Fischer, S. Prestel, M. Ritzmann and P. Skands, *Vincia for Hadron Colliders*, 1605.06142.
- [63] M. Ritzmann, D. A. Kosower and P. Skands, *Antenna Showers with Hadronic Initial States*, *Phys. Lett.* **B718** (2013) 1345–1350, [1210.6345].
- [64] S. Platzer and S. Gieseke, *Coherent Parton Showers with Local Recoils*, *JHEP* **01** (2011) 024, [0909.5593].
- [65] S. Platzer and S. Gieseke, *Dipole Showers and Automated NLO Matching in Herwig++*, *Eur. Phys. J.* **C72** (2012) 2187, [1109.6256].
- [66] S. Gieseke, P. Stephens and B. Webber, *New formalism for QCD parton showers*, *JHEP* **12** (2003) 045, [hep-ph/0310083].
- [67] S. Schumann and F. Krauss, *A Parton shower algorithm based on Catani-Seymour dipole factorisation*, *JHEP* **03** (2008) 038, [0709.1027].
- [68] T. Sjostrand and P. Z. Skands, *Transverse-momentum-ordered showers and interleaved multiple interactions*, *Eur. Phys. J.* **C39** (2005) 129–154, [hep-ph/0408302].
- [69] *Performance of Boosted W Boson Identification with the ATLAS Detector*, Tech. Rep. ATLAS-PHYS-PUB-2014-004, CERN, Geneva, Mar, 2014.
- [70] A. J. Larkoski, I. Moult and D. Neill, *Power Counting to Better Jet Observables*, *JHEP* **12** (2014) 009, [1409.6298].
- [71] A. J. Larkoski, I. Moult and D. Neill, *Analytic Boosted Boson Discrimination*, *JHEP* **05** (2016) 117, [1507.03018].
- [72] M. Spannowsky and M. Stoll, *Tracking New Physics at the LHC and beyond*, *Phys. Rev.* **D92** (2015) 054033, [1505.01921].
- [73] V. Mnih, N. Heess, A. Graves and K. Kavukcuoglu, *Recurrent models of visual attention*, *CoRR* **abs/1406.6247** (2014) .
- [74] D. E. Soper and M. Spannowsky, *Finding physics signals with event deconstruction*, *Phys. Rev.* **D89** (2014) 094005, [1402.1189].