



Leveraging universality of jet taggers through transfer learning

Frédéric A. Dreyer^{1,a}, Radosław Grabarczyk¹, Pier Francesco Monni²

¹ Clarendon Laboratory, Rudolf Peierls Centre for Theoretical Physics, University of Oxford, Parks Road, Oxford OX1 3PU, UK

² Theoretical Physics Department, CERN, 1211 Geneva 23, Switzerland

Received: 23 March 2022 / Accepted: 25 May 2022

© The Author(s) 2022

Abstract A significant challenge in the tagging of boosted objects via machine-learning technology is the prohibitive computational cost associated with training sophisticated models. Nevertheless, the universality of QCD suggests that a large amount of the information learnt in the training is common to different physical signals and experimental setups. In this article, we explore the use of transfer learning techniques to develop fast and data-efficient jet taggers that leverage such universality. We consider the graph neural networks `LundNet` and `ParticleNet`, and introduce two prescriptions to transfer an existing tagger into a new signal based either on fine-tuning all the weights of a model or alternatively on freezing a fraction of them. In the case of W -boson and top-quark tagging, we find that one can obtain reliable taggers using an order of magnitude less data with a corresponding speed-up of the training process. Moreover, while keeping the size of the training data set fixed, we observe a speed-up of the training by up to a factor of three. This offers a promising avenue to facilitate the use of such tools in collider physics experiments.

1 Introduction

The tagging of energetic heavy particles through machine learning methods is one of the key technical challenges at the Large Hadron Collider. Such identification techniques are used to search for new-physics signatures (see e.g. Refs. [1–8]), or to study the properties of Standard Model particles [9], notably to identify boosted electroweak bosons [10–13], the Higgs boson [14–21], or to assign jet flavour [22–38]. The most challenging scenario is the one in which such heavy objects decay into hadronic jets, in which case the ability to identify them from the decay products is seriously challenged by the overwhelming background arising from QCD jets. Provided one has a robust theoretical control over such

background processes, the use of pattern-recognition methods from computer science can help construct novel taggers with a significantly improved performance with respect to analytic discriminants (see e.g. Refs. [39–41] for recent reviews).

A large variety of methods has been proposed in recent years, achieving a remarkable performance in discriminating signal from background in different experimental measurements. These are based on several types of techniques, which span from the use of theory-motivated observables such as energy-flow polynomials [42], convolutional neural networks [10] and graph networks [43] that use four-momenta as input variables. Among the more recent tools, `LundNet` [44] combines the performance of the state of the art graph networks with theory-motivated kinematic input variables, namely the Lund jet plane [45] of emissions within a jet [46].

In the application of machine learning technology to jet tagging, a first challenge is represented by the robust assessment of the theoretical uncertainty in a given model. This is dominated by the dependence of the model's ability to discriminate a given signal from the QCD background on the underlying simulation that is used in the training. A precise control over these effects demands the development of more accurate event generators, a task that is receiving significant attention in the literature (see e.g. Refs. [47–60] and references therein for some recent developments). A second outstanding challenge is reducing the high computational cost of training sophisticated models, which besides the generation of large samples of events also usually requires running on GPUs for several days. In consequence, the training of such taggers requires computational resources not always in the reach of their potential users. Moreover, the models are highly dependent on the experimental signal (e.g. W vs. top-quark tagging) as well as on the choice of experimental cuts which makes the training of a tagger for specific experimental needs from scratch highly inefficient. The physical picture suggests, however, that most of the information learnt by a

^a e-mail: frederic.dreyer@physics.ox.ac.uk (corresponding author)

tagger is related to the description of the QCD splittings that occur within a jet, which simply encode universal properties of QCD rather than features that depend on the underlying experimental signal. It has been demonstrated multiple times in the past that early layers of a deep convolutional network extract general features from the data, and can thus be potentially reused for new tasks [61]. In the present short article we examine the latter aspect of jet tagging, and we tackle the problem by applying inductive *transfer learning* techniques [62] to leverage an existing model for a new application to a different experimental signature. As a result, we will discuss the construction of computationally efficient jet taggers that can achieve high performance also when trained on a small fraction of the original data set, with a significant reduction in the computational complexity associated with the training. The article is structured as follows. In Sect. 2 we briefly review the graph neural network `LundNet` that we adopt for our studies, and discuss the underlying description of jets in terms of Lund jet plane declusterings. In Sect. 3 we then introduce two transfer-learning procedures that allow one to train a new signal starting from an existing model trained on a different tagging problem. These techniques are then applied to the problem of top tagging in Sect. 4, where we study in detail the performance of transfer learning between top taggers with different transverse-momentum cuts and from a W tagger to a top tagger. Subsequently, we present an analysis of the computational advantages of transfer learning procedures over training new models from scratch. In Sect. 5 we discuss our conclusions.

2 Graph neural networks in the Lund plane

The Lund jet plane [46] is a useful theoretical framework to represent the internal kinematics of a jet by means of Lund diagrams [45]. To define it, one starts by constructing the Cambridge–Aachen [63,64] clustering sequence using the constituents of the jet, which carries out a sequential pairwise recombination of the two proto-jets with the smallest angular separation in rapidity–azimuth. One maps this clustering sequence to a tree of Lund declusterings, each of which encodes the kinematic properties of the corresponding clustering step. Each declustering $p_i \rightarrow p_a, p_b$ can be parametrised in terms of the following set of variables:

$$\begin{aligned} \Delta &\equiv \Delta_{ab}, \quad k_t \equiv p_{Tb} \Delta_{ab}, \quad m^2 \equiv (p_a + p_b)^2, \\ z &\equiv \frac{p_{Tb}}{p_{Ta} + p_{Tb}}, \quad \psi \equiv \tan^{-1} \frac{y_b - y_a}{\phi_b - \phi_a}, \end{aligned} \quad (1)$$

where p_a, p_b are the post-branching momenta with their transverse momenta ordered such that $p_{Tb} < p_{Ta}$, $\Delta_{ab} = \sqrt{(y_a - y_b)^2 + (\phi_a - \phi_b)^2}$ (with y and ϕ denoting the rapidity and azimuth, respectively), ψ is an azimuthal angle around




Fig. 1 A graphic representation of the Lund plane for the radiation within a jet. The blue triangle represents the primary Lund plane, with secondary and tertiary Lund planes shown in red and orange, respectively

the subjet axis, and z is the transverse momentum fraction of the branching. The construction of the Lund jet plane can be schematically understood with the help of Fig. 1. The (primary) Lund plane associated with the initial proto-jet represents a two-dimensional parametrisation of the phase space available to further radiation from it. This is indicated by the large (blue) triangle in the $\ln k_t - \ln 1/\Delta$ plane in Fig. 1. Each subsequent primary emission along the hard branch of the tree is shown in red, and it forms a new leaf of the Lund plane, from which secondary emissions will be radiated, indicated by orange leaves. The procedure iterates through all branches of the clustering history, leading to a complete representation of the jet’s substructure. In particular, the structure of the primary Lund jet plane can be computed accurately with perturbative methods [65] and measured experimentally [66].

`LundNet` [67] is a graph neural network which takes the Lund jet plane as input to train efficient and robust jet taggers. The resulting taggers outperform tools with low-level inputs [67] and are relatively resilient to non-perturbative and detector effects given an appropriate choice of cuts in the Lund plane. The jet is mapped into a graph whose nodes represent the declustering steps of the Cambridge–Aachen history, parametrised in terms of tuples $\mathcal{T}^{(i)}$ containing the kinematic variables defined in Eq. (1). In particular, one can define two versions of the `LundNet` network based on the dimensionality of the input tuple, defined as follows:

$$\text{LundNet3} : \mathcal{T}^{(i)} = \{k_t^{(i)}, \Delta^{(i)}, z^{(i)}\}, \quad (2)$$

$$\text{LundNet5} : \mathcal{T}^{(i)} = \{k_t^{(i)}, \Delta^{(i)}, z^{(i)}, m^{(i)}, \psi^{(i)}\}. \quad (3)$$

The edges of the graph correspond to the structure of the Cambridge–Aachen tree.

The LundNet5 network contains more kinematic information for each declustering node, and therefore results in a higher tagging efficiency. Conversely, the LundNet3 network has been shown to be more resilient to non-perturbative and detector effects [67], while having an efficiency similar to state-of-the-art taggers.

The core of the graph architecture relies on an EdgeConv operation [68], which applies a multi-layer perceptron (MLP) to produce a learned edge feature, using combined features of node pairs along each edge as input. This shared MLP consists of two layers, each with a dense network, batch normalisation [69] and ReLU activation [70]. This is followed by an aggregation step which takes an element-wise average of the learned edge features along the edges. The model also includes a shortcut connection [71]. The same MLP is applied to each node, updating all node features while keeping the structure of the graph itself unchanged. The LundNet architecture consists of six successive EdgeConv blocks, with the number of channels for each MLP pair in the block given by (32, 32), (32, 32), (64, 64), (64, 64), (128, 128) and (128, 128). Their final output is concatenated, and processed by a MLP with 384 channels, to which a global average pooling is applied to extract information from all nodes in the graph. This is followed by a fully connected layer with 256 units and a dropout layer with rate 10%. A final softmax output provides the result of the classification. This model is implemented with the Deep Graph Library 0.5.3 [72] and the PyTorch 1.7.1 [73] backend, using an Adam optimiser [74] to minimise the cross entropy loss. The LundNet architecture is summarised in Fig. 2a. Training is performed for 30 epochs using an initial learning rate of 0.001, which is lowered by a factor 10 after the 10th and 20th epochs.

The architecture of LundNet is based on a similar graph neural network, ParticleNet [43], which also provides excellent performance on LHC classification tasks, and we will use it as one of the benchmarks in our study below in comparison to LundNet.

3 Transfer learning

In this section we briefly discuss the application of transfer learning techniques to the design of jet taggers. Transfer learning aims at reusing pre-trained models on new problems, leveraging the knowledge obtained on a similar task to improve the training of a new model, for example by using some or all of the weights of an existing pre-trained neural network as starting point.

Transfer learning has seen a wide range of applications, notably in language processing and computer vision [75, 76]. While earlier data-based approaches to transfer learning

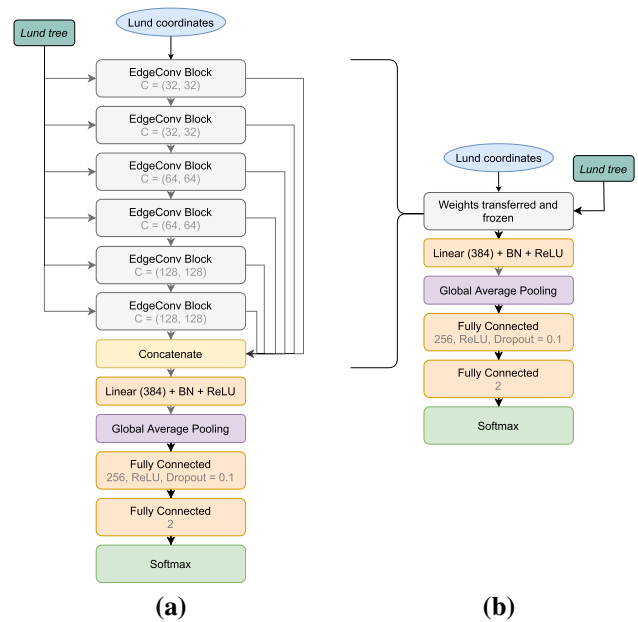


Fig. 2 A flowchart representing the LundNet architecture and the transfer learning procedure employed in this article

focused on domain adaptation [77, 78], there has been a surge of interest in recent years in adapting deep learning models to new tasks [61]. The two main avenues to achieve this goal is either through the retraining of a deep neural network while freezing the weights of its initial layers [79], or through the fine-tuning of the model [80–82]. While most existing applications are based on convolutional or recurrent neural networks, the development of deep learning on graph structured data has also seen advances in transfer learning applied on graph neural networks [83, 84].

In the context of machine learning applications to jet physics, one could expect that different taggers rely on a certain amount of information that is common to different tasks. Concretely, the properties of QCD that define the radiation pattern inside a boosted jet stemming from the QCD background is largely identical within different taggers, and further commonalities can be identified among signals with a similar number of prongs produced by the resonance decay inside a (fat) jet. This suggests that jet physics is an ideal area for the application of transfer learning methods. On the practical side, this would allow for the design of new models/taggers starting from a pre-existing one, not necessarily trained on the same task. The main advantage of transfer learning would then be the considerably reduced computational cost associated with the training of the new model, which does not need to be built from scratch for every new task.

A first important question is the extent to which a network is *transferable*, i.e. whether the transferred model is capable of reaching a performance that is as close as possible to

the fully trained model with just a fraction of the computing resources. A general answer to this question requires a thorough investigation of the features of a given network that are connected to a higher *transferability*, and this goes beyond the scope of this article. Here we instead take a first step in this direction, and consider the two graph neural networks which have been discussed in the previous section. Of these, `ParticleNet` relies on the information carried by the full four momenta of the jet constituents, whereas the `LundNet` models essentially map the Cambridge/Aachen sequence of a jet into its own Lund jet plane. The latter representation encodes the kinematic information at each of the branchings in the fragmentation process, which in turn is to a large extent universal across taggers and depends mainly on the properties of QCD near the soft and/or collinear limits. A second interesting property of the Lund jet plane is that QCD (background) jets roughly have a uniform density of emissions in the Lund jet plane, and hence this structure can be learnt very effectively by a neural network such as `LundNet`. Both of the above properties are expected to facilitate transfer learning in that the input variables on which `LundNet` relies already highlight universal properties of QCD jets and allow the network to distinguish them from those of typical signal jets stemming from the decay of a boosted heavy object. For this reason, our expectation is that transfer learning reaches a rather good performance in the context of this class of models. On the other hand, in the case of `ParticleNet`, the model needs to learn the non-trivial mapping between the information carried by the final-state four momenta used as input, to the physical fragmentation process of the jet. This ends up adding an additional layer of complexity in the training of the network, which is expected to be reflected in a lower performance of the models trained via transfer learning.

In order to explore the *transferability* of the models adopted in this article, we consider two different approaches to inductive transfer learning. Our first transfer-learning approach is a frozen-layer model. Here the weights of the EdgeConv layers of `LundNet` or `ParticleNet` have been pre-trained on a separate jet sample and are kept fixed during the retraining process. The final MLP layers are instead reinitialised to random weights and retrained on a new sample to specialise the tagger to this new pattern recognition task. This procedure is shown in Fig. 2b, and the training on a new data set is performed with the same learning rate and scheduler as the training for the original model. The second approach is a fine-tuning of all the weights in the original tagger. In this case, the learning rate is reduced by a factor ten (or a factor three when transferring from a W to a top tagger), and the tagger is retrained with the same number of epochs and scheduler on a new data set.

The difference in performance between the transferred models and those trained from scratch probes the ability of a network to learn features that are common to different

tasks and therefore its suitability to the application of transfer learning techniques. Moreover, the difference between the frozen-layer and fine-tuning procedures probes how much of the high-level information learnt by a network is extracted from the initial layers, which provides insights on the extrapolating ability of each model.

4 Case study with top tagging

As a case study, we consider the application to top-quark tagging at the LHC. We are interested in discriminating top-quark jets with $p_T > 500$ GeV against the QCD background. We apply and analyse the properties of transfer learning for four different models, `LundNet3` and its transfer from a top-tagger with $p_T > 2$ TeV, `LundNet5` and its transfer from either a top-tagger with $p_T > 2$ TeV or a W -boson tagger with $p_T > 500$ GeV, and finally `ParticleNet` and its transfer from a top-tagger with $p_T > 2$ TeV. All models presented in this section are trained with events generated using Pythia 8.223 [85], considering jets defined according to the anti- k_t algorithm [86] with a jet radius $R = 0.8$ and rapidity $|y| < 2.5$. Signal events are obtained from the simulation of either WW or $t\bar{t}$ production, with W bosons decaying hadronically, while the background sample is obtained from QCD dijet events. The signal and background training data sets consist of 5×10^5 events each. Validation and testing is done on data sets of 5×10^4 events for each of the signal and background process.¹

We start by discussing the reduction in computational complexity that can be achieved through the use of transfer learning techniques, and then provide a phenomenological study of top tagging performance for each model.

4.1 Computational complexity of transferred models

As described in Sect. 3, we consider two transfer learning approaches to retrain an existing jet tagger, a frozen layer model and a fine-tuning model. In this section, we aim to investigate the computational cost of both approaches. For this we consider the construction of a top tagger for 500 GeV jets, starting from a model trained on 2 TeV data. The training time was measured on a NVIDIA GeForce RTX 2080 Ti GPU, training a tagger on either 10^6 or 10^5 total top and QCD jets with an equal number of signal and background events. We measure the time required to train a `LundNet5` and `ParticleNet` model,² given in milliseconds per sample and epoch (which is identical for both the 10^6 and 10^5 samples) as well as the corresponding total training time

¹ The data is available at <https://github.com/JetsGame/data>.

² The training time for both `LundNet5` and `LundNet3` is almost identical.




Fig. 3 Validation accuracy as a function of the number of training epochs

for both data sizes. These measurements are summarised in Table 1. All models were trained for 30 epochs, regardless of their convergence, shown in Fig. 3 for the full data sample. The figure shows the convergence for the two sample cases of LundNet5 (in blue) and ParticleNet (in orange). The solid lines show the evolution for models trained from scratch, the dashed lines refer to the fine-tuning transfer learning setup (trained using a ten-times smaller initial learning rate of 10^{-4}), and finally the dotted lines refer to the frozen-layer transfer learning setup. One observes here that, in practice, the transferred models approach an optimum at a much faster rate, with the fine-tuning setup converging sooner to higher values of the validation accuracy. This implies that such models could be trained for only a few epochs to further reduce the computational cost.

The fine-tuning model does not provide any noticeable speed-up in training time, as it requires the update of all the weights in the network. However, as we will see in Sect. 4.2, it can achieve comparable performance to a full model with only a small fraction of the data. As such, if one uses a tenth of the full data sample, this effectively provides a factor of ten speed-up in training time. Moreover, as already pointed out, the convergence of the model is significantly faster, and requires only a few epochs to converge to an optimal solution, hence providing an opportunity for further optimisation of the training time.

Conversely, the frozen-layer approach has the advantage of reducing the computational cost of the retraining by limiting the update of the weights to the final dense layers, while keeping the EdgeConv blocks unchanged. This results in a reduction of the training time by about a factor of three compared to that of a full model on the same data sample. As for the fine-tuning model, a further reduction can be achieved by reducing the number of epochs and the size of the data set. However, as will be discussed in Sect. 4.2, the frozen layer model requires a larger data sample than the fine-tuning approach to achieve comparable performance.

Table 1 Training time for different taggers. The time was measured when running the models on an NVIDIA GeForce RTX 2080 Ti card. Note that LundNet5 transferred from a top tagger with $p_T > 2$ TeV, when trained on a data set 10 times smaller, still performs better than ParticleNet in terms of AUC (see Table 2) despite the time of training being significantly reduced

| | Training time (ms/sample/epoch) | Total for 10^6 samples (hh:mm:ss) | Total for 10^5 samples (hh:mm:ss) |
|---------------------------|---------------------------------|-------------------------------------|-------------------------------------|
| LundNet5 | 0.46 | 03:48:15 | 00:22:43 |
| LN5 _{frozen} | 0.15 | 01:17:02 | 00:07:36 |
| LN5 _{finetuning} | 0.46 | 03:48:32 | 00:22:45 |
| ParticleNet | 3.60 | 30:09:44 | 02:59:17 |
| PN _{frozen} | 2.16 | 18:13:21 | 01:47:37 |
| PN _{finetuning} | 3.60 | 29:59:46 | 03:01:04 |

4.2 Performance of top taggers

We now study the performance of our top taggers, using the area under the ROC curve (AUC) as an indicator of a model's performance, and summarise our results in Table 2. For each model, we consider the AUCs corresponding to a training from scratch (indicated by AUC in the table), transfer learning with a fine-tuning setup (indicated by AUC_{FT}), and transfer learning with a frozen-layer setup (indicated by AUC_{FR}). For the above three options, we also consider the values of the AUC obtained with 10% of the original training data (denoted by AUC^(10%) in Table 2). We find that, in the case of LundNet5 models trained in the full data set, the fine-tuning setup reproduces exactly the AUC of the model trained from scratch. We also observe that the frozen option, despite being considerably cheaper from a computational viewpoint, leads to AUC values which are extremely close to those of the above models, indicating that LundNet5 models are very suitable for the application of transfer learning as discussed in the previous section. We also consider the more demanding transfer of a LundNet5 W tagger to a top tagger at the same 500 GeV p_T threshold. We can see from the AUC values shown in Table 2 that while a moderate loss of performance is found for the model trained on the reduced data set, we still recover AUC values for the transferred top tagger that are very close to the fully trained LundNet5 model, and significantly better than most state-of-the-art taggers.

In the case of ParticleNet, the fine-tuning setup still performs as well as the model trained from scratch while the frozen setup leads to visibly smaller AUC values. As expected, this indicates that ParticleNet is less suitable for the application of transfer learning. This is due to the fact that it relies on low-level information, such as four momenta, which makes it less easy to identify general properties of the kinematic pattern of QCD already in early layers

Table 2 Benchmarks for top tagging with $p_T > 500$ GeV. The different columns show the AUC for the different transfer learning models considered in the text, where FT denotes the *fine-tuning* option, FR

| | AUC | AUC _{FT} | AUC _{FR} | AUC ^(10%) | AUC _{FT} ^(10%) | AUC _{FR} ^(10%) |
|------------------------------|--------|-------------------|-------------------|----------------------|------------------------------------|------------------------------------|
| LundNet3 (from top 2 TeV) | 0.9820 | 0.9820 | 0.9816 | 0.9773 | 0.9802 | 0.9791 |
| LundNet5 (from top 2 TeV) | 0.9866 | 0.9865 | 0.9863 | 0.9826 | 0.9850 | 0.9845 |
| LundNet5 (from W 500 GeV) | – | 0.9863 | 0.9858 | – | 0.9834 | 0.9832 |
| ParticleNet (from top 2 TeV) | 0.9826 | 0.9826 | 0.9793 | 0.9765 | 0.9795 | 0.9772 |

of the network. For models trained on 10% of the original training data set, we observe that in the case of LundNet3 and LundNet5, the values of AUC obtained with transfer learning models are hardly affected by the reduction in sample size, and they still perform nearly as well as the original models trained on the full data set. For ParticleNet the performance of the models obtained through transfer learning is instead closer to that of the model trained (from scratch) on the reduced data set, in line with our expectation that this class of models is less transferable. We also show the dependence of the AUC on the total (signal plus background) size of the training data set in Fig. 4. The figure shows that transfer learning gives a significant advantage for small sizes of the training data set. For example, the retrained LundNet5 model with the fine-tuning setup and 1.25×10^4 events for signal and background data sets achieves $\text{AUC} = 0.983$, meaning that state-of-the-art performance can be achieved using far smaller data sets than those needed to train a network from scratch, with a considerable speed-up of the process. Concretely, when retraining with 2.5×10^4 samples, the training time is almost two orders of magnitude smaller than that needed to train a similarly performing LundNet model from scratch. Importantly, the difference between the fully trained model and the fine-tuning and frozen-layer transfer learning setups is rather moderate in the case of LundNet5, which indicates that such class of models have rather high *transferability* and they can easily be retrained on a different task. In the case of ParticleNet, we observe that the fine-tuning setup still produces AUC values higher than those of the model fully trained on smaller data sets, although it does not reach the tagging accuracy observed for LundNet5. Moreover, Fig. 4 also shows that the performance of ParticleNet gets significantly worse when using the frozen-layer setup, with the fully trained model outperforming the transfer learning results already for a training done on 10^5 events, while LundNet5 reaches almost the asymptotic values of AUC for this data sample (see also Table 2). Overall, this clearly shows that the use of transfer learning provides a promising avenue to reduce the amount of data required to train new taggers, with certain classes of models such as LundNet being more suitable for the application of these techniques. Whether it is possible to

denotes the *frozen-layer* option, and the (10%) superscript refers to results obtained with just one tenth of the original training data




Fig. 4 Area under the ROC curve as a function of the total signal and background training data set size

define a metric quantifying a priori the ability of a model to be transferred to a different task with reduced computational resources than those needed for a full training, and how to construct better taggers with such features remain interesting open questions.

We now move on to study the ROC curves corresponding to the different models in Fig. 5, showing the background rejection $1/\epsilon_{\text{QCD}}$ versus signal efficiency, ϵ_{Top} . A better performing tagger has a corresponding ROC curve closer to the top-right corner of the figure. The upper panel shows the ROC corresponding to the models LundNet3, LundNet5 and ParticleNet all trained from scratch for a top tagger with $p_T > 500$ GeV. We observe that, as expected, LundNet5 performs better than the other two models, which achieve a very similar performance. This is due to the additional information stored in the tuples associated with each node of the graph (see Eq. (3)). The second panel of Fig. 5 shows the ROC obtained with LundNet5 and different transfer learning options from a top tagger with $p_T > 2$ TeV, divided by the ROC of the model trained from scratch (shown in the upper panel). The dashed blue line corresponds to the fine-tuning setup in which all weights are re-trained on the new task. This option clearly reproduces the performance of the tagger trained from scratch, but as already observed before it does not lead to any reduction of the computational complexity associated with the training. The dotted blue line,

instead, corresponds to the transfer learning obtained with the frozen-layer setup which, as already observed in Table 2, leads to a performance that is very close to that of the original model, with an AUC less than a permille below the full model, and background rejection at intermediate signal efficiencies within 20% of the fully trained tagger. This performance remains far better than most state-of-the-art jet taggers, and orders of magnitude above analytic substructure discriminants.

The remaining three panels in Fig. 5 show a similar comparison in the case of LundNet3 and ParticleNet models transferred from a top tagger with $p_T > 2$ TeV, and LundNet5 models transferred from a W -boson tagger with $p_T > 500$ GeV. For the fine-tuned W , the initial learning rate is set to 3×10^{-4} to allow for a larger perturbation of the pre-trained top model. All of the above four panels also report, in red, the result obtained with a reduced training data set of 10% of the original size, i.e. 10^5 events, with either the fine-tuning (dashed) or frozen-layer (dotted) setup. For LundNet, the plot confirms the conclusions drawn from the AUC study above, showing that these models (both for LundNet3 and LundNet5) still reach the performance of state-of-the-art taggers also in the transfer learning setups, with the frozen-layer setup being only moderately less accurate than the computationally more demanding fine-tuning. While it is clearly easier to transfer a model from a similar tagger trained on a different kinematic regime, we see that transfer learning still reaches highly competitive ROC curves also when the starting model is a W tagger, shown in the last panel of Fig. 5, which demonstrates that the techniques studied in this article can be adopted across wide families of jet taggers. As already observed, ParticleNet, shown in the fourth panel of Fig. 5, performs less well under the transfer learning setups, with a wider gap between the fine-tuning and frozen-layer options.

In general, from Fig. 5 we conclude that the retrained models achieve performances which are extremely close to the models trained from scratch, meaning that the output of EdgeConv operations is a representation of the data which can be efficiently reused for other tasks. The benefits of transfer learning then consist of a significantly shorter training time and a smaller data set size required to converge on an efficient tagger. All models in the above comparison have been trained for 30 epochs. We stress once again that the computational cost can be reduced further by exploiting the fact that transferred LundNet models converge to an optimum with less epochs, as discussed above in Fig. 3.

5 Conclusions

In this article, we have explored the use of transfer learning methods to train efficient jet taggers from existing models.




Fig. 5 QCD rejection vs. top tagging efficiency

With this, we aimed to investigate the ability of a neural network to learn universal features of QCD and to transfer them to a separate task. In practice, we have considered the application of transfer learning to top tagging at different transverse momentum thresholds and to the tagging of two- and three-pronged boosted objects, e.g. W boson and top quark decays. We studied two jet taggers constructed from graph neural networks, LundNet and ParticleNet, and conducted a detailed study of the performance of transferred models as well as of the reduction in computational complexity provided by transfer learning.

We have implemented two transfer-learning procedures. The first one relies on fine-tuning all weights in a model by retraining it on a new data sample with a lower learning rate, while the second freezes the edge convolutions and retrains solely the final dense layers of the network. In the case of LundNet taggers, we find that the fine-tuning approach requires a similar training time per epoch and sample as the fully trained model, but converges to an almost optimal solu-

tion after just a few epochs (compared to tens of epochs for a full model) and requires only a small fraction of the data. Concretely, a model can achieve nearly the same performance using a third of the epochs and a tenth of the original training sample, which leads to a dramatic speed-up of the training process. On the other hand, the frozen-layer method provides a further speed-up in training time by a factor three as only a small fraction of the model weights are updated, but requires a comparatively larger sample size to achieve a similar performance to the fine-tuning approach.

For the two specific LundNet taggers considered (LundNet3 and LundNet5, which differ in the dimensionality of the kinematic inputs associated with each node of the graph), we observe that fine-tuning with a tenth of the data achieves a background rejection moderately lower than that of a fully trained model, with the transferred LundNet3 tagger recovering slightly more of the performance of the baseline model. The frozen-layer approach performs comparably, although in both LundNet3 and LundNet5 it achieves slightly lower background rejection for the same training sample than the fine-tuning method.

The conclusions are somewhat different for ParticleNet, where the frozen-layer method performs noticeably worse than the fine-tuning approach, regardless of the amount of data and number of epochs. Furthermore, the background rejection that can be achieved with a reduced data set is significantly smaller. We attribute this to the fact that ParticleNet relies on kinematic information structured as the four momenta of the jet constituents, which in turn makes it more challenging for the EdgeConv layers to extract general features about the jet fragmentation. In comparison, LundNet uses kinematic information of the sequential clustering steps of the Cambridge/Aachen algorithm as input, which carries denser information about the jet fragmentation dynamics. This is reflected in a larger gap between the fine-tuning and frozen-layer approaches in the ParticleNet case.

Our results show that transfer learning constitutes a promising avenue to build computationally efficient and versatile taggers with state-of-the-art performance. This opens a wide array of possibilities for more wide-spread adoption of machine learning jet-tagging technology for experimental studies at colliders, such as the Large Hadron Collider and future facilities. This article provides a first step towards this goal, and motivates further investigations on the application of these methods to particle phenomenology. In this context, a number of interesting theoretical questions arise. As future directions, it would be informative to study concrete metrics of transferability of a network, and which features of the input variables and choices in the architecture of a model can lead to more transferable designs. Furthermore, it would be interesting to study knowledge transfer in jet taggers from first

principles, and gain analytical insights into the behaviour of transferred models [38, 87].

Acknowledgements We are grateful to Alexander Huss for constructive comments and collaboration in the early stages of this work, and to Stefano Carrazza and Huilin Qu for valuable comments on the manuscript. This work was supported by a Royal Society University Research Fellowship (URF\R1\211294) (FD), and by the CERN's Summer Student Programme (RG).

Data Availability Statement This manuscript has associated data in a data repository. [Authors' comment: The code used is available in [44]. The training data can be obtained from [88].]

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

Funded by SCOAP³.

References

1. J.H. Collins, K. Howe, B. Nachman, Anomaly detection for resonant new physics with machine learning. *Phys. Rev. Lett.* **121**(24), 241803 (2018). <https://doi.org/10.1103/PhysRevLett.121.241803>. arXiv:1805.02664 [hep-ph]
2. T. Heimgel, G. Kasieczka, T. Plehn, J.M. Thompson, QCD or what? *SciPost Phys.* **6**(3), 030 (2019). <https://doi.org/10.21468/SciPostPhys.6.3.030>. arXiv:1808.08979 [hep-ph]
3. A. De Simone, T. Jacques, Guiding new physics searches with unsupervised learning. *Eur. Phys. J. C* **79**(4), 289 (2019). <https://doi.org/10.1140/epjc/s10052-019-6787-3>. arXiv:1807.06038 [hep-ph]
4. R.T. D'Agnolo, A. Wulzer, Learning new physics from a machine. *Phys. Rev. D* **99**(1), 015014 (2019). <https://doi.org/10.1103/PhysRevD.99.015014>. arXiv:1806.02350 [hep-ph]
5. J.H. Collins, K. Howe, B. Nachman, Extending the search for new resonances with machine learning. *Phys. Rev. D* **99**(1), 014038 (2019). <https://doi.org/10.1103/PhysRevD.99.014038>. arXiv:1902.02634 [hep-ph]
6. J.A. Aguilar-Saavedra, F.R. Joaquim, J.F. Seabra, Mass unspecific supervised tagging (MUST) for boosted jets. *JHEP* **03**, 012 (2021). [https://doi.org/10.1007/JHEP03\(2021\)012](https://doi.org/10.1007/JHEP03(2021)012). arXiv:2008.12792 [hep-ph]. [Erratum: *JHEP* **04**, 133 (2021)]
7. G. Kasieczka et al., The LHC Olympics 2020 a community challenge for anomaly detection in high energy physics. *Rep. Prog. Phys.* **84**(12), 124201 (2021). <https://doi.org/10.1088/1361-6633/ac36b9>. arXiv:2101.08320 [hep-ph]
8. G. Karagiorgi, G. Kasieczka, S. Kravitz, B. Nachman, D. Shih, Machine learning in the search for new fundamental physics. arXiv:2112.03769 [hep-ph]
9. A. Radovic, M. Williams, D. Rousseau, M. Kagan, D. Bonaccorsi, A. Himmel, A. Aurisano, K. Terao, T. Wongjirad, Machine learning at the energy and intensity frontiers of particle

- physics. *Nature* **560**(7716), 41–48 (2018). <https://doi.org/10.1038/s41586-018-0361-2>
10. L. de Oliveira, M. Kagan, L. Mackey, B. Nachman, A. Schwartzman, Jet-images—deep learning edition. *JHEP* **07**, 069 (2016). [https://doi.org/10.1007/JHEP07\(2016\)069](https://doi.org/10.1007/JHEP07(2016)069). arXiv:1511.05190 [hep-ph]
 11. J. Barnard, E.N. Dawe, M.J. Dolan, N. Rajcic, Parton shower uncertainties in jet substructure analyses with deep neural networks. *Phys. Rev. D* **95**(1), 014018 (2017). <https://doi.org/10.1103/PhysRevD.95.014018>. arXiv:1609.00607 [hep-ph]
 12. G. Louppe, K. Cho, C. Becot, K. Cranmer, QCD-aware recursive neural networks for jet physics. *JHEP* **01**, 057 (2019). [https://doi.org/10.1007/JHEP01\(2019\)057](https://doi.org/10.1007/JHEP01(2019)057). arXiv:1702.00748 [hep-ph]
 13. C.M.S. Collaboration, A.M. Sirunyan et al., Identification of heavy, energetic, hadronically decaying particles using machine-learning techniques. *JINST* **15**(06), P06005 (2020). <https://doi.org/10.1088/1748-0221/15/06/P06005>. arXiv:2004.08262 [hep-ex]
 14. J. Lin, M. Freytsis, I. Moul, B. Nachman, Boosting $H \rightarrow b\bar{b}$ with machine learning. *JHEP* **10**, 101 (2018). [https://doi.org/10.1007/JHEP10\(2018\)101](https://doi.org/10.1007/JHEP10(2018)101). arXiv:1807.10768 [hep-ph]
 15. S.H. Lim, M.M. Nojiri, Spectral analysis of jet substructure with neural networks: boosted Higgs case. *JHEP* **10**, 181 (2018). [https://doi.org/10.1007/JHEP10\(2018\)181](https://doi.org/10.1007/JHEP10(2018)181). arXiv:1807.03312 [hep-ph]
 16. K. Datta, A. Larkoski, B. Nachman, Automating the construction of jet observables with machine learning. *Phys. Rev. D* **100**(9), 095016 (2019). <https://doi.org/10.1103/PhysRevD.100.095016>. arXiv:1902.07180 [hep-ph]
 17. E.A. Moreno, T.Q. Nguyen, J.-R. Vlimant, O. Cerri, H.B. Newman, A. Periwai, M. Spiropulu, J.M. Duarte, M. Pierini, Interaction networks for the identification of boosted $H \rightarrow b\bar{b}$ decays. *Phys. Rev. D* **102**(1), 012010 (2020). <https://doi.org/10.1103/PhysRevD.102.012010>. arXiv:1909.12285 [hep-ex]
 18. A. Chakraborty, S.H. Lim, M.M. Nojiri, Interpretable deep learning for two-prong jet classification with jet spectra. *JHEP* **07**, 135 (2019). [https://doi.org/10.1007/JHEP07\(2019\)135](https://doi.org/10.1007/JHEP07(2019)135). arXiv:1904.02092 [hep-ph]
 19. C.K. Khosa, S. Marzani, Higgs boson tagging with the Lund jet plane. *Phys. Rev. D* **104**(5), 055043 (2021). <https://doi.org/10.1103/PhysRevD.104.055043>. arXiv:2105.03989 [hep-ph]
 20. L. Cavallini, A. Coccaro, C.K. Khosa, G. Manco, S. Marzani, F. Parodi, D. Rebuffi, A. Rescia, G. Stagnitto, Tagging the Higgs boson decay to bottom quarks with colour-sensitive observables and the Lund jet plane. arXiv:2112.09650 [hep-ph]
 21. H. Qu, C. Li, S. Qian, Particle transformer for jet tagging. arXiv:2202.03772 [hep-ph]
 22. L.G. Almeida, M. Backović, M. Cliche, S.J. Lee, M. Perelstein, Playing tag with ANN: boosted top identification with pattern recognition. *JHEP* **07**, 086 (2015). [https://doi.org/10.1007/JHEP07\(2015\)086](https://doi.org/10.1007/JHEP07(2015)086). arXiv:1501.05968 [hep-ph]
 23. D. Guest, J. Collado, P. Baldi, S.-C. Hsu, G. Urban, D. Whiteson, Jet flavor classification in high-energy physics with deep neural networks. *Phys. Rev. D* **94**(11), 112002 (2016). <https://doi.org/10.1103/PhysRevD.94.112002> arXiv:1607.08633 [hep-ex]
 24. P.T. Komiske, E.M. Metodiev, M.D. Schwartz, Deep learning in color: towards automated quark/gluon jet discrimination. *JHEP* **01**, 110 (2017). [https://doi.org/10.1007/JHEP01\(2017\)110](https://doi.org/10.1007/JHEP01(2017)110) arXiv:1612.01551 [hep-ph]
 25. ATLAS Collaboration, Quark versus gluon jet tagging using jet images with the ATLAS detector
 26. G. Kasieczka, T. Plehn, M. Russell, T. Schell, Deep-learning top taggers or the end of QCD? *JHEP* **05**, 006 (2017). [https://doi.org/10.1007/JHEP05\(2017\)006](https://doi.org/10.1007/JHEP05(2017)006) arXiv:1701.08784 [hep-ph]
 27. A. Butter, G. Kasieczka, T. Plehn, M. Russell, Deep-learned top tagging with a Lorentz layer. *SciPost Phys.* **5**(3), 028 (2018). <https://doi.org/10.21468/SciPostPhys.5.3.028> arXiv:1707.08966 [hep-ph]
 28. ATLAS Collaboration, Identification of jets containing b -hadrons with recurrent neural networks at the ATLAS experiment
 29. CMS Collaboration, A.M. Sirunyan et al., Identification of heavy-flavour jets with the CMS detector in pp collisions at 13 TeV. *JINST* **13**(5), P05011 (2018). <https://doi.org/10.1088/1748-0221/13/05/P05011>. arXiv:1712.07158 [physics.ins-det]
 30. T. Cheng, Recursive neural networks in quark/gluon tagging. *Comput. Softw. Big Sci.* **2**(1), 3 (2018). <https://doi.org/10.1007/s41781-018-0007-y> arXiv:1711.02633 [hep-ph]
 31. S. Macaluso, D. Shih, Pulling out all the tops with computer vision and deep learning. *JHEP* **10**, 121 (2018). [https://doi.org/10.1007/JHEP10\(2018\)121](https://doi.org/10.1007/JHEP10(2018)121) arXiv:1803.00107 [hep-ph]
 32. E.A. Moreno, O. Cerri, J.M. Duarte, H.B. Newman, T.Q. Nguyen, A. Periwai, M. Pierini, A. Serikova, M. Spiropulu, J.-R. Vlimant, JEDI-net: a jet identification algorithm based on interaction networks. *Eur. Phys. J. C* **80**(1), 58 (2020). <https://doi.org/10.1140/epjc/s10052-020-7608-4> arXiv:1908.05318 [hep-ex]
 33. A. Andreassen, B. Nachman, D. Shih, Simulation assisted likelihood-free anomaly detection. *Phys. Rev. D* **101**(9), 095004 (2020). <https://doi.org/10.1103/PhysRevD.101.095004> arXiv:2001.05001 [hep-ph]
 34. ATLAS Collaboration, Deep sets based neural networks for impact parameter flavour tagging in ATLAS
 35. G. Kasieczka, S. Marzani, G. Soyez, G. Stagnitto, Towards machine learning analytics for jet substructure. *JHEP* **09**, 195 (2020). [https://doi.org/10.1007/JHEP09\(2020\)195](https://doi.org/10.1007/JHEP09(2020)195) arXiv:2007.04319 [hep-ph]
 36. R.T. d’Agnolo, G. Grosso, M. Pierini, A. Wulzer, M. Zanetti, Learning new physics from an imperfect machine. *Eur. Phys. J. C* **82**(3), 275 (2022). <https://doi.org/10.1140/epjc/s10052-022-10226-y>. arXiv:2111.13633 [hep-ph]
 37. A. Romero, D. Whiteson, M. Fenton, J. Collado, P. Baldi, Safety of quark/gluon jet classification. arXiv:2103.09103 [hep-ph]
 38. F. Dreyer, G. Soyez, A. Takacs, Quarks and gluons in the Lund plane. arXiv:2112.09140 [hep-ph]
 39. D. Guest, K. Cranmer, D. Whiteson, Deep learning and its application to LHC physics. *Annu. Rev. Nucl. Part. Sci.* **68**(1), 161–181 (2018). <https://doi.org/10.1146/annurev-nucl-101917-021019>
 40. F. Van der Veken, G. Azzopardi, F. Blanc, L. Coyle, E. Fol, M. Giovannozzi, T. Pieloni, S. Redaelli, L. Rivkin, B. Salvachua, M. Schenk, R. Tomas, G. Valentino, Application of machine learning techniques at the CERN Large Hadron Collider 06 (2020). <https://doi.org/10.22323/1.364.0006>
 41. A.J. Larkoski, I. Moul, B. Nachman, Jet substructure at the Large Hadron Collider: a review of recent advances in theory and machine learning. *Phys. Rep.* **841**, 1–63 (2020). <https://doi.org/10.1016/j.physrep.2019.11.001>
 42. P.T. Komiske, E.M. Metodiev, J. Thaler, Energy flow polynomials: a complete linear basis for jet substructure. *JHEP* **04**, 013 (2018). [https://doi.org/10.1007/JHEP04\(2018\)013](https://doi.org/10.1007/JHEP04(2018)013) arXiv:1712.07124 [hep-ph]
 43. H. Qu, L. Gouskos, ParticleNet: jet tagging via particle clouds. *Phys. Rev. D* **101**(5), 056019 (2020). <https://doi.org/10.1103/PhysRevD.101.056019> arXiv:1902.08570 [hep-ph]
 44. F.A. Dreyer, H. Qu, LundNet v1.0.0 (2021). <https://doi.org/10.5281/zenodo.4443152>
 45. B. Andersson, G. Gustafson, L. Lonnblad, U. Pettersson, Coherence effects in deep inelastic scattering. *Z. Phys. C* **43**, 625 (1989). <https://doi.org/10.1007/BF01550942>
 46. F.A. Dreyer, G.P. Salam, G. Soyez, The Lund jet plane. *JHEP* **12**, 064 (2018). [https://doi.org/10.1007/JHEP12\(2018\)064](https://doi.org/10.1007/JHEP12(2018)064) arXiv:1807.04758 [hep-ph]
 47. K. Hamilton, P. Nason, E. Re, G. Zanderighi, NNLOPS simulation of Higgs boson production. *JHEP* **10**, 222 (2013). [https://doi.org/10.1007/JHEP10\(2013\)222](https://doi.org/10.1007/JHEP10(2013)222) arXiv:1309.0017 [hep-ph]
 48. S. Alioli, C.W. Bauer, C. Berggren, F.J. Tackmann, J.R. Walsh, S. Zuberi, Matching fully differential NNLO calculations and

- parton showers. *JHEP* **06**, 089 (2014). [https://doi.org/10.1007/JHEP06\(2014\)089](https://doi.org/10.1007/JHEP06(2014)089) arXiv:1311.0286 [hep-ph]
49. S. Höche, Y. Li, S. Prestel, Drell–Yan lepton pair production at NNLO QCD with parton showers. *Phys. Rev. D* **91**(7), 074015 (2015). <https://doi.org/10.1103/PhysRevD.91.074015> arXiv:1405.3607 [hep-ph]
 50. P.F. Monni, P. Nason, E. Re, M. Wiesemann, G. Zanderighi, MiNNLO_{PS}: a new method to match NNLO QCD to parton showers. *JHEP* **05**, 143 (2020). [https://doi.org/10.1007/JHEP05\(2020\)143](https://doi.org/10.1007/JHEP05(2020)143) arXiv:1908.06987 [hep-ph]
 51. P.F. Monni, E. Re, M. Wiesemann, MiNNLO_{PS}: optimizing 2 → 1 hadronic processes. *Eur. Phys. J. C* **80**(11), 1075 (2020). <https://doi.org/10.1140/epjc/s10052-020-08658-5> arXiv:2006.04133 [hep-ph]
 52. S. Alioli, C.W. Bauer, A. Broggio, A. Gavardi, S. Kallweit, M.A. Lim, R. Nagar, D. Napoletano, L. Rottoli, Matching NNLO predictions to parton showers using N3LL color-singlet transverse momentum resummation in Geneva. *Phys. Rev. D* **104**(9), 094020 (2021). <https://doi.org/10.1103/PhysRevD.104.094020> arXiv:2102.08390 [hep-ph]
 53. J.M. Campbell, S. Höche, H.T. Li, C.T. Preuss, P. Skands, Towards NNLO+PS matching with sector showers. arXiv:2108.07133 [hep-ph]
 54. S. Prestel, Matching N3LO QCD calculations to parton showers. *JHEP* **11**, 041 (2021). [https://doi.org/10.1007/JHEP11\(2021\)041](https://doi.org/10.1007/JHEP11(2021)041) arXiv:2106.03206 [hep-ph]
 55. S. Höche, F. Krauss, S. Prestel, Implementing NLO DGLAP evolution in parton showers. *JHEP* **10**, 093 (2017). [https://doi.org/10.1007/JHEP10\(2017\)093](https://doi.org/10.1007/JHEP10(2017)093) arXiv:1705.00982 [hep-ph]
 56. M. Dasgupta, F.A. Dreyer, K. Hamilton, P.F. Monni, G.P. Salam, Logarithmic accuracy of parton showers: a fixed-order study. *JHEP* **09**, 033 (2018). [https://doi.org/10.1007/JHEP09\(2018\)033](https://doi.org/10.1007/JHEP09(2018)033) arXiv:1805.09327 [hep-ph]
 57. G. Bewick, S. Ferrario Ravasio, P. Richardson, M.H. Seymour, Logarithmic accuracy of angular-ordered parton showers. *JHEP* **04**, 019 (2020). [https://doi.org/10.1007/JHEP04\(2020\)019](https://doi.org/10.1007/JHEP04(2020)019) arXiv:1904.11866 [hep-ph]
 58. M. Dasgupta, F.A. Dreyer, K. Hamilton, P.F. Monni, G.P. Salam, G. Soyez, Parton showers beyond leading logarithmic accuracy. *Phys. Rev. Lett.* **125**(5), 052002 (2020). <https://doi.org/10.1103/PhysRevLett.125.052002> arXiv:2002.11114 [hep-ph]
 59. J.R. Forshaw, J. Holguin, S. Plätzer, Building a consistent parton shower. *JHEP* **09**, 014 (2020). [https://doi.org/10.1007/JHEP09\(2020\)014](https://doi.org/10.1007/JHEP09(2020)014) arXiv:2003.06400 [hep-ph]
 60. Z. Nagy, D.E. Soper, Summations of large logarithms by parton showers. *Phys. Rev. D* **104**(5), 054049 (2021). <https://doi.org/10.1103/PhysRevD.104.054049> arXiv:2011.04773 [hep-ph]
 61. C. Tan, F. Sun, T. Kong, W. Zhang, C. Yang, C. Liu, A survey on deep transfer learning, in *Artificial Neural Networks and Machine Learning—ICANN 2018*, ed. by V. Kůrková, Y. Manolopoulos, B. Hammer, L. Iliadis, I. Maglogiannis (Springer International Publishing, Cham, 2018), pp. 270–279
 62. K. Weiss, T. Khoshgoftaar, D. Wang, A survey of transfer learning. *J. Big Data* **3** (2016). <https://doi.org/10.1186/s40537-016-0043-6>
 63. Y.L. Dokshitzer, G. Leder, S. Moretti, B. Webber, Better jet clustering algorithms. *JHEP* **9708**, 001 (1997). <https://doi.org/10.1088/1126-6708/1997/08/001> arXiv:hep-ph/9707323
 64. M. Wobisch, T. Wengler, Hadronization corrections to jet cross-sections in deep inelastic scattering. arXiv:hep-ph/9907280
 65. A. Lifson, G.P. Salam, G. Soyez, Calculating the primary Lund Jet Plane density. *JHEP* **10**, 170 (2020). [https://doi.org/10.1007/JHEP10\(2020\)170](https://doi.org/10.1007/JHEP10(2020)170) arXiv:2007.06578 [hep-ph]
 66. ATLAS Collaboration, G. Aad et al., Measurement of the Lund jet plane using charged particles in 13 TeV proton–proton collisions with the ATLAS detector. *Phys. Rev. Lett.* **124**(22), 222002 (2020). <https://doi.org/10.1103/PhysRevLett.124.222002> arXiv:2004.03540 [hep-ex]
 67. F.A. Dreyer, H. Qu, Jet tagging in the Lund plane with graph networks. *JHEP* **03**, 052 (2021). [https://doi.org/10.1007/JHEP03\(2021\)052](https://doi.org/10.1007/JHEP03(2021)052) arXiv:2012.08526 [hep-ph]
 68. Y. Wang, Y. Sun, Z. Liu, S.E. Sarma, M.M. Bronstein, J.M. Solomon, Dynamic graph CNN for learning on point clouds. *ACM Trans. Graph.* **38**, 146 (2019). <https://doi.org/10.1145/3326362>
 69. S. Ioffe, C. Szegedy, Batch normalization: accelerating deep network training by reducing internal covariate shift, in *Proceedings of the 32nd International Conference on Machine Learning*, vol. 37. PMLR, Lille, France, 07–09 July 2015 (2015), pp. 448–456. <http://proceedings.mlr.press/v37/ioffe15.html>
 70. X. Glorot, A. Bordes, Y. Bengio, Deep sparse rectifier neural networks, in *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, vol. 15. PMLR, Fort Lauderdale, FL, USA, 11–13 Apr 2011 (2011), pp. 315–323. <http://proceedings.mlr.press/v15/glorot11a.html>
 71. K. He, X. Zhang, S. Ren, J. Sun, Deep residual learning for image recognition, in *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (IEEE, Las Vegas, 2016), pp. 770–778. <https://doi.org/10.1109/CVPR.2016.90>
 72. M. Wang, D. Zheng, Z. Ye, Q. Gan, M. Li, X. Song, J. Zhou, C. Ma, L. Yu, Y. Gai, T. Xiao, T. He, G. Karypis, J. Li, Z. Zhang, Deep graph library: a graph-centric, highly-performant package for graph neural networks (2020). arXiv:1909.01315 [cs.LG]
 73. A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, S. Chintala, Pytorch: an imperative style, high-performance deep learning library, in *Advances in Neural Information Processing Systems*, vol. 32, ed. by H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, R. Garnett (Curran Associates, Inc., Red Hook, 2019), pp. 8024–8035
 74. D.P. Kingma, J. Ba, Adam: a method for stochastic optimization (2014). CoRR. arXiv:1412.6980
 75. S.J. Pan, Q. Yang, A survey on transfer learning. *IEEE Trans. Knowl. Data Eng.* **22**(10), 1345–1359 (2010). <https://doi.org/10.1109/TKDE.2009.191>
 76. F. Zhuang, Z. Qi, K. Duan, D. Xi, Y. Zhu, H. Zhu, H. Xiong, Q. He, A comprehensive survey on transfer learning (2019). CoRR. arXiv:1911.02685
 77. H. Daumé III, Frustratingly easy domain adaptation (2009). CoRR. arXiv:0907.1815
 78. F. Li, S. J. Pan, O. Jin, Q. Yang, X. Zhu, Cross-domain co-extraction of sentiment and topic lexicons, in *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, July 2012 (Association for Computational Linguistics, Jeju Island, 2012), pp. 410–419. <https://aclanthology.org/P12-1043>
 79. J. Yosinski, J. Clune, Y. Bengio, H. Lipson, How transferable are features in deep neural networks? (2014). CoRR. arXiv:1411.1792
 80. R.B. Girshick, J. Donahue, T. Darrell, J. Malik, Rich feature hierarchies for accurate object detection and semantic segmentation (2013). CoRR. arXiv:1311.2524
 81. J. Donahue, Y. Jia, O. Vinyals, J. Hoffman, N. Zhang, E. Tzeng, T. Darrell, Decaf: a deep convolutional activation feature for generic visual recognition, in *Proceedings of the 31st International Conference on International Conference on Machine Learning—Volume 32, ICML'14*, pp. I–647–I–655. JMLR.org (2014)
 82. H.-W. Ng, V.D. Nguyen, V. Vonikakis, S. Winkler, Deep learning for emotion recognition on small datasets using transfer learning, in *Proceedings of the 2015 ACM on International Conference on Multimodal Interaction, ICMI '15* (Association for Comput-

- ing Machinery, New York, 2015), pp. 443–449. <https://doi.org/10.1145/2818346.2830593>
83. W. Hu, B. Liu, J. Gomes, M. Zitnik, P. Liang, V.S. Pande, J. Leskovec, Pre-training graph neural networks (2019). CoRR. [arXiv:1905.12265](https://arxiv.org/abs/1905.12265)
84. J. Lee, H. Kim, J. Lee, S. Yoon, Transfer learning for deep learning on graph-structured data, in *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence, AAAI'17* (AAAI Press, 2017), pp. 2154–2160
85. T. Sjöstrand, S. Ask, J. Christiansen, R. Corke, N. Desai, P. Ilten, S. Mrenna, S. Prestel, C. Rasmussen, P. Skands, An introduction to PYTHIA 8.2. *Comput. Phys. Commun.* **191**, 159–177 (2015). <https://doi.org/10.1016/j.cpc.2015.01.024>
86. M. Cacciari, G.P. Salam, G. Soyez, The anti- k_t jet clustering algorithm. *J. High Energy Phys.* **2008**, 063–063 (2008). <https://doi.org/10.1088/1126-6708/2008/04/063>
87. G. Kasieczka, S. Marzani, G. Soyez, G. Stagnitto, Towards machine learning analytics for jet substructure. *JHEP* **09**, 195 (2020). [https://doi.org/10.1007/JHEP09\(2020\)195](https://doi.org/10.1007/JHEP09(2020)195). [arXiv:2007.04319](https://arxiv.org/abs/2007.04319) [hep-ph]
88. S. Carrazza, F.A. Dreyer, JetsGame/data v1.0.0 (2019). <https://doi.org/10.5281/zenodo.2602515>