

Quarks and gluons in the Lund plane

Frédéric A. Dreyer,^a Gregory Soyez^b and Adam Takacs^c

^a*Rudolf Peierls Centre for Theoretical Physics,
Parks Road, Oxford OX1 3PU, U.K.*

^b*IPhT, Université Paris-Saclay, CNRS UMR 3681, CEA Saclay,
Orme des Merisiers, Bât 774, F-91191 Gif-sur-Yvette, France*

^c*Department of Physics and Technology, University of Bergen,
Allégaten 55, 5007 Bergen, Norway*

E-mail: frederic.dreyer@physics.ox.ac.uk, gregory.soyez@ipht.fr,
adam.takacs@uib.no

ABSTRACT: Discriminating quark and gluon jets is a long-standing topic in collider phenomenology. In this paper, we address this question using the Lund jet plane substructure technique introduced in recent years. We present two complementary approaches: one where the quark/gluon likelihood ratio is computed analytically, to single-logarithmic accuracy, in perturbative QCD, and one where the Lund declusterings are used to train a neural network. For both approaches, we either consider only the primary Lund plane or the full clustering tree. The analytic and machine-learning discriminants are shown to be equivalent on a toy event sample resumming exactly leading collinear single logarithms, where the analytic calculation corresponds to the exact likelihood ratio. On a full Monte Carlo event sample, both approaches show a good discriminating power, with the machine-learning models usually being superior. We carry out a study in the asymptotic limit of large logarithm, allowing us to gain confidence that this superior performance comes from effects that are subleading in our analytic approach. We then compare our approach to other quark-gluon discriminants in the literature. Finally, we study the resilience of our quark-gluon discriminants against the details of the event sample and observe that the analytic and machine-learning approaches show similar behaviour.

KEYWORDS: Jets and Jet Substructure, Resummation

ARXIV EPRINT: [2112.09140](https://arxiv.org/abs/2112.09140)

Contents

1	Introduction	1
2	Lund plane(s) and baseline discriminants	3
3	Analytic discriminants in the Lund-plane with strong angular ordering	5
3.1	Generic considerations	5
3.2	Optimal discriminant for primary Lund declusterings	7
3.3	Extension to the full clustering tree	9
4	Beyond strong angular ordering: including clustering logarithms	10
4.1	Generic considerations	10
4.2	Clustering logarithm with the full Lund tree	11
4.3	Clustering logarithms with primary radiation only	14
5	Machine learning approaches	15
5.1	Primary Lund plane and LSTM	15
5.2	Full Lund tree and Lund-Net	16
6	Validation in a pure-collinear (toy) parton shower	16
7	Full Monte Carlo simulations	20
7.1	Setup	20
7.2	Tagging performance	21
7.3	Resilience	24
7.4	Comparison with other approaches	27
7.5	Effect of clustering logarithms and of azimuthal angles	30
7.6	Asymptotic single-logarithmic limit	31
8	Conclusions	33
A	Sudakov factors with exact splitting functions	35

1 Introduction

Quarks and gluons, the constituents of the proton, are fundamental entities of essential relevance for physics at the Large Hadron Collider (LHC) at CERN. While these particles are ubiquitous at hadron colliders, they are never observed directly, but rather fragment and hadronise immediately into collimated sprays of colourless hadrons. These decay products

are referred to as jets and are generally defined through the application of a sequential recombination algorithm (see e.g. [1, 2]).

As most of the jets observed at collider experiments arise from the fragmentation of a light parton, a detailed understanding of their properties is crucial for experimental analyses. In this context, many experimental studies make use of tools that reliably identify the flavour of jets, e.g. to enhance signals from new physics (decaying predominantly to quarks) from QCD backgrounds (producing predominantly gluon jets). Since quark and gluon branch into one another, it is highly non-trivial to even define what is meant by a “quark jet” or a “gluon jet” (see, for example, the discussion in ref. [3], as well as refs. [4, 5]). As a direct consequence, it is delicate to introduce a properly-defined flavoured jet algorithm [6].

Over the past decade, jet substructure, the study of the internal dynamics of jets, has proven a useful approach to study the decay of heavy particles at and above the electroweak scale, providing a promising avenue to search for signs of new physics beyond the Standard Model (see [2, 7, 8] for recent reviews). While jet substructure has applications in many directions including for example precision measurements in QCD and the study of the quark-gluon plasma produced in heavy-ion collisions, recent years have seen an increasing interest in leveraging progress in deep learning to a range of jet tagging problems [9–18].

Several jet substructure techniques have been introduced to address the question of quark/gluon discrimination. This includes jet-shape based observables like jet angularities [19, 20], energy-energy correlation functions [21] or the jet charge [22–24], counting observables like the charged track multiplicity or the Iterative Soft Drop multiplicity [25], as well as a series of recent deep-learning-based approaches using a range of network architectures and inputs [26–30]. Other techniques, such as jet topics [4, 31, 32], are based on a statistical ensemble of events and are directly meant to obtain separate distributions for quarks and gluons. These are not discussed here as we instead target quark/gluon discriminants working on individual jets.

Recently, the Lund Jet Plane has been introduced [33] as a powerful technique to tackle a wide range of jet substructure applications. For example, the primary Lund plane density has been measured by the ATLAS [34] and ALICE [35] collaborations highlighting, for example, differences between general-purpose Monte-Carlo event generators. This Lund plane density is amenable to precision calculations in perturbative QCD [36], showing an agreement with the ATLAS measurement. Finally, Lund-plane variables can be used as inputs to machine-learning tagger [30, 33].

In this paper, we will use the Lund plane approach to study quark/gluon discrimination. We will do this using both an analytic approach and machine-learning tagging methods. In both cases, we will build two taggers: one based on information from primary Lund declusterings only, and a second based on the full Lund declustering tree. Our analytic approach is based on a resummed calculation of the likelihood ratio at the single-logarithmic accuracy, i.e. matching the logarithmic accuracy obtained in ref. [36] for the primary Lund plane density. We note that likelihood ratios have already been relied upon in the context of boosted-jet discrimination, for example, shower deconstruction [37–39]. Our machine-learning taggers follow the guidelines from refs. [30, 33]. One of the main novelties of this work is that we will aim to gain a first-principles understanding of the

behaviour of the neural network by comparing it with our analytic discriminants in specific limits where the analytic approach is known to be optimal. This provides cross-validation of both approaches, shedding light on the importance of subleading effects in the analytic tagger and providing information on the convergence of deep-learning methods.

The paper is organised as follows. In section 2 we describe the Lund plane which will serve as framework for this study. We describe the analytic strongly angular-ordered Lund-plane discriminant in section 3, both for the primary Lund plane and for the full clustering tree. We discuss the inclusion of clustering logarithms in section 4. The baseline machine learning models used in our comparisons are described in section 5, and we provide a validation of the analytic discriminants against these models using a toy shower in section 6. Finally, we perform an in-depth comparison of the performance and resilience of a wide range of methods on full Monte Carlo simulations in section 7, showing that our approaches are either better or on par with state-of-the-art methods both in terms of discriminating power and in terms of resilience.

The code implementing our analytic quark-gluon discriminant based on Lund declusterings is available at <https://gitlab.com/gsoyez/lund-quark-gluon>.

2 Lund plane(s) and baseline discriminants

In order to fix once and for all the notations to be used throughout this paper, we briefly remind the reader of how the primary Lund plane declusterings are constructed [33]. We also introduce a generalisation beyond the primary plane that instead keeps the full *declustering tree* that we exploit later in this paper.

Primary Lund declusterings. For a given jet, we first recluster its constituents with the Cambridge/Aachen algorithm [40, 41].¹ We then build the list of primary declusterings as follows:

1. start with j being the full reclustered jet;
2. undo the last step of the clustering, $j \rightarrow j_1 + j_2$, giving two subjects j_1 and j_2 . We assume without loss of generality that j_1 is the “harder branch” i.e. that $p_{t1} > p_{t2}$.
3. We define the set of coordinates $\mathcal{T} = \{\Delta, k_t, z, \psi, \dots\}$ for this branching:

$$\begin{aligned} \Delta &= \sqrt{(y_1 - y_2)^2 + (\phi_1 - \phi_2)^2}, & k_t &= p_{t2}\Delta, \\ z &= \frac{p_{t2}}{p_{t1} + p_{t2}}, & \psi &= \tan^{-1} \frac{y_2 - y_1}{\phi_2 - \phi_1}. \end{aligned} \tag{2.1}$$

4. Iterate by going back to step 2 with $j \leftarrow j_1$.

This produces a tuple, ordered from the first declustering to the last,

$$\mathcal{L}_{\text{primary}} = [\mathcal{T}_1, \dots, \mathcal{T}_i, \dots, \mathcal{T}_n] \tag{2.2}$$

that we refer to as the primary Lund declusterings associated with the jet j .

¹The reason why we use the Cambridge/Aachen algorithm instead of other algorithms of the generalised- k_t family [42] algorithm, is discussed in section 2.4 of ref. [33].

The averaged primary Lund plane density is then simply defined as the average number of declusterings for a given $\ln \Delta$ and $\ln k_t$:

$$\rho(\Delta, k_t) = \frac{1}{N_{\text{jets}}} \frac{d^2 N}{d \ln \Delta d \ln k_t}. \quad (2.3)$$

This quantity has been measured by the ATLAS collaboration in ref. [34] (see also ref. [43] for a preliminary ALICE measurement) and studied analytically in ref. [36].

Lund declustering tree. Instead of focusing only on following the hardest branch through the declustering, one can retain the full Cambridge/Aachen tree structure yielding an associated tree of Lund variables:

$$\mathcal{L}_{\text{tree}}(j) = [(\mathcal{T}(j), \mathcal{L}_{\text{tree}}(j_1), \mathcal{L}_{\text{tree}}(j_2))] \quad (2.4)$$

where the tree, $\mathcal{L}_{\text{tree}}(j)$, associated with a jet j has a set $\mathcal{T}(j)$ of Lund coordinates associated with the branching $j \rightarrow j_1 + j_2$, with $p_{t1} > p_{t2}$, as well as sub-trees $\mathcal{L}_{\text{tree}}(j_1)$ and $\mathcal{L}_{\text{tree}}(j_2)$ associated with j_1 and j_2 respectively. Note that this structure can be flattened into a tuple

$$\mathcal{L}_{\text{tree}}(j) = [(\mathcal{T}_1, i_{\text{hard},1}, i_{\text{soft},1}), \dots, (\mathcal{T}_i, i_{\text{hard},i}, i_{\text{soft},i}), \dots, (\mathcal{T}_n, i_{\text{hard},n}, i_{\text{soft},n})], \quad (2.5)$$

where $i_{\text{hard},i}$ (resp. $i_{\text{soft},i}$) indicate the index in the tuple for the next branching along the hard (resp. soft) branch, or 0 in the lack thereof.

Iterated Soft Drop multiplicity. The iterated Soft Drop multiplicity [25] can be straightforwardly defined from the tuple of primary Lund declusterings as the number of declusterings satisfying a given Soft Drop [44, 45] condition $z_i > z_{\text{cut}}(\Delta_i/R)^\beta$ with R the jet radius and z_{cut} and β the Soft Drop parameters, with $\beta < 0$. A standard choice is to take $\beta = -1$ so as to effectively use Soft Drop to impose a cut on $z\Delta$ which is similar to a k_t scale. In this paper, we define the iterated Soft Drop multiplicity as the number of primary declusterings above a fixed k_t cut. This is motivated by the fact that a dimensionful k_t cut is more adequate than a cut on $z\Delta$ to separate between a perturbative region ($k_t > k_{t,\text{cut}}$) and a non-perturbative region ($k_t < k_{t,\text{cut}}$).² The use of a k_t cut will also be used with all the other methods introduced in this paper, allowing for a direct comparison of their performance.

Analytically, one can show that the Iterated Soft Drop multiplicity is the optimal quark-gluon discriminant in the double-logarithmic approximation (see below for a proof).

Baseline discriminant: the average Lund-plane density. The baseline approach we will consider throughout this paper is the one that was introduced in ref. [33]. We first compute the average primary Lund plane densities $\rho_{q,g}(\Delta, k_t)$ separately for the quarks and gluons samples, respectively. For a given jet with Lund declusterings $\{(\Delta_i, k_{t,i})\}$, we then define a likelihood ratio

$$\mathbb{L}_{\text{density}} = \prod_i \frac{\rho_g(\Delta_i, k_{t,i})}{\rho_q(\Delta_i, k_{t,i})}. \quad (2.6)$$

In practice, the average densities $\rho_{q,g}(\Delta, k_t)$ are computed in bins of $\ln \Delta$ and $\ln k_t$.

²The fundamental physics motivation behind this choice is that, in a resummed calculation in perturbative QCD, the scale entering the strong coupling is typically the k_t of the emission.

This approach has already shown to give good results when applied to discriminating boosted W bosons from QCD jets. The application to quark-gluon discrimination considered in this paper is actually simpler as it does not involve a separate isolation and treatment of a hard two-prong decay as was the case in [33] for boosted W tagging.

Eq. (2.6) would be an optimal discriminant if the Lund plane declusterings were independent. We know that this is not the case in practice due to effects such as the energy lost by the leading parton, flavour changes and clustering effects, as discussed in [36]. However, we still expect this approach to yield a better performance than the Iterated Soft Drop multiplicity as it captures some aspects of the quark and gluon radiation patterns beyond the soft-collinear approximation. Throughout this paper, we will consider quark-gluon tagging using the average Lund plane density as a reference approach from which we want to build more performant discriminants.

3 Analytic discriminants in the Lund-plane with strong angular ordering

3.1 Generic considerations

In this section, we introduce a series of quark-gluon discriminants based on a first-principles treatment of the Lund plane declusterings in perturbative QCD. The performance of these new tools will be assessed later in sections 6 and 7.

The core idea is to explicitly compute the likelihood ratio

$$\mathbb{L} = \frac{p_g(\mathcal{L})}{p_q(\mathcal{L})}, \tag{3.1}$$

for a set of Lund declusterings \mathcal{L} — either primary-only or including the full tree — where $p_{q,g}(\mathcal{L})$ denotes the probability to observe the given set of Lund declusterings assuming the jet is either initiated by a quark or by a gluon.³ For this procedure to be infrared-and-collinear safe, we only consider emissions above a given (relative) transverse momentum cut, i.e. require $k_t \geq k_{t,\text{cut}}$.

In the (double-logarithmic) soft-collinear limit, emissions are independent and the single-emission probability for quarks and gluons only differ by the overall colour factor (C_F for a quark, C_A for a gluon). For n primary emissions \mathcal{T}_i , one therefore has

$$\mathbb{L}_{LL} = \prod_i \frac{p_g(\mathcal{T}_i)}{p_q(\mathcal{T}_i)} = \left(\frac{C_A}{C_F}\right)^n. \tag{3.2}$$

This shows that the likelihood ratio is only a (monotonic) function of n and hence that the iterated Soft Drop multiplicity is the optimal discriminant at leading (double) logarithmic accuracy. In this limit, additional, non-primary, declusterings in the full Lund tree all come with a factor C_A and therefore do not contribute to eq. (3.2).

In what follows, we want to extend this result to single-logarithmic accuracy, as what was done in [36] for the average primary Lund plane density. For this, several single-logarithmic effects have to be taken into account: (i) corrections to the running of the

³Throughout this paper, we do not distinguish between quarks and anti-quarks.

strong coupling, (ii) collinear effects stemming either from splittings where the flavour of the leading parton changes or from finite z splittings, (iii) clustering effects where the exact Cambridge/Aachen clustering has to be taken into account for multiple soft emissions at commensurate angles, and (iv) effects of multiple soft emissions at large angles.

In the (soft-collinear) limit where emissions are independent, running coupling effects — (i) in the above list — do not change the double-logarithmic result in eq. (3.2) as both the quark and gluon probabilities are multiplied by α_s taken at the same scale, namely the (relative) k_t of the emission. Running-coupling corrections will nevertheless appear together with the other single-logarithmic contributions and this is discussed below.

Next, the effect of soft-wide-angle radiation — (iv) in the above list — depends on the details of the hard process that underlines the samples of quark and gluon jets. These contributions would, for example, be different in “quark jets” in $qg \rightarrow Zq$ and $qq \rightarrow qq$ events (see e.g. the discussion in refs. [3, 46]). In this paper, we will focus on universal aspects in the collinear limit and therefore neglect these contributions which scale like the square of the jet radius. (For phenomenological applications in section 7, we will therefore use $R = 0.4$.) Beyond the small-radius limit, soft-wide-angle effects would have to be included. In this case, one should consider the approach where one has to discriminate a specific “quark-enriched” signal process from a specific “gluon-enriched” background process. This study goes beyond the scope of this paper. It would nevertheless be interesting, in a follow-up study, to investigate if the analytic techniques developed in this paper could be used to assess the process-dependence of quark-gluon tagging, potentially in combination with the concept of jet topics [31].

In a similar spirit, fixed-order (e.g. NLO) corrections to the underlying hard process would also be process-specific and, as such, fall beyond the scope of this paper. That said, if we were to address a process-specific tagging using our analytic techniques, NLO effects would become relevant in two cases. Firstly, in the context of our resummed approach, they would start contributing one order beyond our single-logarithmic approach. Secondly, exact fixed-order corrections would impact the region of large quark and gluon tagging efficiencies (see e.g. the discussion in ref. [5]).

Clustering effects are delicate to handle in an analytic calculation as even in the large- N_c limit they, in principle, require the full matrix angular dependence for an arbitrary number of emissions strongly-ordered in energy. Since we can expect that collinear effects, and flavour-changing contributions in particular, are numerically dominant in the context of quark-gluon discrimination, we will as a first step neglect clustering effects. In other words, in this section we work in the regime where emissions are strongly ordered in angle and derive a quark-gluon discriminant either using only primary emissions, section 3.2, or using the full clustering tree, section 3.3.

We come back to the question of clustering logarithms in section 4. We will see explicitly in our Monte Carlo simulations in section 7 that clustering effects have a smaller numerical impact on quark-gluon discrimination than the collinear enhancements.

3.2 Optimal discriminant for primary Lund declusterings

We start by considering only the primary Lund plane declusterings $\{(\Delta_i, k_{t,i})\}$ with $k_t \geq k_{t,\text{cut}}$. For these, we want to compute the likelihood ratio

$$\mathbb{L}_{\text{primary}} = \frac{p_g(\{\Delta_i, k_{t,i}, z_i, \dots\})}{p_q(\{\Delta_i, k_{t,i}, z_i, \dots\})}, \quad (3.3)$$

at single logarithmic accuracy in perturbative QCD, in the limit where the emissions are strongly ordered in angle, i.e. that $\Delta_1 \gg \Delta_2 \gg \dots \gg \Delta_n$. In this limit, we should include in eq. (3.3) the contributions associated either with the running of the strong QCD coupling effects, or with any hard-collinear effect.

The quark and gluon probability distributions can be computed iteratively starting from the first (largest-angle) splitting. A key point to take into account is the fact that collinear branchings can change the flavour of the leading branch, either through a $g \rightarrow q\bar{q}$ splitting, or through a $q \rightarrow qg$ splitting where the emitted gluon is harder than the final quark. At each splitting, we should therefore keep track of the flavour of the leading parton as well as of its splitting fraction z and its relative transverse momentum k_t . It is convenient to introduce a matrix

$$p_{ab}^{(i)} \equiv \begin{pmatrix} p(q_i|q_0) & p(q_i|g_0) \\ p(g_i|q_0) & p(g_i|g_0) \end{pmatrix}, \quad (3.4)$$

where $p(b_i|a_0)$ denotes the probability that the harder branch has flavour b after the i^{th} declustering, given that it started (at step “0”) with a jet of flavour a . This matrix is initialised as $p_{ab}^{(0)} = \delta_{ab}$ and is recursively constructed from step $i - 1$ to step i for each of the $i = 1, \dots, n$ Lund declusterings.

Assuming that just before branching i the jet has flavour a , the probability after branching i should include two effects: (i) the probability that the splitting has the observed kinematic properties Δ_i, z_i, \dots , potentially including a change of the leading flavour, and (ii) a Sudakov factor implementing the fact that no emission has occurred between the previous angle Δ_{i-1} and Δ_i (with $\Delta_0 \equiv R$), and with $k_t > k_{t,\text{cut}}$. This Sudakov resums the virtual corrections between Δ_{i-1} and Δ_i . This leads to the recursion

$$p_{ab}^{(i)} = \frac{\alpha_s(k_{ti})}{\pi\Delta_i} \begin{pmatrix} \tilde{P}_{qq}(z_i) & \tilde{P}_{qg}(z_i) \\ \tilde{P}_{gq}(z_i) & \tilde{P}_{gg}(z_i) \end{pmatrix} \begin{pmatrix} S_q^{(i-1,1)} & 0 \\ 0 & S_g^{(i-1,1)} \end{pmatrix} p_{ab}^{(i-1)}. \quad (3.5)$$

In this expression, the splitting kernels \tilde{P}_{ab} are directly related to the Altarelli-Parisi splitting functions with the extra requirement that since the declustering procedure follows the

hardest branch one should impose $z_i < \frac{1}{2}$:⁴

$$\tilde{P}_{qq} = P_{qq}(z)\Theta\left(z < \frac{1}{2}\right) = C_F \frac{1+(1-z)^2}{z}\Theta\left(z < \frac{1}{2}\right), \quad (3.6a)$$

$$\tilde{P}_{gq} = P_{qq}(z)\Theta\left(z < \frac{1}{2}\right) = C_F \frac{1+z^2}{1-z}\Theta\left(z < \frac{1}{2}\right), \quad (3.6b)$$

$$\tilde{P}_{qg} = [P_{qg}(z) + P_{qg}(1-z)]\Theta\left(z < \frac{1}{2}\right) = 2n_f T_R [z^2 + (1-z)^2]\Theta\left(z < \frac{1}{2}\right), \quad (3.6c)$$

$$\tilde{P}_{gg} = [P_{gg}(z) + P_{gg}(1-z)]\Theta\left(z < \frac{1}{2}\right) = 2C_A \left[\frac{1-z}{z} + \frac{z}{1-z} + z(1-z)\right]\Theta\left(z < \frac{1}{2}\right). \quad (3.6d)$$

The Sudakov factors, $S_{q,g}^{(i-1,i)}$, between the angle of the last splitting Δ_{i-1} and the angle of the current splitting Δ_i is computed as

$$S_f^{(i-1,i)} = \exp\left[-\int_{\Delta_i}^{\Delta_{i-1}} \frac{d\Delta}{\Delta} \int dz \frac{\alpha_s(p_{ti}z\Delta)}{\pi} P_f(z_i)\Theta(p_{ti}z\Delta > k_{t,\text{cut}})\right], \quad (3.7)$$

with P_f the total splitting function for a parton of flavour f and p_{ti} the transverse momentum (with respect to the beam) of parton i before splitting. The k_t of the emission is taken as $p_{ti}z\Delta$ which is equivalent to our definition in eq. (2.1) in the collinear limit.⁵ This Sudakov is evaluated at next-to-leading logarithmic (NLL) accuracy with $\Delta_i \ll \Delta_{i-1}$, and we find

$$\begin{aligned} S_f^{(i-1,i)} = \exp\left\{ -\frac{C_f}{2\pi\alpha_s\beta_0^2} \left[(1-\lambda_{i-1}) \ln \frac{1-\lambda_{i-1}}{1-\lambda_{\text{cut}}} - (1-\lambda_i) \ln \frac{1-\lambda_i}{1-\lambda_{\text{cut}}} - \lambda_i + \lambda_{i-1} \right. \right. \\ \left. \left. - \frac{\alpha_s\beta_1}{\beta_0} \left(\frac{1}{2} \ln^2(1-\lambda_i) - \frac{1}{2} \ln^2(1-\lambda_{i-1}) + \frac{\lambda_i - \lambda_{i-1}}{1-\lambda_{\text{cut}}} \ln(1-\lambda_{\text{cut}}) \right) \right. \right. \\ \left. \left. + \left(\frac{\alpha_s K}{2\pi} - \frac{\alpha_s\beta_1}{\beta_0} \right) \left(\frac{\lambda_i - \lambda_{i-1}}{1-\lambda_{\text{cut}}} - \ln \frac{1-\lambda_{i-1}}{1-\lambda_i} \right) \right] \right\}, \quad (3.8) \end{aligned}$$

with $\alpha_s \equiv \alpha_s(p_{t,\text{jet}}R)$,

$$\lambda_{i-1} = 2\alpha_s\beta_0 \left(\ln \frac{R}{x\Delta_{i-1}} - B_f \right), \quad \beta_0 = \frac{11C_A - 4n_f T_R}{12\pi}, \quad (3.9a)$$

$$\lambda_i = 2\alpha_s\beta_0 \left(\ln \frac{R}{x\Delta_i} - B_f \right), \quad \beta_1 = \frac{17C_A^2 - 5C_A n_f - 3C_F n_f}{24\pi^2}, \quad (3.9b)$$

$$\lambda_{\text{cut}} = 2\alpha_s\beta_0 \ln \frac{p_{t,\text{jet}}R}{k_{t,\text{cut}}}, \quad K = \left(\frac{67}{18} - \frac{\pi^2}{6} \right) C_A - \frac{5}{9} n_f, \quad (3.9c)$$

$$B_q = -\frac{3}{4} \quad B_g = -\frac{11C_A - 4n_f T_R}{12C_A}, \quad (3.9d)$$

⁴We have chosen notations where the indices of the \tilde{P} kernels represent the flavour of the hard branch, so as to make the matrix product in eq. (3.5) more obvious. As a consequence, these indices do not always match with the standard indices in the Altarelli-Parisi kernels where the indices instead refer to the flavour of the emitted parton with momentum fraction z . Finally, our probability distributions are taken differentially in Δ_i and z_i . The specific choice of variables is however irrelevant for the problem of quark-gluon classification as it cancels in the likelihood ratio.

⁵Conversely, the value of p_{ti} can be deduced from Δ_i , k_{ti} and z_i using $p_{ti} = k_{ti}/(z_i\Delta_i)$.

and x defined as the momentum fraction of the total jet momentum carried by the subject j before branching i . We note that in the rare occurrences where the Lund declusterings are not ordered in angle — which cannot be ruled out with the Cambridge/Aachen declustering procedure — we set the Sudakov to $S_f = 1$. We also point out that the contribution from hard-collinear splitting to the above expressions have been computed by setting an upper bound e^{B_f} on the z integration in (3.7). This is correct at NLL accuracy. Although it has the drawback to insert uncontrolled subleading corrections — compared to the traditional expression which can be recovered by keeping only the first non-trivial term in B_f — it has the advantage of having a clean endpoint, i.e. $S_f = 1$ for $\lambda_{i-1} \leq \lambda_{\text{cut}}$.

If we introduce the short-hand notations $\tilde{P}_{ab}^{(i)}$ and $S_{ab}^{(i,i-1)} \equiv \delta_{ab} S_a^{(i,i-1)}$ for the full splitting matrix and Sudakov matrix, the probabilities after including all the Lund declusterings takes the form

$$p^{(\text{final})} = S^{(n+1,n)} \tilde{P}^{(n)} S^{(n,n-1)} \dots \tilde{P}^{(i)} S^{(i,i-1)} \dots \tilde{P}^{(1)} S^{(1,0)} p^{(0)}, \quad (3.10)$$

where the leftmost factor in the right-hand side takes into account the fact that there are no more emissions between the angle of the last declustering, Δ_n , and the smallest angle accessible after the last splitting: $\Delta_{n+1} \equiv \Delta_{\text{min}} = k_{t,\text{cut}}/p_{tn}$ with p_{tn} the transverse momentum (with respect to the beam) of the leading parton after the last declustering. Eq. (3.10) has the simple physical interpretation of successive primary branchings, producing the factors $\tilde{P}^{(i)}$, interleaved with Sudakov factors, $S^{(i,i-1)}$, which resum virtual corrections between two primary emissions. Finally, the probabilities associated with an initial quark or gluon jet are given by

$$p_q(\{\Delta_i, k_{t,i}, z_i, \dots\}) = p^{(\text{final})}(q|q_0) + p^{(\text{final})}(g|q_0), \quad (3.11a)$$

$$p_g(\{\Delta_i, k_{t,i}, z_i, \dots\}) = p^{(\text{final})}(q|g_0) + p^{(\text{final})}(g|g_0), \quad (3.11b)$$

translating the fact that we are inclusive over all flavours of the final leading parton.

The probabilities in eqs. (3.11) can be directly inserted in (3.3) to obtain a quark-gluon discriminant. It is, by construction, the optimal discriminant at single-logarithmic accuracy in the limit where the declusterings are strongly ordered in angle. Since the above procedure keeps track of the flavour and momentum fraction x of the leading parton at each step, it takes into account the possible correlations between the different declusterings, hence going beyond the independent-emission assumption used with the average Lund plane density (section 2, eq. (2.6)).

3.3 Extension to the full clustering tree

While non-primary (secondary, tertiary, ...) declusterings have no impact at leading-logarithmic accuracy, they start carrying information at our single-logarithmic accuracy. Generalising the approach from the previous section to the full clustering tree is mostly a technical step. This time, we therefore settle to compute

$$\mathbb{L}_{\text{tree}} = \frac{p_g(\mathcal{L}_{\text{tree}})}{p_q(\mathcal{L}_{\text{tree}})} \quad (3.12)$$

in the strongly angular-ordered limit.

This is again done recursively over the full (de-)clustering tree. For this, consider a declustering $j_{\text{parent}} \rightarrow j_{\text{hard}} + j_{\text{soft}}$, with kinematic variables \mathcal{T} , i.e. with an angle Δ , a soft momentum fraction z and relative transverse momentum k_t . The probabilities associated with the parent jet can be deduced from those of the subjects as follows:

$$p_q(\mathcal{L}_{\text{parent}}) = S_q(\Delta_{\text{prev}}, \Delta) \left[\tilde{P}_{qq}(z)p_q(\mathcal{L}_{\text{hard}})p_g(\mathcal{L}_{\text{soft}}) + \tilde{P}_{gq}(z)p_g(\mathcal{L}_{\text{hard}})p_q(\mathcal{L}_{\text{soft}}) \right] \quad (3.13a)$$

$$p_g(\mathcal{L}_{\text{parent}}) = S_g(\Delta_{\text{prev}}, \Delta) \left[\tilde{P}_{gg}(z)p_g(\mathcal{L}_{\text{hard}})p_g(\mathcal{L}_{\text{soft}}) + \tilde{P}_{qg}(z)p_q(\mathcal{L}_{\text{hard}})p_g(\mathcal{L}_{\text{soft}}) \right] \quad (3.13b)$$

where Δ_{prev} is the angle at which the last declustering before the one under consideration happened (with $\Delta_{\text{prev}} = R$ for the first declustering). As in section 3.2, $S(\Delta_{\text{prev}}, \Delta)$ is a Sudakov factor imposing that no other emission with $k_t > k_{t,\text{cut}}$ occurred since the last declustering at an angle Δ_{prev} , cf. eq. (3.8).⁶ The splitting kernels \tilde{P}_{ab} are the same as in eq. (3.6). These expressions have the same form as eq. (3.5) except that, at each step, they also include the probability for the soft branch.

This recursion is applied until each branch can no longer be declustered in which case, if the last splitting has occurred at an angle Δ_{last} , one then just includes a factor

$$p_q(\mathcal{L} = \emptyset) = S_q(\Delta_{\text{last}}, \Delta_{\text{min}}), \quad (3.14a)$$

$$p_g(\mathcal{L} = \emptyset) = S_g(\Delta_{\text{last}}, \Delta_{\text{min}}), \quad (3.14b)$$

where, as for the primary case, $\Delta_{\text{min}} = k_{t,\text{cut}}/p_t$ for a final branch of momentum p_t .

4 Beyond strong angular ordering: including clustering logarithms

4.1 Generic considerations

We conclude this section on analytic methods by discussing the inclusion of clustering logarithms in our approach. These logarithms arise from situations where we have at least two emissions with commensurate angles and the exact Cambridge/Aachen clustering has to be considered in order to label the emissions as primary, secondary, ternary, etc. When the emissions at commensurate angles are strongly ordered in energy, this leads to single-logarithmic contributions (see e.g. [36]).

In practice, the Cambridge/Aachen clustering can produce clusterings which are not in agreement with the naive physical expectation. Consider for example a quark-initiated jet with two gluon emissions. The harder emission is emitted from the quark and comes with a colour factor C_F . The softer emission can either be seen as emitted from the quark, with a colour factor C_F , or as emitted from the gluon, with a colour factor C_A . When the two emission angles are similar, the actual Cambridge/Aachen clustering will sometimes cluster the second gluon in the C_F^2 contribution with the first gluon, yielding a secondary Lund declustering, or, conversely, cluster the second gluon from the $C_F C_A$ contribution with the hard quark, yielding a primary declustering.

⁶Although we have only made explicit the angular dependence, the Sudakov factors also depend on the prong momenta.

In order to compute these contributions, we need the full angular structure of the matrix elements for an arbitrary number of emissions at commensurate angles. When computing the average Lund plane density, this can be addressed, at least in the large- N_c limit, by a Monte Carlo integration similar to the one used to resum non-global logarithms in [47]. (See [48] for an approach valid at finite N_c .)

In our quark-gluon-discrimination application, one would have to keep track of all the possible colour configuration with which an emission can be radiated by the full set of emissions at larger transverse momentum (or, at our accuracy, at larger energies). This is beyond what can be practically achieved. Instead, we will adopt a simplified approach where we apply a matrix-element correction which only describes correctly situations where (any number of) pairs of emissions are at commensurate angles. This is similar in spirit to the NODS scheme introduced in [49] to implement subleading- N_c corrections in parton showers.⁷

4.2 Clustering logarithm with the full Lund tree

Since clustering logarithms have an explicit dependence on radiation in/from different leaves, we first consider the situation where the quark-gluon tagging is done using the full Lund declustering tree. The case where only primary radiation is considered will be discussed in section 4.3 below.

In our approximation where we only allow for two emissions at commensurate angles, we then consider two declusterings $\mathcal{T}_1 \equiv \{\Delta_1, k_{t1}, z_1, \psi_1\}$ and $\mathcal{T}_2 \equiv \{\Delta_2, k_{t2}, z_2, \psi_2\}$, with $\Delta_1 \sim \Delta_2 \ll 1$. Since clustering corrections happen for two emissions at similar angles and we only aim at describing the configurations where we have only pairs of particles at commensurate angles, we can assume that \mathcal{T}_1 and \mathcal{T}_2 correspond to consecutive Lund declusterings and that all the other emissions are at widely different angles. We can further assume that \mathcal{T}_1 happens before \mathcal{T}_2 in the sequence of declusterings, i.e. $\Delta_1 > \Delta_2$. Our approach is to modify the emission probability in eq. (3.12) for \mathcal{T}_2 to include corrections due to the presence of \mathcal{T}_1 .

There are two main kinematic configurations to consider: either \mathcal{T}_1 and \mathcal{T}_2 are both reconstructed as consecutive “primary” emissions from the same hard branch, or \mathcal{T}_2 is reconstructed as a “secondary” emission from \mathcal{T}_1 .⁸ In the “primary” case, we can either have $z_2 \ll z_1$ ($k_{t2} \ll k_{t1}$) or $z_2 \gg z_1$ ($k_{t2} \gg k_{t1}$), while in the “secondary” case we can assume $z_2 \ll z_1$ ($k_{t2} \ll k_{t1}$). This is illustrated by the Lund diagrams in figure 1. At single-logarithmic accuracy, the clustering correction is computed in the flavour channel where both emissions are gluons. The distinction between the primary and secondary cases is decided by the Cambridge/Aachen clustering. In both cases, if C_R is the colour factor of the common hard branch, the matrix element corresponding to a given clustering

⁷Note the key difference that the NODS method produces the correct behaviour at large- N_c for any number of emissions at commensurate angles. The matrix-element correction only applies to subleading- N_c corrections. In our case, the correct behaviour is only guaranteed for pairs of emissions at commensurate angles even in the large- N_c limit.

⁸Primary and secondary are here counted from the hard branch common to both emissions, even if this one can be anywhere in the Lund tree.




Figure 1. Schematic representation of the three kinematic configurations affected by clustering logarithms. We consider two consecutive declusterings \mathcal{T}_1 and \mathcal{T}_2 with \mathcal{T}_2 occurring after \mathcal{T}_1 in the Lund sequence. If \mathcal{T}_2 is reconstructed as “primary” (i.e. in the same plane as \mathcal{T}_1) it can either be much softer (case (a)) or much harder (case (b)) than \mathcal{T}_1 . Case (c) describes the situation where \mathcal{T}_2 is reconstructed as a secondary emission from \mathcal{T}_1 , and hence can be considered much softer.

sequence will have a contribution proportional to C_R^2 and one proportional to $C_R C_A$. In the strongly-angular-ordered limit, only the first term (C_R^2) contributes to the “primary” clustering and only the second term ($C_R C_A$) to the “secondary” clustering. The gluon-emission diagrams shown in figure 1 represent the two contributions for each clustering case. The first particles to cluster are highlighted in blue.

Let us first handle the case where \mathcal{T}_1 and \mathcal{T}_2 are both “primary” emissions. Say the parent parton has a colour factor C_R . If $z_1 \gg z_2$, the $C_R d^2\Delta_2/\Delta_2^2$ behaviour which corresponds to the collinear limit in section 3.3 should be replaced by the full soft-gluon radiation squared matrix element

$$\left[\frac{C_A}{2} \frac{1}{\Delta_{12}^2} + \frac{C_A}{2} \frac{\Delta_1^2}{\Delta_{12}^2 \Delta_2^2} + \left(C_R - \frac{C_A}{2} \right) \frac{1}{\Delta_2^2} \right] d^2\Delta_2, \quad (4.1)$$

where $\Delta_{12}^2 = \Delta_1^2 + \Delta_2^2 - 2\Delta_1\Delta_2 \cos(\psi_2 - \psi_1)$ is the angle between the two emitted gluons. This means that, in the gluon emission part of eq. (3.5), we should apply a correction factor

$$\Omega_{\text{prim}} = 1 + \frac{C_A}{2C_R} \left(\frac{\Delta_2^2}{\Delta_{12}^2} + \frac{\Delta_1^2}{\Delta_{12}^2} - 1 \right). \quad (4.2)$$

It is straightforward to show that the “primary” case with $z_1 \ll z_2$, gives the same correction Ω_{prim} . As expected, $\Omega_{\text{prim}} \rightarrow 1$ when $\Delta_1 \gg \Delta_2$ (or when $\Delta_1 \ll \Delta_2$) so that the strongly-ordered limit is recovered. Since both emissions are primary, we never have $\Delta_{12} \ll \Delta_1, \Delta_2$.

We now turn to the “secondary” case where \mathcal{T}_2 is emitted from the soft branch of \mathcal{T}_1 . Here, Δ_2 is the angle between the emissions \mathcal{T}_1 and \mathcal{T}_2 , i.e. $\Delta_2 \equiv \Delta_{12}$, and we denote by




Figure 2. Clustering correction factors Ω_{prim} (primary) and Ω_{sec} (secondary) for quark (left) and gluon (right) leading partons. The black regions correspond to the kinematic regions where emissions are not clustered in the indicated Lund plane.

$\Delta_{02} = \Delta_1^2 + \Delta_2^2 - 2\Delta_1\Delta_2 \cos(\psi_2 - \psi_1)$ the angle between \mathcal{T}_2 and the hard branch of \mathcal{T}_1 . In the collinear limit this would correspond to a factor C_A/Δ_2^2 which has to be replaced by the full angular structure

$$\frac{C_A}{2} \frac{1}{\Delta_2^2} + \frac{C_A}{2} \frac{\Delta_1^2}{\Delta_{02}^2 \Delta_2^2} + \left(C_R - \frac{C_A}{2} \right) \frac{1}{\Delta_{02}^2}, \tag{4.3}$$

yielding a correction factor

$$\Omega_{\text{sec}} = \frac{1}{2} \left(1 + \frac{\Delta_1^2}{\Delta_{02}^2} + \frac{\Delta_2^2}{\Delta_{02}^2} \right) + \frac{C_R}{C_A} \frac{\Delta_2^2}{\Delta_{02}^2}. \tag{4.4}$$

Without surprise, $\Omega_{\text{sec}} \rightarrow 1$ if $\Delta_2 \ll \Delta_1$, recovering the strongly-ordered case. The correction factors Ω_{prim} and Ω_{sec} are plotted in figure 2. We see that they are indeed localised around the region where both emissions have commensurate angles. They tend to be larger for quarks than for gluons.

It is interesting to notice that the above correction which accounts for clustering logarithms introduces a dependence on the azimuthal angle ψ . It is the only dependence on ψ at the single logarithmic accuracy.

In principle, the Sudakov factors should also receive single-logarithmic corrections due to clustering effects. Since clustering logarithms only affect flavour-diagonal emissions of two soft gluons (at least within our approximations), it is however relatively straightforward to convince oneself that these corrections only lead to a reshuffling of some contributions between different factors in the overall probability distribution and can therefore be neglected. This can be understood as follows. Suppose one has a parton emitting a gluon “1” at an

angle Δ_1 . In our approach, that the previous emission occurred at an angle Δ_0 and that the following emissions on the hard and soft branches occur at angles $\Delta_{2,\text{hard}}$ and $\Delta_{2,\text{soft}}$ respectively. In the strongly-ordered limit, emission “1” is involved in three contributions to the Sudakov: one between Δ_1 and Δ_0 , one between $\Delta_{2,\text{hard}}$ and Δ_1 , and one between $\Delta_{2,\text{soft}}$ and Δ_1 . Within our approximation, if we want to compute the corrections to the Sudakov at angles commensurate with Δ_1 , we can assume that all the other emissions are at widely separate angles, i.e. $\Delta_{2,\text{hard}}, \Delta_{2,\text{soft}} \ll \Delta_1 \ll \Delta_0$. To compute the overall Sudakov factor, summed over the three regions described above, one should integrate the exact matrix element, including the full angular structure. Within our approximation where we only target correctness for two emissions at commensurate angles, this integral is proportional to $C_R \log(\Delta_0/\Delta_{2,\text{hard}}) + C_A \log(\Delta_1/\Delta_{2,\text{soft}})$, which is the same as the strongly-ordered limit.

In practice, inserting correction factors due to clustering logarithms in eq. (3.13) requires some care as it depends whether the Lund declustering, \mathcal{T} , that it implements comes from following the hard or the soft branch at the previous declustering. First, at our single-logarithmic accuracy, clustering corrections are only non-trivial for two successive gluon emissions. Then, say that the previous branching, happening at an angle Δ_{prev} is denoted by $\mathcal{T}_{\text{prev}}$. If \mathcal{T} follows $\mathcal{T}_{\text{prev}}$ along the harder branch, we only apply a correction for the contributions where $\mathcal{T}_{\text{prev}}$ did not have a flavour change. The correction is then applied only for the flavour-diagonal contribution with a colour factor C_R being C_F or C_A depending on the flavour of the hard parton. Conversely, if \mathcal{T} follows $\mathcal{T}_{\text{prev}}$ along the softer branch, only the flavour-diagonal term in $p_g(\mathcal{L}_{\text{parent}})$ receives a correction with a colour factor C_R given by the flavour of the hard branch at the branching $\mathcal{T}_{\text{prev}}$.

It is interesting to note that the correction factors $\bar{\Omega}_{\text{prim}}$ and $\bar{\Omega}_{\text{sec}}$ explicitly depend on the azimuthal angles of the declusterings, which is new compared to the strongly-angular-ordered case. If we want to consider only the Δ , z and k_t variables for each Lund declusterings (see the discussion in section 7 below), we can integrate out the ϕ dependence, averaging over the domain allowed by the fact that the declustering \mathcal{T}_1 is undone before \mathcal{T}_2 . We denote these azimuthally-averaged correction factors by $\bar{\Omega}_{\text{prim}}$ and $\bar{\Omega}_{\text{sec}}$. They only depend on the ratio $x = \Delta_2/\Delta_1$ and are found to be

$$\bar{\Omega}_{\text{prim}} \stackrel{x < 1/2}{=} 1 + \frac{C_A}{C_R} \frac{x^2}{1-x^2} \tag{4.5a}$$

$$\stackrel{x > 1/2}{=} 1 + \frac{C_A}{2C_R} \left[\frac{1+x^2}{|1-x^2|} \frac{1 - \frac{2}{\pi} \tan^{-1} \left(\frac{1+x}{|1-x|} \sqrt{\frac{2x-1}{2x+1}} \right)}{1 - \frac{2}{\pi} \tan^{-1} \left(\sqrt{\frac{2x-1}{2x+1}} \right)} \right] \tag{4.5b}$$

$$\bar{\Omega}_{\text{sec}} \stackrel{x < 1/2}{=} 1 + \frac{C_R}{C_A} \frac{x^2}{1-x^2} \tag{4.5c}$$

$$\stackrel{x > 1/2}{=} \frac{1}{2} + \left[\frac{\text{sgn}(1-x^2)}{2} + \frac{C_R}{C_A} \frac{x^2}{|1-x^2|} \right] \frac{1 - \frac{2}{\pi} \tan^{-1} \left(\frac{1+x}{|1-x|} \sqrt{\frac{2x-1}{2x+1}} \right)}{1 - \frac{2}{\pi} \tan^{-1} \left(\sqrt{\frac{2x-1}{2x+1}} \right)}. \tag{4.5d}$$

4.3 Clustering logarithms with primary radiation only

Our last analytic step is to include the effect of clustering logarithms in the Lund quark-gluon discriminant which only uses primary declusterings. We do this in an approximation where we only allow for pairs of emissions to be at commensurate angles.

As for the full Lund tree, two types of corrections should be included: corrections to the matrix element for the radiation of two soft gluons at commensurate angles, and potential corrections to the Sudakov factor. Corrections to the real radiation are trivial: the splitting factors $\tilde{P}_{qq}(z_i)$ and $\tilde{P}_{gg}(z_i)$ in eq. (3.5) should be multiplied by a factor Ω_{prim} (eq. (4.2)), respectively with $C_R = C_F$ and $C_R = C_A$.

Since we are no longer including a Sudakov factor for the soft branch, our previous argument saying that the overall Sudakov factor was not affected by clustering effects no longer holds. Let us therefore again consider a gluon emission “1” at an angle Δ_1 and relative transverse momentum k_{t1} . In the soft-gluon limit, the total Sudakov factor at (relative) transverse momentum scales smaller than k_{t1} should use the full matrix element for radiation from the system including both the parent parton and emission “1”, i.e.

$$\begin{aligned}
 -\log S_f^{(i-1,i)} &= \int_{k_{t,\text{cut}}}^{k_{t1}} \frac{\alpha_s(k_t)}{\pi^2} \int d^2\Delta_2 \left[\left(C_R - \frac{C_A}{2} \right) \frac{1}{\Delta_2^2} + \frac{C_A}{2} \frac{1}{\Delta_{12}^2} + \frac{C_A}{2} \frac{\Delta_1^2}{\Delta_2^2 \Delta_{12}^2} \right] \\
 &\times [1 - \Theta(\Delta_{12} < \Delta_1) \Theta(\Delta_{12} > \Delta_2)], \tag{4.6}
 \end{aligned}$$

where the square bracket in the second line imposes that “2” is clustered as a primary emission. This gives a correction compared to the strongly-angular-ordered case which is found to be⁹

$$-\delta \log S_f^{(i-1,i)} = (C_A - C_R) \xi \left[\int_{k_{t,\text{cut}}}^{k_{t1}} \frac{\alpha_s(k_t)}{\pi} \right], \tag{4.7}$$

with $\xi = 0.323006$. Note that this contribution happens to vanish when the parent parton is a gluon ($C_R = C_A$).

5 Machine learning approaches

5.1 Primary Lund plane and LSTM

A natural approach to adopt is the Deep-Learning technique used in the original study of the primary Lund Plane [33], which was already showing an excellent discriminating power in the context of boosted W tagging. Here we only consider the long short-term memory (LSTM) [50] network architecture as it showed the best performance in [33].

In practice, we input the list of Lund declusterings $\{(\ln \Delta_i, \ln k_{ti})\}$ to an LSTM of dimension 128 connected to a dropout layer (with rate 20%), with a final dense layer of dimension two and softmax activation. The network weights are initialised with a He uniform variance scaling initialiser [51], and the training is performed using an Adam optimisation algorithm [52] with a batch size of 128, a learning rate of 0.0005 and a categorical cross-entropy loss function. Our model is implemented using TensorFlow v2.1.0.

The data sample is split into a training sample (80%), a validation sample (10%) and a testing sample (10%). We train over a maximum of 50 epochs, with an early stopping when the performance does not increase over four epochs.

⁹This result is essentially the same as the $\mathcal{O}(\alpha_s^2)$ clustering logarithms contribution to the primary Lund plane density found in [36] (see eq. (3.25) there).

For each configuration, we have run five independent trainings. For the quality measures reported below, the central value is obtained by averaging over the five runs and the uncertainty band is taken as their envelope.

5.2 Full Lund tree and Lund-Net

In order to take full advantage of the information contained in secondary leaves of the Lund plane, we consider the Lund-Net model introduced in ref. [30] and its associated code [53].

As input, we transform the tree of Lund declusterings into a graph, with the kinematic variables \mathcal{T} of a declustering serving as attributes of a node on the graph. The Cambridge/Aachen clustering sequence is used to form bidirectional edges along the nodes connected in the Lund tree.

The graph architecture uses an EdgeConv operation [54], which applies a multi-layer perceptron (MLP) to each incoming edge of a node, using combined features of the node pairs as inputs, producing a learned edge feature. This initial shared MLP consists of two layers, each consisting of a dense network with batch normalisation [55] and ReLU activation [56], which are followed by an aggregation step taking an element-wise average of the learned edge features of the incoming edges as well as a shortcut connection [57]. The same MLP is applied to all nodes, leading to updated node features but keeping the structure of the graph unchanged. The Lund-Net architecture consists of six such EdgeConv blocks stacked together, and the number of channels of the MLPs are (32, 32), (32, 32), (64, 64), (64, 64), (128, 128) and (128, 128). Their output is concatenated for each node, 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 final fully connected layer with 256 units and a dropout layer with rate 10%, with a softmax output giving the result of the classification. The Lund-Net model is implemented with the Deep Graph Library 0.5.3 [58] using the PyTorch 1.7.1 [59] backend, and training is performed for 30 epochs, using an Adam optimiser [52] to minimise the cross entropy loss. An initial learning rate of 0.001 is used, which is lowered by a factor 10 after the 10th and 20th epochs. As for the LSTM approach, the data sample is randomly split in 80/10/10% training/validation/testing samples, and we take the average and envelope of five runs.

In this paper, the inputs for each Lund declustering include, by default, $\ln \Delta$, $\ln k_t$, $\ln z$ and ψ . In section 6 which probes the collinear limit of our discriminants, the azimuthal angle is irrelevant and therefore not included in any of our approaches. Furthermore, in section 7.4 we discuss the effect of adding particle-ID information to the inputs, and in section 7.5 we discuss the effect of the azimuthal angle ψ . When imposing a cut on the (relative) transverse momentum, only the Lund declusterings with k_t above the cut are included in the data sample.

6 Validation in a pure-collinear (toy) parton shower

Before turning to a full Monte Carlo-based assessment of the discriminating performance of the tools introduced in the previous sections, we provide a cross-validation between the analytic and deep-learning approaches. To do this, we use a setup in which our analytic

approach in sections 3.2 and 3.3 corresponds to the exact likelihood-ratio discriminant. This is achieved by generating events directly in the strong-angular-ordered limit, where our analytic approach from section 3 becomes exact: for simplicity, we use a fixed-coupling approximation, a fixed initial jet p_t of 1 TeV (with $R=1$)¹⁰ and a fixed cut on emissions $k_{t,\text{cut}} = 1$ GeV. We generate pure-quark and pure-gluon samples by starting the simulation with either an initial hard quark or an initial hard gluon. We simulate branchings using the full Altarelli-Parisi splitting functions, keeping track at each emission of the angle Δ and energy fraction z of the emission. In the strict collinear limit, a parton of momentum p_t branches in two partons of momenta $(1-z)p_t$ and zp_t , so the transverse momentum of each parton in the cascade — or, equivalently, its fraction of the initial jet p_t — can be deduced from the angles (Δ_i) and momentum fractions (z_i) at each branching. In practice, we have used (a slightly adapted version of)¹¹ the `microjet` code [60] to simulate events with strong angular ordering. The Lund declusterings are taken directly from the event trees, without any reclustering with the Cambridge/Aachen jet algorithm. This guarantees the absence of clustering logarithms.

Our analytic approaches to quark-gluon discrimination are applied as described in sections 3.2 and 3.3 except for two details: (i) they have been adapted to use a fixed-coupling approximation, and (ii) the Sudakov factor in eq. (3.7) has been computed keeping the full splitting function so as to guarantee that the resulting probability distributions match exactly that of the generated sample, including corrections strictly beyond our single-logarithmic approximation. With a fixed-coupling approximation, the calculation of the Sudakov exponent is relatively straightforward and expressions are given in appendix A for completeness.

With this setup in mind, we want to check that the machine learning (ML) approach using an LSTM network trained on primary Lund declusterings (section 5.1) converges to the same performance as what is given by the analytic approach in section 3.2. Similarly, we expect that the Lund-Net approach from section 5.2, trained on full Lund trees, gives the same performance as that of the analytic discriminant based on the full Lund tree in section 3.3. We also want to check that these new tools offer a better discriminating power than what is obtained using either the Iterated Soft Drop multiplicity or the average primary Lund plane density (see section 2).

In practice, we use a sample of 10^6 events generated with our adapted version of the `microjet` code, with α_s fixed either to 0.1 or to 0.5. These samples are either used to compute the analytic discriminant,¹² eqs. (3.3) or (3.12) or as inputs to train/validate/test our neural-network-based models. For the methods using machine-learning, the event sample is split in 80/10/10% training/validation/testing samples. This is repeated five

¹⁰Since we work in the pure collinear limit, the jet radius just plays the role of a reference scale for the logarithms of the angles, i.e. large-angle corrections scaling like powers of R are absent. The precise value of R is therefore irrelevant and we chose $R = 1$ for simplicity.

¹¹Our adaptation compared to the original work in [60] mostly consists in imposing a k_t cutoff (instead of a small cut on z , as well as to keep the full tree of the generated cascade rather than just the final particles.

¹²For $\alpha_s = 0.1$, our analytic results have been obtained with a sample size of 10^7 events instead of the default sample of 10^6 events. This shows no visually observable differences on the results presented here.

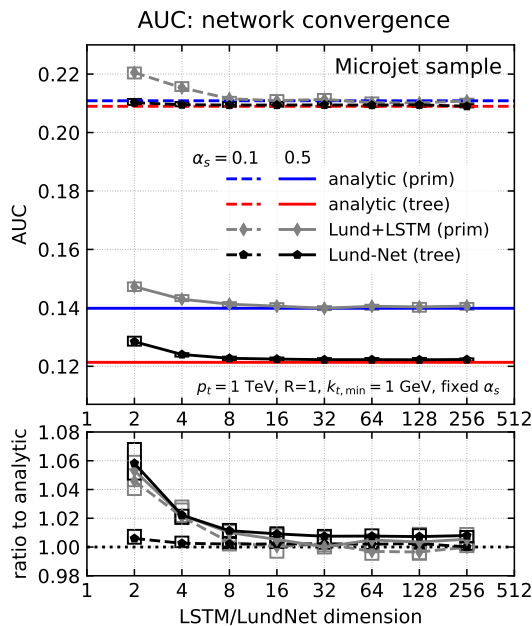


Figure 3. Area under the ROC curve (AUC) as a function of dimension of the network. The networks converge to the optimal (analytic) discriminant above ≈ 32 nodes, proving that the networks learned all the features of the microjet samples.

times for 5 different subdivisions and we take the average and envelope of these five runs respectively as an estimate of the performance and of the associated uncertainty. The quark and gluon efficiencies, ε_q and ε_g , are defined as the fraction of the events in the quark and gluon samples, respectively, which pass a cut on the analytic likelihood ratio or on the network discriminant.

Since the analytic models reproduce the exact likelihood ratio for these collinear samples, they are expected to provide the optimal discriminants. We first study how the area under the ROC curve (AUC) evolves as a function of the dimension of the LSTM or EdgeConv block in our machine-learning setup, varied between 2 and 256 nodes,¹³ compared to the expected exact analytic result. This is shown in figure 3. Here, the AUC is defined as the area under the $(\varepsilon_q, \varepsilon_g)$ curve, meaning that a lower AUC means a better discriminating power. It is remarkable that for a network dimension of 32 or above, the neural network is able to reproduce the expected optimal discriminant to within at most 1%, for both values of α_s . If we look at the full ROC curves, figure 4, we see again the same level of agreement. The larger uncertainty at smaller quark efficiencies is expectable as only a fraction of the background events pass the tagger. Based on figures 3 and 4, we note a hierarchy between the classifiers with Lund-based methods performing better than the Iterated SoftDrop multiplicity and, among the Lund-based methods, the ones using the full tree information performing better than the ones using only primary declusterings. Improving the (logarithmic) accuracy of the analytic approach and exploiting more jet substructure

¹³For Lund-Net, the dimension refers to the size of the first MLP in the initial EdgeConv block, keeping the scaling of the successive layers identical to the one in section 5.2.

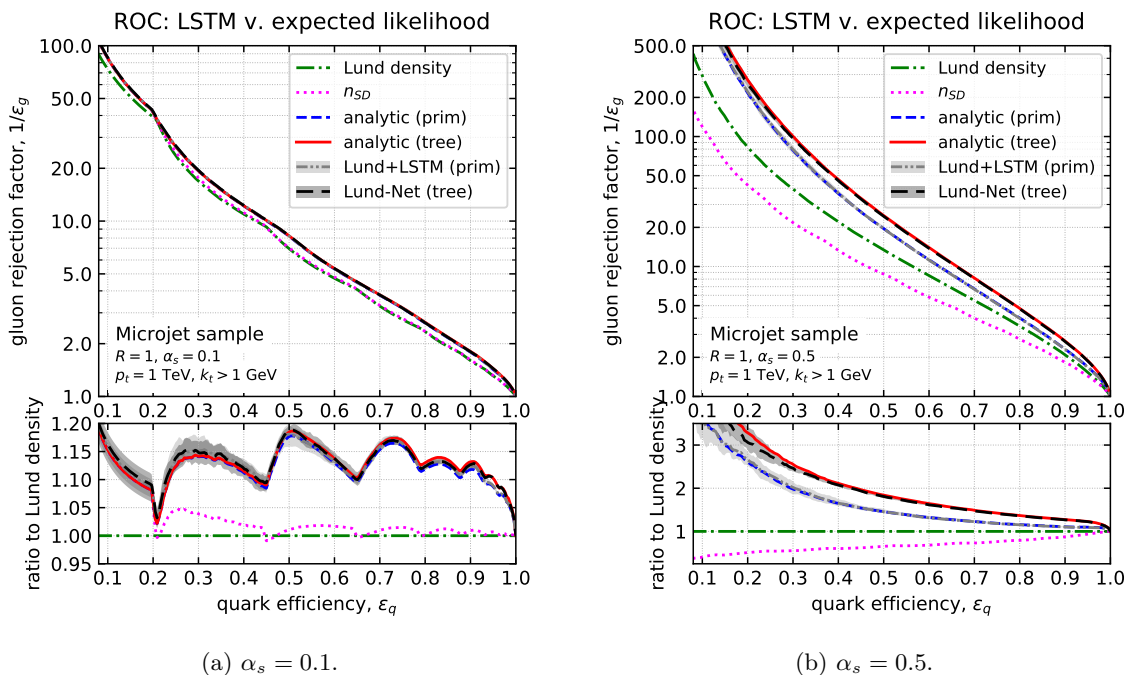


Figure 4. The ROC curves for the analytic and ML classifiers for the microjet sample. For this sample, our analytic approach is exact, showing that the neural networks capture the full (single-logarithmic) information in the training sample.

information both lead to a performance increase. Furthermore, the performance differences are enhanced as one opens up the phase space to include more emissions (or increase α_s).

To further investigate the $\lesssim 1\%$ difference between the Lund-Net and analytic results for $\alpha_s = 0.5$, we show in figure 5 the performance this time as a function of the size of the training sample (keeping the size of the validation and testing samples to 10^5 jets). The shaded band around the analytic expectation represent the statistical fluctuations obtained by splitting the full 10^6 event sample in 10 subsamples of 10^5 events, running our analytic discriminant independently on each subsample. For a testing sample size of 10^5 jets and $\epsilon_q = 0.1$, we only keep $\sim 0.1\%$, i.e. ~ 100 , of the gluon jets which is compatible with the $\sim 10\%$ observed statistical uncertainty. One sees that within the statistical uncertainties, the performance of the Lund-Net approach matches that of the analytic expectation for a training sample of 4×10^5 events or more. This is seen both for the AUC, figure 5a and for the ROC curves, figure 5b. In the latter case, the convergence is slightly slower at small $\epsilon_q = 0.1$, as one could have expected.

Before we close this section, we note that additional tests targetting the asymptotic single-logarithmic limit of full Monte Carlo simulations will be carried on in section 7.6.

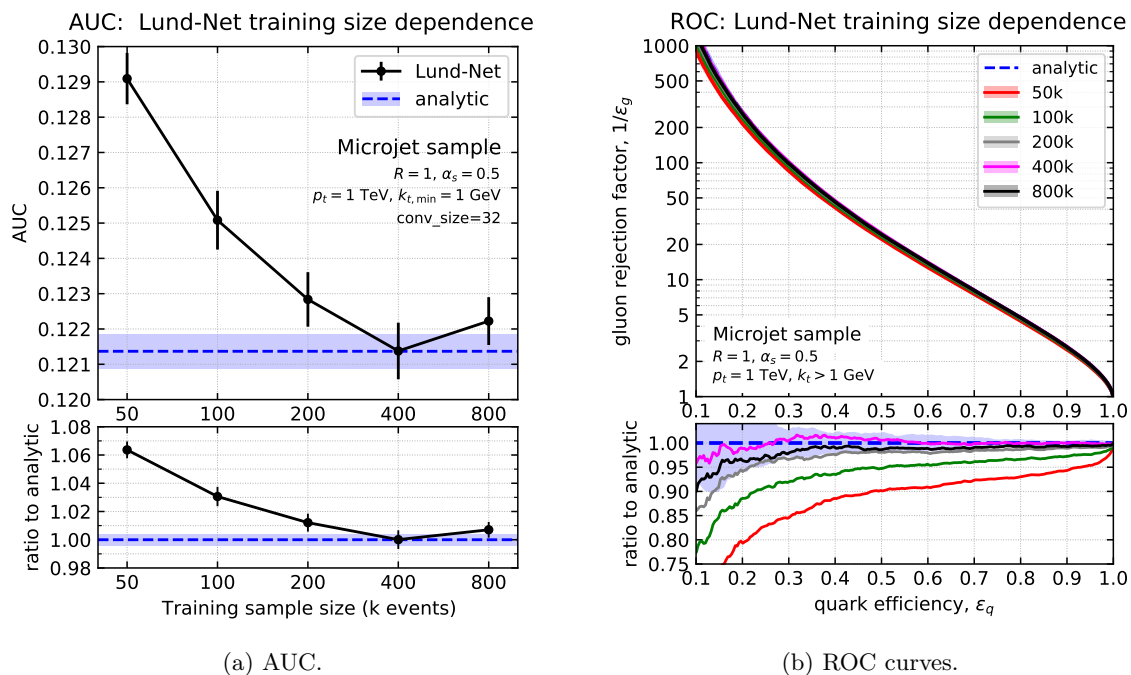


Figure 5. Convergence of the ML-based Lund taggers as a function of the training sample size for the AUC (a) and ROC curve (b). The ML testing phase is always performed on the same 100k events and the shaded band around the analytic results is the fluctuations across different samples of 100k events. The optimal performance is reached for training sample sizes of 400k events and above.

7 Full Monte Carlo simulations

7.1 Setup

We now move to testing the performance of Lund-plane-based quark-gluon discriminants with full Monte Carlo samples. For our reference quark and gluon samples, we simulated respectively $Z+q$ and $Z+g$ events with Pythia v8.24 [61, 62] with multi-parton interactions enabled with the Monash13 tune [63]. For the $Z+q$ sample, only light quark flavours (u , d and s) have been included. The Z boson is set to decay to invisible neutrinos. Jets are then reconstructed on the remaining final-state using the anti- k_t algorithm [42] with $R = 0.4$, as implemented in FastJet [64, 65]. We select at most the two hardest jets within $|y| < 2.5$ and keep only the ones with $500 < p_t < 550$ GeV. For each selected jet, we recluster its constituents with the Cambridge/Aachen algorithm and we construct the Lund declusterings following the recipe described in section 2. The studies described below are performed with quarks and gluon samples of 10^6 jets each.

To probe the resilience of our quark-gluon discriminants against various effects we have generated additional event samples. The first one uses the same setup as above with hadronisation and multi-parton interactions switched off, hence probing the influence of non-perturbative effects. The second uses dijet events, $qq \rightarrow qq$ with light quark flavours and $gg \rightarrow gg$, and is meant to probe the dependence on the hard process. The third one uses the same setup as the reference sample (Z +jet with non-perturbative effects enabled),

this time generated with Herwig v7.2.0 [66, 67] so as to probe the dependence on the Monte Carlo generator.

Since our goal in this paper is to address “universal” aspects of quark-gluon tagging (recall the discussion in section 2), one can consider that varying the Monte Carlo generator and the hard process yielding the quark and gluon samples provide a picture of how much freedom there is beyond these universal aspects. For a practical analysis where one would consider specific signal and background processes, this discussion would have to be revisited and it would be interesting to include samples including full NLO matrix elements matched with parton shower, using tools such as MC@NLO [68], Powheg [69–71] or Sherpa [72, 73].

We test a total of six quark-gluon discriminants: the Iterated Soft Drop multiplicity (n_{SD}) and the discriminant based on the average Lund plane density (Lund density), both described in section 2, our new analytic discriminants using either the primary declusterings only (analytic(prim), sections 3.2 and 4.3) or the full declustering tree (analytic(tree), sections 3.3 and 4.2), and the deep-learning approaches using either only the primary Lund declusterings (Lund+LSTM (prim)) or the full Lund tree (Lund-Net (tree)) both described in section 5.

For the analytic models, clustering contributions are included with their dependence on the azimuthal angle ψ . We further discuss the influence of clustering logarithms and of the azimuthal angle ψ in section 7.5 below. Our analytic models are only considered in the presence of a k_t cut on the Lund declusterings, guaranteeing infrared-and-collinear safety.

As in the previous section, for the methods using machine learning, the event sample is subdivided into 8×10^5 training jets, 10^5 evaluation jets and 10^5 testing jets. We use five different subdivisions of the full sample to assess the uncertainties on the performance. For the discriminant based on the average Lund plane density, we use the first 9×10^5 events to build a (binned) estimate of $\rho_{q,g}(\Delta, k_t)$ and the 10^5 remaining events as a testing sample.

7.2 Tagging performance

We first look at the performance of our taggers. In this section we use our reference Monte Carlo sample, i.e. Z +jet events generated with Pythia with hadronisation and multi-parton interactions enabled. Figure 6 shows the area under the ROC curve as a function of the k_t cut applied on Lund declusterings, and figure 7 shows the ROC curves themselves for two specific choices of the cut: no cut (figure 7a) or $k_t > 1$ GeV (figure 7b).

Leaving aside for now the methods based on deep learning, we see the expected pattern. First, the average Lund density brings a small improvement compared to the Iterated SoftDrop approach. It is interesting to notice that while the performance of the Lund density approach flattens as the k_t cut is lowered, that of ISD gets worse at small k_t cuts. Since the Iterated SoftDrop multiplicity and our analytic approaches are based on perturbative QCD arguments, one might have anticipated their performance degradation for low values of the k_t cut. In particular, since our analytic models include the running of the strong coupling with k_t , they become unstable as we approach the Landau pole. The average Lund density approach however directly uses the Pythia sample to estimate the likelihood and is therefore free of these effects.




Figure 6. Area under the ROC curve (AUC) obtained using Z +jet events simulated with Pythia. All the Lund-plane based quark-gluon discriminants studied in the paper are shown: the Iterated Soft Drop multiplicity (n_{SD}), the likelihood based on the average primary Lund-plane density, and our new analytic and ML-based discriminants using either primary Lund branchings of the full Lund declustering tree. The AUC is plotted as a function of the minimum transverse momentum, k_t , cut imposed on the Lund declusterings.

Let us now focus on our analytic discriminants. Compared to the average density we see a $\sim 5\%$ improvement in the AUC when using only the primary declusterings, reaching 9% for $k_{t,cut} = 1$ GeV. As visible in figure 7b, this improvement increases towards smaller quark efficiency where it can reach 30–50% for ε_q in the 0.2–0.5 range. Adding the information from the full clustering tree, this improvement in AUC increases slightly, reaches e.g. $\sim 12\%$ for $k_{t,cut} = 1$ GeV. Looking at the ROC curve, this improvement is seen mostly at large ε_q with limited impact at smaller ε_q . As for ISD, the performance of our analytic models worsens for small k_t cuts, below 1 GeV. This is most likely due to a breakdown of the perturbative approach.

If we now turn to the Lund methods using deep learning, we see a clear improvement in discriminating power for all k_t cuts and across all values of the quark efficiency. The AUC is reduced (i.e. improved) by 20–40% for a cut on k_t below 1 GeV and the gluon rejection factor is improved by a factor between 2 and 3 for ε_q in the 0.2–0.5 range.

A striking feature of the machine-learning-based approaches in figures 6 and 7, especially compared to the results shown in the collinear sample in section 6, is that they show a substantial performance improvement compared to the analytic models. There can be several explanations for this. Of course, since our analytic approach is purely perturbative, differences can be of non-perturbative origin. This is certainly the case at very small

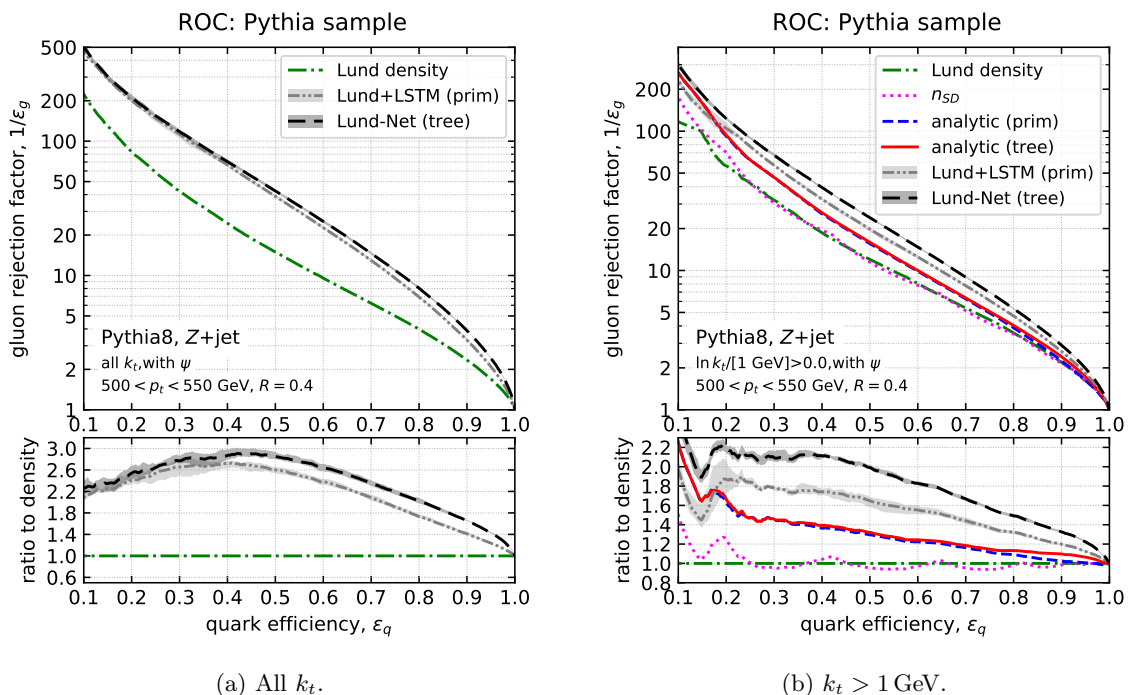


Figure 7. ROC curves corresponding to the AUC showed in figure 6. The left plot corresponds to all Lund declustering being included while for the right plot a k_t cut of 1 GeV has been applied. Since the methods anchored in perturbative QCD are meant to be effective in the region where perturbative QCD applies, they are only shown on the right plot.

values of $k_{t,\text{cut}}$ where our perturbative approach breaks down when the performance of the machine-learning-based approaches keeps improving. However, the gain is already visible at values of the k_t cut close to 10 GeV, where non-perturbative corrections are relatively small. From a pure perturbative perspective, there are at least three possible explanations as to why the machine learning approaches may outperform our analytic discriminant.

First, our treatment of clustering logarithms is only correct for pairs of emissions at commensurate angles so we should expect corrections even in the single-logarithmic limit.

Secondly, our analytic discriminant works in the limit of small angles, where quark-gluon discrimination can be thought of as universal (at least within our single-logarithmic approximation). The deep-learning methods will learn additional information, starting at the single-logarithmic accuracy, from radiation at large angle. Since this information is process-dependent, one should expect this gain in performance to come at the expense of an enlarged sensitivity to the hard process. We will come back to this point in section 7.3.

Lastly, there can be effects of subleading perturbative corrections that are not included in our analytic approach. These can either be subleading logarithmic corrections beyond single logarithms, or finite, fixed-order, corrections which would induce additional correlations between the Lund declusterings that are neglected at our analytic accuracy but that the neural network training would pick. In section 7.6, we show that if we take a more asymptotic regime, the gap between the analytic and deep-learning approaches shrink,

strongly suggesting that the differences seen in figures 6 and 7 is not dominated by our simplified treatment of clustering logarithms.

We should also point out that subleading logarithmic corrections, or fixed-order corrections, are not fully included in Monte-Carlo event generators like Pythia. The improvement seen with deep-learning approaches should therefore be taken with caution.¹⁴

Of course, to this list of perturbative effects, one should also add non-perturbative corrections which, while beyond the reach of our analytic approach, are captured by the neural networks.

7.3 Resilience

The discriminating power of a quark-gluon tagger is not necessarily the only quality feature we may want to require. Indeed, our extraction of the tagging performance is obtained for a specific event sample which can have its own limitations or simply be different from the event sample used in later practical applications.

Ideally, one would want a tagger to be *resilient*, i.e. to show a degree of insensitivity to potential mismodelling aspects or to specific details of an event sample. For example, if we want to be able to describe a tagger from first-principles perturbative QCD, we would want to limit its sensitivity to the modelling of non-perturbative effects. In a similar spirit, we want to limit the sensitivity of a tagger to the details of the event generator used to obtain the event sample. More specifically, in the context of quark-gluon discrimination, we want our taggers to be insensitive to the details of the (hard) processes contributing to the signal and background(s) we try to separate. This last point is intimately related to the intrinsic ill-defined nature of quark-gluon tagging. In this context, resilience can be seen as a measure of universality.¹⁵

In this section, we therefore investigate the resilience of our tagger against the three effects listed above: (i) non-perturbative effects, (ii) the choice of the hard process and, (iii) the choice of a Monte Carlo event generator. The first is probed by comparing our reference sample to a sample generated at parton level, i.e. with hadronisation and multiparton interactions switched off (see section 7.1 for details). For the second we use either our reference Z + jet sample or a sample of dijet event, and for the third, we compare our default Pythia8 sample to a Herwig7 sample.

For the analytic models, we apply them directly to the different event samples, obtaining in each case the quark and gluon efficiency as a function of the cut on the model's output, i.e. either the Iterated Soft Drop multiplicity, or the analytic likelihood ratio for the “Lund density” approach or for our new primary or full analytic discriminants. For

¹⁴It is tempting to argue that Monte Carlo event generators implement a more precise kinematics than the approximate one used in our analytic approach. For example, the Sudakov factor in the analytic calculations only retains the contributions up to single logarithms. For the fixed-coupling toy microjet sample used in section 6, we had instead kept the full z dependence of the splitting function in the Sudakov factor. We have checked that the effect of keeping the full splitting instead of keeping only the terms relevant at the single-logarithmic accuracy is, at most, 0.5%. This is clearly insufficient to explain the differences between the analytic and machine-learning approaches observed here.

¹⁵The idea of being resilient against details of the hard process however extends to tagging applications beyond quark-gluon discrimination.

machine-learning approaches, we have trained the networks on our reference Pythia8 Z +jet sample with hadronisation and multi-parton interactions, and applied the resulting network to the other event samples.

For a given fully-specified tagger, i.e. a tagging method and cut on its output (a.k.a. a working point), one obtains quark and gluon efficiencies $\varepsilon_{q,g}^{(\text{ref})}$ and $\varepsilon_{q,g}^{(\text{alt})}$, respectively for the reference and alternative event samples. The resilience of the taggers is then simply defined as the inverse of the relative change between the two samples:

$$\zeta = \left[\left(\frac{2(\varepsilon_q^{(\text{alt})} - \varepsilon_q^{(\text{ref})})}{\varepsilon_q^{(\text{alt})} + \varepsilon_q^{(\text{ref})}} \right)^2 + \left(\frac{2(\varepsilon_g^{(\text{alt})} - \varepsilon_g^{(\text{ref})})}{\varepsilon_g^{(\text{alt})} + \varepsilon_g^{(\text{ref})}} \right)^2 \right]^{-1}. \quad (7.1)$$

With this choice, a bigger ζ corresponds to a more resilient tagger. We should then select a working point at which we evaluate the performance and resilience of a tagger. In practice, we take the point at which the tagging performance, defined as the significance $\varepsilon_q/\sqrt{\varepsilon_g}$, is maximal for the reference sample. This is typically realised for $\varepsilon \sim 0.3$ – 0.5 . The resulting performance (significance) and resilience are denoted by Π_{best} and ζ_{best} , respectively. Ideally, we therefore seek for a tagger with large Π_{best} and ζ_{best} . We have tested that selecting instead a fixed quark efficiency, e.g. $\varepsilon_q = 0.5$, produces similar results.¹⁶ That said, while quantitative arguments can be made about the relative discriminating performance of our taggers, it is more delicate to reach such a precise quantitative discussion of resilience. The discussion below therefore tries to remain mostly at a qualitative level, i.e. noting that taggers with larger resilience are likely to have less modelling uncertainties. It would be interesting — and clearly beyond the scope of this paper — to perform a dedicated study of resilience.

We want to study how resilience and performance behave for our discriminants, varying the k_t cut on Lund declusterings. Our results are presented in figure 8, for the three types of resilience we want to investigate: resilience against the specifics of the hard process (figure 8a), resilience against non-perturbative effects (figure 8b), and resilience against the choice of the event generator (figure 8c). To guide the eye, the results corresponding to a k_t cut of 1 GeV are represented with filled symbols, with all the other results using open symbols. For all three resiliences, the usual trade-off is observed: as we increase the k_t cut, performance decreases and resilience increases. Overall, our analytic models and our Deep Learning results show a similar behaviour, although a given performance-resilience point is achieved for a different k_t cut for different taggers. Our analytic models however appear as slightly more resilient to non-perturbative effects than their machine-learning equivalents.

Compared to the Iterative Soft Drop and Lund density approaches, one sees that the analytic model typically bring a gain in performance without sacrificing in resilience.

Focusing on the results with a k_t cut of 1 GeV, it is interesting to see that the machine-learning-based techniques reach a larger performance, as already seen in figures 6 and 7, at the expense of having a smaller resilience. This hints towards the interpretation that this

¹⁶In general, one can argue that the lower ε_q values should be ignored because they are subject to large statistical fluctuations. The large and low ε_q values are also impractical because they do not yield a large discriminating power.

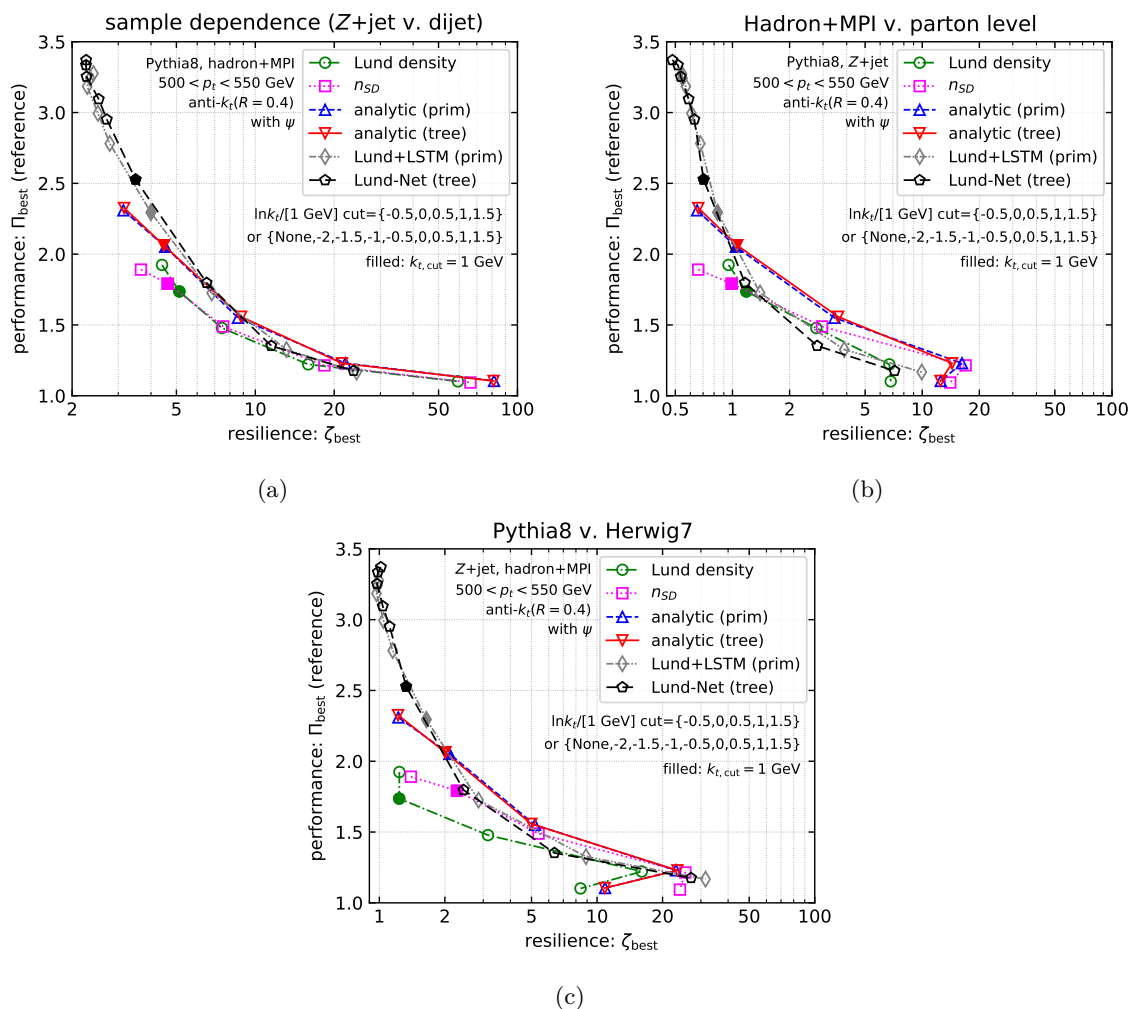


Figure 8. Plots of performance as a function of resilience for different discriminants. The curves are obtained by scanning over a range of accessible $k_{t,cut}$ values. The filled symbols correspond to a k_t cut of 1 GeV. Different plots correspond to different probes of resilience: (a) probes the sample dependence (replacing the Z +jet sample with a dijet sample), (b) probes non-perturbative effects (using parton-level simulations instead of full simulations with hadronisation and multi-parton interactions) and (c) probes the effect of the Monte Carlo generator (using Herwig 7 instead of Pythia 8).

gain in performance is obtained by the neural networks exploiting information (i) going beyond the “universal” collinear behaviour (worse resilience against the choice of hard process), (ii) in non-perturbative effects (worse resilience against hadronisation and MPI), and (iii) specific to the modelling of the events (worse resilience against the choice of event generator). In all three cases, increasing the k_t cut by a few hundred MeVs would result in a behaviour very similar to the one of the analytic model, both in terms of performance and in terms of resilience.

Finally, if all one cares about is performance, machine-learning discriminants using the full information in the Lund tree show the best result, albeit at the expense of a poor resilience. This should at least be kept in mind when using a quark-gluon discriminant for potentially different applications, or when assessing uncertainties associated with a tagger.

7.4 Comparison with other approaches

In this section, we compare the performance of our Lund-plane-based taggers to that of other existing taggers.

We therefore select our taggers based on the full Lund tree: the Lund-Net tagger with no k_t cut and the analytic discriminant (“analytic(tree)” in previous figures) with a k_t cut of 1 GeV (referred to as “Lund NLL” in this section), and compare them to a series of pre-existing discriminants. We first consider benchmark jet shapes:

- angularities [19, 20], defined as the following sum over the jet constituents $\lambda_\alpha = (\sum_{i \in \text{jet}} p_{t,i} \Delta R_{i,\text{jet}}^\alpha) / (R^\alpha \sum_{i \in \text{jet}} p_{t,i})$. We work with $\alpha = 1$, sometimes referred to as width or girth.
- energy-energy correlation functions [21], defined as the following sum over pairs of jet constituents $\text{EEC}_\beta = (\sum_{i,j \in \text{jet}} p_{t,i} p_{t,j} \Delta R_{ij}^\beta) / [R^\beta (\sum_{i \in \text{jet}} p_{t,i})^2]$. In this case, we will set $\beta = 1/2$.
- to probe the effect of a k_t cut similar to the one we introduce in the Lund plane techniques, we have considered the case where the angularities and energy-correlation function are defined on the Lund declusterings (primary and secondary) above a given k_t cut. We recall that in this case, we expect that the Iterative Soft Drop multiplicity and our analytic Lund-tree discriminant are respectively optimal at leading and next-to-leading logarithmic accuracy in QCD.

In all cases, a cut is applied on one of these jet shapes and the full ROC curve is obtained by varying the cut. We then consider a series of recent machine-learning-based quark-gluon discriminants (for which we also apply a cut on the network output):

- Particle-Net described in ref. [28], based on point clouds. In practice, we have directly used the `ParticleNet` code available from [74], modifying the provided `keras` example to use our event sample. We have used a batch size of 1000 and kept the best model over a training of 50 epochs. Note that this model includes the particle ID in the network inputs.
- Particle-Flow Networks (PFN) from ref. [27]. This includes the rapidity, azimuth (both relative to the jet axis) and transverse momentum information of each jet constituent. Each particle is mapped into a per-particle latent space. The sum over all particles of these spaces is then mapped onto a final discriminating variable. We have also considered the so-called PFN-ID approach where the particle IDs are also included. In practice, we have used the code provided in the `EnergyFlow` package [75], modifying the examples to use our event samples and training over 60 epochs.
- Energy-Flow Networks (EFN) also from [27] and again only adapted from the example given in the `EnergyFlow` package to our needs. The approach is similar to that of the PFN above except that the latent space uses IRC-safe information through a weighting proportional to the p_t of each particle.

Finally, we also consider the Lund-Net approach, labelled “Lund-Net(+ID)” where each Lund tuple $(\ln \Delta, \ln k_t, \ln z, \Psi)$ is supplemented by one additional integers for each of the two subjects j_1 and j_2 and determined as follows: if the subjet has a single constituent we use the PDG ID of the constituent, otherwise we set the integer to 0. The idea is similar to having particle identification added from the PFN to the PFN-ID approach.¹⁷

Our findings are presented in figure 9, for the signal significance (top row), and for the trade-off between significance and resilience against the choice of Monte-Carlo event generator (bottom row). Focusing first on the signal significance for analytic discriminants, figure 9a, we see relatively different patterns between the shape-based observables λ_1 and $\text{EEC}_{0.5}$ and the Lund-based observables, with the performance peaking at larger values of the quark efficiency in the latter case, with shape-based observables reaching a better overall performance. However, computing the EEC and λ using the Lund declusterings with a k_t above 1 GeV, i.e. with the same input information in all cases, we see that our analytic Lund discriminant shows indeed an improvement at all quark efficiencies compared to the shape-based discriminants. A similar pattern is seen in the resilience plot on figure 9c with the Lund analytic model with a 1-GeV k_t cut being intermediate between the jet shape with a 1-GeV k_t cut and with no k_t cut. It is interesting to notice that the addition of particle ID information to the Lund-Net approach improves the performance at low k_t cut, or with no k_t cut at all, but changes neither the performance nor the resilience once a larger k_t cut is applied. This is most likely due to the fact that all the input subjects have more than one constituent and hence the ID information is 0. This contrasts with the findings in ref. [30], where the addition of the jet mass had a negative impact on resilience at large k_t cuts. At low resilience (large significance), the jet shapes give a slightly better performance vs. significance behaviour than our Lund-plane approach. Focusing instead on the AUC — the bottom-right table in figure 9 — we see that our analytic Lund-tree approach does a better job than the other jet shapes (i.e. a lower AUC), including jet shapes computed on the full set of constituents.

Moving to machine-learning-based models, figure 9b, we see a significance pattern mostly similar across different models. The performance of Lund-based models is on par with the one obtained from Particle-Net. Compared to energy/particle flow approaches, our Lund-based results show a slightly better performance than the EFN and PFN results, but fall slightly lower than the performance of PFN-ID. Adding the particle ID information using our Lund-Net(+ID) approach recovers a performance similar to the PFN-ID approach, although with a marginally smaller average peak performance. If we instead look at the AUC, we see that the Lund-Net(+ID) reaches the best performance (lowest AUC), marginally better than the PFN-ID and Lund-Net approaches, then followed by the Particle-Net model. While the PFN-ID method shows a small performance improvement at mid signal efficiency, the Lund-Net(+ID) setup has a small advantage at small and large signal efficiencies.¹⁸ The observed differences are however of a size similar to the statisti-

¹⁷In an experimental context, complete particle-ID information would not be available. One could however separate charged tracks from neutral energy deposits in calorimeters, with potential additional information such as heavy quarks or electromagnetic v. hadronic calorimeters.

¹⁸Including the particle ID information in a more coherent way, e.g. as a separate information that is fed to the final dense layers, one might be able to make up the difference with PFN-ID at mid quark efficiencies.

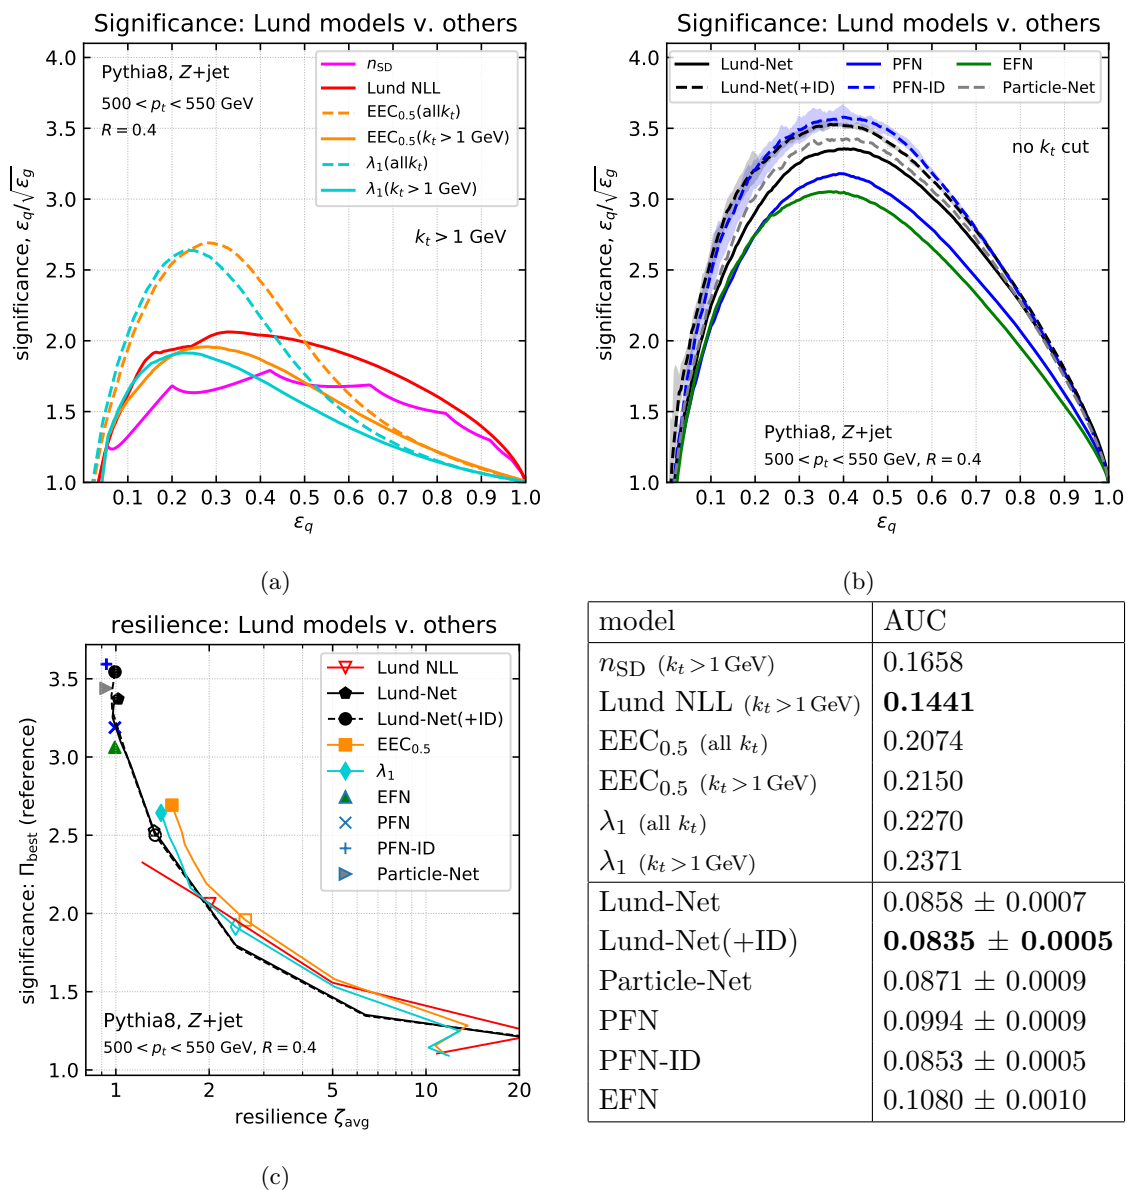


Figure 9. Comparison of the Lund-plane-based approaches with other models. Explicit plots of the signal significance $\varepsilon_q/\sqrt{\varepsilon_g}$ are shown in the upper plots, first for analytic discriminants, figure (a), then for machine-learning-based approaches, figure (b). The bottom panel, figure (c), shows the corresponding performance v. resilience plot, where the resilience is measured with respect to the choice of Monte Carlo generator (cf. section 7.3). The table on the bottom-right corner gives the area under the ROC curve (AUC) for the different models (lower is better). For the ML-based models, the uncertainty is half the difference between the minimal and maximal values obtained over 5 different runs.




Figure 10. Plots showing the ratio of the ROC curve with and without azimuthal-angle dependence and/or clustering logarithms. The plots of the left column include only primary declusterings while the plots of the right column include the full declustering tree. The bottom plots include declusterings above a 1-GeV k_t cut, while the top plots include all relative transverse momenta. In all cases, the ROC curve is normalised to the ML Results including azimuthal-angle dependence.

cal fluctuations observed in our simulations (only shown, on the significance plot, for the Lund-Net(+ID) and PFN-ID for the sake of readability).

Finally, from figure 9c, we see a similar degree of resilience for all machine-learning-based approaches. Again, it would be interesting to train energy/particle flow networks or the ParticleNet network on Lund declusterings above a certain k_t cut (or using another cut definition) to study the performance versus resilience trade-off in a broader perspective.

7.5 Effect of clustering logarithms and of azimuthal angles

All the results presented so far have included the dependence on the Lund azimuthal angles ψ_i . Since it is known that these are not properly described at single-logarithmic accuracy by the standard dipole showers (including Pythia8) [76], we want to briefly investigate their impact on discriminating power. Additionally, from our analytic perspective, the dependence on the Lund azimuthal angles ψ_i only comes in through the clustering logarithms. This therefore gives us an explicit opportunity to investigate the numerical impact of including clustering logarithms into our analytic discriminants.

In figure 10, we show the background rejection, $1/\varepsilon_g$, including either only primary declusterings (left column) or the full declustering tree (right column), as obtained using

our reference Pythia 8 Z +jet sample with hadronisation and multi-parton interactions. The top row, with only machine-learning results, corresponds to the case without a k_t cut while the bottom row, with both ML and analytic results, includes only the declusterings with $k_t \geq 1$ GeV. To increase readability, we show in all cases, the ratio relative to what is obtained with the corresponding ML model — LSTM or Lund depending on whether only the primary or all the declusterings are used — including the ψ dependence.

We first discuss the ML results, presented in figure 10 either with (solid, black) or without (dashed, green) ψ information. We see that including the ψ information brings a 5–15% performance gain, mostly at intermediate quark efficiency. This gain is larger at lower k_t where non-perturbative effects are larger. When a 1-GeV k_t cut is imposed, we also show the results of our analytic quark-gluon taggers, again using either primary-only information (left column), or using the full declustering tree (right column). In each case, three results are given: including the ψ angles (solid, blue), not including the ψ angles but including the (ψ -averaged) clustering logarithms from eq. (4.5) (dashed, red), or including neither the ψ angles, nor the clustering logarithms (dotted, magenta). We again see a $\sim 10\%$ increase in performance brought by the inclusion of azimuthal angles. The fact that these performance gains are of similar magnitude in the deep-learning and analytic approaches indicates that our simplified treatment of clustering logarithms is a decent approximation.

Finally, when the azimuthal angles are not included, we see that the influence of clustering logarithms is small.

7.6 Asymptotic single-logarithmic limit

In this final study, we want to further study the differences between the analytic and ML results. Our aim is here to take a limit where subleading effects decrease. Since our analytic approach technically resums double and single logarithms of $\log(p_t R/k_{t,\text{cut}})$, we want to proceed in a similar way as for the NLL-accuracy tests in [76], i.e, take the limit $\alpha_s(p_t R) \rightarrow 0$, $\log(p_t R/k_{t,\text{cut}}) \rightarrow \infty$ while keeping $\alpha_s(p_t R) \log(p_t R/k_{t,\text{cut}})$ constant. The main idea behind this limit is that subleading-logarithmic contributions as well as fixed-order contributions are suppressed as $\alpha_s \rightarrow 0$.

Generating and analysing events over an exponentially increasing range of scales poses a series of numerical challenges which, in practice, make it unreachable for standard Monte Carlo event generators like Pythia8. We have therefore used instead the PanScales e^+e^- code developed precisely to overcome these challenges in ref. [76]. We have therefore generated $e^+e^- \rightarrow Z \rightarrow q\bar{q}$ (quark) events and $e^+e^- \rightarrow H \rightarrow gg$ (gluon) events with a centre-of-mass energy Q , fixing $\alpha_s(Q) \log(Q/k_{t,\text{cut}}) = 0.32$ and taking $\alpha_s(Q)$ to be either 0.04, 0.02, 0.01, corresponding to $L = \log(Q/k_{t,\text{cut}})$ of either 8, 16, or 32. In all cases, we have used the PanLocal shower in its antenna variant with the β parameter set to 1/2. Subleading colour corrections are included using the Nested Ordered Double-Soft (NODS) scheme as described in [49]. This produces event samples with a single-logarithmic accuracy with the exception of the subleading- N_c corrections for which the NODS method only guarantees the correct behaviour for (any number of) pairs of emissions at commensurate angles. We reconstruct the e^+e^- Lund declusterings in each of the event hemispheres. The e^+e^- reconstruction follows an almost-trivial adaptation of the hadronic collisions

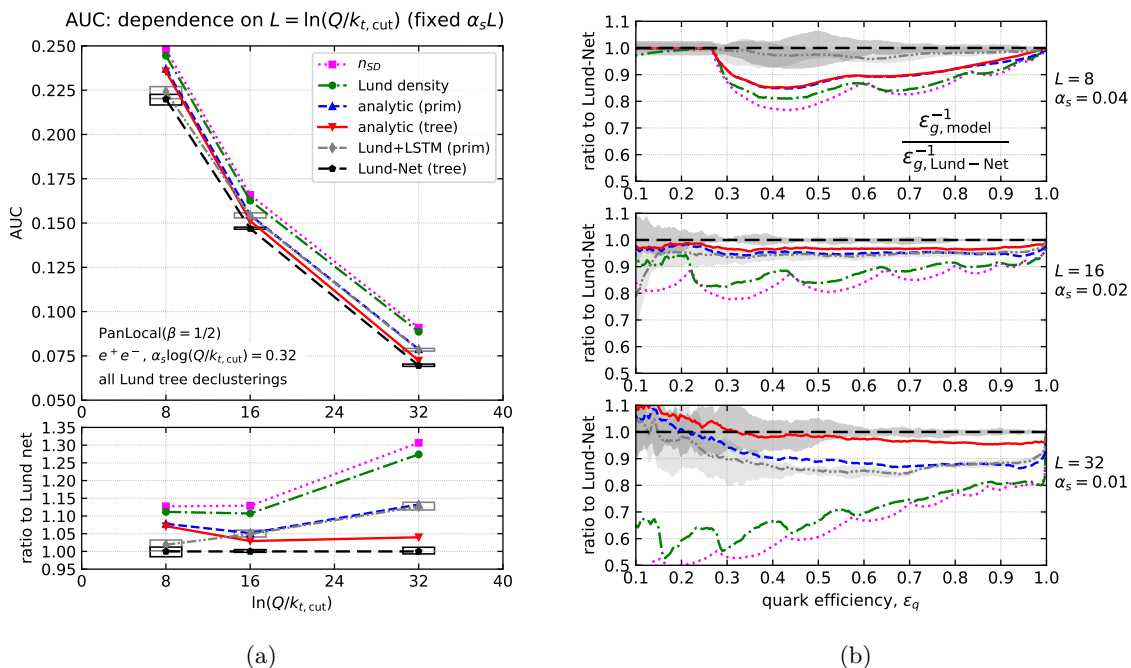


Figure 11. Dependence of (a) the AUC, and (b) the ROC curve on $L = \log(Q/k_{t,cut})$ fixing $\alpha_s L = 0.32$. The results are obtained with an e^+e^- parton-level setup using the PanLocal PanScales shower which is NLL-accurate. Different curves show different methods and we see that the analytic and ML models converge to one another as we take the asymptotic limit $L \rightarrow \infty$, $\alpha_s \rightarrow 0$, $\alpha_s L = \text{constant}$. For the plots on the right column, we take the ratio of the ROC curves to the Lund-Net results.

technique described in section 2, except perhaps for the reconstruction of the azimuthal angle ψ which is described in details in [76].

Both our analytic methods and the approaches based on Machine Learning can be straightforwardly applied to these new sets of events. We therefore study the same set of quark-gluon discriminants as in section 7.2. The contribution from clustering logarithms is included in our analytic models and information from the Lund azimuthal angles is included both in the analytic and ML approaches.

Our results are presented in figure 11 for the AUC, figure 11a and for the ROC curves, figure 11b. For the latter, we have normalised the gluon rejection rate ϵ_g^{-1} to the Lund-Net rejection rate. One can see from these plots that the difference between the analytic and Machine-Learning-based methods decrease when increasing L (decreasing α_s) at fixed $\alpha_s L$. This is true separately for the approaches using only primary declusterings (“analytic(prim)” and “Lund+LSTM”) and for the approaches using the full declustering tree (“analytic(tree)” and “Lund-Net”). At the same time, the performance gain compared to the Iterated SoftDrop multiplicity increases.

Before closing this section, we want to address a last point about the azimuthal angle dependence and single-logarithmic accuracy. It has been pointed out in ref. [76] that, due to non-physical recoil effects, dipole showers such as Pythia or Dire, would generate a spurious dependence on ψ , potentially biasing the assessment of quark-gluon classification. Using

the machinery described in this section, we have studied potential differences between the Pythia shower and the PanScales showers which are free of this effect [76]. Within the few-percent accuracy of our studies we have not been able to isolate a clearly-visible impact of this effect.

8 Conclusions

This paper addresses the question of quark/gluon discrimination using the Lund-plane approach to characterise the substructure of jets. Our main result is that it is possible to compute the quark-gluon likelihood ratio from the first principles in QCD. The calculation is done at the single-logarithmic accuracy, including all collinear contributions as well as clustering effects for any number of pairs of emissions at commensurate angles. This automatically provides us with an optimal quark/gluon tagger at the same accuracy.

As expectable, this tagger shows an improved performance either compared to using the average Lund plane density to build the likelihood ratio, or compared to the Iterated Soft Drop multiplicity which corresponds to the optimal quark/gluon discriminant at leading (double) logarithmic accuracy. Most of the improvement ($\lesssim 10\%$ for the AUC compared to n_{SD}) is already captured when including only primary declusterings, but the effect of additional declustering in subsidiary Lund planes ($\sim 3\%$) is clearly visible, especially at larger quark efficiencies. The gain in performance can be attributed to the better treatment of the kinematics of each emission, e.g. through the full Altarelli-Parisi splitting functions, through the full antenna pattern for emissions at commensurate angles, and to a better treatment of the correlations between emissions, e.g. by taking into account the energy of the emitting parton or by including clustering effects. Furthermore, this gain in performance is accompanied by a gain in resilience against effects beyond our perturbative calculation. In this context, we have studied three specific effects: the dependence against non-perturbative effects, the dependence against the specific choice of quark/gluon enriched samples used as benchmarks, and the choice of Monte Carlo event generator.

In section 7.4, we have compared our Lund-based approach to other typical quark/gluon taggers using jet substructure, like angularities or energy correlation functions. Focusing for simplicity on the case where all the taggers are applied to Lund declusterings with a k_t above 1 GeV to reduce non-perturbative effects, we see that the Lund-based likelihood approach gives a gain in performance, especially at large quark efficiency, while maintaining a similar degree of resilience.¹⁹

The second set of results in this paper is the extensive study of quark/gluon taggers using deep-learning techniques combined with Lund declustering information. When applied to the full set of declusterings in a jet, our ML-based tagger reaches a performance (and resilience) comparable to that obtained with Particle-Flow networks [27], and marginally better than what is achieved by Particle-Net [28]. Compared to the analytic results, our ML-tagger gives a clearly visible performance gain, even when considering declusterings above a given k_t cut. This gain in performance, however, comes at a price in all the forms

¹⁹If the jet shapes are computed on the full jet, they yield a larger significance at smaller quark efficiency at the expense of a reduced resilience against non-perturbative effects.

of resilience we have studied, especially the sensitivity to hadronisation and multi-parton interactions.

One of the key points of this paper is the direct comparison between the analytic and deep-learning approaches. Since our tagger targets the optimal discriminant in the single-logarithmic approximation one can directly compare its performance with that of the deep-learning models. We first did that in the strongly angular-ordered limit where our analytic calculation is exact. The results in section 6 indicate a convergence of the deep-learning taggers to the optimal performance as long as the size of the network is taken large enough. Beyond the collinear limit, where our analytic treatment is only approximate, the deep-learning approach shows a better performance. However, if we move progressively to the single-logarithmic asymptotic limit ($\alpha_s \rightarrow 0$ at fixed $\alpha_s \log(Q/k_{t,\text{cut}})$) the difference between the two approaches drastically reduces as we showed in section 7.6. At the same time, the gain in performance compared to leading-logarithmic-accurate taggers — i.e. the average Lund density and the Iterated Soft Drop multiplicity — increases.

The above observations strongly suggest that the gain in performance observed for deep-learning taggers (in addition to our analytic tagger) in phenomenological Monte Carlo applications come predominantly either from subleading effects (beyond single logarithms), from large-angle soft emissions (not included in our analytic calculation), or from non-perturbative effects.

This points towards several interesting physics considerations. First, subleading logarithmic effects, albeit present in data, are not properly included in any parton shower Monte Carlo generator today. Conclusions regarding subdominant logarithmic effects should therefore be taken carefully. In this context, it would be interesting in the future to further investigate potential differences between standard dipole showers (like Pythia) which are known to have failures at the single-logarithmic accuracy, or even at leading-log for subleading colour effects, and the PanScales showers which are NLL-accurate.

Then, large-angle soft emissions are process-dependent and should therefore be treated carefully when applied outside the configurations where they have been tested and calibrated. In the future, it would be interesting to see if an analytic treatment similar to the one adopted in this paper could allow for quantitative assessment of the process-dependence of quark/gluon discrimination (see also ref. [46] for a Monte-Carlo-based study).

Finally, non-perturbative effects come with non-negligible modelling uncertainties and should therefore also be taken carefully. The ability to progressively reduce non-perturbative effects by increasing the k_t cut-off on Lund declusterings could help further investigating the impact of non-perturbative effects, and the associated systematic uncertainties, in a practical context.

In conclusion, we have seen that Lund-plane declusterings were useful to define a variety of quark/gluon discriminants, bridging regions targetting high discriminating performance and regions where a high-precision degree of control can be reached from first-principles QCD.

Acknowledgments

G.S. is especially grateful to Ben Nachman and Eric Metodiev for stimulating discussions. G.S. and F.D. are also grateful to Gavin Salam for earlier work on the Lund jet plane without which this paper would never have been possible. A.T. wishes to thank the Institut de Physique Theorique for the hospitality. This work has been supported in part by the French Agence Nationale de la Recherche, under grant ANR-15-CE31-0016 (G.S. and A.T.), by the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement No. 788223, PanScales) (G.S.), by the Starting Grant from Trond Mohn Foundation (BFS2018REK01) (A.T.), and by a Royal Society University Research Fellowship (URF\R1\211294) (F.D.).

A Sudakov factors with exact splitting functions

In section 6, we have used an event sample generated in the strong-angular-ordered limit to compare the performance of our analytic discriminants (also in the limit of strong angular ordering, see section 3 to that of deep-learning-based discriminants. To guarantee that the analytic approach reproduces the exact likelihood ratio, we have kept the full Altarelli-Parisi splitting functions in the Sudakov (using a fixed-coupling approximation). If we have a hard parton of momentum xp_t (with p_t the initial transverse momentum of the jet) and flavour f , and compute the Sudakov factor between an angle Δ_{i-1} and Δ_i , we find:

$$-\log S_f^{(i-1,i)} = \frac{2\alpha_s C_f}{p_i} \left[\frac{\log^2 x_1}{2} - \frac{\log^2 x_2}{2} + B_f \log \frac{x_2}{x_1} + \text{Li}_2(x_1) - \text{Li}_2(x_2) + \delta R_f \right], \quad (\text{A.1})$$

with

$$x_1 = \frac{k_{t,\text{cut}}}{x\Delta_{i-1}p_t}, \quad B_q = -\frac{3}{4}, \quad (\text{A.2})$$

$$x_2 = \frac{k_{t,\text{cut}}}{x\Delta_i p_t}, \quad B_g = -\frac{11C_A - 4n_f T_R}{12C_A}, \quad (\text{A.3})$$

and

$$\delta R_q = \frac{3}{2}(x_2 - x_1), \quad (\text{A.4})$$

$$\delta R_g = \frac{3}{2}(x_2 - x_1) + \left(\frac{1}{2} - \frac{n_f T_R}{C_A} \right) \left[(x_2 - x_1) - \frac{1}{2}(x_2^2 - x_1^2) + \frac{2}{9}(x_2^3 - x_1^3) \right]. \quad (\text{A.5})$$

Open Access. This article is distributed under the terms of the Creative Commons Attribution License ([CC-BY 4.0](https://creativecommons.org/licenses/by/4.0/)), which permits any use, distribution and reproduction in any medium, provided the original author(s) and source are credited. SCOAP³ supports the goals of the International Year of Basic Sciences for Sustainable Development.

References

- [1] G.P. Salam, *Towards jetography*, *Eur. Phys. J. C* **67** (2010) 637 [[arXiv:0906.1833](#)] [[INSPIRE](#)].
- [2] S. Marzani, G. Soyez and M. Spannowsky, *Looking inside jets: an introduction to jet substructure and boosted-object phenomenology*, *Lect. Notes Phys.* **958** (2019) 1 [[arXiv:1901.10342](#)] [[INSPIRE](#)].
- [3] P. Gras et al., *Systematics of quark/gluon tagging*, *JHEP* **07** (2017) 091 [[arXiv:1704.03878](#)] [[INSPIRE](#)].
- [4] P.T. Komiske, E.M. Metodiev and J. Thaler, *An operational definition of quark and gluon jets*, *JHEP* **11** (2018) 059 [[arXiv:1809.01140](#)] [[INSPIRE](#)].
- [5] A.J. Larkoski and E.M. Metodiev, *A theory of quark vs. gluon discrimination*, *JHEP* **10** (2019) 014 [[arXiv:1906.01639](#)] [[INSPIRE](#)].
- [6] A. Banfi, G.P. Salam and G. Zanderighi, *Infrared safe definition of jet flavor*, *Eur. Phys. J. C* **47** (2006) 113 [[hep-ph/0601139](#)] [[INSPIRE](#)].
- [7] A.J. Larkoski, I. Moult and B. Nachman, *Jet substructure at the Large Hadron Collider: a review of recent advances in theory and machine learning*, *Phys. Rept.* **841** (2020) 1 [[arXiv:1709.04464](#)] [[INSPIRE](#)].
- [8] R. Kogler et al., *Jet substructure at the Large Hadron Collider: experimental review*, *Rev. Mod. Phys.* **91** (2019) 045003 [[arXiv:1803.06991](#)] [[INSPIRE](#)].
- [9] L. de Oliveira, M. Kagan, L. Mackey, B. Nachman and A. Schwartzman, *Jet-images — deep learning edition*, *JHEP* **07** (2016) 069 [[arXiv:1511.05190](#)] [[INSPIRE](#)].
- [10] P. Baldi, K. Bauer, C. Eng, P. Sadowski and D. Whiteson, *Jet substructure classification in high-energy physics with deep neural networks*, *Phys. Rev. D* **93** (2016) 094034 [[arXiv:1603.09349](#)] [[INSPIRE](#)].
- [11] G. Kasieczka, T. Plehn, M. Russell and T. Schell, *Deep-learning top taggers or the end of QCD?*, *JHEP* **05** (2017) 006 [[arXiv:1701.08784](#)] [[INSPIRE](#)].
- [12] A. Butter et al., *The machine learning landscape of top taggers*, *SciPost Phys.* **7** (2019) 014 [[arXiv:1902.09914](#)] [[INSPIRE](#)].
- [13] H. Qu and L. Gouskos, *ParticleNet: jet tagging via particle clouds*, *Phys. Rev. D* **101** (2020) 056019 [[arXiv:1902.08570](#)] [[INSPIRE](#)].
- [14] E.A. Moreno et al., *JEDI-net: a jet identification algorithm based on interaction networks*, *Eur. Phys. J. C* **80** (2020) 58 [[arXiv:1908.05318](#)] [[INSPIRE](#)].
- [15] CMS collaboration, *Performance of the DeepJet b tagging algorithm using 41.9 fb⁻¹ of data from proton-proton collisions at 13 TeV with phase 1 CMS detector*, CMS-DP-2018-058, CERN, Geneva, Switzerland (2018).
- [16] Y.-L. Du, D. Pablos and K. Tywoniuk, *Jet tomography in heavy-ion collisions with deep learning*, *Phys. Rev. Lett.* **128** (2022) 012301 [[arXiv:2106.11271](#)] [[INSPIRE](#)].
- [17] Y.-L. Du, D. Pablos and K. Tywoniuk, *Deep learning jet modifications in heavy-ion collisions*, *JHEP* **03** (2021) 206 [[arXiv:2012.07797](#)] [[INSPIRE](#)].
- [18] L. Apolinário, N.F. Castro, M. Crispim Romão, J.G. Milhano, R. Pedro and F.C.R. Peres, *Deep learning for the classification of quenched jets*, *JHEP* **11** (2021) 219 [[arXiv:2106.08869](#)] [[INSPIRE](#)].

- [19] C.F. Berger, T. Kucs and G.F. Sterman, *Event shape/energy flow correlations*, *Phys. Rev. D* **68** (2003) 014012 [[hep-ph/0303051](#)] [[INSPIRE](#)].
- [20] L.G. Almeida, S.J. Lee, G. Perez, G.F. Sterman, I. Sung and J. Virzi, *Substructure of high- p_T jets at the LHC*, *Phys. Rev. D* **79** (2009) 074017 [[arXiv:0807.0234](#)] [[INSPIRE](#)].
- [21] A.J. Larkoski, G.P. Salam and J. Thaler, *Energy correlation functions for jet substructure*, *JHEP* **06** (2013) 108 [[arXiv:1305.0007](#)] [[INSPIRE](#)].
- [22] R.D. Field and R.P. Feynman, *A parametrization of the properties of quark jets*, *Nucl. Phys. B* **136** (1978) 1 [[INSPIRE](#)].
- [23] D. Krohn, M.D. Schwartz, T. Lin and W.J. Waalewijn, *Jet charge at the LHC*, *Phys. Rev. Lett.* **110** (2013) 212001 [[arXiv:1209.2421](#)] [[INSPIRE](#)].
- [24] Z.-B. Kang, X. Liu, S. Mantry, M.C. Spraker and T. Wilson, *Dynamic jet charge*, *Phys. Rev. D* **103** (2021) 074028 [[arXiv:2101.04304](#)] [[INSPIRE](#)].
- [25] C. Frye, A.J. Larkoski, J. Thaler and K. Zhou, *Casimir meets Poisson: improved quark/gluon discrimination with counting observables*, *JHEP* **09** (2017) 083 [[arXiv:1704.06266](#)] [[INSPIRE](#)].
- [26] G. Louppe, K. Cho, C. Becot and K. Cranmer, *QCD-aware recursive neural networks for jet physics*, *JHEP* **01** (2019) 057 [[arXiv:1702.00748](#)] [[INSPIRE](#)].
- [27] P.T. Komiske, E.M. Metodiev and J. Thaler, *Energy flow networks: deep sets for particle jets*, *JHEP* **01** (2019) 121 [[arXiv:1810.05165](#)] [[INSPIRE](#)].
- [28] H. Qu and L. Gouskos, *ParticleNet: jet tagging via particle clouds*, *Phys. Rev. D* **101** (2020) 056019 [[arXiv:1902.08570](#)] [[INSPIRE](#)].
- [29] J.S.H. Lee, S.M. Lee, Y. Lee, I. Park, I.J. Watson and S. Yang, *Quark gluon jet discrimination with weakly supervised learning*, *J. Korean Phys. Soc.* **75** (2019) 652 [[arXiv:2012.02540](#)] [[INSPIRE](#)].
- [30] F.A. Dreyer and H. Qu, *Jet tagging in the Lund plane with graph networks*, *JHEP* **03** (2021) 052 [[arXiv:2012.08526](#)] [[INSPIRE](#)].
- [31] E.M. Metodiev and J. Thaler, *Jet topics: disentangling quarks and gluons at colliders*, *Phys. Rev. Lett.* **120** (2018) 241602 [[arXiv:1802.00008](#)] [[INSPIRE](#)].
- [32] J. Brewer, J. Thaler and A.P. Turner, *Data-driven quark and gluon jet modification in heavy-ion collisions*, *Phys. Rev. C* **103** (2021) L021901 [[arXiv:2008.08596](#)] [[INSPIRE](#)].
- [33] F.A. Dreyer, G.P. Salam and G. Soyez, *The Lund jet plane*, *JHEP* **12** (2018) 064 [[arXiv:1807.04758](#)] [[INSPIRE](#)].
- [34] ATLAS collaboration, *Measurement of the Lund jet plane using charged particles in 13 TeV proton-proton collisions with the ATLAS detector*, *Phys. Rev. Lett.* **124** (2020) 222002 [[arXiv:2004.03540](#)] [[INSPIRE](#)].
- [35] ALICE collaboration, *Physics preliminary summary: measurement of the primary Lund plane density in pp collisions at $\sqrt{s} = 13$ TeV with ALICE*, ALICE-PUBLIC-2021-002, CERN, Geneva, Switzerland (2021).
- [36] A. Lifson, G.P. Salam and G. Soyez, *Calculating the primary Lund jet plane density*, *JHEP* **10** (2020) 170 [[arXiv:2007.06578](#)] [[INSPIRE](#)].
- [37] D.E. Soper and M. Spannowsky, *Finding physics signals with shower deconstruction*, *Phys. Rev. D* **84** (2011) 074002 [[arXiv:1102.3480](#)] [[INSPIRE](#)].

- [38] D.E. Soper and M. Spannowsky, *Finding top quarks with shower deconstruction*, *Phys. Rev. D* **87** (2013) 054012 [[arXiv:1211.3140](#)] [[INSPIRE](#)].
- [39] D. Ferreira de Lima, P. Petrov, D. Soper and M. Spannowsky, *Quark-gluon tagging with shower deconstruction: unearthing dark matter and Higgs couplings*, *Phys. Rev. D* **95** (2017) 034001 [[arXiv:1607.06031](#)] [[INSPIRE](#)].
- [40] Y.L. Dokshitzer, G.D. Leder, S. Moretti and B.R. Webber, *Better jet clustering algorithms*, *JHEP* **08** (1997) 001 [[hep-ph/9707323](#)] [[INSPIRE](#)].
- [41] M. Wobisch and T. Wengler, *Hadronization corrections to jet cross-sections in deep inelastic scattering*, in *Workshop on Monte Carlo generators for HERA physics (plenary starting meeting)*, (1998), p. 270 [[hep-ph/9907280](#)] [[INSPIRE](#)].
- [42] M. Cacciari, G.P. Salam and G. Soyez, *The anti- k_t jet clustering algorithm*, *JHEP* **04** (2008) 063 [[arXiv:0802.1189](#)] [[INSPIRE](#)].
- [43] ALICE collaboration, *QCD dynamics studied with jets in ALICE*, in *55th rencontres de Moriond on QCD and high energy interactions*, (2021) [[arXiv:2105.10523](#)] [[INSPIRE](#)].
- [44] A.J. Larkoski, S. Marzani, G. Soyez and J. Thaler, *Soft drop*, *JHEP* **05** (2014) 146 [[arXiv:1402.2657](#)] [[INSPIRE](#)].
- [45] M. Dasgupta, A. Fregoso, S. Marzani and G.P. Salam, *Towards an understanding of jet substructure*, *JHEP* **09** (2013) 029 [[arXiv:1307.0007](#)] [[INSPIRE](#)].
- [46] S. Bright-Thonney and B. Nachman, *Investigating the topology dependence of quark and gluon jets*, *JHEP* **03** (2019) 098 [[arXiv:1810.05653](#)] [[INSPIRE](#)].
- [47] M. Dasgupta and G.P. Salam, *Resummation of nonglobal QCD observables*, *Phys. Lett. B* **512** (2001) 323 [[hep-ph/0104277](#)] [[INSPIRE](#)].
- [48] Y. Hatta and T. Ueda, *Resummation of non-global logarithms at finite N_c* , *Nucl. Phys. B* **874** (2013) 808 [[arXiv:1304.6930](#)] [[INSPIRE](#)].
- [49] K. Hamilton, R. Medves, G.P. Salam, L. Scyboz and G. Soyez, *Colour and logarithmic accuracy in final-state parton showers*, *JHEP* **03** (2021) 041 [[arXiv:2011.10054](#)] [[INSPIRE](#)].
- [50] S. Hochreiter and J. Schmidhuber, *Long short-term memory*, *Neural Comput.* **9** (1997) 1735.
- [51] K. He, X. Zhang, S. Ren and J. Sun, *Delving deep into rectifiers: surpassing human-level performance on ImageNet classification*, [arXiv:1502.01852](#) [[INSPIRE](#)].
- [52] D.P. Kingma and J. Ba, *Adam: a method for stochastic optimization*, 12, 2014 [[arXiv:1412.6980](#)] [[INSPIRE](#)].
- [53] F.A. Dreyer and H. Qu, *fdreyer/LundNet: v1.0.0*, [Zenodo](#), (2021).
- [54] Y. Wang, Y. Sun, Z. Liu, S.E. Sarma, M.M. Bronstein and J.M. Solomon, *Dynamic graph CNN for learning on point clouds*, *ACM Trans. Graph.* **38** (2019) 1.
- [55] S. Ioffe and C. Szegedy, *Batch normalization: accelerating deep network training by reducing internal covariate shift*, in *Proceedings of the 32nd international conference on machine learning*, Lille, France, 7–9 July 2015, *Proc. Machine Learn. Res.* **37** (2015) 448.
- [56] X. Glorot, A. Bordes and Y. Bengio, *Deep sparse rectifier neural networks*, in *Proceedings of the 14th international conference on artificial intelligence and statistics*, Fort Lauderdale, FL, U.S.A., 11–13 April 2011, *Proc. Machine Learn. Res.* **15** (2011) 315.
- [57] K. He, X. Zhang, S. Ren and J. Sun, *Deep residual learning for image recognition*, in *2016 IEEE conference on Computer Vision and Pattern Recognition (CVPR)*, Las Vegas, NV, U.S.A., [IEEE](#), (2016), p. 770 [[arXiv:1512.03385](#)] [[INSPIRE](#)].

- [58] M. Wang et al., *Deep graph library: a graph-centric, highly-performant package for graph neural networks*, [arXiv:1909.01315](#).
- [59] A. Paszke et al., *Pytorch: an imperative style, high-performance deep learning library*, in *Advances in neural information processing systems* 32, H. Wallach et al. eds., [Curran Associates Inc.](#), (2019), p. 8024.
- [60] M. Dasgupta, F. Dreyer, G.P. Salam and G. Soyez, *Small-radius jets to all orders in QCD*, *JHEP* **04** (2015) 039 [[arXiv:1411.5182](#)] [[INSPIRE](#)].
- [61] T. Sjöstrand, S. Mrenna and P.Z. Skands, *PYTHIA 6.4 physics and manual*, *JHEP* **05** (2006) 026 [[hep-ph/0603175](#)] [[INSPIRE](#)].
- [62] T. Sjöstrand et al., *An introduction to PYTHIA 8.2*, *Comput. Phys. Commun.* **191** (2015) 159 [[arXiv:1410.3012](#)] [[INSPIRE](#)].
- [63] P. Skands, S. Carrazza and J. Rojo, *Tuning PYTHIA 8.1: the Monash 2013 tune*, *Eur. Phys. J. C* **74** (2014) 3024 [[arXiv:1404.5630](#)] [[INSPIRE](#)].
- [64] M. Cacciari and G.P. Salam, *Dispelling the N^3 myth for the k_t jet-finder*, *Phys. Lett. B* **641** (2006) 57 [[hep-ph/0512210](#)] [[INSPIRE](#)].
- [65] M. Cacciari, G.P. Salam and G. Soyez, *FastJet user manual*, *Eur. Phys. J. C* **72** (2012) 1896 [[arXiv:1111.6097](#)] [[INSPIRE](#)].
- [66] M. Bahr et al., *HERWIG++ physics and manual*, *Eur. Phys. J. C* **58** (2008) 639 [[arXiv:0803.0883](#)] [[INSPIRE](#)].
- [67] J. Bellm et al., *HERWIG 7.0/HERWIG++ 3.0 release note*, *Eur. Phys. J. C* **76** (2016) 196 [[arXiv:1512.01178](#)] [[INSPIRE](#)].
- [68] S. Frixione and B.R. Webber, *Matching NLO QCD computations and parton shower simulations*, *JHEP* **06** (2002) 029 [[hep-ph/0204244](#)] [[INSPIRE](#)].
- [69] P. Nason, *A new method for combining NLO QCD with shower Monte Carlo algorithms*, *JHEP* **11** (2004) 040 [[hep-ph/0409146](#)] [[INSPIRE](#)].
- [70] S. Frixione, P. Nason and C. Oleari, *Matching NLO QCD computations with parton shower simulations: the POWHEG method*, *JHEP* **11** (2007) 070 [[arXiv:0709.2092](#)] [[INSPIRE](#)].
- [71] S. Alioli, P. Nason, C. Oleari and E. Re, *A general framework for implementing NLO calculations in shower Monte Carlo programs: the POWHEG BOX*, *JHEP* **06** (2010) 043 [[arXiv:1002.2581](#)] [[INSPIRE](#)].
- [72] E. Gerwick, S. Hoeche, S. Marzani and S. Schumann, *Soft evolution of multi-jet final states*, *JHEP* **02** (2015) 106 [[arXiv:1411.7325](#)] [[INSPIRE](#)].
- [73] SHERPA collaboration, *Event generation with Sherpa 2.2*, *SciPost Phys.* **7** (2019) 034 [[arXiv:1905.09127](#)] [[INSPIRE](#)].
- [74] H. Qu and L. Gouskos, *Particle-net*, <https://github.com/hqucms/ParticleNet>, accessed 21 June 2021.
- [75] P. Komiske, E. Metodiev and J. Thaler, *Energy flow*, <https://energyflow.network>, accessed 21 June 2021.
- [76] M. Dasgupta, F.A. Dreyer, K. Hamilton, P.F. Monni, G.P. Salam and G. Soyez, *Parton showers beyond leading logarithmic accuracy*, *Phys. Rev. Lett.* **125** (2020) 052002 [[arXiv:2002.11114](#)] [[INSPIRE](#)].