

The Lund jet plane

Frédéric A. Dreyer,^{a,b} Gavin P. Salam^{c,b,d,1} and Grégory Soyez^e

^aCenter for Theoretical Physics, Massachusetts Institute of Technology,
Cambridge, MA 02139, U.S.A.

^bRudolf Peierls Centre for Theoretical Physics, Clarendon Laboratory,
Parks Road, Oxford OX1 3PU, U.K.

^cCERN, Theoretical Physics Department,
CH-1211 Geneva 23, Switzerland

^dAll Souls College,
Oxford OX1 4AL, U.K.

^eIPhT, Université Paris-Saclay, CNRS UMR 3681,
CEA Saclay, F-91191 Gif-sur-Yvette, France

E-mail: frederic.dreyer@physics.ox.ac.uk,
gavin.salam@physics.ox.ac.uk, gregory.soyez@ipht.fr

ABSTRACT: Lund diagrams, a theoretical representation of the phase space within jets, have long been used in discussing parton showers and resummations. We point out that they can be created for individual jets through repeated Cambridge/Aachen declustering, providing a powerful visual representation of the radiation within any given jet. Concentrating here on the primary Lund plane, we outline some of its analytical properties, highlight its scope for constraining Monte Carlo simulations and comment on its relation with existing observables such as the z_g variable and the iterated soft-drop multiplicity. We then examine its use for boosted electroweak boson tagging at high momenta. It provides good performance when used as an input to machine learning. Much of this performance can be reproduced also within a transparent log-likelihood method, whose underlying assumption is that different regions of the primary Lund plane are largely decorrelated. This suggests a potential for unique insight and experimental validation of the features being used by machine-learning approaches.

KEYWORDS: Jets, QCD Phenomenology

ARXIV EPRINT: [1807.04758](https://arxiv.org/abs/1807.04758)

¹On leave from CNRS, UMR 7589, LPTHE, F-75005, Paris, France.

Contents

1	Introduction	1
2	Lund diagrams and the primary Lund plane	2
2.1	Construction of the primary Lund plane	4
2.2	Averaged Lund plane density and basic analytical properties	5
2.3	Use for measurements and constraints on Monte-Carlo generators	7
2.4	Declustering other jet-algorithm sequences and higher-order effects	8
2.5	Relations with other observables	12
3	Application to boosted-W tagging	12
3.1	Log-likelihood use of Lund Plane	13
3.2	Machine-learning use of Lund plane	16
3.3	Jet-shape discriminant	17
3.4	Simulation, detector effects and reconstruction	18
3.5	Results	19
3.6	Resilience to non-perturbative effects	20
4	Conclusions	23
A	The Lund plane for (C/A-reclustered) anti-k_t jets	25
B	Moderate energy jets and secondary Lund planes	27
C	Detector effects and the subjet-particle rescaling algorithm (SPRA)	30

1 Introduction

Jets, the collimated bunches of hadrons that result from the fragmentation of energetic quarks and gluons, are among the most fascinating objects that are used at colliders. The study of the internal structure of jets has become a prominent area of research at CERN's Large Hadron Collider, both theoretically [1] and experimentally [2]. This is a reflection of its power to probe the Higgs sector of the Standard Model (SM) and to search for physics beyond the Standard Model (BSM), but also of the considerable scope for learning more about the quantum chromodynamics (QCD) associated with the development of jets.

Theoretical and phenomenological work on jet substructure has taken two main directions in the past years. On one hand there has been extensive effort to manually construct observables (see e.g. refs. [3–16]) that can help distinguish between different origins of jets: e.g. those stemming from the hadronic decay of a boosted $W/Z/H$ boson or top quark (signals), versus those from the fragmentation of a quark or gluon (background). That effort has been accompanied by extensive calculations of the properties of those observables within perturbative QCD (e.g. refs. [17–30]) and experimental measurements of their distributions (e.g. refs. [31–35]).

In addition to the study of specific manually crafted observables, several groups have highlighted the power of machine-learning (ML) approaches to exploit jet-substructure information, using a variety of ML architectures. As inputs they have mainly considered discretised images of the particles inside a jet [36–39], clustering histories from sequential recombination jet algorithms [40–42], or a basis of substructure observables [43–45]. The performances that they obtain for signal versus background discrimination are often substantially better than those based on manually constructed observables. This good performance comes, however, at the price of limited clarity as to what jet substructure features are actually being exploited. One consequence is that it is difficult to establish to what extent widely-used modelling tools, e.g. parton-shower Monte Carlo generators and detector simulations, reliably predict those features, an aspect that is critical for the quantitative interpretation of collider searches.¹

The purpose of this article is to introduce a representation of the internal structure of jets that helps bridge the fault-line between manually constructed observables and ML approaches. In particular, we ask whether it is possible to organise the information within a jet such that (a) it can be straightforwardly measured and understood in data (b) it can be manually organised into transparent and physically motivated discrimination observables (without ML) and (c) it can serve as an input to ML for signal/background discrimination, specifically one whose main discriminating characteristics can be clearly identified and understood.

The representation that we use is inspired by Lund diagrams [48], which serve as a theoretical representation of the phase-space within jets and are often used in discussions of Monte Carlo parton shower algorithms and resummation of logarithmically enhanced terms in perturbation theory. In a Lund diagram the available phase-space is mapped to a triangle in a two dimensional (logarithmic) plane that shows the transverse momentum and the angle of any given emission with respect to its emitter. Each given emission creates new phase space (a triangular leaf) for further emissions. One of the key observations of this paper is that Lund diagrams need not merely be a construct for theoretical calculations. They can be constructed for individual jets, essentially by following the clustering tree of the Cambridge/Aachen [49, 50] jet algorithm. The pattern of emissions, notably within the first triangular phase-space region, the primary Lund plane, carries considerable information about the jet.

2 Lund diagrams and the primary Lund plane

To help understand how the primary Lund plane is constructed, figure 1 shows three representations for each of two jets.

The top representation shows the set of particles in the jet, with the direction and length of each line segment schematically representing the direction and scalar momentum of the corresponding particle. The black particle (*a*) is the primary particle, i.e. the one that initiated the jet. Particles (*b*) and (*c*) are emissions inside the jet.

¹Note that some approaches attempt to circumvent the use of modelling tools [46, 47].

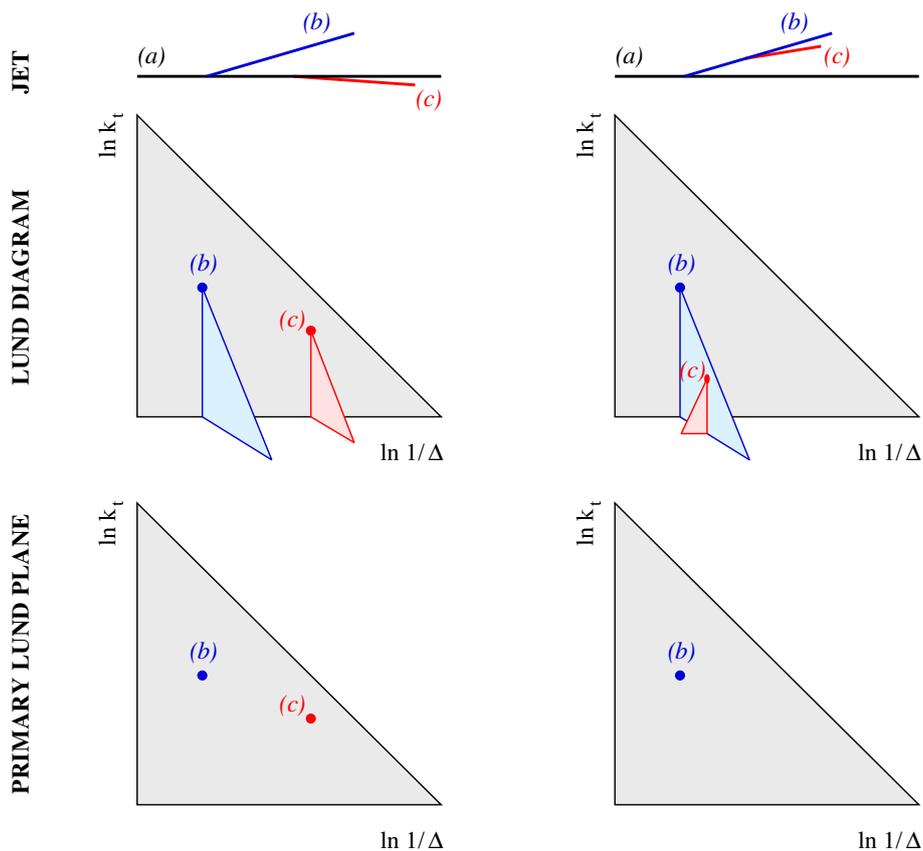


Figure 1. Different representations for two jets. Top: the particles inside the jet. Middle: the full Lund diagram. Bottom: the primary Lund plane. See text for further details.

The middle representation gives the full Lund diagrams for each of the two jets. The phase-space for emission from each particle is represented as a triangle in a $\ln \Delta$ and $\ln k_t$ plane, where Δ and k_t are respectively the angle and transverse momentum of an emission with respect to its emitter. The triangles are colour-coded to match the colours of the particles in the upper row. The black triangle represents the primary phase space, i.e. emission from (a) (our classification of which particle emits which other ones is based on the concept of angular ordering of emissions). Considering the left-hand jet, the blue particle (b) in the jet is represented as a blue point at the appropriate (Δ, k_t) coordinate on the (black) triangle associated with its emitter (a). The blue particle has its own phase-space region, the blue triangle, which is known as a secondary Lund triangle, or “leaf” where the particle could have, but in this case didn’t, emit. Similarly for the red particle, (c), which is also emitted from (a). In contrast, for the right-hand jet, (c) was emitted from (b) and so its point appears on the (secondary) blue triangle associated with particle (b), while its red phase-space triangle emerges as a tertiary triangle, or leaf, off (b)’s triangle.

Finally, the bottom diagram shows the primary Lund plane, which contains just the positions of the emissions from (a), but no information about what further secondary emissions may have been produced. It is this simpler representation that we will use throughout most of the article.

2.1 Construction of the primary Lund plane

Our starting point for constructing the primary Lund plane is to (re-)cluster a jet's constituents with the Cambridge-Aachen (C/A) algorithm [49, 50], which has significant advantages over other members of the generalised- k_t [51] family (see section 2.4).² The C/A algorithm identifies the pair of particles i and j closest in rapidity ($y = \ln \frac{E+p_z}{E-p_z}$, with E and p_z the particle's energy and longitudinal momentum with respect to the colliding beams) and azimuth ϕ , i.e. with the minimal value of $\Delta_{ij}^2 = (y_i - y_j)^2 + (\phi_i - \phi_j)^2$. It then recombines them into a “pseudojet” with momentum $p = p_i + p_j$. This procedure is repeated until all particles (and pseudojets) have been recombined, or are separated by Δ_{ij} larger than some parameter R .

To create a primary Lund plane representation of a jet we then work backwards through the C/A clustering. One starts with the full jet and then proceeds as follows:

1. Decluster the current object to produce two pseudojets, p_a and p_b , labelled such that $p_{ta} > p_{tb}$, where p_{ti} is the transverse momentum of i with respect to the colliding beams. We will consider p_b to be the emission and $p_a + p_b$ to be the emitter. In the limit where p_b carries little momentum relative to p_a , $p_a + p_b$ and p_a can be thought of being the same particle, simply differing through the loss of a small amount of momentum by the radiation of a gluon p_b .
2. Determine a number of variables associated with the declustering, e.g.

$$\Delta \equiv \Delta_{ab}, \quad k_t \equiv p_{tb} \Delta_{ab}, \quad m^2 \equiv (p_a + p_b)^2, \quad (2.1a)$$

$$z \equiv \frac{p_{tb}}{p_{ta} + p_{tb}}, \quad \kappa \equiv z \Delta, \quad \psi \equiv \tan^{-1} \frac{y_b - y_a}{\phi_b - \phi_a}, \quad (2.1b)$$

In the limit $p_{tb} \ll p_{ta}$ and $\Delta \ll 1$, k_t is the transverse momentum of particle b (the emission) relative to its emitter, ψ is an azimuthal angle around the (sub)jet axis, and z is the momentum fraction of the branching. In our default definition of the Lund plane, the coordinates associated with this declustering will be $\ln \Delta$ and $\ln k_t$. One may also, however, make other choices of coordinates, such as for example $\ln \Delta$ and $\ln \kappa$, or $\ln \Delta$ and $\ln k_t / p_{t,\text{jet}}$ (with $p_{t,\text{jet}}$ the jet transverse momentum). We will denote the variables as a tuple $\mathcal{T}^{(i)} = \{k_t^{(i)}, \Delta^{(i)}, \dots\}$ for the i^{th} iteration of this step.

3. Repeat the procedure by going to step 1 for the harder branch, p_a .

This procedure gives an ordered list of tuples of variables

$$\mathcal{L}_{\text{primary}} = \left[\mathcal{T}^{(i)}, \dots, \mathcal{T}^{(n)} \right] \quad (2.2)$$

containing the kinematic variables for each of the primary branchings off the main emitter. The k_t and Δ elements of the tuples (specifically their logarithms) can be interpreted as set of coordinates of points in the Lund plane, corresponding to the full set of primary

²Throughout this paper, we also use the C/A algorithm for the initial jet finding. The case where jets are clustered with the anti- k_t algorithm (and re-clustered with the C/A algorithm) is discussed in appendix A.

branchings, as in the lower row of figure 1. The tuple elements other than k_t and Δ provide complementary information for each point.

One could additionally follow the lower p_t branch at each declustering. This would effectively create secondary, tertiary, etc., Lund planes (or triangles), i.e. one for each emission, giving the full Lund diagram as in the middle row of figure 1. We postpone the study of full Lund diagrams to future work, although a brief discussion of the use of a secondary Lund plane is given in appendix B.

A last point in this subsection is to consider infrared and collinear safety. The full list of tuples produced by the primary declustering procedure is not infrared and collinear safe. For example, if one adds an infinitesimally soft emission, it may lead to an additional primary declustering, and so an extra tuple in the list. However if one considers the subset of tuples in some specified finite region of the Lund plane (for example all tuples with k_t larger than some cut, or all tuples in a given two-dimension bin of k_t and Δ), then one recovers infrared and collinear safety. Specifically, the extra soft emission will not modify that subset, because either it will be declustered as a primary emission that is outside that finite region, or it will be clustered with a harder emission inside that region, but will not modify the kinematic variables of the tuple associated with that harder emission's own primary declustering. In practice only the pattern of declusterings with $k_t \gg \Lambda_{\text{QCD}}$ is amenable to perturbative calculation.

2.2 Averaged Lund plane density and basic analytical properties

The simplest analysis of the Lund plane is to examine the average density of points per jet and per unit area in the $\ln k_t - \ln \Delta$ plane, which we denote

$$\rho(\Delta, k_t) = \frac{1}{N_{\text{jet}}} \frac{dn_{\text{emission}}}{d \ln k_t d \ln 1/\Delta}. \quad (2.3)$$

One can also define a density in terms of dimensionless variables, e.g.

$$\bar{\rho}(\Delta, \kappa) = \frac{1}{N_{\text{jet}}} \frac{dn_{\text{emission}}}{d \ln \kappa d \ln 1/\Delta}. \quad (2.4)$$

The quantity $\rho(\Delta, k_t)$ is represented in figure 2a for a sample of ($C/A, R = 1$) jets with $p_t > 2$ TeV, simulated using the dijet process in Pythia 8.230 [52] with the Monash13 tune [53]. For the case of a quark-initiated jet (about 80% of the jets in the sample figure 2a), to leading order in perturbative QCD and for $\Delta \ll 1$, one expects

$$\rho \simeq \frac{\alpha_s(k_t) C_F}{\pi} \bar{z} (p_{gq}(\bar{z}) + p_{gq}(1 - \bar{z})), \quad \bar{z} = \frac{k_t}{p_{t,\text{jet}} \Delta}, \quad (\Delta \ll 1), \quad (2.5)$$

where $C_F = \frac{4}{3}$, $p_{gq}(z) = \frac{1+(1-z)^2}{z}$ ($0 < z < 1$); \bar{z} is an effective momentum fraction and coincides with z in eq. (2.1) when there is a single emission. For $\bar{z} \ll 1$ the \bar{z} -dependent factor in ρ is equal to 2 and so the density of primary Lund emissions is just proportional to the strong coupling,

$$\rho \simeq \frac{2\alpha_s(k_t) C_F}{\pi}, \quad (\Delta \ll 1, \bar{z} \ll 1). \quad (2.6)$$

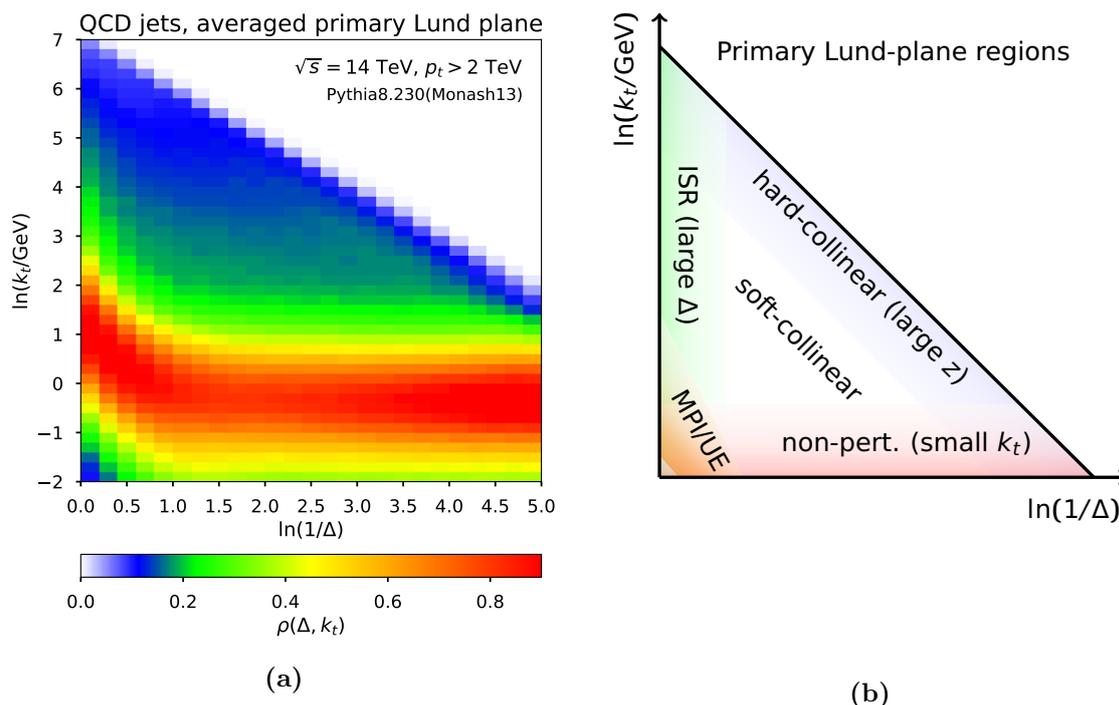


Figure 2. (a) The average primary Lund plane density, ρ , for jets clustered with the C/A algorithm and $R = 1$ having $p_t > 2$ TeV and $|y| < 2.5$, in a simulated QCD dijet sample. (b) Schematic representation of the different regions of the Lund plane.

The upper diagonal edge in the figure is a consequence of the kinematic limit, $k_t < \frac{1}{2} p_{t,\text{jet}} \Delta$. At low scales $\alpha_s(k_t)$ gets large, which accounts for the bright red band around $k_t = 1$ GeV. In this region the Lund plane density is not amenable to perturbative calculation. Equivalently eq. (2.5) receives large corrections from non-perturbative terms proportional to powers of k_t/Λ_{QCD} . At values of $\Delta \sim 1$, initial state radiation (ISR) and multi-parton interactions (MPI/UE) contribute to increasing the density, which is reflected in the contours of constant colour bending upwards to the left. The different regions are outlined schematically in figure 2b.

Beyond leading perturbative order, several further physical effects contribute to the structure of the Lund plane. The upper boundary gets smeared out because of degradation of the leading subjet energy as one declusters the jet.³ The leading subjet can also change flavour as one moves down the clustering tree, in particular when there is an emission close to the upper, kinematic boundary. This can then alter the density of emissions at smaller angles, i.e. subsequent declusterings. The underlying physics of these two effects is closely connected with small- R resummations, cf. refs. [54, 55]. Non-global [56] and clustering [57, 58] logarithms introduce correlations between regions of the Lund plane at

³This smearing does not occur if one examines $\bar{\rho}(\Delta, \kappa)$, from eq. (2.4), since κ is defined in terms of the local z fraction of the emission, which does not depend on earlier splittings at larger angles (while k_t does). However, instead the non-perturbative boundary gets smeared, as does the relation between a given location on the plane and the invariant mass of the pair being declustered.

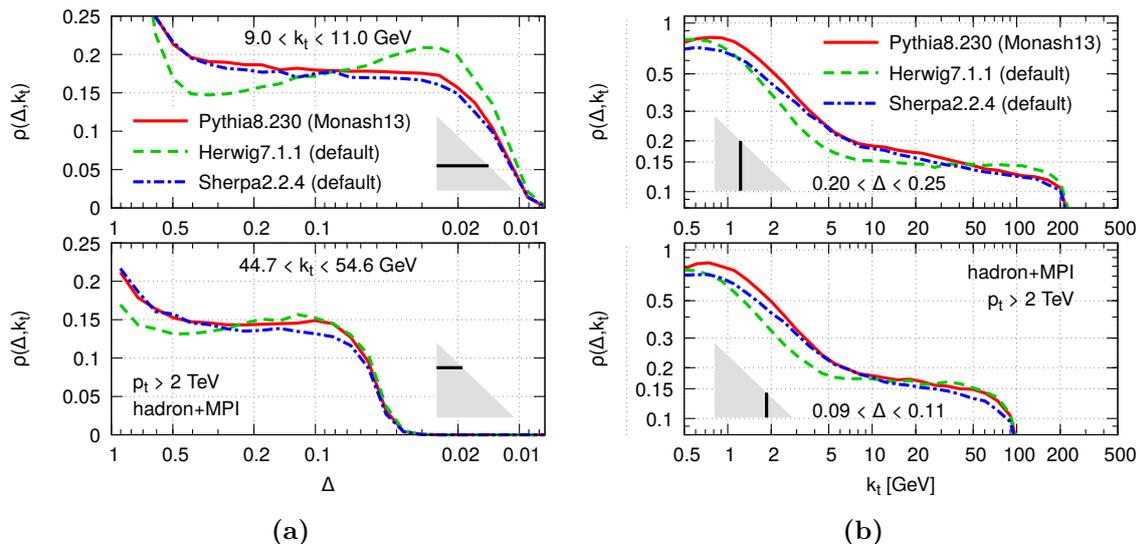


Figure 3. Emission density along slices of the Lund plane, at fixed k_t (top) and Δ (bottom), comparing three event generators.

similar Δ values but different k_t 's. For each effect that introduces a correlation, there is typically also an impact on the average Lund density beyond leading order. We leave the detailed study of these contributions to future work. Note that we expect them to contribute generically at an accuracy $\alpha_s^n L^{n-1}$, where L is some combination of $\ln k_t$ and $\ln \Delta$. This is the same logarithmic order as the running coupling effects that we have explicitly highlighted in eqs. (2.5) and (2.6). Below, in section 2.4, we will see evidence that in the bulk of the Lund plane running effects dominate numerically over these other effects. Finally, one should keep in mind that most of these subleading effects are, at least to some extent, included in modern Monte Carlo programs. (The extent to which there are differences between Monte Carlo programs is discussed next.) In that respect these effects are expected to have only a modest impact on the tagging studies carried out in section 3.

2.3 Use for measurements and constraints on Monte-Carlo generators

The Lund jet plane density ρ in eq. (2.3) can be directly measured experimentally and compared to analytic predictions and parton-shower Monte-Carlo simulations. Here we concentrate on the latter. For such quantitative studies it is convenient to examine slices of the Lund plane density at fixed k_t and fixed Δ . Two of each are shown in figure 3, illustrating the potential of the Lund plane for providing insight into event generators. The figure compares the output of three different generators, Pythia 8.230 (Monash13 tune), Sherpa 2.2.4 [59] and Herwig 7.1.1 [60] (angular-ordered shower). The slices at fixed $\ln k_t$ illustrate a somewhat different trend between the angular-ordered generator and the other two, which both have transverse-momentum ordered showers. The differences span the $\pm 20\text{--}30\%$ range and we have found that they are robust against non-perturbative effects (detector effects are discussed in appendix C). The slices at fixed Δ illustrate the coverage of both the high- k_t , perturbative region, where the density is an infrared and

collinear safe quantity, and the low- k_t , non-perturbative region. In the latter, for k_t below a few GeV, one also sees differences between generators of about 15%. The ability to clearly identify separate perturbative and non-perturbative regions provides a powerful advantage relative to quantities such as jet-shapes that have been measured in the past, and whose distributions tend, to some extent, to mix perturbative and non-perturbative sensitivity.

Returning to more general considerations about the interest of measuring the Lund plane density, an additional remark here is that at low k_t , the Lund plane density could be seen as providing an effective constraint on the strong coupling in the infrared, which one might also be able to relate to the α_0 parameter of refs. [61, 62].

2.4 Declustering other jet-algorithm sequences and higher-order effects

The choice of the C/A algorithm to create the clustering sequence is related both to physical properties of the C/A algorithm and associated higher-order perturbative structures that appear when one calculates ρ and $\bar{\rho}$. To illustrate this, we write $\bar{\rho}$ as an expansion in powers of α_s .

$$\bar{\rho}(\Delta, \kappa) = \sum_{n=1} \left(\frac{\alpha_s}{2\pi}\right)^n \bar{\rho}_n(\Delta, \kappa). \quad (2.7)$$

The first order term, for a quark-induced jet, is given by an expression similar to eq. (2.5),

$$\bar{\rho}_1(\Delta, \kappa) = 2C_F \bar{z} (p_{gq}(\bar{z}) + p_{gq}(1 - \bar{z})), \quad \bar{z} = \frac{\kappa}{\Delta}, \quad \left(\bar{z} < \frac{1}{2}, \Delta \ll 1\right), \quad (2.8)$$

for all algorithms of the generalised- k_t family. For small κ this reduces to $4C_F$.

At higher orders, $n > 1$, one expects that there may be logarithmic enhancements, i.e. terms of the form $\alpha_s^n L^m$ where L may generically be either $\ln 1/\Delta$ or $\ln 1/\kappa$. Ignoring potential subleading- N_C factorisation violation issues, or equivalently super-leading logarithms [63, 64], there are strong reasons to believe that with the C/A algorithm the highest logarithmic enhancement will correspond to $m = n - 1$, i.e. at most a single logarithmic correction factor relative to the leading order result. Such terms will arise from: the running of the coupling, which is naturally single logarithmic; flavour-changing hard-collinear splittings, which have only (single) collinear logarithms; and non-global and clustering logarithms, which have only (single) soft logarithms. In particular the clustering logarithms are single logarithmic because the angular-ordered nature of the algorithm matches the underlying angular ordered pattern of soft-collinear radiation.

In contrast, if one uses the anti- k_t or k_t algorithms, the clustering results in double-logarithmic enhancements, i.e. terms $\alpha_s(\alpha_s L^2)^{n-1}$, or equivalently $m = 2(n - 1)$. We show this explicitly at order α_s^2 . First consider the k_t algorithm, and a configuration with two primary emissions, with $\theta_{1q} \ll \theta_{2q} \ll 1$, $k_{t2} \ll k_{t1}$, which implies $z_2 \ll z_1$. Here θ_{ij} is the angle between particles i and j , and q represents the (leading) quark. As originally pointed out in the article that proposed the Cambridge algorithm [49], when the emissions are on the same side of the quark, then $\theta_{12} < \theta_{2q}$ and so emission 2 clusters with emission 1 rather than with the quark, cf. figure 4a. The resulting pseudojet retains the kinematics of emission 1. With the C/A algorithm, emission 2 would have formed its own independent

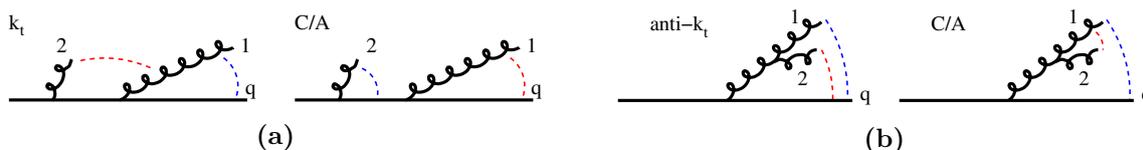


Figure 4. Configurations that lead to terms $\alpha_s^2 L^2$ for the Lund plane density for algorithms other than C/A, showing with a red (blue) dashed line the clustering that occurs first (second). (a) For the k_t -declustered Lund plane: emission 2 clusters with emission 1 rather than with the quark and so fails to appear as part of the primary Lund plane. (b) For the anti- k_t -declustered Lund plane: emission 2 gets declustered as a primary emission, even though physically it belongs to the secondary Lund plane associated with emission 1, resulting in a spurious enhancement of the primary Lund plane density.

primary declustering, but with the k_t algorithm it does not do so.⁴ This leads to a deficit of primary declusterings. For a given choice of emission-2 kinematics, the region of emission-1 kinematics where this occurs is proportional to $\ln^2 \theta_2/\kappa_2$, and results in the following double logarithmic suppression relative to the leading-order result,

$$\begin{aligned} \bar{\rho}_2^{(k_t)}(\Delta, \kappa) &\simeq -(4C_F)^2 \int \frac{d\theta_{2q}}{\theta_{2q}} \int \frac{d\kappa_2}{\kappa_2} \int_{\kappa_2}^{\theta_2} \frac{d\kappa_1}{\kappa_1} \int_{\kappa_1}^{\theta_2} \frac{d\theta_{1q}}{\theta_{1q}} \int_{-\pi/2}^{\pi/2} \frac{d\phi_{12}}{2\pi} \delta\left(\ln \frac{\kappa_2}{\kappa}\right) \delta\left(\ln \frac{\theta_{2q}}{\Delta}\right) \\ &= -4C_F^2 \ln^2 \frac{\Delta}{\kappa} + \mathcal{O}(L). \end{aligned} \quad (2.9)$$

The minus sign accounts for the fact that in this region the primary Lund plane contribution from emission 2 is lost because 2 clusters with 1. The limits for the ϕ_{12} integration account for the fact that both emissions have to be on the same side of the quark in order for 2 to cluster with 1.

For the anti- k_t algorithm, the issue that arises is that secondary splittings can end up being categorised as primary in terms of the (de)clustering sequence. Consider emission 1 with an angle θ_{1q} with respect to the quark and momentum fraction z_1 . It can emit a soft gluon 2 at an angle $\theta_{12} \ll \theta_{1q}$, carrying a momentum fraction z_2 , defined relative to the quark momentum, and satisfying $z_2 \ll z_1$, cf. figure 4b. With the C/A algorithm, the condition on the angles would ensure that emission 2 is always clustered with 1, before 1 clusters with the quark, and hence emission 2 will never on its own be considered as a primary declustering. For the anti- k_t algorithm, emission 2 will cluster with emission 1 only if $\theta_{12}^2/z_1^2 < \theta_{2q}^2 \simeq \theta_{1q}^2$. In the remaining region, $\theta_{12}^2 > z_1^2 \theta_{1q}^2$, emission 2 will be clustered directly with the quark, and hence will be considered as a primary declustering. Fixing $\theta_{1q} \simeq \theta_{2q}$ and z_2 , but integrating over the θ_{12} and z_1 , one finds a double logarithmic enhancement,

$$\begin{aligned} \bar{\rho}_2^{(\text{anti-}k_t)}(\Delta, \kappa) &\simeq 16C_F C_A \int \frac{d\theta_{1q}}{\theta_{1q}} \int_{\kappa}^{\theta_{1q}} \frac{d\kappa_1}{\kappa_1} \int_{\frac{\kappa_1}{\theta_{1q}}}^{\frac{\kappa_1}{z_2}} \frac{dz_2}{z_2} \int_{\kappa_1}^{\theta_{1q}} \frac{d\theta_{12}}{\theta_{12}} \delta\left(\ln \frac{z_2 \theta_{1q}}{\kappa}\right) \delta\left(\ln \frac{\theta_{1q}}{\Delta}\right) \\ &= +8C_F C_A \ln^2 \frac{\Delta}{\kappa} + \mathcal{O}(L). \end{aligned} \quad (2.10)$$

⁴This was part of the motivation in ref. [49] for inventing the Cambridge algorithm.

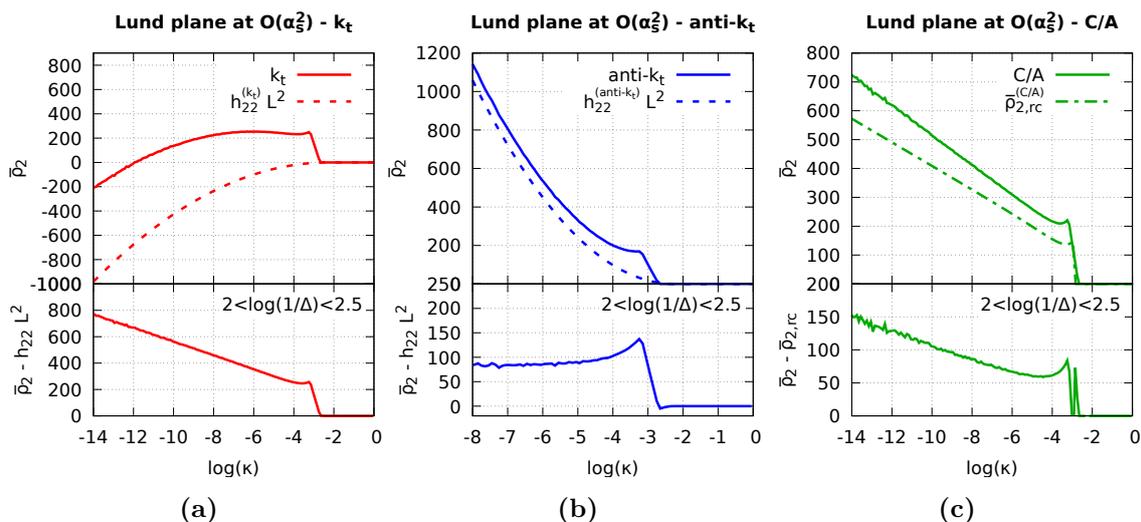


Figure 5. Evaluations with Event2 of the second-order contribution to the Lund plane, in a bin of $\ln 1/\Delta$, as a function of κ , for (de)clustering sequences obtained with the k_t , anti- k_t and C/A jet algorithms. In (a) and (b) the dashed line corresponds to the analytic expectations, eqs. (2.9) and (2.10) for clustering-induced double-logarithms in the k_t and anti- k_t algorithms. In (c), for the C/A algorithm, which is seen here to be free of double logarithms, the dot-dashed line corresponds to the (single-logarithmic) running coupling correction, eq. (2.11), illustrating that it dominates the second-order correction.

In setting the lower limit of the θ_{12} integral, we have made use of the condition $\theta_{12} > z_1 \theta_{1q} = \kappa_1$. The upper bound on the z_2 integration comes from the constraint $z_2 < z_1 = \kappa_1/\theta_{1q}$ and, for physical values of κ and Δ , the solution of the δ -function constraint on z_2 is always below that bound.

To verify our calculations we have used the Event2 program [65, 66] to evaluate the exact result for $\bar{\rho}_2(\Delta, \kappa)$ using e^+e^- versions of the k_t , anti- k_t and C/A algorithms.⁵ Figure 5, in the upper panels, shows $\bar{\rho}_2$ for the k_t , anti- k_t and C/A algorithms, compared in the first two cases also to the double logarithmic calculations of eqs. (2.9) and (2.10). The quantity h_{22} denotes the coefficient of $L^2 = \ln^2 \frac{\Delta}{\kappa}$ in those equations. The lower panels for the k_t and anti- k_t algorithms show the result after subtraction of the double-logarithmic terms. One sees clearly that at most single logarithmic terms remain. Figure 5 therefore validates our calculation of the double logarithmic coefficients for the k_t and anti- k_t algorithms as well as our prediction of an absence of double logarithmic corrections for the C/A algorithm.

One might ask why the presence or absence of double logarithmic corrections should have any impact on the choice of reclustering algorithm. The answer lies in the origin of

⁵Specifically, the Event2 program generates $e^+e^- \rightarrow$ jets events. We used it in 3-jet NLO mode. We cluster each event with the e^+e^- k_t (Durham) algorithm [67] so as to obtain exactly two jets, and then recluster each of those jets with the e^+e^- version of the k_t , C/A or anti- k_t algorithm, as defined by FastJet’s `ee_genkt_algorithm` [68] with $p = 1, 0, -1$ respectively. The resulting sequence is then used to determine the Lund-plane density, following the steps of section 2.1 (but not first reclustering with C/A). For an ij declustering we then define $\ln 1/\Delta = -\ln \tan \theta_{ij}/2$ and $\kappa = \min(E_i, E_j)/(E_i + E_j)\Delta$. This definition ensures that eq. (2.8) remains valid even for Δ of order 1.

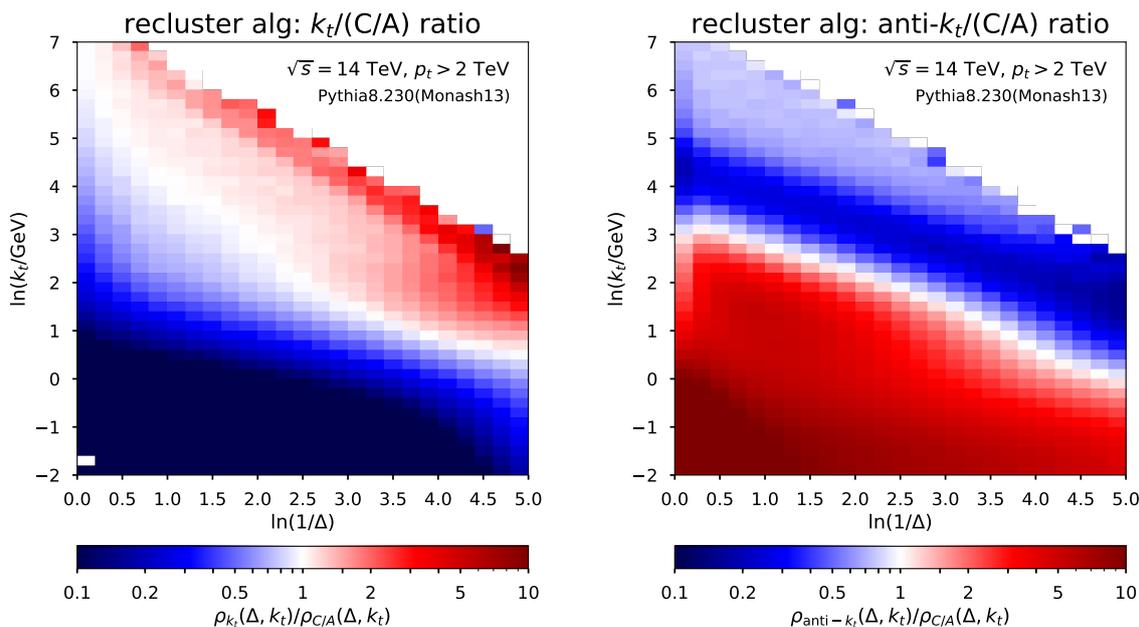


Figure 6. The $\rho(\Delta, k_t)$ results as obtained with k_t (left) and anti- k_t (right) declustering, normalised to the result for C/A declustering.

those double logarithms: for both k_t and anti- k_t , they come about because the reclustering correlates a given part of the Lund plane with a large (double-logarithmic) other region of the Lund plane, e.g. in the case of the k_t algorithm anywhere with smaller Δ and larger k_t .

In contrast, with the C/A algorithm, the correlations are only within single-logarithmic regions, for example, correlations between a given point and the region of similar Δ , or between a given point and emissions in the band along the hard collinear boundary. The more limited correlations for C/A reclustering should make resummation calculations for the Lund plane more straightforward. They are also expected to result in a cleaner input for tagging applications, and to simplify the interpretations of comparisons between Monte Carlo simulations and data. Note that the single logarithmic slope observed in the case of the C/A algorithm also contains running coupling (rc) corrections, which have the form

$$\bar{\rho}_{2,rc}^{(C/A)}(\Delta, \kappa) = \bar{\rho}_1(\Delta, \kappa) 4\pi b_0 \ln \frac{1}{\kappa} + \mathcal{O}(1), \quad (2.11)$$

where $b_0 = \frac{11C_A - 2n_f}{12\pi}$. One sees from figure 5c that this contribution accounts for the bulk of the single-logarithmic slope seen at second order for the C/A algorithm. The residual single-logarithmic corrections from clustering effects are therefore a small component of the overall single-logarithmic contributions.

The above calculations of double logarithmic coefficients may appear to be somewhat abstract. Additionally therefore, it can be instructive to examine how the use of the k_t and anti- k_t algorithms for (de)clustering modifies the structure of the Lund plane in Pythia simulations. Figure 6 show the ratio of the averaged primary Lund plane density as obtained with the k_t and anti- k_t (de)clustering to that with our default C/A choice.

One sees large modifications, especially at lower k_t values, where the double logarithmic effects are largest. In those regions, the modifications are in line with the expectations from the calculations above: a suppression for the k_t algorithm and a (larger) enhancement for the anti- k_t algorithm. There are also features that go beyond our second-order double-logarithmic calculation. For example for the k_t algorithm there is a strong enhancement around $\ln 1/\Delta = 5$ and $\ln k_t/\text{GeV} = 2$. For the anti- k_t algorithm there is a line of deficit sloping downwards from $\ln 1/\Delta = 0$, $\ln k_t/\text{GeV} \simeq 4.5$. We do not have systematic analytical explanations for these features, and their presence provides further reasons for choosing the C/A algorithm for the declustering.

2.5 Relations with other observables

The declustering sequence that produces the primary Lund plane is closely connected with a range of other jet observables. The simplest case is that of the average multiplicity of iterated soft-drop steps [69], N_{SD} , which relates to the averaged $\bar{\rho}$ density through

$$\langle N_{\text{SD}} \rangle = \int_0^\infty \frac{d\Delta'}{\Delta'} \int \frac{d\kappa}{\kappa} \bar{\rho}(\Delta', \kappa) \left[\Theta\left(\frac{\Delta'}{2} - \kappa\right) - \Theta\left(\kappa - z_{\text{cut}}(\Delta')^{1+\beta}\right) \right]. \quad (2.12)$$

This relation applies to soft drop with a generic value of β [70], assuming a jet radius of 1 for simplicity. This is an exact relation, and it holds because the iterated soft drop procedure simply follows the same set of declustering steps as the primary Lund plane and counts those that satisfy the kinematic condition that is represented in the second of the Θ -functions in eq. (2.12). The first Θ -function just represents the kinematic boundary induced by the condition $\kappa < \max(z, 1-z)\Delta$. The counting for N_{SD} is inclusive over all primary splittings and so the average over events that produces $\bar{\rho}$, also gives the average number of iterated soft-drop steps, $\langle N_{\text{SD}} \rangle$.

Further exact relations exist between various soft-drop observables and the tuple of declustering variables defined in section 2.1. For example the soft-drop mass and z_g [71] variables are given by $m^{(i)}$ and $z^{(i)}$ (cf. eq. (2.1)) from the first of the entries in the primary declustering sequence, $\mathcal{L}_{\text{primary}}$ of eq. (2.2), that satisfies $z^{(i)} \geq z_{\text{cut}}(\Delta^{(i)})^\beta$. This is because ignoring the earlier declusterings with $z^{(i)} < z_{\text{cut}}(\Delta^{(i)})^\beta$ is functionally identical to the procedure of discarding (i.e. grooming away) the softer branch in the soft-drop procedure.⁶

3 Application to boosted- W tagging

We now turn to the potential of the Lund plane for identifying hadronically decaying boosted electroweak bosons, concentrating specifically on the example of W identification. Figure 7 shows the averaged (primary) Lund plane for hadronic W decays ($p_t > 2$ TeV), to be compared to figure 2 for dijets. Two main differences are clearly visible to the human eye. One is the diagonally oriented patch in the W case, around $\ln 1/\Delta = 2.5$ and $\ln k_t/\text{GeV} = 4$, which is connected with the fixed-mass two-pronged structure of the

⁶In v1 of the arXiv version of this paper, we had stated connections between $\bar{\rho}(\Delta, \kappa)$ and the distributions of jet observables such as the soft-drop z_g and the jet broadening. Those statements were correct only in a context with just primary soft branchings and so not hold in full QCD.

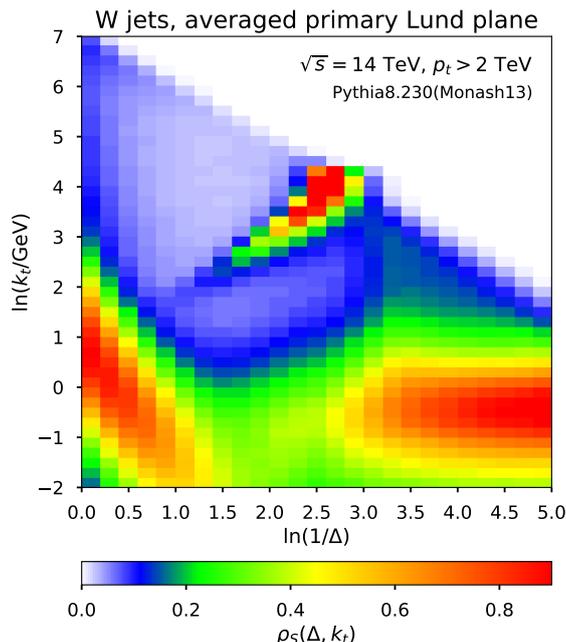


Figure 7. Lund-plane emission density, $\rho_S(\Delta, k_t)$, for hadronically decaying boosted W bosons, in WW events, using the same jet-clustering and selection as in figure 2.

W : lines of constant mass in the Lund plane are up-right going diagonals. The other important feature in the W case is the considerable depletion of emissions in the upper-left region and below the W -mass structure. The depletion is principally a consequence of the colour-singlet nature of the W . It extends also into the non-perturbative region (but not substantially affecting structures of the Lund plane associated with the underlying event).

We investigate two broad approaches to making use of the information in the Lund plane. One is a log-likelihood type approach, while the other will be to use machine learning.

3.1 Log-likelihood use of Lund Plane

The log-likelihood approach uses two main inputs: the first requires the identification of the “leading” emission, (ℓ) , which in the W case is likely to be associated with the two-prong decay. We take this leading emission to be the first emission in the Lund declustering sequence that satisfies $z > z_{\text{cut}}$ with $z_{\text{cut}} = 0.025$, which corresponds to the emission that would be selected by the mMDT tagger [19] with the same z_{cut} or equivalently by the Soft-Drop (SD) procedure [70] with $\beta = 0$ and that z_{cut} . We define a \mathcal{L}_ℓ log likelihood function

$$\mathcal{L}_\ell(m^{(\ell)}, z^{(\ell)}) = \ln \left(\frac{1}{N_S} \frac{dN_S}{dm^{(\ell)} dz^{(\ell)}} \bigg/ \frac{1}{N_B} \frac{dN_B}{dm^{(\ell)} dz^{(\ell)}} \right), \quad (3.1)$$

using the ratio of $dN_X/dm^{(\ell)} dz^{(\ell)}$ ($X = S, B$), the differential distribution in the mass and z variables of the leading emission $(m^{(\ell)}, z^{(\ell)})$ for a simulated signal sample S (W bosons) with N_S jets, and the analogous quantity for a background (QCD dijet) sample B . In practice we bin logarithmically in $m^{(\ell)}$ and $z^{(\ell)}$ to construct a discretised approximation to $\mathcal{L}_\ell(m^{(\ell)}, z^{(\ell)})$.

The second likelihood input is designed to bring sensitivity to the pattern of non-leading ($n\ell$) emissions, i.e. the pattern of additional radiation, within the primary Lund plane, that decorates the basic two-prong structure. It involves a function

$$\mathcal{L}_{n\ell}(\Delta, k_t; \Delta^{(\ell)}) = \ln \left(\rho_S^{(n\ell)} / \rho_B^{(n\ell)} \right), \quad (3.2)$$

where $\rho_X^{(n\ell)}$ is determined just over the non-leading emissions,

$$\rho_X^{(n\ell)}(\Delta, k_t; \Delta^{(\ell)}) = \frac{dn_{\text{emission},X}^{(n\ell)}}{d \ln k_t d \ln 1/\Delta d \ln \Delta^{(\ell)}} \bigg/ \frac{dN_X}{d \ln \Delta^{(\ell)}}, \quad (3.3)$$

as a function of the angle $\Delta^{(\ell)}$ of the leading emission, with $X = S, B$ corresponding either to the W signal ($X = S$) or to the QCD background sample ($X = B$). Our overall log-likelihood signal-background discriminator for a given jet is then given by

$$\mathcal{L}_{\text{tot}} = \mathcal{L}_\ell(m^{(\ell)}, z^{(\ell)}) + \sum_{i \neq \ell} \mathcal{L}_{n\ell}(\Delta^{(i)}, k_t^{(i)}; \Delta^{(\ell)}) + \mathcal{N}(\Delta^{(\ell)}), \quad (3.4)$$

where the normalisation term \mathcal{N} is

$$\mathcal{N}(\Delta^{(\ell)}) = - \int d \ln \Delta d \ln k_t \left(\rho_S^{(n\ell)} - \rho_B^{(n\ell)} \right), \quad (3.5)$$

up to an overall constant. In the sum over non-leading emissions, $i \neq \ell$ in eq. (3.4), each non-leading emission i contributes information (through the $\mathcal{L}_{n\ell}(\Delta^{(i)}, k_t^{(i)}; \Delta^{(\ell)})$ term) about whether its corresponding region of the Lund plane tends to be more populated by signal or background emissions. The normalisation term \mathcal{N} accounts for the average difference in the number of non-leading emissions between signal and background jets.

It is instructive to think about the conditions under which eq. (3.4) would be the optimal discriminator that can be constructed from the sequence in primary Lund-plane declusterings: (1) the identification of the leading emission associated with the W 's two-prong structure should be correct; (2) non-leading emissions in the Lund plane should effectively be independent of each other, which is the basis of the sum over $i \neq \ell$ in eq. (3.4); (3) that pattern of independent emission should depend on $\Delta^{(\ell)}$ but not on $m^{(\ell)}$. Each of these approximations has its imperfections, but none is expected to be particularly badly violated.

To help illustrate how the log-likelihood approach works in practice, we show the leading-emission distribution density, $\frac{1}{N_X} \frac{dN_X}{d \ln m^{(\ell)} d \ln z^{(\ell)}}$, for background ($X = B$) and signal ($X = S$) jets in figure 8. The background is diffuse, while the signal is peaked around $m^{(\ell)} = M_W$ and concentrated at larger $z^{(\ell)}$ values, as one would expect. The non-leading emission density, $\rho_B^{(n\ell)}$, is shown for background (dijets) in figure 9a, for jets where the leading emission has $1.5 < \ln 1/\Delta^{(\ell)} < 2$ and $z^{(\ell)} > z_{\text{cut}}$ (roughly $\ln k_t/\text{GeV} > 2.5$). For similar rapidities and lower k_t 's there is a modest depletion in the number of emissions. This is a partial shadow cast by the leading emission: non-leading emissions with similar Δ and ψ to the leading emission will be clustered with it. The other main feature of note in figure 9a is the empty area in the upper-left region of the plot: given that the emission

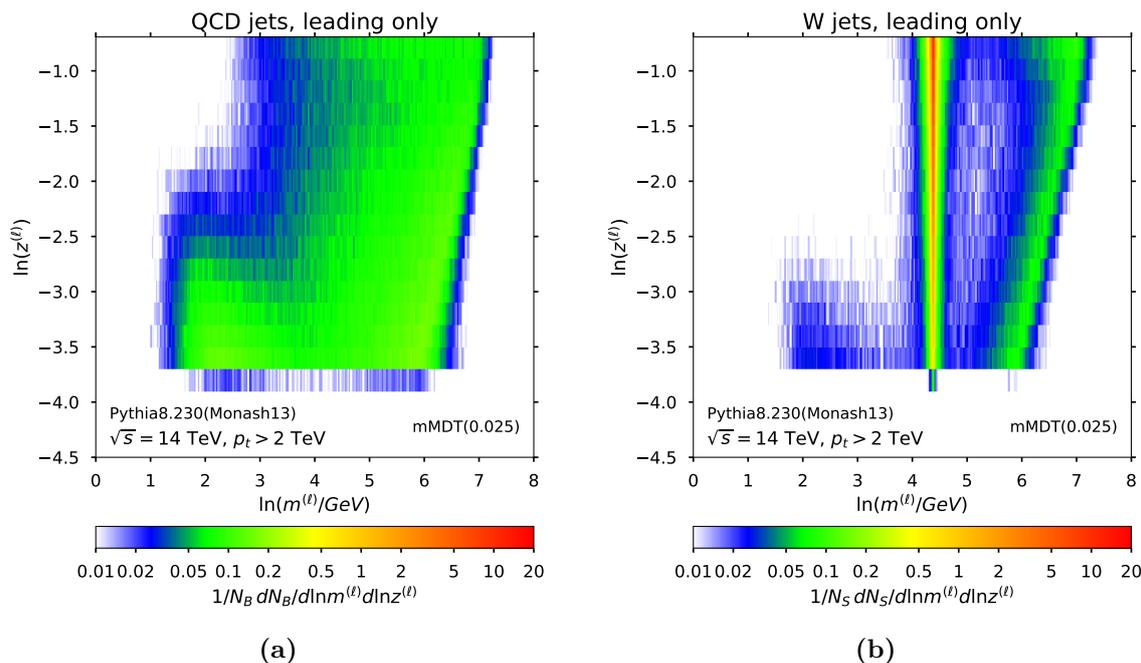


Figure 8. Distribution of the leading emission $\frac{1}{N_X} \frac{dN_X}{d \ln m^{(\ell)} d \ln z^{(\ell)}}$ for (a) background QCD jets ($X = B$) and (b) signal W jets ($X = S$). Note the mass peak around $\ln m^{(\ell)} = \ln M_W$ in the case of W jets and the enhancement at larger values of $\ln z^{(\ell)}$.

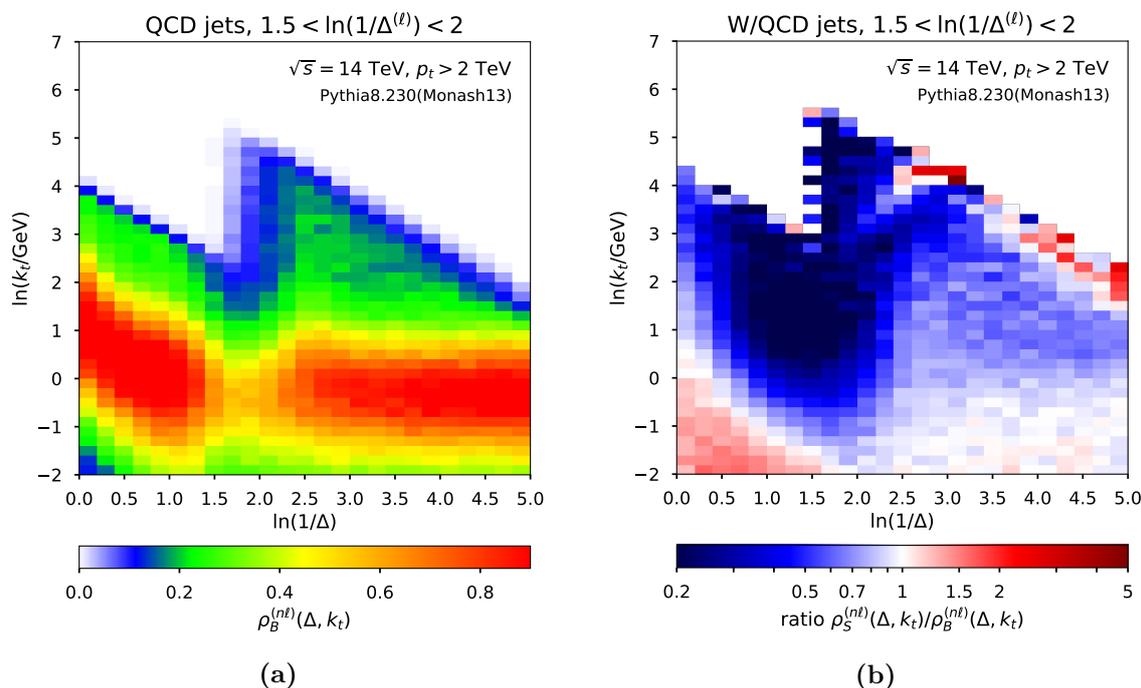


Figure 9. (a) Non-leading Lund-plane density for dijets, $\rho_B^{(nl)}(\Delta, k_t; \Delta^{(\ell)})$, eq. (3.3), i.e. after removing the leading emission, specifically the first emission that passes an mMDT selection procedure. (b) Ratio of the W to QCD Lund-plane $\rho_{S,B}^{(nl)}(\Delta, k_t; \Delta^{(\ell)})$ densities, whose logarithm is equal to the likelihood function \mathcal{L}_{nl} of eq. (3.2). Both plots correspond to the choice $1.5 < \ln 1/\Delta^{(\ell)} < 2$.

classified as leading had $1.5 < \ln 1/\Delta^{(\ell)} < 2$, there cannot have been any emissions with a small Δ and $z > z_{\text{cut}}$.

Figure 9b shows the $\mathcal{L}_{n\ell}$ likelihood function. Most of its discriminating power will come from the extensive dark blue region. For each emission that is present in this region, one gets a negative $\mathcal{L}_{n\ell}$ contribution to \mathcal{L}_{tot} , which drives \mathcal{L}_{tot} to be more background like (i.e. negative). Instead if there are few emissions, the positive contribution of the \mathcal{N} term results in a more signal-like (positive) final \mathcal{L}_{tot} value. Note that the dark blue region of figure 9b stretches down to k_t 's below a few GeV, and since that region tends to contain a significant number of emissions in the background case (cf. figure 9a), one expects that some of the sensitivity in W v. QCD discrimination will come from low- k_t non-perturbative effects. This highlights the importance of direct experimental measurements of such regions.⁷ One can explicitly check the influence of the low- k_t region on the tagging performance by imposing a minimum k_t cut in the construction of the Lund plane. This is discussed further in section 3.5 below.

In our practical implementation of the log-likelihood approach, we will use $\ln m$ bins of size 0.025, $\ln z$ bins of size 0.2 for \mathcal{L}_ℓ ; for $\mathcal{L}_{n\ell}$ we will take bins in $\ln k_t$ and $\ln \Delta$ of 0.2 and bins in $\ln \Delta^{(\ell)}$ of 0.5. The likelihood functions will be calculated using 500,000 simulated signal and background jets, while performance will be evaluated on an independent sample of 200,000 signal and background jets.

3.2 Machine-learning use of Lund plane

Our second approach to using the Lund-plane information for W tagging is to provide it as an input to a variety of machine learning (ML) methods.

The input can be provided in the form of a sequence of $\{\ln 1/\Delta, \ln k_t\}$ pairs; we use this kind of sequence with dense (DNN) and Long Short-Term Memory (LSTM) neural networks [73].⁸ In practice the sequence is zero-padded to form a 60×2 dimensional matrix. The DNN consists of four layers of size 200 with ReLU activation, and a final two-dimensional layer with softmax activation. For the LSTM network, we use a cell with 128-dimensional output connected to a dropout layer with rate 20%, with a final dense layer of dimension two and softmax activation.⁹ In addition to the $\ln 1/\Delta$ and $\ln k_t$ variables, one could add variables such as ψ , the logarithm of the subjet mass ($\ln m$), or the particle multiplicity in the subjet. With realistic experimental resolutions, we have found that they bring a small additional gain in background rejection, in the 10 – 25% range.¹⁰ Keeping in mind that the variability in performance associated with different training choices is also in

⁷It also highlights an important difference between our log-likelihood approach and the shower-deconstruction approach [72]. Both are partially analytical multi-variate log-likelihood approaches. Shower-deconstruction exploits far more detailed information on correlations than our log-likelihood method can. However since the underlying shower-deconstruction likelihoods must be calculated perturbatively, one loses access to the substantial information that is contained in the non-perturbative region.

⁸LSTMs for jet substructure were investigated recently also in ref. [41].

⁹Similar performance can also be achieved using a Gated Recurrent Unit [74].

¹⁰In contrast, at truth particle level, the gain is much more significant, of the order of a factor of two, mostly from the mass information.

the 10% range, we chose not to include these extra variables for our final results. However they could be further investigated in future work.

Alternatively one can create a 2-dimensional Lund image for each event (in which only a few pixels are turned on) and provide it as an input to a convolutional neural network (CNN), where additional information such as the azimuthal angle can be encoded through the pixel intensity by adding new channels. Each jet is represented as a 50×50 pixel image. These images are used to train a neural network consisting of three two-dimensional convolutional layers with ReLU activation and 128 output filters each, which are each connected to a max pooling layer and a spatial dropout layer with rate 5%. The first convolution window is of size 10×10 pixels, with the following two layers having windows of size 4×4 . The last convolutional layer is connected to a dense network with 256 neurons and another dropout layer leading to a final two-dimensional output layer with softmax activation. This network can also be trained on jet images, where each pixel corresponds to a bin in (y, ϕ) -space around the jet axis. In this case the pixel intensity is given by the normalised scalar p_t sum of particles within that phase space region.

Our machine learning is implemented in `Keras` 2.0.8 [75], using `TensorFlow` 1.2.1 [76] as the backend. All model weights are initialised with a He uniform variance scaling initialiser [77], and each training is performed using a batch size of 128, with Adam as optimisation algorithm [78] and a categorical cross-entropy loss function. The parameters for the machine learning are similar, where relevant, to those in refs. [37, 38], though we used a greater number of pixels for the images and larger networks. Training is carried out on a sample of 500 000 jets for each of the W and background samples. During each epoch 80% of the sample is used for adapting the network weights. At the end of each epoch performance is tested on a validation sample consisting of 10% of the events (that were not used during the epoch's training). If that performance has not improved over the past 4 epochs training halts. The maximum number of epochs is 15 and training typically halts at epoch 10–15. The final performance of the network is then evaluated using a further independent 10% of the sample.

Other recent work has also made use of declustering sequences with machine learning tools. Ref. [40] has used recursive neural networks on the complete declustering tree (with various clustering algorithms), using the momenta of the subjects at each stage as inputs. Ref. [41], has used the anti- k_t clustering sequence as a way of ordering all constituent particle momenta, which are then provided in that order to an LSTM. The Lund plane will not necessarily lead to more powerful discriminants than these approaches, however it offers a degree of physical insight and a scope for direct measurements of average densities and correlations. Both of these aspects are potentially valuable for understanding and correctly calibrating machine learning approaches to jet tagging. The LSTM that we have used with the Lund plane is also arguably more straightforward to use with standard machine-learning tools such as `Keras` than the RNN used in ref. [40].

3.3 Jet-shape discriminator

In addition to the log-likelihood and machine-learning based approaches, we will also include a comparison with an optimised choice of jet-shape discriminator. For this purpose,

we apply the SoftDrop algorithm ($\beta = 2$, $z_{\text{cut}} = 0.05$) and use the resulting groomed jet to calculate both the jet mass and the D_2 observable [13] with $\beta = 2$ (itself very similar to C_2 [11] for any given jet mass). The jet mass and D_2 are then given as an input to a boosted-decision tree (BDT) in the TMVA [79] package. This pair of observables, used with just a mass cut, not a BDT,¹¹ was found to be close to optimal in terms of background rejection among a comparison of 88 shape-mass combinations in the recent Les Houches (LH) study [80] (figure III.29), for a fixed signal efficiency of 0.4. Its performance is substantially above that of the default ATLAS and CMS jet-shape discriminant choices. We refer to it as $D_2^{\text{[loose]}}$.¹²

We could also have chosen the dichroic D_2 variable, used as a benchmark in [80], whose performance is only slightly worse but is more resilient against non-perturbative effects. We will return to the question of resilience below in section 3.5.

3.4 Simulation, detector effects and reconstruction

In evaluating the performances of different methods, an important consideration concerns the inclusion of realistic experimental resolutions. This is relevant both for reconstructing the W mass and as regards the radiation part in the rest of the Lund plane. The baseline that we adopt for our comparisons is to use the Delphes [81] fast detector simulation, version 3.4.1, with the `delphes_card_CMS_NoFastJet.tcl` card to simulate both detector effects and particle flow reconstruction. The particle flow outputs have artefacts on angular scales associated with the hadronic and electromagnetic calorimeter, and these have an adverse effect on performance, both in terms of mass resolution and availability of Lund plane information. Accordingly we use a “subjet-particle rescaling algorithm” (SPRA1) at the level of small-radius ($R_h = 0.12$) subjets to retain angular information from electromagnetic calorimeter deposits and charged tracks at small angular scales while retaining full hadronic calorimetry energy information at larger scales. The SPRA algorithm is closely related to earlier methods that proposed jet-wide, subjet and hadronic calorimeter rescaling of charged tracks or electromagnetic calorimeter deposits [82–89]. The details of the Delphes particle flow effect on the Lund plane, of the SPRA algorithm and of the improvements it brings are given in appendix C.

Our results will use simulated dijet events as the background and WW events as the signal, selecting jets, clustered with the C/A algorithm with radius $R = 1$ using `FastJet3.3.0` [68], with $p_t > 2$ TeV and $|y| < 2.5$ at a centre-of-mass energy of $\sqrt{s} = 14$ TeV. We use Pythia 8.230 with its Monash 13 tune to generate the events. The samples coincide with those used in the recent Les Houches study [80]. We choose to concentrate on a high- p_t sample for two reasons: (1) the LHC is increasingly focusing on this region and (2) one expects that high- p_t jets contain the most information, because it is at high- p_t

¹¹The Lund-likelihood approach can be sensibly used with a fixed mass window, simply by discretising the \mathcal{L}_ℓ likelihood ratio to have a single mass bin. However it is not so straightforward to force machine-learning methods to use a fixed mass window. Therefore to obtain meaningful comparisons across all types of methods we must use the D_2 variable in combination with the full mass information.

¹²In the LH report, this observable is denoted as $D_2^{(2)}[l \otimes \frac{l}{l}]$.

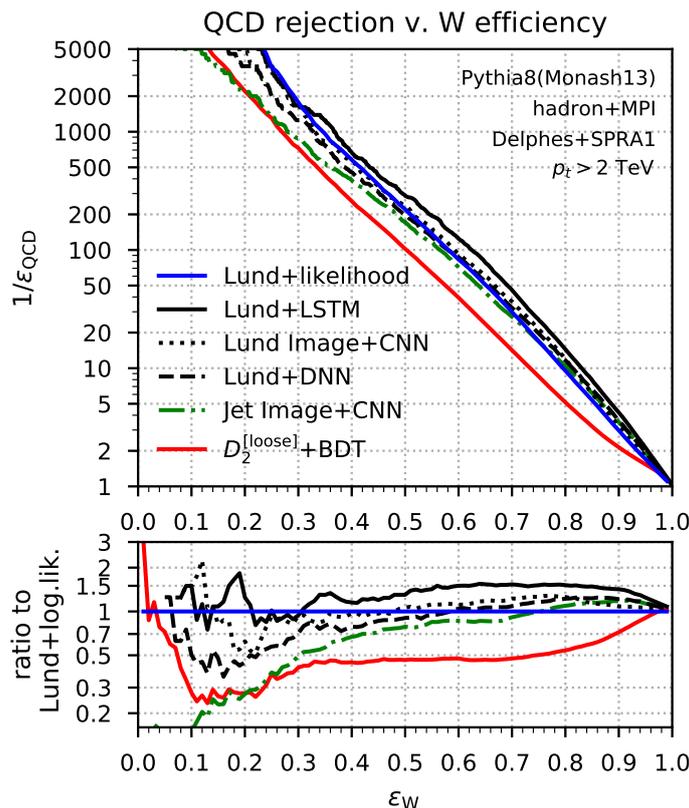


Figure 10. Background rejection ($1/\epsilon_{\text{QCD}}$) versus signal efficiency (ϵ_W), per jet, for different W -tagging methods. The lower panel shows the ratio to the Lund+likelihood method.

that the colour-singlet nature of the boson has the most impact on the radiation pattern. Results at lower p_t are given in appendix B.

Note that we have not included pileup in our simulations. We therefore work within the assumption that methods for pileup mitigation such as SoftKiller [90], PUPPI [91], constituent subtractor [92] or machine learning approaches [93] can successfully remove the contamination in the regions that are critical for discrimination. It is also possible to use area subtraction [94, 95], given that at each stage of the declustering one has (sub)jets with a well-defined area. However, area subtraction is more likely to be susceptible to fluctuations than other methods. A further possibility is to supplement pileup mitigation with methods such as filtering [3], trimming [96] or recursive soft-drop [97], applied only to larger subjets, in order to reduce their contamination from pileup.

3.5 Results

Results for the performance of the different tagging methods of sections 3.1–3.3 are shown in figure 10. The upper panel shows the background rejection factor for each method as a function of signal tagging efficiency. The lower panel shows the ratio of that rejection factor to the one obtained for the Lund-likelihood method (the blue line). Four of the methods are based on machine learning, as discussed in section 3.2: three of them use Lund plane inputs (Lund-image (with CNN), Lund+DNN, Lund+LSTM), while the other

uses a normal jet image (with CNN). Finally the plot also includes the jet-shape plus mass approach, labelled $D_2^{\text{[loose]}}$.

Overall, the LSTM approach with Lund inputs performs best across nearly the full range of signal efficiencies. Taking $\epsilon_W = 0.4$ as a reference, the Lund-likelihood is within a factor of 0.7–0.8 of the LSTM performance, while the other machine learning methods are slightly worse than the Lund-likelihood method. The D_2 -based shape discrimination is a factor of two worse than the Lund-likelihood method.¹³ At higher (lower) signal efficiencies, the machine learning approaches appear to perform relatively slightly better (worse).

The pattern of performances is fairly insensitive to the details both of the Lund-likelihood procedure and the machine learning. For example, in the Lund-likelihood approach, using the subjet that gives a mass closest to the W mass, rather than the first one to pass the mMDT z_{cut} condition, affects signal rejection performance only at the $\sim 10\%$ level (making it better at high efficiencies, worse at lower efficiencies). Similarly, as mentioned in section 3.2, adding mass and azimuth (ψ) information to the LSTM has only a modest 10–25% effect after accounting for detector effects (after detector simulation, this gain is present only when one uses SPRA), which does not appear to be particularly significant relative to other training uncertainties.

We also note that the ROC curves becomes noisy at small signal efficiencies. This can at least in part be attributed to statistical uncertainties associated with the finite size of our training/testing samples, in particular when estimating the background rejection factor, where only a small fraction of the events pass the tagger. For example, the Lund-likelihood method uses a sample of 200 000 events for testing. This corresponds to a statistical uncertainty of about 30% at $\epsilon_W = 0.2$. Correspondingly, the LSTM uses 50 000 events for the testing phase, meaning an expected statistical uncertainty that is twice as large (given similar background rejection).

3.6 Resilience to non-perturbative effects

One can argue that performance is not the only feature one may request from a boosted object tagger. In particular, one may require that the tagger remain relatively insensitive to model-dependent non-perturbative effects. Such insensitivity could translate into a reduced uncertainty on the determination of the tagger’s signal efficiency and background rejection rates. It could also allow for the possibility of understanding the tagger’s behaviour with first-principles perturbative QCD calculations.

To carry out studies of sensitivity to non-perturbative effects, we will compare performance between parton and hadron level. Parton-level results cannot be sensibly passed through a detector simulation, so the study must be carried out with actual particles (i.e. partons or hadrons). However, as we discussed in section 3.4 and appendix C, real detector effects can have a significant impact on the mass resolution in particular, which can affect the conclusions of any multivariate study that uses the mass. Accordingly we carry out two sets of studies in parallel. In the first set of studies, we classify jets in terms of whether

¹³Note that D_2 also has access to information that is being discarded in the Lund-plane variables, specifically that associated with the substructure of secondary leaves off the primary plane.

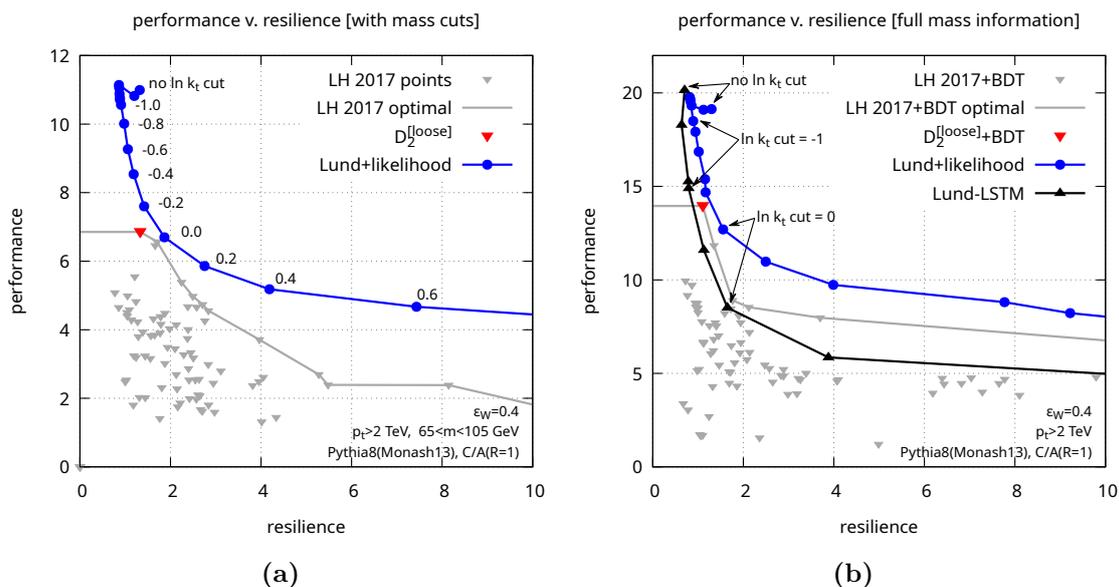


Figure 11. Plots of performance ($\epsilon_W / \sqrt{\epsilon_{\text{QCD}}}$), for fixed (hadron+MPI) signal efficiency $\epsilon_W = 0.4$, versus resilience to non-perturbative effects, (ζ of eq. (3.6)). Grey triangles (and the red one) correspond to the full range of shape observables studied in the LH 2017 study [80]. The blue circles and black triangles correspond to the Lund-likelihood and Lund-LSTM methods respectively, with each point along a line corresponding to different lower cuts on the Lund-plane $\ln k_t/\text{GeV}$ value, below which declusterings are ignored (i.e. not passed to the LSTM or likelihood method; training is repeated for each different k_t cut). In (a) the shape observables are used together with a cut on a mass variable, $65 < m < 105$ GeV (the mass may be groomed, or ungroomed, depending on the point); for the Lund likelihood, eq. (3.1) is evaluated with a single (groomed) mass bin, covering the same mass range as for the shapes, plus an outflow mass bin. In (b) shape variables are combined with the full particle-level resolution mass information through a boosted decision tree (BDT) and the cut that defines $\epsilon_W = 0.4$ is placed on the BDT output; for the Lund-likelihood and LSTM methods, full resolution Lund-plane information is used (including the mass for the likelihood method). For the Lund+likelihood method, the values of the $\ln k_t$ cuts are the same as for (a), i.e. spaced every 0.2 units of $\ln k_t$. For the Lund+LSTM curves points are separated by 0.5 units of $\ln k_t$.

they satisfy a loose requirement on the (possibly groomed) jet mass, $65 < m < 105$ GeV, and then do not further use any mass information. Such a study is fairly realistic in terms of how much mass information is accessible in a detector, but cannot be performed with machine learning, because the latter is likely to “cheat” and learn the mass information from other variables in the jet. To be able to also examine machine learning, we therefore carry out a second set of studies, in which full particle-level information is available, allowing reconstruction of the W mass peak. All methods then exploit the unrealistically good particle-level mass resolution on that W mass peak.

In figure 11, we show the performance achieved by the different tagging approaches versus their resilience to underlying event and hadronisation corrections. This is calculated following the procedure introduced in section III.2 of the 2017 Les Houches proceedings [80].

The performance, $\epsilon_W/\sqrt{\epsilon_{\text{QCD}}}$, is plotted versus the resilience ζ , which is calculated using both hadron+MPI-level efficiencies ϵ and parton-level efficiencies ϵ' (all computed for a set of cuts on a shape variable, or multi-variate tagger output, that gives a hadron+MPI-level signal efficiency $\epsilon_W = 0.4$),

$$\zeta = \left(\frac{\Delta\epsilon_W^2}{\langle\epsilon\rangle_W^2} + \frac{\Delta\epsilon_{\text{QCD}}^2}{\langle\epsilon\rangle_{\text{QCD}}^2} \right)^{-\frac{1}{2}}, \quad (3.6)$$

where $\Delta\epsilon = \epsilon - \epsilon'$ and $\langle\epsilon\rangle = \frac{1}{2}(\epsilon + \epsilon')$. The left-hand plot shows the results obtained in a specific mass-bin, comparing our likelihood method with the results from the LH report [80].¹⁴ The right-hand plot shows the results with full mass information, and includes results with machine learning. Both parton-level and hadron+MPI-level efficiencies are calculated using a discriminator determined/trained using hadron+MPI-level events (this statement holds for all likelihood, LSTM and BDT-based results).

Figure 11 shows grey triangles for each of the 88 combinations of a single shape variable and mass used in the LH 2017 report [80] (the shape and mass being combined via a BDT in the right-hand plot). The grey line is the upper envelope of those points. The specific $D_2^{\text{[loose]}}$ variant discussed in section 3.3 is highlighted in red and one can see that it has the best performance among all shape+mass taggers.

For methods that use the Lund plane information one can impose a lower limit, $k_{t,\text{cut}}$, on the value of k_t for which Lund-plane declusterings are considered. Declusterings with lower k_t values are simply ignored, both at the training stage and subsequently when evaluating performance and resilience. For the Lund-LSTM method, the tagger is trained separately for each $k_{t,\text{cut}}$ value. Larger values of $k_{t,\text{cut}}$ are expected to yield taggers that are more resilient to non-perturbative effects. The results for the Lund-likelihood and Lund-LSTM methods are shown as blue and black points respectively (linked by lines) in figure 11, each point corresponding to a specific value of the k_t cut.

Without a k_t cut for the Lund-based taggers, performances qualitatively mirror those in figure 10 at the corresponding value of $\epsilon_W = 0.4$: the Lund-LSTM method performs best, then comes the Lund-likelihood method, followed by $D_2^{\text{[loose]}}$. Quantitative differences relative to figure 10 are a consequence of the lack of detector simulation and the use of a broad mass bin (figure 11a) or full mass resolution (figure 11b). The quantitative differences are especially large in the latter case, as one would expect (e.g. for the LSTM, $\epsilon_W/\sqrt{\epsilon_{\text{QCD}}} \simeq 20$ at $\epsilon_W = 0.4$ translates to $1/\epsilon_{\text{QCD}} \sim 2500$, compared to $1/\epsilon_{\text{QCD}} \simeq 700$ in figure 10).

For the Lund-likelihood method, imposing a low k_t cut, $\ln k_{t,\text{cut}} < -1$, has little impact on the performance or resilience relative to the situation without any cut.¹⁵ Further raising

¹⁴Albeit with jets obtained with an initial C/A clustering rather than an anti- k_t clustering as in the original LH study.

¹⁵There are small variations in both performance and resilience as one lowers the cut below $\ln k_{t,\text{cut}} = -1$. Further work would be needed to establish whether these variations are robust with respect to statistical fluctuations. Genuine structure in that region could be generated by a variety of effects, for example related to the breakdown of the independent-emission picture that underlies the log likelihood construction. Non-perturbative effects can also induce significant smearing, which complicate the understanding of that region.

the value of $k_{t,\text{cut}}$ initially leads to a rapid loss in performance and modest improvement in resilience. This suggests that there is information in the non-perturbative region when discriminating boosted W jets from QCD jets. For $k_{t,\text{cut}} = 1$ GeV ($\ln k_{t,\text{cut}} = 0$), performance is slightly better than the best shape variable at comparable resilience. Only for yet higher values of $k_{t,\text{cut}}$ does resilience improve substantially, and then the Lund-likelihood performance remains above that of the shape variables (well above for figure 11a). Thus it appears that the Lund-likelihood method performs well not just in terms of raw performance, but also, with a k_t cut, in terms of performance for a given degree of sensitivity to non-perturbative effects.

For the Lund-LSTM method, even a small cut on k_t rapidly leads to a loss of performance. For $\ln k_{t,\text{cut}} \gtrsim -1$, its performance falls below that of the Lund-likelihood method and that remains the case as $k_{t,\text{cut}}$ is further increased. In fact, for $\ln k_{t,\text{cut}} \gtrsim -0.5$, the performance of the Lund-LSTM method even starts to fall slightly below the most optimal shape variables. This is somewhat puzzling and hints at potential fragility of machine-learning approaches.

Overall we see that while an ML based approach can achieve substantially better performance, the models obtained are not particularly resilient to non-perturbative corrections. We note however that other training methods, e.g. based on adversarial networks [98–100], could improve the robustness of the derived taggers to specific effects such as hadronisation, MPI and pileup.

While we have focused here on resilience to corrections from hadronisation and MPI, one could similarly study the resilience of the methods against pileup or detector effects.

4 Conclusions

The Lund plane offers a powerful new way to study and exploit the internal structure of jets. In contrast to traditional shape observables it connects much more directly to individual regions of phase space. This makes it useful across a range of applications in jet physics. It also brings many declustering based jet observables, such as the iterated soft-drop multiplicity and z_g into a single unified framework.

One way of studying the Lund plane is in terms of its average density, as a function of angle and transverse momentum. This density is amenable to calculation within both resummed and fixed-order perturbative methods. We limited our discussion of such a calculation to first order, section 2.2, and identified a number of the contributions that would become relevant at higher orders. Experimentally, we believe that much of the Lund plane phase-space can be reliably determined. This conclusion is based on Delphes fast-detector simulations in conjunction with subjet-particle rescaling type algorithms (SPRA, appendix C), to recover information at small angles that might otherwise be obscured by finite calorimeter resolution. This offers a clear potential for carrying out experimental measurements of the pattern of radiation in both the perturbative and non-perturbative regimes. One application of such measurements would be to constrain Monte Carlo simulation programs, which as we saw in section 2.3 show up to 30% differences in their predictions of the Lund plane density from one program to another. Another application

would be to directly identify which kinematic regions of a jet’s radiation pattern are modified in heavy-ion collisions, thus shedding light on the mechanisms of partonic energy loss in a hot, dense medium.¹⁶

A use case for the Lund plane that we have explicitly examined is for tagging boosted electroweak bosons. Compared to the jet-image type inputs that have been the mainstay of “visual” machine-learning approaches to jet substructure tagging so far (e.g. boosted- W tagging), many of the features that can be exploited are immediately visible to the human eye. With certain machine-learning methods (notably LSTM’s) the Lund-plane inputs appear to yield superior W -tagging performance as compared to jet images. This is despite the fact that by discarding information about secondary leaves of the Lund diagram, we are actually providing less information to the machine learning methods than comes with jet images. We note that for reliable comparisons of the relative quantitative performance of different methods it was important to take into account detector effects.

The fundamental information that is contained in the Lund plane, i.e. the kinematics of declustering sequences, has been used in other recent work on machine learning [40, 41]. However the Lund plane as a visualisation provides powerful insight into the physical structure of that information and into how that information differs according to the origin of the jet, cf. figure 7 for W jets versus figure 2 for QCD jets. In particular, the Lund plane’s simplicity, and the relatively moderate degree of correlation between different parts of the plane, have the consequence that much of the performance obtained by machine-learning algorithms can be reproduced using conceptually simple log-likelihood approaches. This opens the prospect for a substantial degree of experimental and theoretical understanding of the robustness of the Lund plane information for tagging. That understanding may be useful also in terms of the construction of high-performance decorrelated taggers [104].

There is a potential for a range of other applications, including top-tagging, quark-gluon discrimination, further improvements of boosted electroweak boson tagging, or extensions to the recently proposed soft-drop photon isolation approach [105]. Furthermore, Lund-plane type studies need not necessarily be restricted to the study of final-state jets, but may also be informative for the initial state, for example to help discriminate different mechanisms of Higgs-boson production.

Finally, while we have restricted most of our discussion here to the primary Lund plane (other than a brief discussion in appendix B), one cannot help but wonder about the potential benefits to be had from exploiting the structure of the full Lund diagram for jets, cf. the middle row of figure 1. This may be relevant both for developing our generic understanding of the structure of jets, and for certain tagging applications, for example with recursive neural networks (as in ref. [40]) or tree-LSTM architectures [106] to capture the full clustering tree in the machine-learning training.

¹⁶This provided the original motivation for introducing the Lund plane as a measurable quantity, in the context of the “Novel tools and observables for jet physics in heavy-ion collisions / 5th Heavy Ion Jet Workshop” [101], and has in the meantime also been studied in refs. [102, 103]. See also <https://gitlab.cern.ch/gsalam/2017-lund-from-MC> for a corresponding implementation. In heavy-ion collisions, background (i.e. UE) contamination appears to be a non-negligible issue, as it may be also in high-pileup pp collisions, depending on the precise pileup-mitigation scheme being used. Various potential approaches to address this were highlighted in section 3.4.

We therefore look forward to a wide range of studies with Lund-diagram related observables in future work. To facilitate such studies the `LundPlane` library is available as part of the `FastJet contrib` package.

Acknowledgments

We are grateful to the referee for helpful comments on the manuscript. We thank Michele Selvaggi for making it straightforward to interface an external FastJet instance to Delphes, and Jesse Thaler for discussions on log-likelihood approaches. We are also grateful to CERN’s Techlab for computing time on a Tesla P100, and to the NVIDIA Corporation for the donation of a Titan Xp GPU used for this research. Our introduction of the Lund plane as a measurable observable was originally motivated by discussions in the context of the “Novel tools and observables for jet physics in heavy-ion collisions / 5th Heavy Ion Jet Workshop”. A study of the Lund plane in terms of its sensitivity to jet quenching has appeared as part of the workshop proceedings [101] and has also been examined in refs. [102, 103]. F.D. is supported by the SNF grant P2SKP2_165039 and by the Office of High Energy Physics of the U.S. Department of Energy (DOE) under grant DE-SC-0012567. F.D. thanks the University of Zurich and the Pauli Center for Theoretical Studies, G.P.S. thanks the IPhT Saclay, and G.S. thanks CERN for hospitality while this work was being completed. G.S. is supported in part by the French Agence Nationale de la Recherche, under grant ANR-15-CE31-0016.

A The Lund plane for (C/A-reclustered) anti- k_t jets

Throughout the main text we have determined the Lund plane for jets obtained with an initial Cambridge/Aachen clustering. This choice has the advantage of procedural simplicity and also makes it relatively straightforward to interpret the structures that appear in different regions of the Lund plane. One could, however, instead use the anti- k_t [51] algorithm to find the initial jets, and then recluster their constituents with the C/A algorithm in order to obtain the Lund plane. We expect that this might be the preferred experimental approach given that much effort into experimental jet calibration gets directed to the anti- k_t algorithm. The averaged primary Lund plane obtained for C/A-reclustered anti- k_t jets is shown in figure 12. It is almost identical to figure 2, except near $\Delta = 0$ and $\ln k_t = 0$, where the anti- k_t jets appear to have an additional structure: an up-right going diagonal structure for $\ln k_t/\text{GeV} \sim 0$ and $\ln R/\Delta \lesssim 0.75$.

A reasonable hypothesis is that this structure is associated with the clustering of soft (mostly underlying-event) particles near the edge of the jet. To help understand this in more detail, figure 13 shows the rapidity-azimuth distribution of particles in a single C/A jet (upper row) versus a C/A-reclustered anti- k_t jet (lower row), at various stages of the declustering. At each stage, the softer subjet is shown in red. The declustering steps are aligned such that the last step shown corresponds to a similar pair of subjets in the two sequences. At the earlier stages, for C/A, one sees that large values of Δ are associated with large-area softer (red) subjets. In contrast, with reclustered anti- k_t jets, the softer subjets

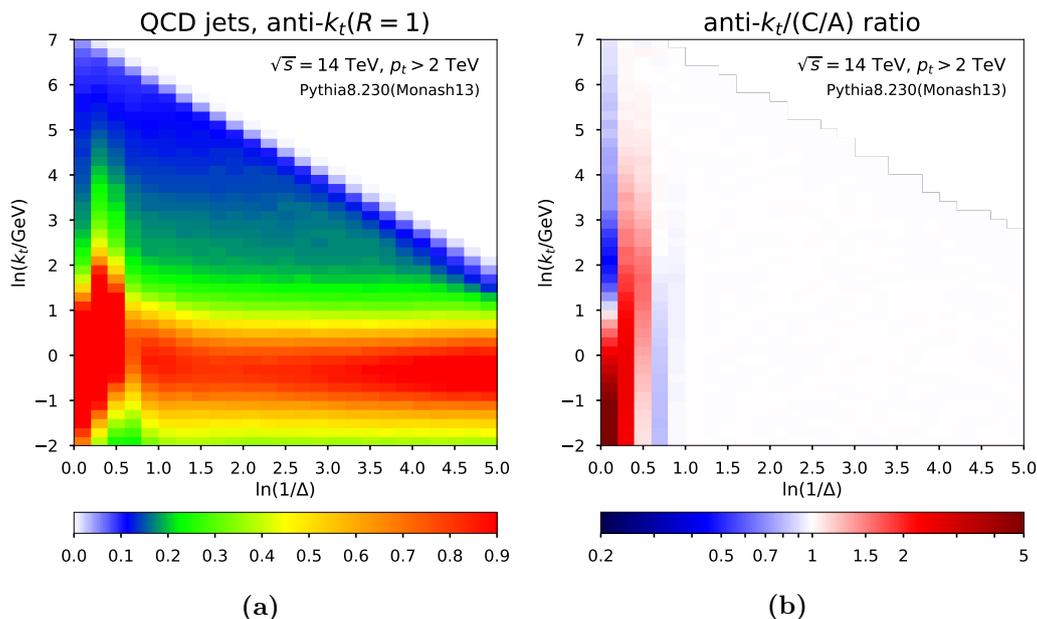


Figure 12. The averaged primary Lund plane density (a) for jets initially obtained with anti- k_t clustering (whose constituents are then reclustered with the Cambridge/Aachen algorithm) and its ratio (b) to the averaged Lund plane density for jets originally obtained with Cambridge/Aachen clustering (figure 2). Note the structure around $\Delta = 0$ and $\ln k_t = 0$ that is present here and not in figure 2.

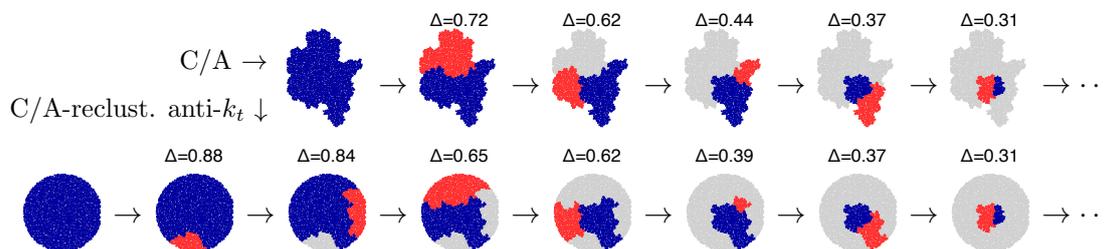


Figure 13. Illustration of declustering sequences for a Cambridge/Aachen jet (upper row) and the same neighbourhood of an event clustered with the anti- k_t algorithm and subsequently reclustered with Cambridge/Aachen (lower row). The jets include ghost particles [94, 95] so as to illustrate the area of the jets involved at each declustering stage. The plot shows each stage of the declustering, with the softer subjet shown in red (b in the notation of section 2.1) and the harder subjet shown in blue (a).

for large values of Δ tend to have smaller areas, constrained by the circular boundary of the original anti- k_t jet. Smaller areas imply a smaller amount of p_t coming from the underlying event. Therefore in the Lund plane the peak of subjets at large Δ (small $\ln R/\Delta$) should come at lower k_t for the reclustered anti- k_t jets than for the C/A jets.

It is natural to ask whether the pattern of softer-subjet area versus Δ seen in figure 13 holds beyond the case of a single jet. To answer this question, figure 14 shows the average subjet area as a function of $\ln R/\Delta$ for C/A jets (in blue) and C/A-reclustered anti- k_t jets

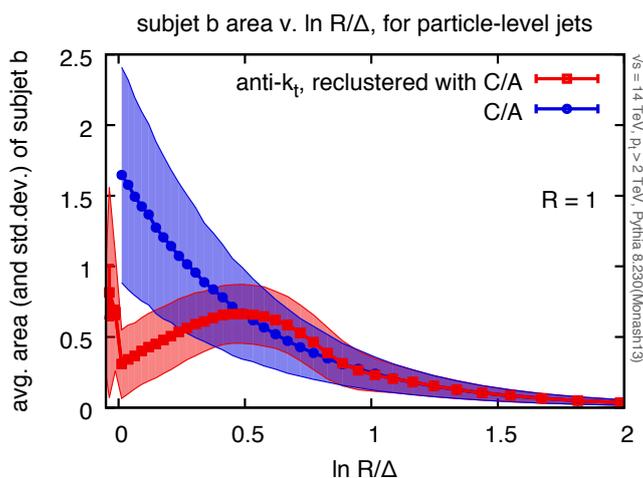


Figure 14. Average area (points), and standard deviation (band), of the softer subjet (subjet b) in the Lund plane declusterings, shown as a function of Δ for C/A jets (blue) and for C/A-reclustered anti- k_t jets (red).

(in red). The bands represent the standard-deviation of the jet areas. The event sample and jet selection are identical to those used in figure 12. In the C/A case, the softer subjet area increases for smaller $\ln R/\Delta$, i.e. for increasing Δ , and is consistent with an area A_b that scales as $A_b \sim \Delta^2$. If the density of p_t per unit area from the underlying event is ρ , then one expects the k_t of typical softer subjets to go as $k_t \sim \rho A_b \Delta \sim \Delta^3$. In contrast, for $\ln R/\Delta < 0.5$, the typical area for reclustered anti- k_t softer subjets tends to decrease as $\ln R/\Delta$ decreases. The scaling near $\Delta = 0$ is found to be roughly $A_b \sim \Delta^{-2.6}$, which would lead to $k_t \sim \Delta^{-1.6}$. Note, however that the area scaling behaviours given here include a component where the subjet is moderately hard and so the scaling behaviours for pure underlying-event jets may differ in their details.

A point to note about reclustered anti- k_t jets is that it is possible to have $\Delta > R$, cf. the points at negative $\ln R/\Delta$ in figure 14. This occurs only rarely and tends to be driven by specific configurations of hard particles in the jet.

A final comment is that since the difference in structure between C/A and C/A-reclustered anti- k_t jets' Lund planes is in a region that is anyway dominated by soft particles from the underlying event, we expect it to have little impact on discrimination power if one uses reclustered anti- k_t rather than C/A jets. Explicit studies with the Lund-likelihood method bear out this expectation.

B Moderate energy jets and secondary Lund planes

For most of this article, jets have been considered with a $p_t > 2$ TeV cut on the transverse momentum. However, when considering lower energy jets, the peak in the primary Lund plane associated with the W splitting, cf. figure 7, moves to the left and the shadow region to its left, associated with the colour-singlet nature of the W , become less visible. This is expected to reduce the performance achieved by taggers based just on the primary Lund

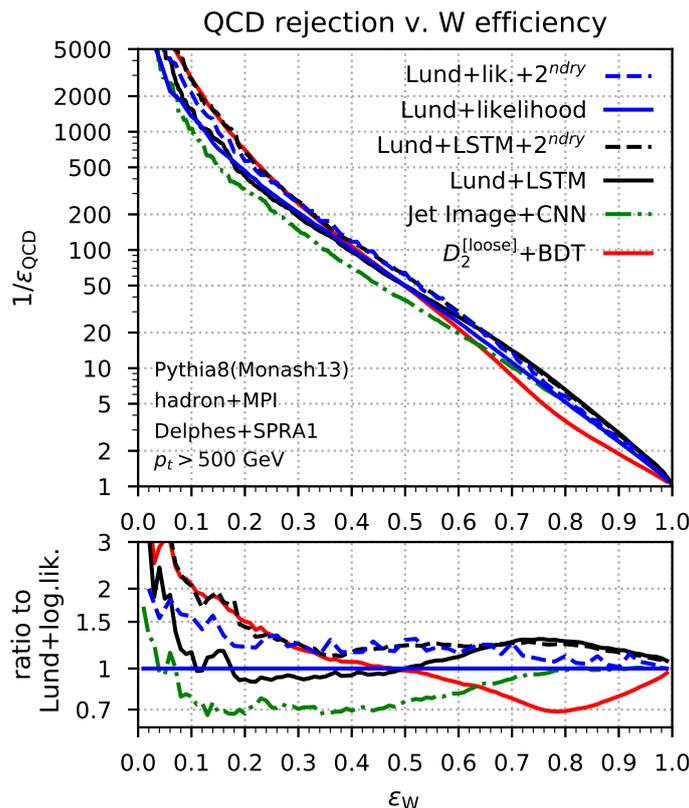


Figure 15. Background rejection ($1/\epsilon_{\text{QCD}}$) versus signal efficiency (ϵ_W), with a transverse momentum cut on the jets of 500 GeV. The lower panel shows the ratio to the Lund+likelihood method. The solid curves are to be compared to the corresponding ones in figure 10 for jets with $p_t > 2$ TeV (note the different scale in the lower panel).

plane variables (though the larger fraction of gluon-induced background jets at lower p_t may partially counteract this).

The reduced performance at lower p_t from primary-Lund plane methods is visible in figure 15, which shows the signal efficiency against the background rejection for $p_t > 500$ GeV jets, using the primary Lund-plane log-likelihood method, the primary Lund-plane LSTM method and also the D_2^{loose} observable (sensitive also to emissions beyond primary Lund-plane ones). Those curves are to be compared to the corresponding ones in figure 10 for $p_t > 2$ TeV jets. Using $\epsilon_W = 0.4$ as a reference point, one sees that all methods are worse in background rejection at lower p_t . The loss is a factor of 5 – 7 for the primary-Lund-plane based methods, while it’s only a little more than a factor of 2 for the D_2^{loose} observable, with the result that for $\epsilon_W = 0.4$ D_2^{loose} ’s performance at moderate- p_t is comparable to that of the primary Lund-plane methods. At higher (lower) signal efficiencies, D_2^{loose} does worse (better) than the primary-Lund-plane methods.

The D_2 observable effectively takes into account information not just from the primary but also secondary Lund planes, information that is discarded by the primary Lund-plane log likelihood and LSTM methods. Figure 16 shows the secondary Lund planes for the leading primary emission, defined as in section 3.1 as the first emission in the C/A declus-

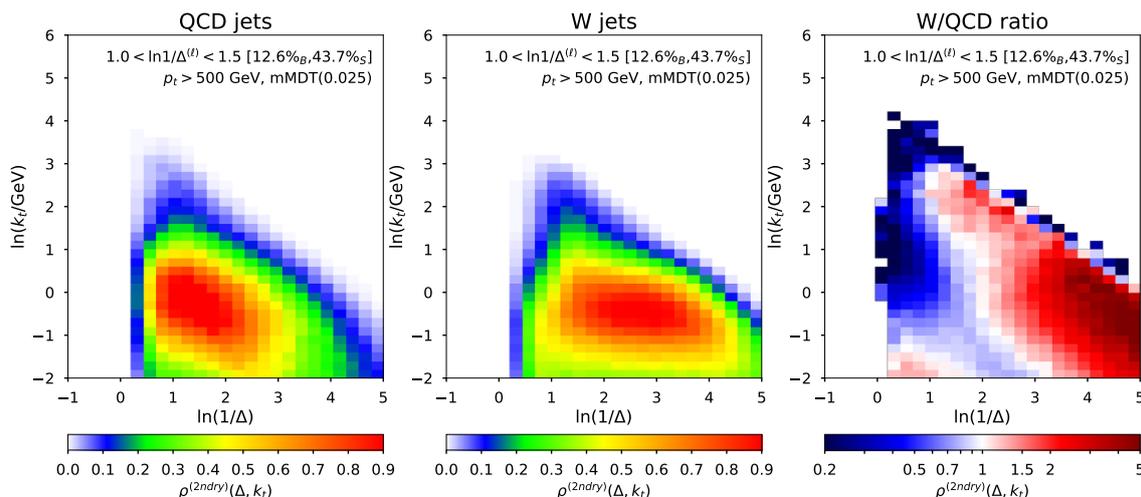


Figure 16. Averaged density, $\rho^{(2\text{ndry})}(\Delta, k_t)$, of the secondary Lund plane associated with the leading (mMDT(0.025)) emission, for jets in which the leading emission’s angle satisfied $1 < \ln 1/\Delta^{(\ell)} < 1.5$. From left to right: for jets in dijet events, for jets in WW events, and the ratio of the two. The percentages in square brackets indicate the fractions of jets for which the leading emission is in the given bin, for background and signal respectively.

tering sequence that satisfies $z > 0.025$ (cf. the “mMDT(0.025)” label in the figures). The left-hand plot corresponds to QCD jets, the middle one to W jets. The right-hand plot, which shows the W/QCD secondary Lund-plane ratio, helps illustrate the nature of the discriminating information contained in the secondary Lund plane. In particular, the leading emission in QCD jets will tend to have more large-angle radiation, while for W jets the emissions are more likely to be relatively hard and collinear.

To help test the hypothesis that D_2 is doing well effectively because of the secondary information, we have attempted to explicitly add secondary Lund-plane information also to the log-likelihood and LSTM methods.

For the log-likelihood method we found that the best performance came with two steps: replacing the mMDT(0.025) identification of the leading emission with an identification based on finding the primary emission that had the smallest value of $|\ln(p_{ta}p_{tb}\Delta^2/m_W^2)\ln z|$; the additional likelihood for the secondary plane, similar to the non-leading primary plane likelihood, is combined with the primary \mathcal{L}_{tot} from eq. (3.4) not by direct addition, but via a 2-dimensional map of the two likelihoods.¹⁷ The fact that this last step was needed suggests that there may be scope for better understanding correlations between the primary and secondary Lund planes.

For the LSTM method, we start by identifying the secondary Lund plane using the same method as in the previous log-likelihood approach. The primary and secondary Lund plane declusterings are then given as input to two separate LSTMs with a dropout layer,

¹⁷A two-dimensional map of the two likelihoods helps take into account the correlations between the primary and secondary likelihoods. These correlations can, for example, come from radiation from the leading parton which gets clustered as a secondary emission or vice-versa.

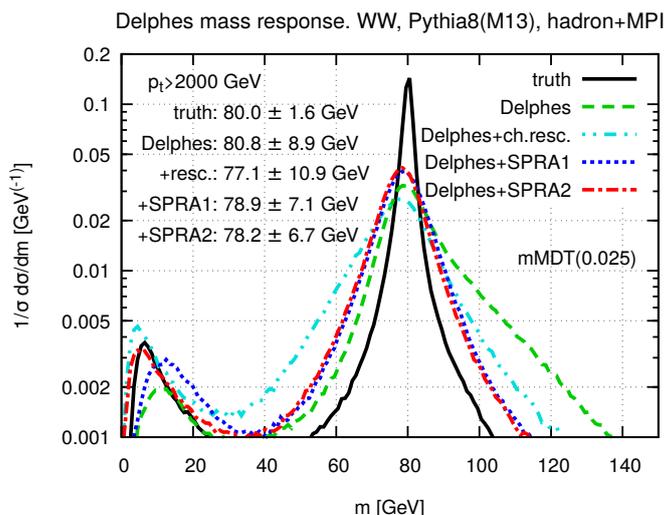


Figure 17. Reconstructed mMDT ($z_{\text{cut}} = 0.025$) W mass peak at particle level (“truth”), with Delphes particle flow [81], and additionally with the rescaling algorithms described in the text to improve resolution. The average and standard deviation for the reconstructed W peak at the different levels are shown based on a fit of a Gaussian distribution between 50 and 110 GeV. The generation and selection of jets is as described in section 3.4, using the WW process, in particular selecting jets with the requirement $p_t > 2$ TeV, $|y| < 2.5$.

with 128 units for the primary plane and 64 for the secondary one. The output of the LSTMs are then concatenated and passed to a dense layer with 100 dimensions, and then given to a final two-dimensional layer with softmax activation.

The performances of the methods including the leading emission’s secondary Lund plane are shown as dashed lines in figure 15, and are labelled “+2nddry”. They improve the background rejection by 20 – 30% in the $\epsilon_W \sim 0.3 - 0.6$ range, such that our Lund-based tagging including the secondary plane now outperforms D_2 down to about $\epsilon_W \sim 0.3$. With the LSTM approach the performance is matched also at lower values of ϵ_W . Note that at high p_t we did not find a significant benefit in including secondary Lund plane information.

We leave a more extensive study of the secondary Lund plane and its impact on jet tagging for future work, keeping in mind also that today’s parton showers may not always correctly reproduce the patterns of correlations between emissions [107].

C Detector effects and the subjet-particle rescaling algorithm (SPRA)

At particle (“truth”) level, in a W sample, with $p_{t,\text{jet}} > 2$ TeV, the mass of the (sub)jet selected by the mMDT procedure is very sharply peaked around the true W mass, with an effective resolution of about 1.5 GeV. In contrast if one uses the Delphes fast detector simulation [81], with particle flow (PF) and the `delphes_card_CMS_NoFastJet.tcl` card, one finds an mMDT mass resolution of about 9 GeV. This is illustrated in figure 17 (we return to the SPRA curves below). Such a large difference in resolution between particle and detector-level can have a big impact on conclusions about performance, especially given

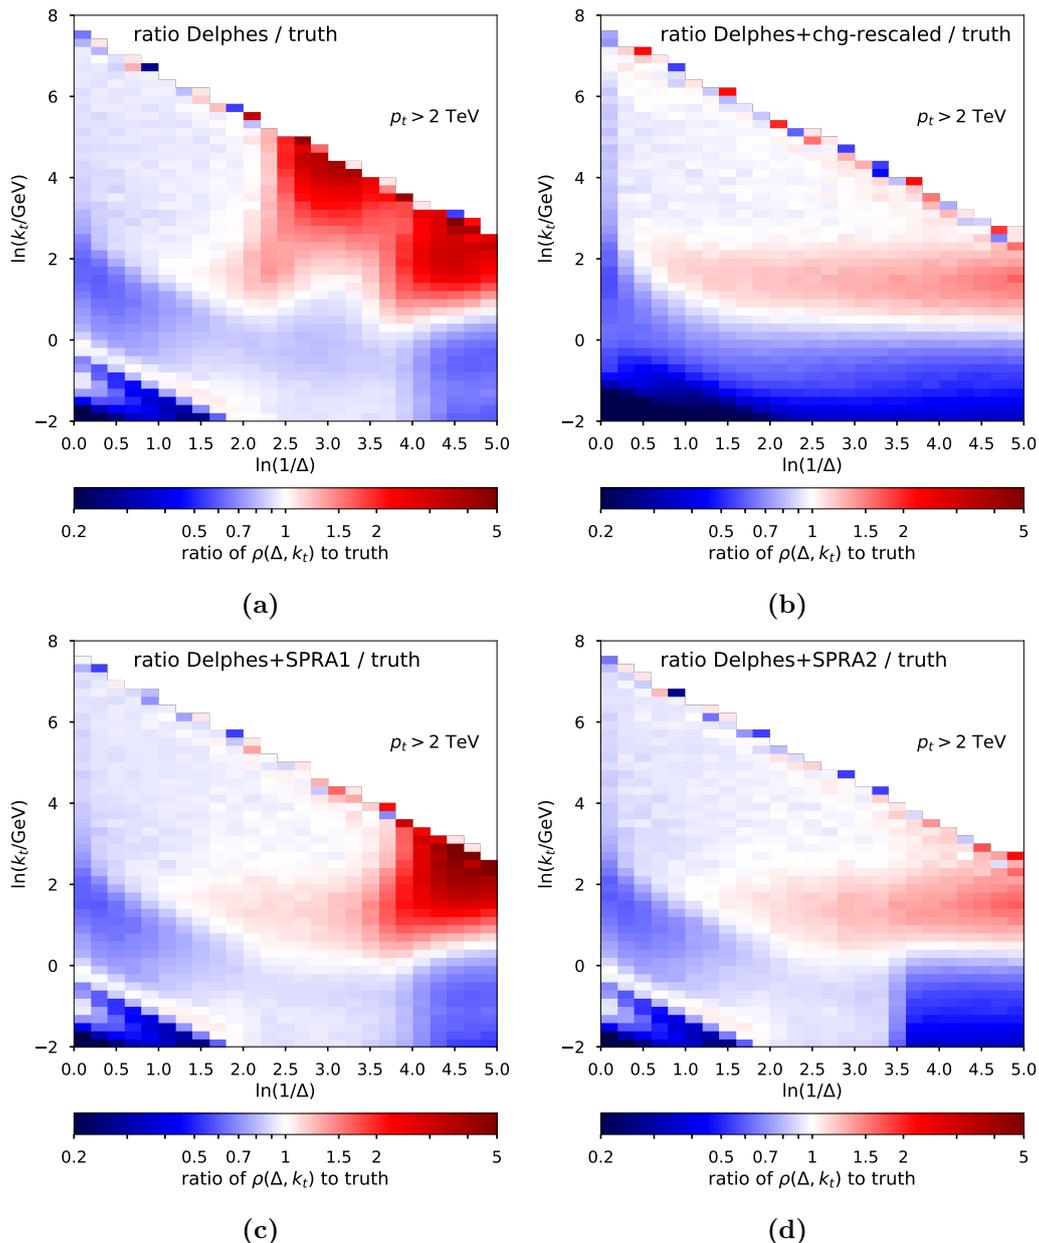


Figure 18. Ratio of detector-level Lund-plane densities to the truth density. (a) With Delphes’ particle flow algorithm. (b) Delphes particle flow supplemented with charged rescaling (after rescaling, neutral particles are discarded). (c,d) Delphes particle flow supplemented with SPRA1, SPRA2 respectively. The generation and selection of jets is as described in section 3.4, using the dijet process, selecting jets with the requirement $p_t > 2$ TeV, $|y| < 2.5$.

that both ML and likelihood-based approaches tend to derive considerable discrimination power from the mass variable. This is true even if the mass is not directly passed as an input to a ML approach, because the mass can quite effectively be deduced from the Lund-plane k_t and Δ variables.

Detector effects can also have a significant impact on the Lund plane at angular scales commensurate with the hadronic calorimeter spacing and lower. This is illustrated in figure 18a, which shows the ratio of the Lund plane for dijets as obtained with Delphes particle flow relative to the particle-level truth. In the lower-left corner there is a prominent dark blue region indicating missing Lund plane subjects at detector level: this can be interpreted as a consequence of finite p_t thresholds (a given p_t maps onto a downwards-right going diagonal). In the right-hand part of the Lund plane, just below the kinematic limit, there are two prominent enhanced regions (in red), with corresponding deficits at lower k_t : their positions in $\ln 1/\Delta$ coincide with the intrinsic angular resolution scales of the hadronic (HCal) and electromagnetic calorimeters (ECal), which in the central part of the detector have η, ϕ segmentations of 0.087×0.087 and 0.0174×0.0174 respectively. Those segmentations translate to $\ln 1/\Delta$ values of about 2.4 and 4.1.

The origin of the enhancements is relatively straightforward to understand. Consider the structure associated with the hadronic calorimeter scale, Δ_{HCal} . On average, about 10% of particles in jets are undecayed neutral hadrons (for example K_L). Schematically, the particle flow algorithm can identify the energy deposit from such particles as the difference between the energy in a given hadronic calorimeter tower and that observed in the charged tracks that enter the tower. The Delphes PF implementation assigns that energy difference to a point in η, ϕ that is randomly distributed over the calorimeter tower area. If the neutral hadron has a true separation $\Delta_{\text{true}} \ll \Delta_{\text{HCal}}$ and transverse momentum $k_{t,\text{true}}$ relative to the jet core, the reconstructed separation and transverse momentum will be

$$\Delta_{\text{reco}} \sim \Delta_{\text{HCal}} \gg \Delta_{\text{true}}, \quad k_{t,\text{reco}} \sim \frac{\Delta_{\text{HCal}}}{\Delta_{\text{reco}}} k_{t,\text{true}} \gg k_{t,\text{true}}, \quad (\text{C.1})$$

where the scaling of the transverse momentum, k_t , relative to the jet core, follows from the assumption that the neutral hadron’s transverse momentum p_t relative to the beam direction is correctly determined ($k_t = p_t \Delta$, cf. eq. (2.1a)). For a jet core containing $\mathcal{O}(10\text{--}20)$ particles, there will typically be at least one neutral hadron, which will be reconstructed with an angular scale Δ_{HCal} and a transverse momentum that is larger than its true transverse momentum. It is this that creates the red bump around $\ln 1/\Delta \simeq \ln 1/\Delta_{\text{HCal}} \simeq 2.4$: with the particle flow algorithm there is nearly always something in that region, whereas at truth level there isn’t. This is arguably also the origin of the long tail to high masses for the Delphes curve in figure 17. An analogous phenomenon occurs on the electromagnetic calorimeter granularity scale. The depletions at lower k_t values may be shadows induced by the enhanced regions. (Given an emission at some Δ , a fraction of emissions at lower k_t and similar Δ and ψ get clustered with it and so are assigned to its secondary Lund plane).

The limitations induced by detector granularity have been addressed in a range of past work. Katz, Son and Tweedie [82] were the first to document the issue in the context of substructure and they proposed a solution, whereby 5×5 groups of ECal cells associated with a single HCal cell were rescaled to match the total HCal + ECal energy. Ref. [83] extended the procedure, applying the scaling within minijets. Nowadays, CMS has an approach referred to as “split PF photons+neutrals” [88], which is conceptually similar

insofar as it distributes neutral-hadron energy across the ECal cells (it also includes tracking improvements for high-energy jets). Schaetzel and Spannowsky investigated rescaling the charged tracks in jet to match the jet’s total calorimeter energy (“track-flow” in the nomenclature of ref. [87]; here we will call it charge-rescaling). A procedure that is functionally equivalent, track-assisted mass, has been studied by ATLAS [89] to improve its mass resolution. Other studies of the question include refs. [85–87, 108].

To mitigate the impact of detector granularity, we adopt a subjet particle rescaling algorithm (SPRA) that is similar in spirit to that of ref. [83]. We have two variants, SPRA1 and SPRA2. For SPRA1, we take a jet and recluster its Delphes particle-flow objects into subjets using the C/A algorithm with radius $R_h = 0.12$, commensurate with the hadronic calorimeter granularity.¹⁸ Taking each subjet in turn we scale each PF charged-particle (h^\pm) and photon (γ) candidate that it contains by a factor f_1 ,

$$f_1 = \frac{\sum_{i \in \text{subjet}} p_{t,i}}{\sum_{i \in \text{subjet}(h^\pm, \gamma)} p_{t,i}}, \tag{C.2}$$

and discard the other particles (i.e. neutral hadron candidates). If the subjet does not contain any photon or charged-particle candidates, we instead retain all of the subjet’s particles with their original momenta. After having applied this procedure to each subjet, we then recluster the full set of resulting particles, i.e. from all subjets, into a single large jet and evaluate the mass and Lund plane on that set of particles.

The SPRA2 variant is a similar but carries out two levels of rescaling. The motivation behind the double rescaling is to reduce the effects of the granularity of both the hadronic and electromagnetic calorimeters. After applying the SPRA1 algorithm, we recluster the resulting particles with a radius $R_e = 0.03$, commensurate with the electromagnetic calorimeter granularity. Taking each (R_e) subjet in turn, we scale each PF charged-particle (h^\pm) candidate that it contains by a factor f_2 ,

$$f_2 = \frac{\sum_{i \in \text{subjet}} p_{t,i}}{\sum_{i \in \text{subjet}(h^\pm)} p_{t,i}}, \tag{C.3}$$

and discard the other particles (i.e. photons and possibly some neutral hadron candidates). If the subjet does not contain any charged-particle candidates, we instead retain all of the subjet’s particles with their original momenta. Again, after having applied the procedure to each subjet, we take all the particles and produce a single large-radius jet from them.

Figure 17 shows that there is some gain in mass resolution from the SPRA1 algorithm, from about 9 GeV to 7 GeV. There is, however, only limited gain in going to the double rescaling, SPRA2, algorithm.

If we now consider the Lund plane reconstruction, figure 18, we see that SPRA1 (figure 18c) eliminates most of the structure associated with the hadronic calorimeter scale, while SPRA2 (figure 18d) also alleviates some of the structure associated with the electromagnetic calorimeter scale. Overall, the conclusion is that with the help of the SPRA

¹⁸As pointed out in ref. [83], strictly one should choose $\sqrt{2} < R_h/\Delta_{\text{HCal}} < 2$, but we found little difference between a value in that range and our choice that is slightly below the lower edge of this range.

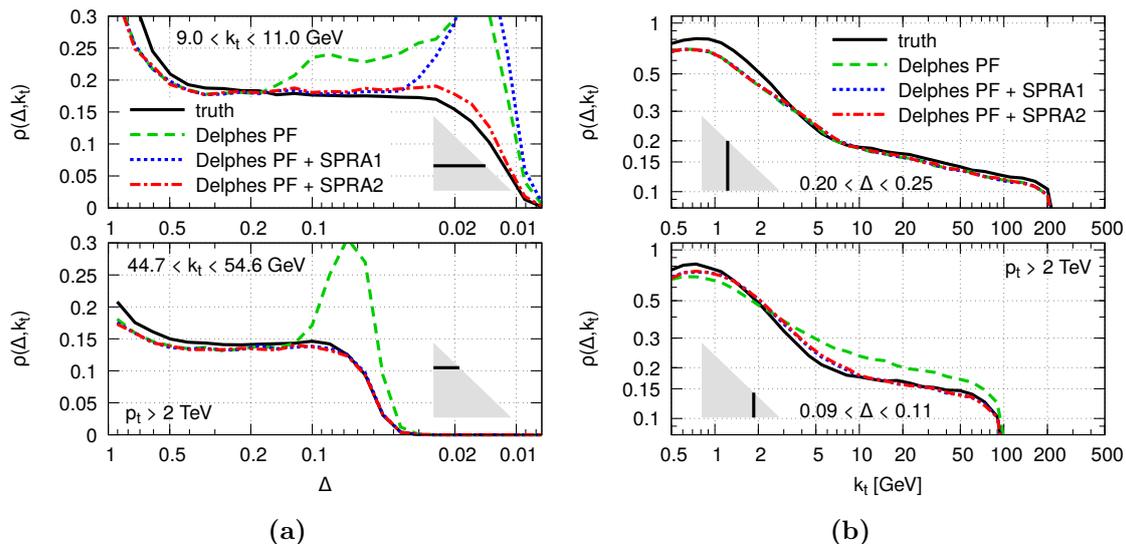


Figure 19. Lund-plane slices comparing the truth result to Delphes PF with and without the SPRA rescalings. The slices are shown at fixed k_t as a function of Δ (left) and at fixed Δ versus k_t (right). The artefacts visible in the top-left plot at scales of the hadronic ($\Delta \sim 0.087$) and electromagnetic ($\Delta \sim 0.0174$) calorimeters are well brought under control by the SPRA1 and SPRA2 rescalings respectively. The generation and selection of jets is as described in section 3.4, using the dijet process, selecting jets with the requirement $p_t > 2$ TeV, $|y| < 2.5$.

algorithms, a large part of the Lund plane can be measured fairly reliably, with detector effects that remain limited to within 20–30%. This is visible also in the plots of Lund-plane slices in figure 19.

An even simpler approach is to adopt a jet-wide track rescaling, where every charged track is multiplied by a factor

$$f_{\text{chg}} = \frac{p_{t,\text{jet}}}{\sum_{i \in \text{jet}(h^\pm)} p_{t,i}}, \quad (\text{C.4})$$

and only those rescaled tracks are used as an input. This performs fairly well, cf. the Lund-plane density ratio in figure 18b.

A final test of the SPRA algorithms is shown in figure 20, which compares the background rejection power of the log-likelihood W -tagger, as a function of signal efficiency, for truth particles and for Delphes PF with and without SPRA. A first observation is that for a signal efficiency of 0.5, the background rejection is about six times worse with Delphes PF as compared to truth particles (the factor is even larger with machine learning taggers). The SPRA1 algorithm brings about a factor of two improvement relative to plain PF. The further gain from the SPRA2 algorithm is limited (and perhaps not statistically significant). Accordingly for our main W -tagger performance results in section 3.5 we use SPRA1, which is arguably also the most similar to the procedure used by CMS in ref. [88]. However at higher p_t 's or for measurements of the Lund plane, it is probably advantageous to use SPRA2.

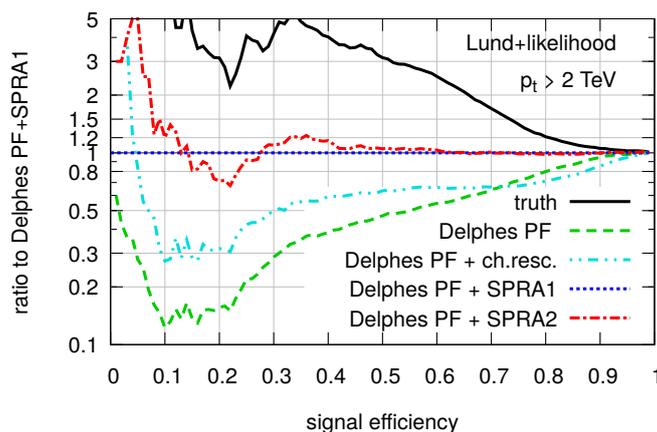


Figure 20. Performance of the Lund-likelihood discriminator, at truth level and with Delphes PF, with and without the SPRA algorithms, and also with charge rescaling. As a function of signal efficiency, the plot shows the ratio of background rejection relative to that obtained for Delphes PF+SPRA1. The log-likelihood maps have been determined separately for each setup (truth, Delphes PF, etc.). The generation and selection of jets is as described in section 3.4, using the dijet process, selecting jets with the requirement $p_t > 2$ TeV, $|y| < 2.5$.

Figure 20 also shows charge-rescaling, which performs less well than the SPRA approaches, though still better than Delphes particle flow alone. Our understanding of the worse performance relative to the SPRA approaches is as follows: the performance observed in figure 20 combines the performance for the Lund plane as seen in figure 18 with the performance for the jet mass that is observed in figure 17. While charge rescaling does well in figure 18, it performs somewhat worse than either of the SPRAs in figure 17. Note also that the region of figure 18 where charge rescaling performs better than SPRA1, $\ln 1/\Delta \gtrsim 3.5$, is a region that does not contribute dominantly to the discrimination power, cf. figure 9b.

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.

References

- [1] 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*, [arXiv:1709.04464](https://arxiv.org/abs/1709.04464) [[INSPIRE](https://inspirehep.net/literature/1709044)].
- [2] L. Asquith et al., *Jet Substructure at the Large Hadron Collider: Experimental Review*, [arXiv:1803.06991](https://arxiv.org/abs/1803.06991) [[INSPIRE](https://inspirehep.net/literature/1803069)].
- [3] J.M. Butterworth, A.R. Davison, M. Rubin and G.P. Salam, *Jet substructure as a new Higgs search channel at the LHC*, *Phys. Rev. Lett.* **100** (2008) 242001 [[arXiv:0802.2470](https://arxiv.org/abs/0802.2470)] [[INSPIRE](https://inspirehep.net/literature/770432)].
- [4] J. Thaler and L.-T. Wang, *Strategies to Identify Boosted Tops*, *JHEP* **07** (2008) 092 [[arXiv:0806.0023](https://arxiv.org/abs/0806.0023)] [[INSPIRE](https://inspirehep.net/literature/770432)].

- [5] D.E. Kaplan, K. Rehermann, M.D. Schwartz and B. Tweedie, *Top Tagging: A Method for Identifying Boosted Hadronically Decaying Top Quarks*, *Phys. Rev. Lett.* **101** (2008) 142001 [[arXiv:0806.0848](#)] [[INSPIRE](#)].
- [6] S.D. Ellis, C.K. Vermilion and J.R. Walsh, *Techniques for improved heavy particle searches with jet substructure*, *Phys. Rev. D* **80** (2009) 051501 [[arXiv:0903.5081](#)] [[INSPIRE](#)].
- [7] S.D. Ellis, C.K. Vermilion and J.R. Walsh, *Recombination Algorithms and Jet Substructure: Pruning as a Tool for Heavy Particle Searches*, *Phys. Rev. D* **81** (2010) 094023 [[arXiv:0912.0033](#)] [[INSPIRE](#)].
- [8] T. Plehn, G.P. Salam and M. Spannowsky, *Fat Jets for a Light Higgs*, *Phys. Rev. Lett.* **104** (2010) 111801 [[arXiv:0910.5472](#)] [[INSPIRE](#)].
- [9] J. Thaler and K. Van Tilburg, *Identifying Boosted Objects with N -subjettiness*, *JHEP* **03** (2011) 015 [[arXiv:1011.2268](#)] [[INSPIRE](#)].
- [10] J. Thaler and K. Van Tilburg, *Maximizing Boosted Top Identification by Minimizing N -subjettiness*, *JHEP* **02** (2012) 093 [[arXiv:1108.2701](#)] [[INSPIRE](#)].
- [11] A.J. Larkoski, G.P. Salam and J. Thaler, *Energy Correlation Functions for Jet Substructure*, *JHEP* **06** (2013) 108 [[arXiv:1305.0007](#)] [[INSPIRE](#)].
- [12] Y.-T. Chien, *Telescoping jets: Probing hadronic event structure with multiple R 's*, *Phys. Rev. D* **90** (2014) 054008 [[arXiv:1304.5240](#)] [[INSPIRE](#)].
- [13] A.J. Larkoski, I. Moult and D. Neill, *Power Counting to Better Jet Observables*, *JHEP* **12** (2014) 009 [[arXiv:1409.6298](#)] [[INSPIRE](#)].
- [14] A.J. Larkoski, J. Thaler and W.J. Waalewijn, *Gaining (Mutual) Information about Quark/Gluon Discrimination*, *JHEP* **11** (2014) 129 [[arXiv:1408.3122](#)] [[INSPIRE](#)].
- [15] I. Moult, L. Necib and J. Thaler, *New Angles on Energy Correlation Functions*, *JHEP* **12** (2016) 153 [[arXiv:1609.07483](#)] [[INSPIRE](#)].
- [16] G.P. Salam, L. Schunk and G. Soyez, *Dichroic subjettiness ratios to distinguish colour flows in boosted boson tagging*, *JHEP* **03** (2017) 022 [[arXiv:1612.03917](#)] [[INSPIRE](#)].
- [17] M.H. Seymour, *Jet shapes in hadron collisions: Higher orders, resummation and hadronization*, *Nucl. Phys. B* **513** (1998) 269 [[hep-ph/9707338](#)] [[INSPIRE](#)].
- [18] I. Feige, M.D. Schwartz, I.W. Stewart and J. Thaler, *Precision Jet Substructure from Boosted Event Shapes*, *Phys. Rev. Lett.* **109** (2012) 092001 [[arXiv:1204.3898](#)] [[INSPIRE](#)].
- [19] M. Dasgupta, A. Fregoso, S. Marzani and G.P. Salam, *Towards an understanding of jet substructure*, *JHEP* **09** (2013) 029 [[arXiv:1307.0007](#)] [[INSPIRE](#)].
- [20] M. Dasgupta, A. Fregoso, S. Marzani and A. Powling, *Jet substructure with analytical methods*, *Eur. Phys. J. C* **73** (2013) 2623 [[arXiv:1307.0013](#)] [[INSPIRE](#)].
- [21] Y.-T. Chien and I. Vitev, *Jet Shape Resummation Using Soft-Collinear Effective Theory*, *JHEP* **12** (2014) 061 [[arXiv:1405.4293](#)] [[INSPIRE](#)].
- [22] D. Bertolini, J. Thaler and J.R. Walsh, *The First Calculation of Fractional Jets*, *JHEP* **05** (2015) 008 [[arXiv:1501.01965](#)] [[INSPIRE](#)].
- [23] M. Dasgupta, L. Schunk and G. Soyez, *Jet shapes for boosted jet two-prong decays from first-principles*, *JHEP* **04** (2016) 166 [[arXiv:1512.00516](#)] [[INSPIRE](#)].

- [24] C. Frye, A.J. Larkoski, M.D. Schwartz and K. Yan, *Precision physics with pile-up insensitive observables*, [arXiv:1603.06375](#) [INSPIRE].
- [25] C. Frye, A.J. Larkoski, M.D. Schwartz and K. Yan, *Factorization for groomed jet substructure beyond the next-to-leading logarithm*, *JHEP* **07** (2016) 064 [[arXiv:1603.09338](#)] [INSPIRE].
- [26] M. Dasgupta, A. Powling, L. Schunk and G. Soyez, *Improved jet substructure methods: Y-splitter and variants with grooming*, *JHEP* **12** (2016) 079 [[arXiv:1609.07149](#)] [INSPIRE].
- [27] S. Marzani, L. Schunk and G. Soyez, *The jet mass distribution after Soft Drop*, *Eur. Phys. J. C* **78** (2018) 96 [[arXiv:1712.05105](#)] [INSPIRE].
- [28] S. Marzani, L. Schunk and G. Soyez, *A study of jet mass distributions with grooming*, *JHEP* **07** (2017) 132 [[arXiv:1704.02210](#)] [INSPIRE].
- [29] A.J. Larkoski, I. Moult and D. Neill, *Factorization and Resummation for Groomed Multi-Prong Jet Shapes*, *JHEP* **02** (2018) 144 [[arXiv:1710.00014](#)] [INSPIRE].
- [30] I. Moult, B. Nachman and D. Neill, *Convolved Substructure: Analytically Decorrelating Jet Substructure Observables*, *JHEP* **05** (2018) 002 [[arXiv:1710.06859](#)] [INSPIRE].
- [31] CMS collaboration, *Identification techniques for highly boosted W bosons that decay into hadrons*, *JHEP* **12** (2014) 017 [[arXiv:1410.4227](#)] [INSPIRE].
- [32] ATLAS collaboration, *Identification of boosted, hadronically decaying W bosons and comparisons with ATLAS data taken at $\sqrt{s} = 8$ TeV*, *Eur. Phys. J. C* **76** (2016) 154 [[arXiv:1510.05821](#)] [INSPIRE].
- [33] ATLAS collaboration, *Search for high-mass diboson resonances with boson-tagged jets in proton-proton collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector*, *JHEP* **12** (2015) 055 [[arXiv:1506.00962](#)] [INSPIRE].
- [34] ATLAS collaboration, *Measurement of the Soft-Drop Jet Mass in pp Collisions at $\sqrt{s} = 13$ TeV with the ATLAS Detector*, *Phys. Rev. Lett.* **121** (2018) 092001 [[arXiv:1711.08341](#)] [INSPIRE].
- [35] CMS collaboration, *Measurement of jet substructure observables in $t\bar{t}$ events from pp collisions at $\sqrt{s} = 13$ TeV*, [CMS-PAS-TOP-17-013](#) (2017).
- [36] J. Cogan, M. Kagan, E. Strauss and A. Schwartzman, *Jet-Images: Computer Vision Inspired Techniques for Jet Tagging*, *JHEP* **02** (2015) 118 [[arXiv:1407.5675](#)] [INSPIRE].
- [37] 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].
- [38] P.T. Komiske, E.M. Metodiev and M.D. Schwartz, *Deep learning in color: towards automated quark/gluon jet discrimination*, *JHEP* **01** (2017) 110 [[arXiv:1612.01551](#)] [INSPIRE].
- [39] 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].
- [40] G. Louppe, K. Cho, C. Becot and K. Cranmer, *QCD-Aware Recursive Neural Networks for Jet Physics*, [arXiv:1702.00748](#) [INSPIRE].
- [41] S. Egan, W. Fedorko, A. Lister, J. Pearkes and C. Gay, *Long Short-Term Memory (LSTM) networks with jet constituents for boosted top tagging at the LHC*, [arXiv:1711.09059](#) [INSPIRE].

- [42] A. Andreassen, I. Feige, C. Frye and M.D. Schwartz, *JUNIPR: a Framework for Unsupervised Machine Learning in Particle Physics*, [arXiv:1804.09720](#) [INSPIRE].
- [43] K. Datta and A. Larkoski, *How Much Information is in a Jet?*, *JHEP* **06** (2017) 073 [[arXiv:1704.08249](#)] [INSPIRE].
- [44] K. Datta and A.J. Larkoski, *Novel Jet Observables from Machine Learning*, *JHEP* **03** (2018) 086 [[arXiv:1710.01305](#)] [INSPIRE].
- [45] P.T. Komiske, E.M. Metodiev and J. Thaler, *Energy flow polynomials: A complete linear basis for jet substructure*, *JHEP* **04** (2018) 013 [[arXiv:1712.07124](#)] [INSPIRE].
- [46] E.M. Metodiev, B. Nachman and J. Thaler, *Classification without labels: Learning from mixed samples in high energy physics*, *JHEP* **10** (2017) 174 [[arXiv:1708.02949](#)] [INSPIRE].
- [47] L.M. Dery, B. Nachman, F. Rubbo and A. Schwartzman, *Weakly Supervised Classification in High Energy Physics*, *JHEP* **05** (2017) 145 [[arXiv:1702.00414](#)] [INSPIRE].
- [48] B. Andersson, G. Gustafson, L. Lönnblad and U. Pettersson, *Coherence Effects in Deep Inelastic Scattering*, *Z. Phys. C* **43** (1989) 625 [INSPIRE].
- [49] Y.L. Dokshitzer, G.D. Leder, S. Moretti and B.R. Webber, *Better jet clustering algorithms*, *JHEP* **08** (1997) 001 [[hep-ph/9707323](#)] [INSPIRE].
- [50] M. Wobisch and T. Wengler, *Hadronization corrections to jet cross-sections in deep inelastic scattering*, in *Monte Carlo generators for HERA physics. Proceedings, Workshop, Hamburg, Germany, 1998–1999*, pp. 270–279, [hep-ph/9907280](#) [INSPIRE].
- [51] M. Cacciari, G.P. Salam and G. Soyez, *The anti- k_t jet clustering algorithm*, *JHEP* **04** (2008) 063 [[arXiv:0802.1189](#)] [INSPIRE].
- [52] T. Sjöstrand et al., *An Introduction to PYTHIA 8.2*, *Comput. Phys. Commun.* **191** (2015) 159 [[arXiv:1410.3012](#)] [INSPIRE].
- [53] 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].
- [54] 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].
- [55] M. Dasgupta, F.A. Dreyer, G.P. Salam and G. Soyez, *Inclusive jet spectrum for small-radius jets*, *JHEP* **06** (2016) 057 [[arXiv:1602.01110](#)] [INSPIRE].
- [56] M. Dasgupta and G.P. Salam, *Resummation of nonglobal QCD observables*, *Phys. Lett. B* **512** (2001) 323 [[hep-ph/0104277](#)] [INSPIRE].
- [57] R.B. Appleby and M.H. Seymour, *Nonglobal logarithms in interjet energy flow with kt clustering requirement*, *JHEP* **12** (2002) 063 [[hep-ph/0211426](#)] [INSPIRE].
- [58] Y. Delenda, R. Appleby, M. Dasgupta and A. Banfi, *On QCD resummation with $k(t)$ clustering*, *JHEP* **12** (2006) 044 [[hep-ph/0610242](#)] [INSPIRE].
- [59] T. Gleisberg et al., *Event generation with SHERPA 1.1*, *JHEP* **02** (2009) 007 [[arXiv:0811.4622](#)] [INSPIRE].
- [60] J. Bellm et al., *HERWIG 7.1 Release Note*, [arXiv:1705.06919](#) [INSPIRE].
- [61] Y.L. Dokshitzer and B.R. Webber, *Calculation of power corrections to hadronic event shapes*, *Phys. Lett. B* **352** (1995) 451 [[hep-ph/9504219](#)] [INSPIRE].

- [62] Y.L. Dokshitzer, G. Marchesini and B.R. Webber, *Dispersive approach to power behaved contributions in QCD hard processes*, *Nucl. Phys. B* **469** (1996) 93 [[hep-ph/9512336](#)] [[INSPIRE](#)].
- [63] J.R. Forshaw, A. Kyrieleis and M.H. Seymour, *Super-leading logarithms in non-global observables in QCD*, *JHEP* **08** (2006) 059 [[hep-ph/0604094](#)] [[INSPIRE](#)].
- [64] S. Catani, D. de Florian and G. Rodrigo, *Space-like (versus time-like) collinear limits in QCD: Is factorization violated?*, *JHEP* **07** (2012) 026 [[arXiv:1112.4405](#)] [[INSPIRE](#)].
- [65] S. Catani and M.H. Seymour, *The dipole formalism for the calculation of QCD jet cross-sections at next-to-leading order*, *Phys. Lett. B* **378** (1996) 287 [[hep-ph/9602277](#)] [[INSPIRE](#)].
- [66] S. Catani and M.H. Seymour, *A general algorithm for calculating jet cross-sections in NLO QCD*, *Nucl. Phys. B* **485** (1997) 291 [Erratum *ibid.* **B 510** (1998) 503] [[hep-ph/9605323](#)] [[INSPIRE](#)].
- [67] S. Catani, Y.L. Dokshitzer, M. Olsson, G. Turnock and B.R. Webber, *New clustering algorithm for multi-jet cross-sections in e^+e^- annihilation*, *Phys. Lett. B* **269** (1991) 432 [[INSPIRE](#)].
- [68] M. Cacciari, G.P. Salam and G. Soyez, *FastJet User Manual*, *Eur. Phys. J. C* **72** (2012) 1896 [[arXiv:1111.6097](#)] [[INSPIRE](#)].
- [69] 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](#)].
- [70] A.J. Larkoski, S. Marzani, G. Soyez and J. Thaler, *Soft Drop*, *JHEP* **05** (2014) 146 [[arXiv:1402.2657](#)] [[INSPIRE](#)].
- [71] A. Larkoski, S. Marzani, J. Thaler, A. Tripathi and W. Xue, *Exposing the QCD Splitting Function with CMS Open Data*, *Phys. Rev. Lett.* **119** (2017) 132003 [[arXiv:1704.05066](#)] [[INSPIRE](#)].
- [72] D.E. Soper and M. Spannowsky, *Finding physics signals with shower deconstruction*, *Phys. Rev. D* **84** (2011) 074002 [[arXiv:1102.3480](#)] [[INSPIRE](#)].
- [73] S. Hochreiter and J. Schmidhuber, *Long short-term memory*, *Neural Comput.* **9** (1997) 1735.
- [74] K. Cho, B. van Merriënboer, Ç. Gülçehre, F. Bougares, H. Schwenk and Y. Bengio, *Learning phrase representations using RNN encoder-decoder for statistical machine translation*, [arXiv:1406.1078](#).
- [75] F. Chollet, *Keras*, <https://keras.io>, (2015).
- [76] M. Abadi et al., *TensorFlow: Large-scale machine learning on heterogeneous systems*, (2015), software available from <https://www.tensorflow.org/>.
- [77] K. He, X. Zhang, S. Ren and J. Sun, *Delving Deep into Rectifiers: Surpassing Human-Level Performance on ImageNet Classification*, [arXiv:1502.01852](#) [[INSPIRE](#)].
- [78] D.P. Kingma and J. Ba, *Adam: A Method for Stochastic Optimization*, [arXiv:1412.6980](#) [[INSPIRE](#)].
- [79] A. Hocker et al., *TMVA — Toolkit for Multivariate Data Analysis*, *PoS(ACAT)040* [[physics/0703039](#)] [[INSPIRE](#)].

- [80] J.R. Andersen et al., *Les Houches 2017: Physics at TeV Colliders Standard Model Working Group Report*, [arXiv:1803.07977](#) [INSPIRE].
- [81] DELPHES 3 collaboration, J. de Favereau et al., *DELPHES 3, A modular framework for fast simulation of a generic collider experiment*, *JHEP* **02** (2014) 057 [[arXiv:1307.6346](#)] [INSPIRE].
- [82] A. Katz, M. Son and B. Tweedie, *Jet Substructure and the Search for Neutral Spin-One Resonances in Electroweak Boson Channels*, *JHEP* **03** (2011) 011 [[arXiv:1010.5253](#)] [INSPIRE].
- [83] M. Son, C. Spethmann and B. Tweedie, *Diboson-Jets and the Search for Resonant Zh Production*, *JHEP* **08** (2012) 160 [[arXiv:1204.0525](#)] [INSPIRE].
- [84] S. Schaetzel and M. Spannowsky, *Tagging highly boosted top quarks*, *Phys. Rev. D* **89** (2014) 014007 [[arXiv:1308.0540](#)] [INSPIRE].
- [85] A.J. Larkoski, F. Maltoni and M. Selvaggi, *Tracking down hyper-boosted top quarks*, *JHEP* **06** (2015) 032 [[arXiv:1503.03347](#)] [INSPIRE].
- [86] S. Bressler, T. Flacke, Y. Kats, S.J. Lee and G. Perez, *Hadronic Calorimeter Shower Size: Challenges and Opportunities for Jet Substructure in the Superboosted Regime*, *Phys. Lett. B* **756** (2016) 137 [[arXiv:1506.02656](#)] [INSPIRE].
- [87] Z. Han, M. Son and B. Tweedie, *Top-Tagging at the Energy Frontier*, *Phys. Rev. D* **97** (2018) 036023 [[arXiv:1707.06741](#)] [INSPIRE].
- [88] CMS collaboration, *V Tagging Observables and Correlations*, CMS-PAS-JME-14-002.
- [89] ATLAS Collaboration, *Jet mass reconstruction with the ATLAS Detector in early Run 2 data*, ATLAS-CONF-2016-035.
- [90] M. Cacciari, G.P. Salam and G. Soyez, *SoftKiller, a particle-level pileup removal method*, *Eur. Phys. J. C* **75** (2015) 59 [[arXiv:1407.0408](#)] [INSPIRE].
- [91] D. Bertolini, P. Harris, M. Low and N. Tran, *Pileup Per Particle Identification*, *JHEP* **10** (2014) 059 [[arXiv:1407.6013](#)] [INSPIRE].
- [92] P. Berta, M. Spousta, D.W. Miller and R. Leitner, *Particle-level pileup subtraction for jets and jet shapes*, *JHEP* **06** (2014) 092 [[arXiv:1403.3108](#)] [INSPIRE].
- [93] P.T. Komiske, E.M. Metodiev, B. Nachman and M.D. Schwartz, *Pileup Mitigation with Machine Learning (PUMML)*, *JHEP* **12** (2017) 051 [[arXiv:1707.08600](#)] [INSPIRE].
- [94] M. Cacciari and G.P. Salam, *Pileup subtraction using jet areas*, *Phys. Lett. B* **659** (2008) 119 [[arXiv:0707.1378](#)] [INSPIRE].
- [95] M. Cacciari, G.P. Salam and G. Soyez, *The Catchment Area of Jets*, *JHEP* **04** (2008) 005 [[arXiv:0802.1188](#)] [INSPIRE].
- [96] D. Krohn, J. Thaler and L.-T. Wang, *Jet Trimming*, *JHEP* **02** (2010) 084 [[arXiv:0912.1342](#)] [INSPIRE].
- [97] F.A. Dreyer, L. Necib, G. Soyez and J. Thaler, *Recursive Soft Drop*, *JHEP* **06** (2018) 093 [[arXiv:1804.03657](#)] [INSPIRE].
- [98] I.J. Goodfellow et al., *Generative Adversarial Networks*, *Adv. Neural Inf. Process. Syst.* **27** (2014) 2672 [[arXiv:1406.2661](#)] [INSPIRE].

- [99] G. Louppe, M. Kagan and K. Cranmer, *Learning to Pivot with Adversarial Networks*, [arXiv:1611.01046](#) [INSPIRE].
- [100] C. Shimmin et al., *Decorrelated Jet Substructure Tagging using Adversarial Neural Networks*, *Phys. Rev. D* **96** (2017) 074034 [[arXiv:1703.03507](#)] [INSPIRE].
- [101] H.A. Andrews et al., *Novel tools and observables for jet physics in heavy-ion collisions*, [arXiv:1808.03689](#) [INSPIRE].
- [102] Y.-T. Chien and R. Kunnawalkam Elayavalli, *Probing heavy ion collisions using quark and gluon jet substructure*, [arXiv:1803.03589](#) [INSPIRE].
- [103] ALICE collaboration, H. Andrews, *Exploring phase space of jet splittings at alice using grooming and recursive techniques*, (2018). Talk at *Quark Matter 2018*, Venice, Italy, https://indico.cern.ch/event/656452/contributions/2869941/attachments/1649044/2636550/HarryAndrews_QuarkMatter18Final.pdf.
- [104] J. Dolen, P. Harris, S. Marzani, S. Rappoccio and N. Tran, *Thinking outside the ROCs: Designing Decorrelated Taggers (DDT) for jet substructure*, *JHEP* **05** (2016) 156 [[arXiv:1603.00027](#)] [INSPIRE].
- [105] Z. Hall and J. Thaler, *Photon isolation and jet substructure*, *JHEP* **09** (2018) 164 [[arXiv:1805.11622](#)] [INSPIRE].
- [106] K.S. Tai, R. Socher and C.D. Manning, *Improved semantic representations from tree-structured long short-term memory networks*, [[arXiv:1503.00075](#)].
- [107] M. Dasgupta, F.A. Dreyer, K. Hamilton, P.F. Monni and G.P. Salam, *Logarithmic accuracy of parton showers: a fixed-order study*, *JHEP* **09** (2018) 033 [[arXiv:1805.09327](#)] [INSPIRE].
- [108] B.T. Elder and J. Thaler, *Aspects of Track-Assisted Mass*, [arXiv:1805.11109](#) [INSPIRE].