Casimir meets Poisson: improved quark/gluon discrimination with counting observables

Christopher Frye,[a] **Andrew J. Larkoski,**[b] **Jesse Thaler**[c] **and Kevin Zhou**[c]

[a]*Department of Physics, Harvard University,*
17 Oxford Street, Cambridge, MA 02138, U.S.A.

[b]*Physics Department, Reed College,*
3203 SE Woodstock Blvd, Portland, OR 97202, U.S.A.

[c]*Center for Theoretical Physics, Massachusetts Institute of Technology,*
77 Massachusetts Avenue, Cambridge, MA 02139, U.S.A.

E-mail: frye@physics.harvard.edu, larkoski@reed.edu, jthaler@mit.edu, knzhou@mit.edu

Abstract: Charged track multiplicity is among the most powerful observables for discriminating quark- from gluon-initiated jets. Despite its utility, it is not infrared and collinear (IRC) safe, so perturbative calculations are limited to studying the energy evolution of multiplicity moments. While IRC-safe observables, like jet mass, are perturbatively calculable, their distributions often exhibit Casimir scaling, such that their quark/gluon discrimination power is limited by the ratio of quark to gluon color factors. In this paper, we introduce new IRC-safe counting observables whose discrimination performance exceeds that of jet mass and approaches that of track multiplicity. The key observation is that track multiplicity is approximately Poisson distributed, with more suppressed tails than the Sudakov peak structure from jet mass. By using an iterated version of the soft drop jet grooming algorithm, we can define a "soft drop multiplicity" which is Poisson distributed at leading-logarithmic accuracy. In addition, we calculate the next-to-leading-logarithmic corrections to this Poisson structure. If we allow the soft drop groomer to proceed to the end of the jet branching history, we can define a collinear-unsafe (but still infrared-safe) counting observable. Exploiting the universality of the collinear limit, we define generalized fragmentation functions to study the perturbative energy evolution of collinear-unsafe multiplicity.

Contents

1 Introduction

The fantastic jet reconstruction performance of ATLAS and CMS [1, 2] — along with increasingly sophisticated tools to predict jet properties from first principles [3–10] — has led to significant advances in the field of jet substructure [11–14]. A key goal in jet substructure is to robustly discriminate quark-initiated jets from gluon-initiated jets [15–23], with many applications to new physics searches at the Large Hadron Collider (LHC) (see e.g. [24–26]). In the eikonal limit, quarks and gluons differ only by their respective

color charges, $C_F = 4/3$ versus $C_A = 3$, such that gluon jets emit more soft gluon radiation than quark jets. At this order, the difference between quark and gluon radiation patterns is controlled entirely by the Casimir ratio $C_A/C_F = 9/4$, which drives (and limits) the expected separation power between quark and gluon jets.

One of the most powerful quark/gluon discriminants is hadron multiplicity, or its charged-particle-only variant, track multiplicity n_{tr} [15, 16, 18, 27–29]. This is an effective discriminant because the average track multiplicity within quark and gluon jets scales approximately as (see e.g. [30, 31])

$$\frac{\langle n_{tr} \rangle_g}{\langle n_{tr} \rangle_q} \simeq \frac{C_A}{C_F}. \tag{1.1}$$

Since multiplicity is not infrared and collinear (IRC) safe, though, it is difficult to predict its discrimination performance from first principles.[1] On the other hand, IRC-safe observables like jet mass and jet width are analytically tractable [33–35], but they exhibit worse quark/gluon performance than multiplicity. The reason is that these discriminants are dominated by a single emission at leading-logarithmic (LL) accuracy, giving rise to Casimir scaling of the quark/gluon discrimination power,

$$(\text{gluon mistag rate}) \simeq (\text{quark efficiency})^{C_A/C_F}, \tag{1.2}$$

and therefore relatively weak separation between quark and gluon jets. This Casimir scaling behavior holds for any observable with a Sudakov form factor at LL accuracy, including a wide range of IRC-safe additive observables [17]. While one can try to interpolate between the IRC-unsafe and IRC-safe regimes using generalized angularities [18], track multiplicity remains one of the best performing — yet analytically puzzling — quark/gluon discriminants.

In this paper, we introduce a new class of "counting observables" that are IRC safe, yet yield comparable quark/gluon performance to track multiplicity. Unlike additive observables, which are only sensitive to a single emission at LL order, these counting observables are directly sensitive to multiple emissions at LL, allowing them to exceed the performance estimate in eq. (1.2). Crucially, the quark/gluon performance of counting observables still depends on the color factors C_A and C_F, but instead of being described by Sudakov form factors, these observables are described by Poisson distributions; this allows their discrimination power to improve as more emissions are included. These counting observables not only clarify the underlying reason why track multiplicity performs so well, but they also demonstrate the new kinds of analytic structures possible from IRC-safe but non-additive observables.[2]

The counting observables we study are based on an iterated variant of soft drop declustering [40]. As a grooming procedure, soft drop starts at the trunk of an angular-ordered

---

[1]It is possible to calculate the evolution with energy of the multiplicity moments; see, e.g., ref. [32] for a review.

[2]An alternative counting method was proposed in ref. [19], which considers associated subjets outside of the jet boundary. Additionally, there has been interest in understanding the scaling of the cross section at high jet multiplicity [36–39]. Here, we focus on counting subjets within the jet of interest.

clustering tree [41, 42] and sequentially removes soft branches with small momentum fraction z_{ij} until a hard branching is found. At a step in the clustering tree where branches i and j split, the splitting is retained in the groomed jet if the momentum fraction satisfies

$$z_{ij} > z_{\text{cut}} \left(\frac{\theta_{ij}}{R_0} \right)^{\beta} , \tag{1.3}$$

where θ_{ij} is an appropriately defined relative angle between branches i and j, and R_0 is the jet radius. For appropriate choices of the soft drop parameters z_{cut} and β, observables defined on the groomed jet are automatically infrared (but not necessarily collinear) safe. While the original soft drop procedure terminates once it finds a hard $1 \rightarrow 2$ splitting, the iterated variant we employ in this paper continues, following the hardest branch (the "trunk") through multiple levels until an angular cutoff scale θ_{cut} is reached.

The simplest counting observable we can define using iterated soft drop (ISD) is just the total number of emissions from the trunk of the clustering tree that ISD records. In particular, this includes all emissions $n \in [1, n_{\text{max}}]$ that satisfy the soft drop condition and lie outside the θ_{cut} cone. We call this observable "soft drop multiplicity",

$$n_{\text{SD}}(z_{\text{cut}}, \beta, \theta_{\text{cut}}) = \sum_n 1, \tag{1.4}$$

which depends on the choice of ISD parameters. It is complementary to the "soft drop level" observable $L_{\text{SD}}(\beta)$ introduced in ref. [43], which also iteratively applies the soft drop condition, but changes the z_{cut} scale. As long as $z_{\text{cut}} > 0$, soft drop multiplicity is infrared safe.

With $\theta_{\text{cut}} > 0$ or $\beta < 0$, n_{SD} is collinear safe as well, so we can use analytic resummation tools to predict its discrimination power. We do this to resum large logarithms of z_{cut} and θ_{cut}, which are of soft and collinear origin, respectively, and which lead to a double-logarithmic observable. The analysis at LL order is straightforward, yielding a Poisson distribution whose average value is set by the phase space "area" of counted emissions. This leads to quark/gluon discrimination power which approaches that of track multiplicity, particularly in the case of $\beta = -1$. Moving from LL to next-to-leading-logarithmic (NLL) order, one finds a slight decrease in discrimination power, due in part to the jet-flavor mixing that appears at this accuracy. We implement the NLL calculation through a set of evolution equations that have a similar form to parton evolution.

With $\theta_{\text{cut}} = 0$ and $\beta \geq 0$, the soft drop multiplicity n_{SD} is no longer collinear safe, so we cannot predict its absolute discrimination power. That said, for the special case of $\beta = 0$ (which was initially introduced as the modified mass drop tagger [3, 44]), we can use renormalization group (RG) techniques to predict the *evolution* of its discrimination power. When $\beta = 0$, soft drop multiplicity has purely collinear divergences, which can be absorbed into a generalized fragmentation function (GFF) that depends on the RG scale μ [45]. After extracting this GFF at low scales (either from LHC data[3] or parton shower

---

[3] Just as for parton distribution functions and ordinary fragmentation functions, extracting GFFs involves matching to fixed-order calculations, as described in ref. [45]. These fixed-order calculations involve a mixture of quark and gluon final-state partons, so multiple event samples with different quark/gluon fractions are required to disentangle the contributions from quark and gluon GFFs.

simulations), one can use a perturbative DGLAP-like evolution equation to predict the discrimination power achievable at higher scales. Intriguingly, in the limit of pure Yang-Mills, one can show that at lowest order, the soft drop multiplicity asymptotes to a true Poisson distribution at large values of μ, such that it behaves like an idealized counting observable (albeit in a theory with only gluons).

The remainder of this paper is organized as follows. In section 2, we define the ISD procedure, introduce soft drop multiplicity, and take a first look at its distribution using parton shower generators. In section 3, we perform an LL analysis, focusing on the contrast between soft drop multiplicity's Poisson behavior and the more familiar Sudakov-peak behavior of additive observables. We extend our analytic calculations to NLL order in section 4 and compare our analytic distributions to those obtained from various parton showers. We consider the collinear-unsafe case of $\theta_{\rm cut} = 0$ and $\beta = 0$ in section 5, deriving the corresponding RG evolution equations and presenting numerical results based on parton shower inputs. We present our conclusions in section 6.

In an appendix, we demonstrate that our analytical tools can also be used to study more general ISD observables, in particular the weighted multiplicity $\sum_n (z_n)^\kappa$ which weights each counted emission according to its momentum fraction z_n. Soft drop multiplicity is a special case ($\kappa = 0$) of this more general observable, and the one most useful for quark/gluon discrimination.

2 Counting observables from soft drop declustering

2.1 Iterated soft drop

Our counting observables are defined using an iterated variant of the soft drop declustering algorithm. We briefly review soft drop here for convenience and to establish conventions.

The soft drop grooming procedure can be applied to any jet found using a standard jet algorithm of characteristic radius R_0. After reclustering the jet using the Cambridge/Aachen (C/A) algorithm [41, 42], soft drop involves sequentially undoing the cluster history to remove wide-angle soft radiation and identify hard 2-prong substructure. For each C/A branching into subjets i and j, there are quantities z_{ij} and θ_{ij}, which are defined differently for different collider environments:

$$e^+ e^- \text{ collisions:} \qquad z_{ij} = \frac{\min(E_i, E_j)}{E_i + E_j}, \qquad \theta_{ij} = \text{angle between } i, j, \qquad (2.1)$$

$$pp \text{ collisions:} \qquad z_{ij} = \frac{\min(p_{Ti}, p_{Tj})}{p_{Ti} + p_{Tj}}, \qquad \theta_{ij} = \Delta R_{ij}, \qquad (2.2)$$

where ΔR represents distance in the rapidity-azimuth plane. The soft drop grooming algorithm can be summarized as follows:

1. Traverse the C/A clustering tree, beginning at the trunk and sequentially examining each branching.

2. Upon arriving at a branching into subjects i and j, check whether the soft drop condition is satisfied:

$$z_{ij} > z_{\text{cut}} \left(\frac{\theta_{ij}}{R_0} \right)^{\beta} , \tag{2.3}$$

where z_{cut} and β are fixed parameters of the algorithm. If so, the algorithm terminates; stop grooming and return the jet as is.

3. If the branching fails this condition, remove the softer of the two subjects (i or j) from the groomed jet and return to Step 2 on the next branching in the remaining clustering tree.

Our analysis is based on ISD where the soft drop algorithm is iterated. In this case, the procedure does not terminate when a hard branching is found, but is instead iteratively applied to the harder of the two subjects. This continues until an angular cutoff is reached, so in addition to z_{cut} and β, ISD depends on an additional parameter θ_{cut}. While ISD could be used as a grooming procedure in its own right, the primary purpose of ISD in this paper is to determine which set of (z_{ij}, θ_{ij}) branchings contribute to the observables we define below. For this purpose, the ISD algorithm proceeds as follows:

1'. Set the counter n equal to 1. Traverse the C/A clustering tree, beginning at the trunk and sequentially examining each branching.

2'. Upon arriving at a branching into subjects i and j, check whether the branching angle satisfies

$$\theta_{ij} > \theta_{\text{cut}} . \tag{2.4}$$

If not, the algorithm terminates.

3'. If $\theta_{ij} > \theta_{\text{cut}}$, then check whether the soft drop condition is satisfied:

$$z_{ij} > z_{\text{cut}} \left(\frac{\theta_{ij}}{R_0} \right)^{\beta} . \tag{2.5}$$

If not, return to Step 2' on the harder of subjects i and j.

4'. If the soft drop condition is satisfied, define

$$z_n \equiv z_{ij} , \qquad \theta_n \equiv \theta_{ij} . \tag{2.6}$$

Then increment $n \to n + 1$ and return to Step 2' on the harder of subjects i and j.

Because we recurse to the harder subjet at each junction, we think of each (z_n, θ_n) splitting as an emission from the "hard core" of the jet and refer to the above procedure as traversing the "trunk" of the clustering tree. A schematic of this procedure is shown in figure 1.

To emphasize, we are not using ISD as an alternative grooming technique to soft drop. In fact, we have found no need to refer to the ISD-groomed jet explicitly in our analysis. Instead, we employ ISD simply as a method to obtain an IRC-safe set of (z_n, θ_n) values to define our counting observables. Of course, the specific values of (z_n, θ_n) depend on the
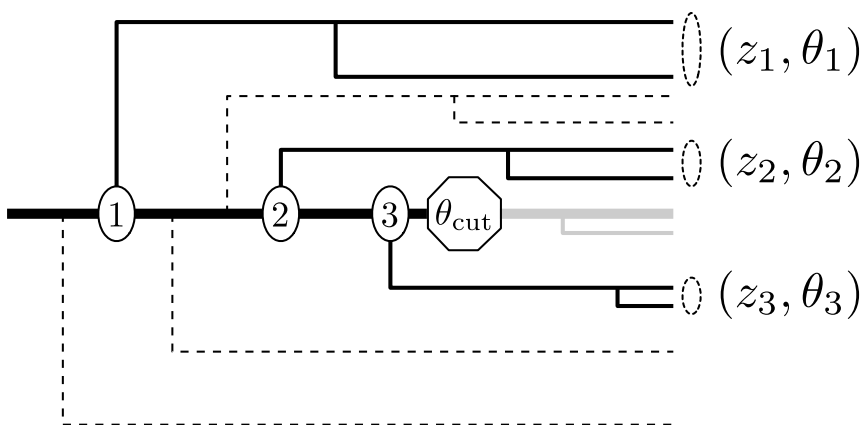
Figure 1. Illustration of the ISD procedure. A C/A tree is declustered from the trunk (thick line), defined by the hardest p_T branches. If a node fails the soft drop condition, it is removed from consideration (dashed lines). If a node passes the soft drop condition after n iterations, this defines the value of (z_n, θ_n). The declustering stops at an angular scale of θ_{cut}, and subsequent nodes are not considered further (gray lines).

precise choice of ISD procedure. In this paper, we focus on the soft drop multiplicity, which counts emissions from the trunk of the clustering tree, and have defined ISD accordingly. In section 2.3, we consider variants of soft drop multiplicity, with corresponding variants to the ISD procedure.

To demonstrate the qualitative behavior of observables defined below in this section, we present results from parton shower simulations. We separately generate $pp \rightarrow Z + q$ and $pp \rightarrow Z + g$ events at center-of-mass energy 13 TeV using MADGRAPH 2.4.0 and let the Z decay to neutrinos for simplicity. We then shower the events through VINCIA 2.0.01 [46, 47], a plug-in to PYTHIA 8.215 [48], with default tuning parameters.[4] Jet are identified using the anti-k_t algorithm [49] with radius $R_0 = 0.6$ in FASTJET 3.1.3 [50]. We use a sample of events in which the hardest jet with $|\eta| < 2.5$ has p_T between 450 and 550 GeV. We recluster and measure our observables on the hardest jet from each event using FASTJET. Because ISD is sufficiently different from ordinary soft drop, we do not use the `RecursiveTools` FJCONTRIB [51], but rather directly traverse the C/A tree in our analysis. We plan to make our code available publicly in a future release of FJCONTRIB.

2.2 Soft drop multiplicity

The (z_n, θ_n) values from ISD allow us to define a variety of interesting jet observables. Here, we focus on soft drop multiplicity n_{SD}, which is simply the total count of the recorded (z_n, θ_n) pairs. This observable, defined already in eq. (1.4), depends implicitly on the ISD parameters z_{cut}, β, and θ_{cut}. Among all of the observables we tested, n_{SD} appears to perform the best for quark/gluon discrimination. We discuss more general observables in section 2.3 and appendix A.

---

[4]In section 4, we show results from four different parton shower generators. Here, we use VINCIA as a representative example since it makes predictions which are intermediate relative to the other generators.
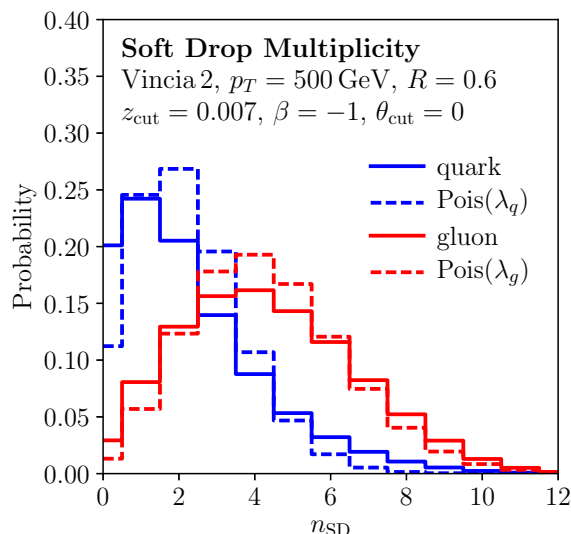
Figure 2. Distribution of the soft drop multiplicity n_{SD} from VINCIA 2.0.01. Using the IRC-safe benchmark parameters in eq. (2.7), we find good quark/gluon discrimination power. For comparison, we show Poisson distributions with the same means as dashed curves ($\lambda_q = 2.2$, $\lambda_g = 4.3$).

As defined above, ISD only follows the harder branch (i.e. the trunk) at each junction of the clustering tree. Therefore, n_{SD} effectively counts emissions from the hard core of the jet, down to the angular resolution scale θ_{cut}. When $z_{\mathrm{cut}} = \theta_{\mathrm{cut}} = 0$, n_{SD} is simply the depth of the trunk of the C/A tree.

When $z_{\mathrm{cut}} > 0$, the soft drop multiplicity is infrared safe, as all soft emissions at finite angles fail the soft drop condition in eq. (2.5). When $\theta_{\mathrm{cut}} > 0$, soft drop multiplicity is also collinear safe, since an exactly collinear splitting along the trunk does not satisfy eq. (2.4). Alternatively, $\beta < 0$ also gives collinear-safe distributions, since an exactly collinear splitting along the trunk does not satisfy eq. (2.5). The borderline case of $\theta_{\mathrm{cut}} = 0$ and $\beta = 0$ is collinear unsafe, but it can be handled using RG methods, as shown in section 5.

In figure 2, we show the soft drop multiplicity distributions for quark and gluon jets as extracted from VINCIA. Results are given using the benchmark parameters

$$z_{\mathrm{cut}} = 0.007, \qquad \beta = -1, \qquad \theta_{\mathrm{cut}} = 0. \tag{2.7}$$

This benchmark is chosen to maximize quark/gluon discrimination power while retaining perturbative calculability, as discussed in section 3. The distributions are approximately Poisson and yield good quark/gluon discrimination power.

2.3 Multiplicity variants

While the focus of this paper is on soft drop multiplicity n_{SD}, many other observables could be defined using the (z_n, θ_n) values recorded by ISD. For example, the techniques developed in this paper can be directly applied to the weighted soft drop multiplicity,

$$n_{\mathrm{SD}}^{(\kappa)} = \sum_n z_n^{\kappa}. \tag{2.8}$$

Note that soft drop multiplicity is a special case ($\kappa = 0$) of this more general observable, with the same criteria for IRC safety. We study the weighted multiplicity in detail in appendix A, but find its quark/gluon discrimination power to be inferior to the discrete $\kappa = 0$ case. In fact, LL reasoning leads one to expect the soft drop multiplicity n_{SD} to have the best discrimination power of any observable defined on the (z_n, θ_n) values; see the end of section 3.3 for a short discussion.

Nevertheless, several other promising variants of soft drop multiplicity might prove useful:

- The weighted soft drop multiplicity in eq. (2.8) only refers to the momentum fractions z_n in the sum over emissions. One could also consider an angle-weighted variant

$$\sum_n z_n^\kappa \theta_n^\alpha,\qquad(2.9)$$

or indeed any function of z_n and θ_n. The potential advantage of including θ_n information is that even for $\theta_{\mathrm{cut}} = 0$, such observables would be collinear safe for $\alpha > 0$.

- Instead of counting emissions only from the trunk of the C/A tree, we could extend the sum to include all branchings down to the angular resolution θ_{cut}. This multiplicity variant would require a modification of the ISD algorithm: in step 4′, the recursion would be applied to both subjets i and j, not just the harder one. This is a step in similarity towards full hadron multiplicity, reducing to it exactly when $z_{\mathrm{cut}} = \theta_{\mathrm{cut}} = 0$. This variant of soft drop multiplicity is more difficult to study analytically, however, due to the nonlinear structure of the recursion. Moreover, it is not clear that this variant would provide a performance advantage over n_{SD}. While gluons emitted from the hard core of a quark (gluon) jet give rise to factors of C_F (C_A), subsequent emissions from those gluons give rise to factors of C_A regardless of the jet flavor; this might wash out quark/gluon discrimination power.

- The original soft drop algorithm uses a C/A tree to mimic the angular-ordered structure of the parton shower. One could also study variants based on reclustering with the generalized-k_t algorithm with exponent p [49, 50]. The C/A algorithm used above corresponds to $p = 0$, while the k_t algorithm uses $p = 1$. For this variant, it would make sense to replace the angular cut θ_{cut} with a cut d_{cut} on the generalized distance measure d_{ij}.

This last k_t variant is of particular interest, given the discussion below in section 3.3. Nonperturbative physics typically dominates when $k_t \simeq \Lambda_{\mathrm{QCD}}$, so it makes sense to use a clustering algorithm where the clustering scale is "parallel" to the nonperturbative scale. This variant of n_{SD} would then allow the nonperturbative phase space to be clearly separated from the perturbative region and avoided. This would open up as much perturbative phase space for measured emissions as possible. We note that it is possible to mimic some of the LL structure of the k_t variant by using ISD with $\beta = -1$, though there would be differences going to NLL order.

We defer an analysis of these variants to future work, anticipating that many of the analytic tools from this paper can be translated to these generalized contexts. Experimentally, one might want to measure a track-based version of n_{SD}, trading collinear safety for improved robustness to pileup, which could be studied with the help of track functions [45, 52, 53].

3 Leading-logarithmic analysis

At LL order, the only difference between quarks and gluons is encoded in the color factors C_F and C_A, so Casimir scaling is a generic feature of many quark/gluon discriminants. Here, we review the case of additive observables (and close variants), where Casimir scaling of the Sudakov form factor yields a universal discrimination power at LL that depends only on C_A/C_F. We then show that the soft drop multiplicity is Poisson distributed, with its mean and variance satisfying Casimir scaling.

In general, any observable that is sensitive to multiple emissions at LL is "Poisson-like" distributed, in the sense that its variance σ^2 and mean μ both scale with the number n of emissions counted, i.e. $\sigma^2 = \mathcal{O}(\mu)$. In the limit of many emissions, all such observables converge to a normal distribution with decreasing relative width $w_{\text{rel}} \sim \sigma/\mu \sim 1/\sqrt{n}$. Then as more emissions are counted, the discrimination power is not a universal function of C_A/C_F, but instead improves as μ increases and the quark/gluon distributions separate.

In this section, we illustrate this behavior for soft drop multiplicity with distributions extracted from VINCIA, using the setup described in section 2.1. We extract ROC (receiver operating characteristic) curves of the quark efficiency versus the gluon mistag rate, and explain their qualitative behavior. In appendix A, we consider weighted soft drop multiplicity, with behavior that interpolates between that of Poisson- and Sudakov-distributed observables.

3.1 Review of additive observables

A generic jet observable is defined on the momenta p_i and quantum numbers q_i of particles within a jet. An additive IRC-safe observable f is one that reduces to the form

$$f\left(\{p_i, q_i\}\right) = \sum_{i \in \text{jet}} f(p_i) \tag{3.1}$$

in the soft/collinear limit, so that the observable depends on a simple sum over the jet constituents, independent of q_i.[5] The function $f(p_i)$ can depend on global properties of the jet (e.g. its p_T), but not on its substructure. Collinear safety implies that $f(p_i)$ is linear in the particle energies E_i. Examples of additive observables include the jet mass [54–56], the radial moments [57], and the angularities [34, 58, 59], among many others.

---

[5] One could consider additive but IRC-unsafe observables which do depend on q_i.
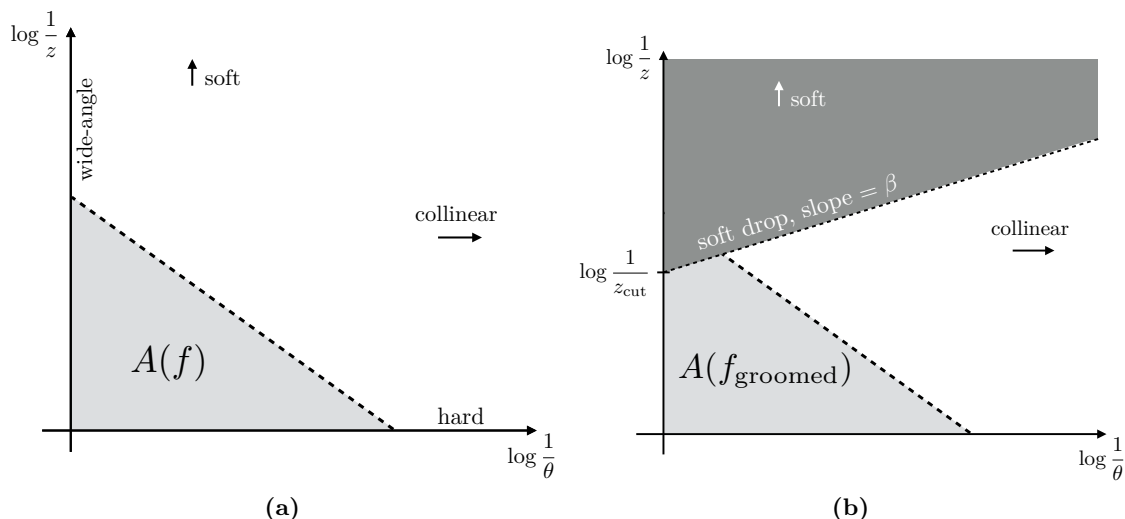
Figure 3. Illustration of emission phase space, which is a close analog of the Lund diagram [60], where gluon emissions are uniformly distributed in the $(\log 1/\theta, \log 1/z)$ plane. (a) The measurement of an additive observable f imposes a Sudakov veto on the phase space area $A(f)$. (b) For groomed observables, the measurement of a quasi-additive observable f_{groomed} also imposes a Sudakov veto.

We now review the Casimir scaling of additive observables at LL order, as discussed in ref. [17].[6] For simplicity of the discussion below, we let α_s be a fixed coupling so that the expressions are more compact, but it is straightforward to include a running coupling at LL order. At this order, we need only consider gluon emissions from the jet core that are both soft and collinear, described by the most singular terms in the splitting function. Parametrizing emissions by their angle θ and energy (or p_T) fraction z, real emissions are uniformly distributed in the $(\log 1/\theta, \log 1/z)$ plane. The density in this emission phase space is

$$\rho_i = \frac{2\alpha_s C_i}{\pi}, \tag{3.2}$$

where C_i is the appropriate color factor, equal to $C_F = 4/3$ for quarks and $C_A = 3$ for gluons. The structure of emission phase space is shown in figure 3a. Virtual emissions are encoded in the boundaries of the emission phase space, where $\log(1/\theta), \log(1/z) \to \infty$, such that the total emission probability at each α_s order is zero to maintain the normalization of the probability distribution.

Applying the strongly-ordered limit and the fact that $f(p_i)$ is linear in E_i, only a single dominant emission contributes to the observable at lowest order:

$$\sum_{i \in \text{jet}} f(p_i) \quad \overset{\text{LL}}{\implies} \quad \max_{i \in \text{jet}} f(p_i). \tag{3.3}$$

---

[6]Casimir scaling of additive observables at LL is identical to the statement of Casimir scaling of the cusp anomalous dimension in QCD, which has a long history in QCD [61, 62]. Casimir scaling is known to hold through three loops [63] in the cusp anomalous dimension, but is not expected to hold exactly [64]. At NLL and beyond, Casimir scaling is broken by the appearance of the non-cusp anomalous dimension.

Therefore, the probability that the observable f is less than some value $f_{\max}$ is equal to the probability that there are no emissions in the region where $f(p_i) > f_{\max}$. This implies a cumulative distribution function

$$\int_0^{f_{\max}} \mathrm{d}f \, p(f) \equiv \Sigma_i(f_{\max}) = e^{-\rho_i A(f_{\max})} \,, \tag{3.4}$$

where $A(f_{\max})$ is the forbidden area of emission phase space, shown in figure 3a:

$$A(f_{\max}) = \int_{f(z,\theta) > f_{\max}} \frac{\mathrm{d}\theta}{\theta} \frac{\mathrm{d}z}{z}. \tag{3.5}$$

Note that the cumulative distributions for quarks and gluons are related by

$$\Sigma_g(f_{\max}) = \left[\Sigma_q(f_{\max}) \right]^{C_A/C_F} \,, \tag{3.6}$$

where $C_A/C_F = 9/4$. That is, the Sudakov form factors for f are related by Casimir scaling. As a result, the ROC curve for quark/gluon discrimination, which simply plots $\Sigma_q(f)$ versus $\Sigma_g(f)$, takes the universal form of eq. (1.2).

From this logic, it is clear that the above analysis also extends to certain non-additive observables. For example, jet observables defined on groomed jets are not additive, since the grooming procedure removes emissions that would otherwise contribute to the sum in eq. (3.1). But groomed observables of the quasi-additive form

$$f_{\mathrm{groomed}}\left(\{p_i, q_i\}\right) = \sum_{i \in \mathrm{groomed\ jet}} f(p_i) \tag{3.7}$$

still exhibit Casimir scaling, since the measured value of f_{groomed} forbid emissions in the region $A(f_{\mathrm{groomed}})$ shown in figure 3b. More generally, Casimir scaling arises whenever the value of the measurement actively forbids emissions from some region of phase space. This vetoed phase space region builds up a Sudakov form factor which in turn controls the discrimination power achievable at LL.

Beyond LL order, different Sudakov-distributed observables will exhibit different discrimination power due to higher-order or nonperturbative effects, but eq. (3.6) is still a representative benchmark. In figure 4, we show ROC curves for jet mass m, the soft-dropped jet mass m_{SD}, and the groomed jet radius R_g, which all roughly follow the prediction from Casimir scaling. We also show track multiplicity n_{tr}, which exhibits substantially better performance and provides a useful discrimination target.

3.2 Soft drop multiplicity

Soft drop multiplicity is not an additive observable, nor does the measured value of n_{SD} actively forbid emissions in any region of phase space. As a result, n_{SD} does not exhibit Sudakov behavior and it instead satisfies a fundamentally different scaling relation. Physically, this is because all emissions that pass the soft drop condition are weighted equally, so n_{SD} depends on multiple emissions even at leading accuracy. These emissions occur in the region of phase space passing the soft drop and angular cuts, shown in figure 5.
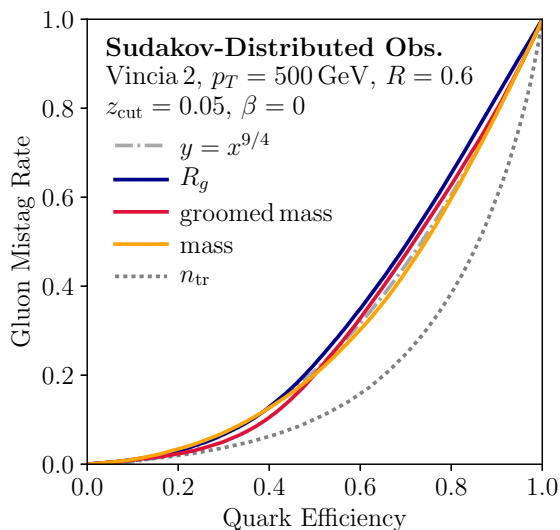
Figure 4. Comparison of the quark/gluon ROC curves for various Sudakov-distributed observables to the $y = x^{9/4}$ prediction from Casimir scaling. Shown are the groomed jet radius, groomed jet mass, and ordinary jet mass. As a useful benchmark, we also show the performance of track multiplicity $n_{\rm tr}$, which is known to be a very strong discriminant.
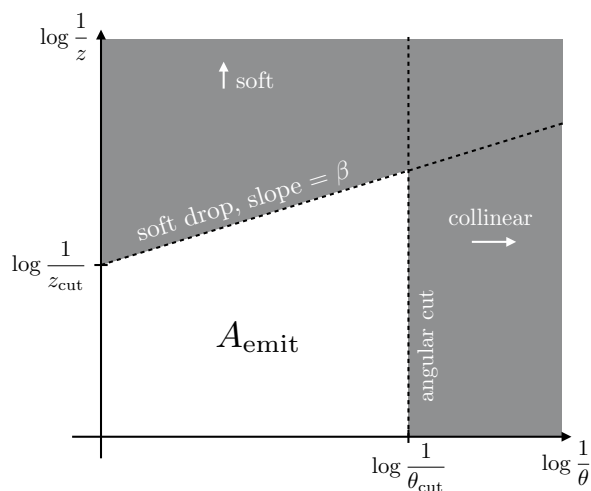
Figure 5. Same as figure 3b, but now highlighting the allowed emission region $A_{\rm emit}$ that is counted by soft drop multiplicity.

Restricting to the IRC safe case with $\theta_{\rm cut} > 0$, the measured region has finite area in the emission plane,

$$A_{\rm emit} = \log \frac{R_0}{\theta_{\rm cut}} \left(\log \frac{1}{2z_{\rm cut}} + \frac{\beta}{2} \log \frac{R_0}{\theta_{\rm cut}} \right), \tag{3.8}$$

and soft drop multiplicity simply counts the number of real emissions in this area. This expression actually holds for all $\beta \in (-\infty, \infty)$ as long as the angular cut $\theta_{\rm cut}$ imposes a non-trivial constraint on emissions. Since real emissions occur independently with uniform probability, they are described by a Poisson process, and the soft drop multiplicity is

Poisson distributed at LL order:[7]

$$\mathcal{P}_i(n_{\mathrm{SD}}) = \mathrm{Pois}(\lambda_i)[n_{\mathrm{SD}}], \qquad \lambda_i = \rho_i A_{\mathrm{emit}}. \tag{3.9}$$

For reference, the Poisson distribution with mean λ is

$$\mathrm{Pois}(\lambda)[n] = \frac{\lambda^n e^{-\lambda}}{n!}. \tag{3.10}$$

Since the variance of a Poisson distribution is also equal to λ, the means and variances of n_{SD} both satisfy Casimir scaling

$$\frac{\langle n_{\mathrm{SD}} \rangle_g}{\langle n_{\mathrm{SD}} \rangle_q} \simeq \frac{C_A}{C_F}, \qquad \frac{\mathrm{Var}(n_{\mathrm{SD}})_g}{\mathrm{Var}(n_{\mathrm{SD}})_q} \simeq \frac{C_A}{C_F}, \tag{3.11}$$

mirroring the behavior of track multiplicity in eq. (1.1), but for an IRC-safe observable.

To be clear, in defining our resummation accuracy, we count large logarithms of z_{cut} and θ_{cut} in the mean/variance of the n_{SD} distribution. That is, we define LL and NLL exactly as for more familiar additive observables, with LL including all terms of the form $\alpha_s^n \log^{n+1}$ that appear in the exponent of the n_{SD} distribution, and NLL including those terms of the form $\alpha_s^n \log^n$. With this definition, eq. (3.8) then shows that n_{SD} is indeed a double-logarithmic observable. In this section, we study this observable's general properties with fixed coupling, i.e. in the double-logarithmic approximation, for purposes of illustration. In section 4, LL and NLL results are computed using the appropriate running coupling.

The above analysis provides several concrete predictions. Our most salient result is that, since the soft drop multiplicity is Poisson distributed at LL, we expect the ratio of the variance to the mean to be close to 1, as shown in figure 6a. We also predict that the mean and variance satisfy the Casimir scaling relations in eq. (3.11), as shown in figure 6b. Though not shown here, we also checked the prediction that for $\beta = 0$, the mean soft drop multiplicity scales as

$$\lambda_i \propto \log \frac{1}{z_{\mathrm{cut}}} \log \frac{1}{\theta_{\mathrm{cut}}}. \tag{3.12}$$

In general, we find good agreement for these predictions at large values of z_{cut}, even out to $z_{\mathrm{cut}} \simeq 0.4$ where $\log z_{\mathrm{cut}}$ is not so large. For lower cut values, nonperturbative and higher-order effects cause these LL results to break down. In section 3.3, we demonstrate how to choose parameters so that nonperturbative effects can be avoided, and in section 4.2, we compute the NLL corrections to the perturbative predictions discussed here.

3.3 Optimal discrimination power

As a direct result of the properties exhibited in section 3.2, the discrimination power of soft drop multiplicity improves as the means $\lambda_i = \rho_i A_{\mathrm{emit}}$ increase. This is because the mean of each distribution is proportional to the Casimir C_i, while the standard deviation is equal to the square root of the mean. The overlap of the distributions is characterized by the relative width

$$w_{\mathrm{rel}} \equiv \frac{\sqrt{\mathrm{Var}(n_{\mathrm{SD}})_i}}{\langle n_{\mathrm{SD}} \rangle_i} = \frac{1}{\sqrt{\lambda_i}}. \tag{3.13}$$

---

[7]Note that at this order, we do not account for color correlations, so the emissions are effectively Abelian.
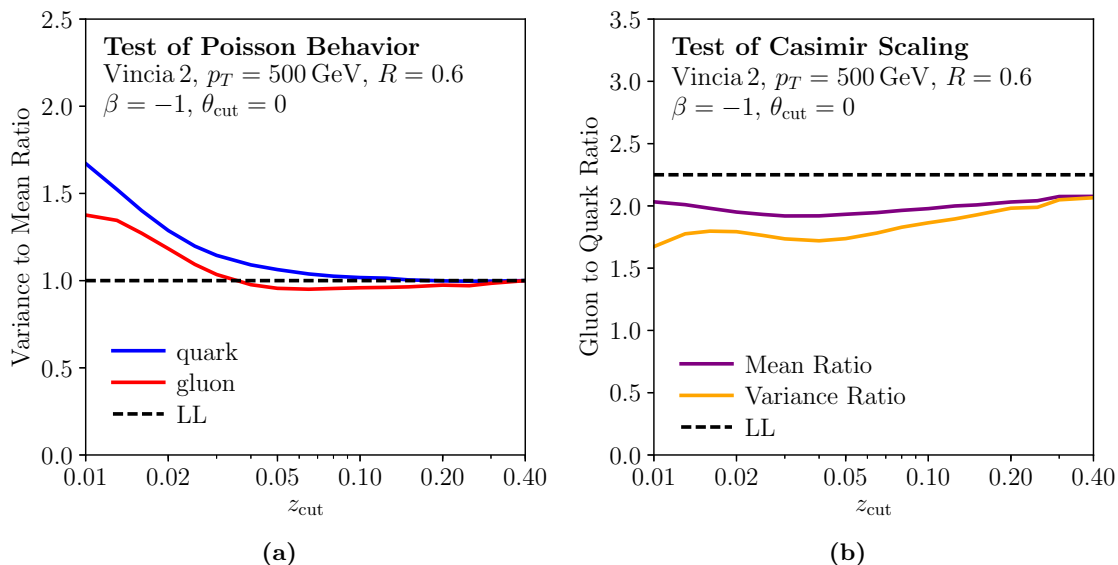
Figure 6. (a) Variance to mean ratio of the soft drop multiplicity as a function of $z_{\rm cut}$. The parameters β and $\theta_{\rm cut}$ are set to the benchmark values in eq. (2.7), and the LL prediction of equal mean and variance is shown as a dashed line. (b) Gluon to quark mean ratios and variance ratios, with the prediction of Casimir scaling shown as a dashed line. In both cases, we see qualitative agreement between VINCIA and the LL predictions down to $z_{\rm cut} = 0.02$.
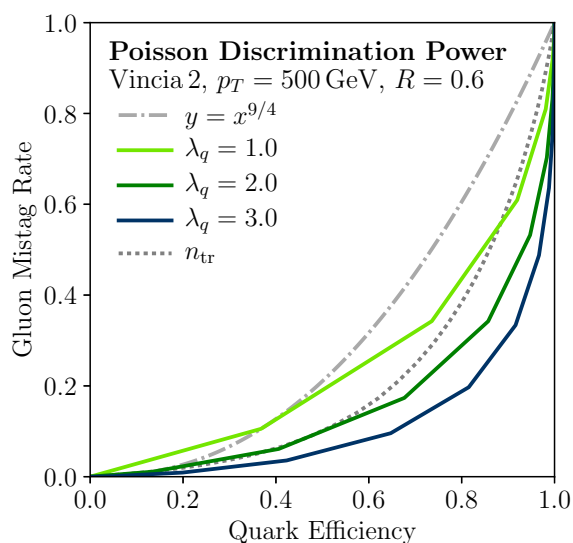
Figure 7. Expected quark/gluon discrimination power for Poisson-distributed observables. The mean observable value for quarks is λ_q, and we assume the mean for gluons is given by Casimir scaling $\lambda_g = (C_A/C_F)\lambda_q$. For reference, we show the $y = x^{9/4}$ curve for additive observables with Casimir scaling, as well as track multiplicity $n_{\rm tr}$ extracted from VINCIA. For mean quark values $\lambda_q \gtrsim 2$, a Poisson-like observable satisfying Casimir scaling would be competitive with track multiplicity. The ROC curves are piecewise linear since the observable takes on discrete integer values.

– 14 –
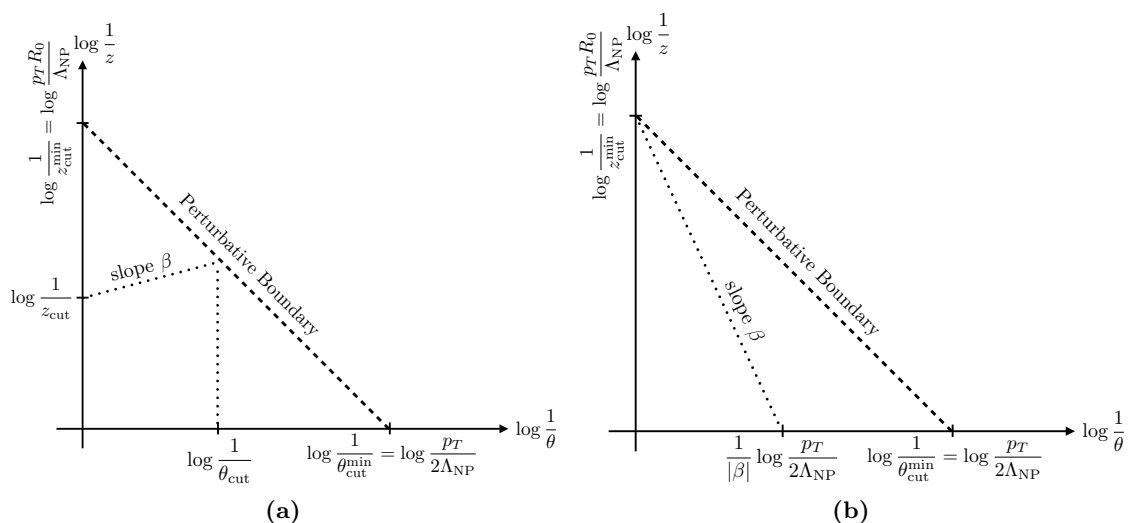
Figure 8. Illustration of the optimal phase space configuration consistent with a perturbative analysis. The dashed line with slope -1 separates perturbative and nonperturbative emissions. (a) For $\beta > -1$, the value of θ_{cut} has to be chosen to avoid allowed emissions above the nonperturbative boundary. (b) For $\beta < -1$, θ_{cut} can be set to zero, with z_{cut} pushed to the nonperturbative boundary. To maximize the allowed perturbative phase space, one should take $\beta = -1$ and z_{cut} set to the optimal value in eq. (3.20).

Indeed, in the many-emission limit where the distributions are approximately Gaussian, have equal mean and variance, and satisfy Casimir scaling, the discrimination power is solely determined by the relative width. As the cuts z_{cut} and θ_{cut} are lowered, the means increase, causing the relative widths to narrow, reducing the overlap between the quark and gluon distributions, and improving the discrimination power.

For reference, the discrimination power of Poisson distributions with different means is shown in figure 7, from which we see that track multiplicity has comparable discrimination power to a $\lambda_q \simeq 2$ observable.

To maximize the quark/gluon discrimination power, one should maximize the mean of the soft drop multiplicity distributions, which corresponds to taking z_{cut} and θ_{cut} as small as possible, for a given exponent β. The validity of this analysis, however, is restricted to perturbation theory, so we must ensure that the values of the chosen parameters do not allow for distributions that are dominated by nonperturbative emissions. We can determine the parameters that enforce perturbative emissions by restricting the minimum relative k_t appropriately.

To enforce that an emission is perturbative, we require that the relative k_t of the emission is larger than a perturbative cutoff scale Λ_{NP}, i.e.

$$z\,\theta \gtrsim \frac{\Lambda_{\mathrm{NP}}}{p_T}\,, \tag{3.14}$$

where z and θ are the energy fraction and splitting angle of the emission, and p_T is the transverse momentum of the jet. Below, we take $\Lambda_{\mathrm{NP}} = 2\,\mathrm{GeV}$ unless otherwise noted.

For an emission that just passes soft drop, and therefore contributes to the soft drop multiplicity, we have

$$z \gtrsim z_{\text{cut}} \frac{\theta^\beta}{R_0^\beta} . \tag{3.15}$$

There are two regimes to consider. For $\beta > -1$ as in figure 8a, we can find the intersection of eqs. (3.14) and (3.15). Setting $\theta \to \theta_{\text{cut}}$, we find a restriction on θ_{cut} to be perturbative:

$$\theta_{\text{cut}} \gtrsim \left(\frac{\Lambda_{\text{NP}}}{z_{\text{cut}} p_T R_0} \right)^{\frac{1}{1+\beta}} R_0 . \tag{3.16}$$

To determine the optimal choice of z_{cut} while enforcing perturbativity, we set θ_{cut} to saturate this inequality and insert it into the double-log expression for the average soft drop multiplicity, eq. (3.8). Maximizing this quantity, we find the optimal ISD parameters to be

$$z_{\text{cut}}|_{\text{optimal}} = \frac{1}{2} \left(\frac{2\Lambda_{\text{NP}}}{p_T R_0} \right)^{\frac{1}{2+\beta}} , \tag{3.17}$$

$$\theta_{\text{cut}}|_{\text{optimal}} = \left(\frac{2\Lambda_{\text{NP}}}{p_T R_0} \right)^{\frac{1}{2+\beta}} R_0 . \tag{3.18}$$

The factors of two arise because the energy fraction of the softer emission is (by definition) less than $1/2$. Inserting these results into the expression for the average soft drop multiplicity, we find the largest perturbative value for the mean soft drop multiplicity to be

$$\langle n_{\text{SD}} \rangle_{\text{optimal}}^{\beta>-1} \simeq \frac{\alpha_s}{\pi} \frac{C_i}{2+\beta} \log^2 \left(\frac{2\Lambda_{\text{NP}}}{p_T R_0} \right) . \tag{3.19}$$

For $\beta < -1$, one can see from the $(\log 1/\theta, \log 1/z)$ phase space in figure 8b that an angular cutoff is not needed to avoid the nonperturbative region, so we can set $\theta_{\text{cut}} = 0$. In this case, z_{cut} saturates the bound eq. (3.14) for $\theta \to R_0$, yielding

$$z_{\text{cut}}|_{\text{optimal}} = \frac{\Lambda_{\text{NP}}}{p_T R_0}, \tag{3.20}$$

and the average soft drop multiplicity is

$$\langle n_{\text{SD}} \rangle_{\text{optimal}}^{\beta<-1} \simeq \frac{\alpha_s}{\pi} \frac{C_i}{|\beta|} \log^2 \left(\frac{2\Lambda_{\text{NP}}}{p_T R_0} \right) . \tag{3.21}$$

Combining these regions for all $\beta \in (-\infty, \infty)$, the maximum attainable mean soft drop multiplicity with perturbative parameters is

$$\langle n_{\text{SD}} \rangle_{\text{optimal}} \simeq \frac{\alpha_s C_i}{\pi} \min \left[\frac{1}{|\beta|}, \frac{1}{|2+\beta|} \right] \log^2 \left(\frac{2\Lambda_{\text{NP}}}{p_T R_0} \right) . \tag{3.22}$$

In particular, the mean is maximized for $\beta = -1$, giving the optimal perturbative discrimination power in this double-log approximation. This result can be understood directly from figure 8, which shows that soft drop multiplicity with $\beta = -1$ can capture all of the perturbative emissions in phase space.
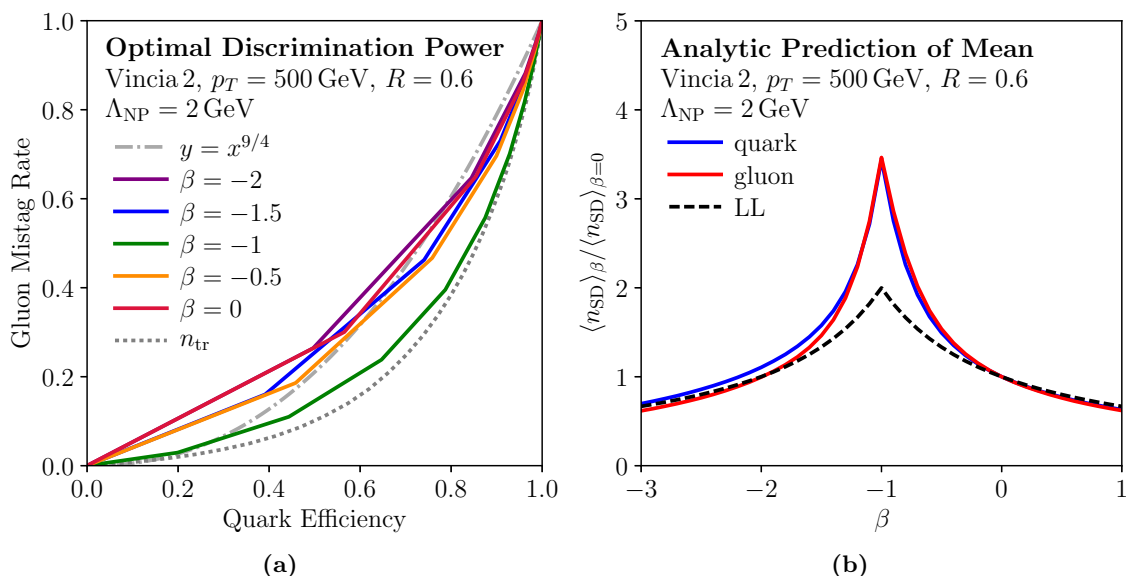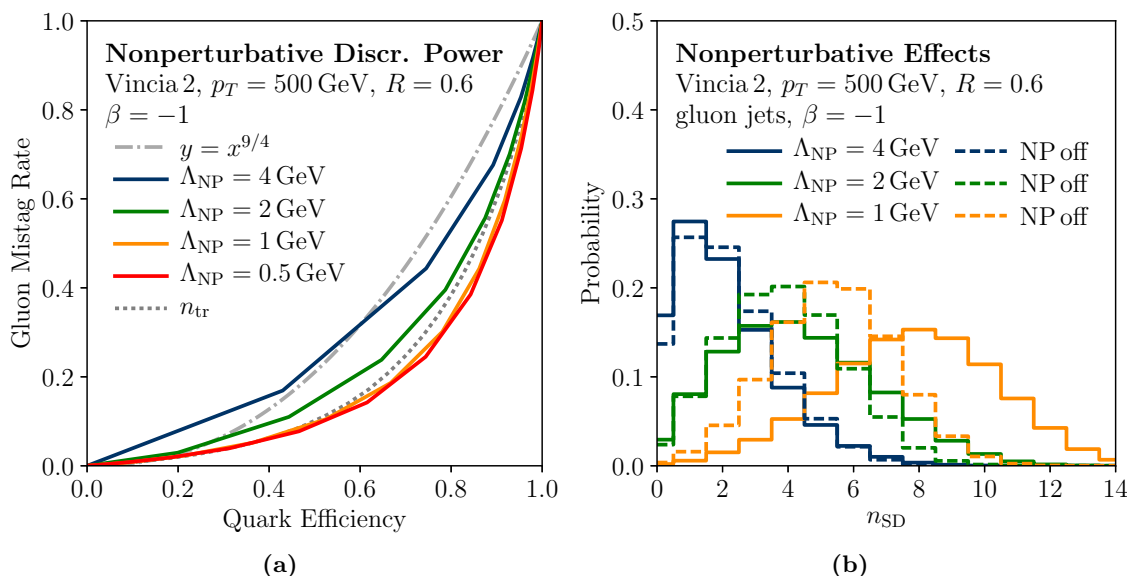
Figure 9. (a) Discrimination power of soft drop multiplicity as a function of β, with the optimal (perturbative) values of $z_{\rm cut}$ and $\theta_{\rm cut}$ computed from eqs. (3.17), (3.18) and (3.20) using $\Lambda_{\rm NP} = 2\,{\rm GeV}$. (b) Ratio of mean $n_{\rm SD}$ as a function of β to mean $n_{\rm SD}$ at $\beta = 0$. The VINCIA results for quarks and gluons agree with the double log prediction from eq. (3.22), except near $\beta = -1$ where nonperturbative effects become important.

We can directly test this double-log prediction in parton shower generators. In figure 9a, we show the quark/gluon ROC curve for soft drop multiplicity with the optimal perturbative soft drop parameters, sweeping through β. The best discrimination power found in VINCIA is indeed observed near $\beta = -1$. For a more quantitative test, eq. (3.22) predicts that the ratio of the optimal soft drop multiplicity for a given value of β to the optimal soft drop multiplicity at $\beta = 0$ is

$$\frac{\langle n_{\rm SD}\rangle_{\rm optimal}}{\langle n_{\rm SD}\rangle_{\rm optimal}^{\beta=0}} = \min\left[\frac{2}{|\beta|}, \frac{2}{|2+\beta|}\right]. \tag{3.23}$$

In figure 9b, we compare this ratio to distributions extracted from VINCIA and find good agreement away from $\beta = -1$. Note that when $\beta = -1$, the counted and nonperturbative regions share a boundary, while in all other cases the two regions only meet at a single point. This explains why nonperturbative sensitivity should be amplified when β nears -1. This extra sensitivity could of course be mitigated by using a more conservative value of $\Lambda_{\rm NP}$, but there is a tradeoff between reducing nonperturbative effects and increasing discrimination power.

In figure 10a, we show the effect that decreasing $\Lambda_{\rm NP}$ (and thus decreasing $z_{\rm cut}$ and $\theta_{\rm cut}$) has on the discrimination power, holding $\beta = -1$ fixed. Note that $n_{\rm SD}$ rivals $n_{\rm tr}$ for $\Lambda_{\rm NP} = 1\,{\rm GeV}$, but that there is no gain in performance when $\Lambda_{\rm NP}$ is taken smaller. In figure 10b we show the shift in gluon $n_{\rm SD}$ distributions from switched off hadronization and underlying event in VINCIA. We take this as an indicator of nonperturbative sensitivity

Figure 10. (a) Discrimination power of soft drop multiplicity as a function of Λ_{NP} with $\beta = -1$, $\theta_{\mathrm{cut}} = 0$, and z_{cut} computed from eq. (3.20). (b) Impact of hadronization and underlying event in VINCIA on gluon distributions.

in the distributions. One can see that perturbative control is lost for $\Lambda_{\mathrm{NP}} < 2\,\mathrm{GeV}$. For $p_T = 500\,\mathrm{GeV}$, $\Lambda_{\mathrm{NP}} = 2\,\mathrm{GeV}$ gives the benchmark parameters in eq. (2.7).

Our perturbative analysis here was restricted to LL order and fixed coupling, and the inclusion of higher-order effects will affect the discrimination power of soft drop multiplicity. In particular, at NLL order, quark and gluon jet flavors can mix, so we expect that higher-order effects in general decrease the discrimination power from the LL prediction. We perform NLL calculations and compare our results to parton showers in section 4. Beyond these higher-order effects, we have restricted the analysis to perturbative parameters. Allowing nonperturbative emissions to contribute to the soft drop multiplicity should improve the discrimination power, however, at the expense of loss of predictivity. We discuss in section 5 how to restore some of this predictive power in the nonperturbative regime with GFFs.

One might wonder if the discrimination power could be further improved by weighting the emissions, e.g. by their energy, as in the weighted soft drop multiplicity of eq. (2.8). At LL order, however, the soft drop multiplicity is provably the most powerful discriminant that can be defined on the (z_n, θ_n) values.[8] To see this, note that the normalized distribution of emissions in the $(\log 1/\theta, \log 1/z)$ plane is identical for quark and gluon jets at LL order, even including running coupling effects. Therefore, once the value of n_{SD} is known for a given jet, no additional discriminatory information can be gleaned from the (z_n, θ_n) values. Nevertheless, weighted soft drop multiplicity provides an example of a more general

---

[8]We thank Ben Nachman for discussions on this point. Specifically, he demonstrated that the quark/gluon likelihood ratio is a monotonic function of n_{SD}, with no other non-trivial (z_n, θ_n) dependence, thus providing further confirmation that n_{SD} is the optimal discriminant one can construct.

observable that can be effectively studied with our analytic tools; we demonstrate this in appendix A.[9]

4 Calculations for IRC-safe soft drop multiplicity

We now demonstrate that the LL predictions of the previous section can be reproduced by a set of perturbative evolution equations. These equations describes how soft drop multiplicity evolves with decreasing θ_{cut}, similar to traditional parton evolution [65]. This approach also admits a generalization to NLL, which we use to make precise predictions for comparison to parton showers.

When talking about the resummation of large logarithms at LL and NLL accuracy, we are specifically referring to factors of $\log z_{\text{cut}}$ and $\log \theta_{\text{cut}}$, not to any logarithms associated with the n_{SD} observable (which is an integer). As we already saw in eq. (3.8), these logarithms control the size of the emission phase space, which in turn control the expected mean value of n_{SD}, so their resummation is essential for predicting the distribution of n_{SD}.

4.1 Leading-logarithmic evolution equations

We begin by analyzing the soft drop multiplicity to LL accuracy. This case is simple enough to keep the structure of the θ_{cut} evolution transparent; the generalization to NLL just requires keeping track of more details. To achieve LL accuracy, we need only consider soft-collinear gluons emitted from the hard core of a jet; flavor-changing effects are not present at this order. Furthermore, the trunk of the clustering tree retains all but an $\mathcal{O}(z_{\text{cut}})$ fraction of the original jet's energy, so for $z_{\text{cut}} \ll 1$, energy losses are negligible at this order as well.

Let $p_n^i(\theta_{\text{cut}})$ denote the probability that, given a jet of flavor i and ISD parameter θ_{cut}, its soft drop multiplicity $n_{\text{SD}}(\theta_{\text{cut}})$ is measured to be n. Here, we leave the dependence on z_{cut} and β implicit, since they do not participate directly in the evolution equations. Since n_{SD} is a discrete counting observable, $p_n^i(\theta_{\text{cut}})$ is finite and should satisfy the normalization condition $\sum_{n=0}^{\infty} p_n^i(\theta_{\text{cut}}) = 1$ for each flavor i.

We can compute the distribution for $p_n^i(\theta_{\text{cut}})$ by solving a set of evolution equations. Consider decreasing the resolution angle from θ_{cut} to $\theta_{\text{cut}} - \delta\theta_{\text{cut}}$. The value of n_{SD} will increase by one if there is an emission in the interval $[\theta_{\text{cut}} - \delta\theta_{\text{cut}}, \theta_{\text{cut}}]$ that passes soft drop; otherwise n_{SD} will remain unchanged. That is,

$$p_n^i(\theta_{\text{cut}} - \delta\theta_{\text{cut}}) = p_{n-1}^i(\theta_{\text{cut}}) \frac{\delta\theta_{\text{cut}}}{\theta_{\text{cut}}} \int_0^{1/2} \mathrm{d}z \, \frac{\alpha_s(z\,\theta_{\text{cut}}\,p_T)}{\pi} \, P_{i \to i}(z) \, \Theta_{\text{SD}}(z, \theta_{\text{cut}}) \tag{4.1}$$
$$+ p_n^i(\theta_{\text{cut}}) \left(1 - \frac{\delta\theta_{\text{cut}}}{\theta_{\text{cut}}} \int_0^{1/2} \mathrm{d}z \, \frac{\alpha_s(z\,\theta_{\text{cut}}\,p_T)}{\pi} \, P_{i \to i}(z) \, \Theta_{\text{SD}}(z, \theta_{\text{cut}}) \right).$$

---

[9]One might be attracted to weighted soft drop multiplicity because it reduces sensitivity to soft emissions. Presumably, the value of Λ_{NP} could be reduced somewhat without introducing significant nonperturbative effects. One cannot increase perturbative discrimination power in this way, however, since any gain in discrimination power from reducing Λ_{NP} must necessarily come with comparable nonperturbative sensitivity.

Here, $P_{i \to i}(z)$ is the splitting function for the hard parton i to emit a collinear gluon of energy fraction z (and remain as flavor i), and $\Theta_{\mathrm{SD}}(z, \theta)$ imposes the soft drop condition,

$$\Theta_{\mathrm{SD}}(z, \theta) \equiv \Theta \left(z - z_{\mathrm{cut}} \frac{\theta^{\beta}}{R_0^{\beta}} \right). \tag{4.2}$$

At LL, $\alpha_s(z\,\theta_{\mathrm{cut}}\,p_T)$ runs with the 1-loop β function.

Using eq. (4.1), we can derive the linear first-order differential equation in θ_{cut},

$$\frac{\mathrm{d}p_n^i}{\mathrm{d}\theta_{\mathrm{cut}}} = \frac{p_n^i(\theta_{\mathrm{cut}}) - p_{n-1}^i(\theta_{\mathrm{cut}})}{\theta_{\mathrm{cut}}} \int_0^{1/2} \mathrm{d}z\, \frac{\alpha_s(z\,\theta_{\mathrm{cut}}\,p_T)}{\pi} P_{i \to i}(z)\, \Theta_{\mathrm{SD}}(z, \theta_{\mathrm{cut}}). \tag{4.3}$$

Because no emissions are recorded outside the jet radius R_0, there is a boundary condition $p_n^i(R_0) = \delta_{n,0}$. With this boundary condition, the solution to eq. (4.3) is

$$p_0^i(\theta_{\mathrm{cut}}) = e^{-I_{i \to i}(\theta_{\mathrm{cut}}, R_0)}, \tag{4.4}$$

$$p_{n \geq 1}^i(\theta_{\mathrm{cut}}) = \int_{\theta_{\mathrm{cut}}}^{R_0} \frac{\mathrm{d}\theta}{\theta}\, e^{-I_{i \to i}(\theta_{\mathrm{cut}}, \theta)} \left(\int_0^{1/2} \mathrm{d}z \frac{\alpha_s(z\theta p_T)}{\pi} P_{i \to i}(z) \Theta_{\mathrm{SD}}(z, \theta) \right) p_{n-1}^i(\theta), \tag{4.5}$$

where

$$I_{i \to i}(\theta_1, \theta_2) = \int_{\theta_1}^{\theta_2} \frac{\mathrm{d}\theta}{\theta} \int_0^{1/2} \mathrm{d}z\, \frac{\alpha_s(z\,\theta\,p_T)}{\pi} P_{i \to i}(z)\, \Theta_{\mathrm{SD}}(z, \theta). \tag{4.6}$$

The expression in eq. (4.4) corresponds to the case of no emissions between R_0 and θ_{cut}. The expression in eq. (4.5) computes the probability that ISD records $n - 1$ emissions in the interval $[\theta, R_0]$, one final emission at θ, then zero emissions in the interval $[\theta_{\mathrm{cut}}, \theta]$, with an integral over the angle θ where the final counted emission occurs.

We can interpret eq. (4.5) as a recursion relation in n with eq. (4.4) as the initial condition. The first step in the recursion ($n = 1$) gives

$$p_1^i(\theta_{\mathrm{cut}}) = \int_{\theta_{\mathrm{cut}}}^{R_0} \frac{\mathrm{d}\theta}{\theta}\, e^{-I_{i \to i}(\theta_{\mathrm{cut}}, \theta)} \left(\int_0^{1/2} \mathrm{d}z\, \frac{\alpha_s(z\,\theta\,p_T)}{\pi} P_{i \to i}(z)\, \Theta_{\mathrm{SD}}(z, \theta) \right) e^{-I_{i \to i}(\theta, R_0)}$$

$$= e^{-I_{i \to i}(\theta_{\mathrm{cut}}, R_0)} I_{i \to i}(\theta_{\mathrm{cut}}, R_0). \tag{4.7}$$

A similar simplification occurs for each value of n, and we recognize the Poisson distribution we found in eq. (3.9):

$$p_n^i(\theta_{\mathrm{cut}}) = \frac{1}{n!} \left[I_{i \to i}(\theta_{\mathrm{cut}}, R_0) \right]^n e^{-I_{i \to i}(\theta_{\mathrm{cut}}, R_0)}. \tag{4.8}$$

At LL, the soft drop multiplicity n_{SD} is thus Poisson distributed with mean $\lambda_i = I_{i \to i}(\theta_{\mathrm{cut}}, R_0)$. With fixed coupling, the mean value agrees exactly with $\lambda_i = \rho_i A_{\mathrm{emit}}$ found before (see eqs. (3.2) and (3.8)):

$$I_{i \to i}(\theta_{\mathrm{cut}}, R_0)|_{\text{fixed } \alpha_s} = \frac{2\alpha_s C_i}{\pi} \log \frac{R_0}{\theta_{\mathrm{cut}}} \left(\log \frac{1}{2z_{\mathrm{cut}}} + \frac{\beta}{2} \log \frac{R_0}{\theta_{\mathrm{cut}}} \right). \tag{4.9}$$

4.2 Next-to-leading-logarithmic corrections

The next-to-leading logarithms take the form $\alpha_s^n \log^n z_{\text{cut}}$ and $\alpha_s^n \log^n \theta_{\text{cut}}$ in the logarithm of $p_n^i(\theta_{\text{cut}})$. To resum these, we must consider emitted partons that are not necessarily soft and that can be either quarks or gluons. This requires us to take energy losses and flavor changes into account at this accuracy. It is convenient to compute $p_n^i(\theta_{\text{cut}})$ by expressing it as

$$p_n^i(\theta_{\text{cut}}) = \sum_{j=q,g} \int_{1/2^n}^{1} \mathrm{d}Z \, p_n^{i \to j(Z)}(\theta_{\text{cut}}) . \tag{4.10}$$

Here, $\mathrm{d}Z \, p_n^{i \to j(Z)}(\theta_{\text{cut}})$ is the differential probability that, given a jet of flavor i, ISD counts n emissions from its hard core that result in a flavor change from i to j, and a remaining energy fraction in the interval $[Z, Z + \mathrm{d}Z]$.[10] These more differential distributions evolve with θ_{cut} as

$$
\begin{aligned}
&p_n^{i \to j(Z)}(\theta_{\text{cut}} - \delta\theta_{\text{cut}}) \\
&= p_n^{i \to j(Z)}(\theta_{\text{cut}}) \left(1 - \frac{\delta\theta_{\text{cut}}}{\theta_{\text{cut}}} \int_0^{1/2} \mathrm{d}z \, \frac{\alpha_s(z\theta_{\text{cut}} Z p_T)}{\pi} P_{j \to \text{any}}(z) \Theta_{\text{SD}}(z, \theta_{\text{cut}}) \right) \\
&\quad + \sum_k \frac{\delta\theta_{\text{cut}}}{\theta_{\text{cut}}} \int_0^{1/2} \mathrm{d}z \, \frac{\alpha_s\left(z\theta \frac{Z}{1-z} p_T\right)}{\pi} P_{k \to j}(z) \Theta_{\text{SD}}(z, \theta) p_{n-1}^{i \to k[Z/(1-z)]}(\theta_{\text{cut}}) \frac{1}{1-z} , \quad (4.11)
\end{aligned}
$$

where

$$P_{i \to \text{any}}(z) = \sum_j P_{i \to j}(z) . \tag{4.12}$$

The middle line of eq. (4.11) is the probability that n emissions are counted at resolution θ_{cut}, and that only virtual or soft-dropped emissions (neither of which have an impact on energy fractions, up to z_{cut} corrections) occur in the interval $[\theta_{\text{cut}} - \delta\theta_{\text{cut}}, \theta_{\text{cut}}]$. The second line is the probability that $n - 1$ emissions are counted at resolution θ_{cut} and result in a flavor conversion $i \to k$, and that an additional counted emission causing further conversion $k \to j$ occurs in $[\theta_{\text{cut}} - \delta\theta_{\text{cut}}, \theta_{\text{cut}}]$.

We now justify that these evolution equations do indeed resum large logarithms to NLL, with one caveat. As is necessary for NLL resummation, these evolution equations contain NLO information about the jet's substructure. To achieve NLL accuracy, we need to properly include the following double-emissions structures: collinear plus collinear (C+C), soft plus collinear (S+C), soft plus soft (S+S), and hard plus soft-collinear (H+SC). Since ISD is an angular-ordered algorithm, collinear emissions factorize in the cross section, so our evolution equations correctly include C+C and S+C double emissions. The S+S case is included as well by letting α_s run with the 2-loop β function in the CMW scheme [66]. The one caveat is that we do not describe H+SC double emissions correctly at NLO, since we use splitting functions instead of full matrix elements.[11] Thus, our approximation

---

[10]The probability for the hard core to be left with energy fraction between $Z/(1-z)$ and $(Z+\mathrm{d}Z)/(1-z)$ is then $p_n^{i \to j[Z/(1-z)]}(\theta_{\text{cut}}) \, \mathrm{d}Z/(1-z)$. This is used in eq. (4.11).

[11]Besides this caveat, though, note that our use of $1 \to 2$ (as opposed to $1 \to 3$) splitting functions is sufficient at NLL, since n_{SD} is a double-logarithmic observable.

should become more accurate as the jet radius R_0 becomes smaller, forcing hard emissions in the jet to become collinear. We also ignore the effects of logarithms of $z_{\rm cut}$ that arise from nonglobal radiation [67], and so do not describe emissions in the jet from secondary radiation from outside of the jet.

Despite the extra complications at NLL order, eq. (4.11) is still a linear first-order differential equation, just as in section 4.1. The solution is

$$p_0^{i \to j(Z)}(\theta_{\rm cut}) = \delta_{ji}\,\delta(Z-1)\,\exp\left[-I_{i \to {\rm any}}(\theta_{\rm cut}, R_0)\right],\tag{4.13}$$

$$p_{n \geq 1}^{i \to j(Z)}(\theta_{\rm cut}) = \sum_k \int_{\theta_{\rm cut}}^{R_0} \frac{{\rm d}\theta}{\theta} \int_0^{1/2} {\rm d}z \,\exp\left[-I_{j(Z) \to {\rm any}}(\theta_{\rm cut}, \theta)\right]$$

$$\times \frac{\alpha_s\left(z\,\theta\,\frac{Z}{1-z}\,p_T\right)}{\pi}\,P_{k \to j}(z)\,\Theta_{\rm SD}(z, \theta)\,p_{n-1}^{i \to k[Z/(1-z)]}(\theta)\,\frac{1}{1-z},\tag{4.14}$$

where

$$I_{j(Z) \to {\rm any}}(\theta_1, \theta_2) = \int_{\theta_1}^{\theta_2} \frac{{\rm d}\theta}{\theta} \int_0^{1/2} {\rm d}z\,\frac{\alpha_s(z\,\theta\,Z\,p_T)}{\pi}\,P_{j \to {\rm any}}(z)\,\Theta_{\rm SD}(z, \theta)\,.\tag{4.15}$$

Note that $p_n^{i \to j(Z)}$ vanishes for $Z < 1/2^n$. The same manipulations that led to eq. (4.7) and the Poisson distribution at LL do not go through at NLL, so we cannot write $p_n^{i \to j(Z)}$ or p_n^i in closed form at this order. Nonetheless, the integrals in eq. (4.14) can be performed numerically by first computing $p_1^{i \to j(Z)}(\theta_{\rm cut})$, then computing $p_2^{i \to j(Z)}(\theta_{\rm cut})$, and so on until p_n^j is negligible. In practice, the probability saturates for n of order 10.

The $n_{\rm SD}$ distributions and ROC curves at LL and NLL accuracy are displayed in figure 11. The uncertainties in the NLL calculation come from varying the α_s scale up and down by a factor of 2. (Scale variation in the LL calculation does not give a reliable estimate of the uncertainty, since flavor-changing processes are absent at LL; we therefore omit bands around the LL predictions.) The fact that the uncertainties are abnormally small in one bin is an artifact of this one-dimensional variation procedure, which leaves the scale-varied distributions properly normalized. Also, the uncertainties in the ROC curve are substantially smaller than the uncertainties in the NLL distributions, since the way we implement the scale variation affects quarks and gluons in a correlated way. We show both $\beta = -1$ and $\beta = -0.5$ with $z_{\rm cut}$ and $\theta_{\rm cut}$ chosen to be "optimal" according to eqs. (3.17), (3.18) and (3.20) with $\Lambda_{\rm NP} = 2\,{\rm GeV}$. One can see that NLL corrections result in a slight decrease in discrimination power compared to LL, due in part to the flavor changes that occur at this order.

4.3 Comparison to parton showers

It is instructive to compare our NLL calculation of the soft drop multiplicity $n_{\rm SD}$ with results obtained from parton shower generators. In addition to the Vincia setup described in section 2.1, we obtained alternative event samples by showering the hard events through Pythia 8.219 [68, 69], Herwig 7.0.1 [70, 71], and Sherpa 2.2.0 [72], interfaced to their default hadronization and underlying event models.

First, to validate the reliability of our NLL calculation, we want to explore the impact of nonperturbative effects on the parton showers. In section 3.3 we noted that hadronization
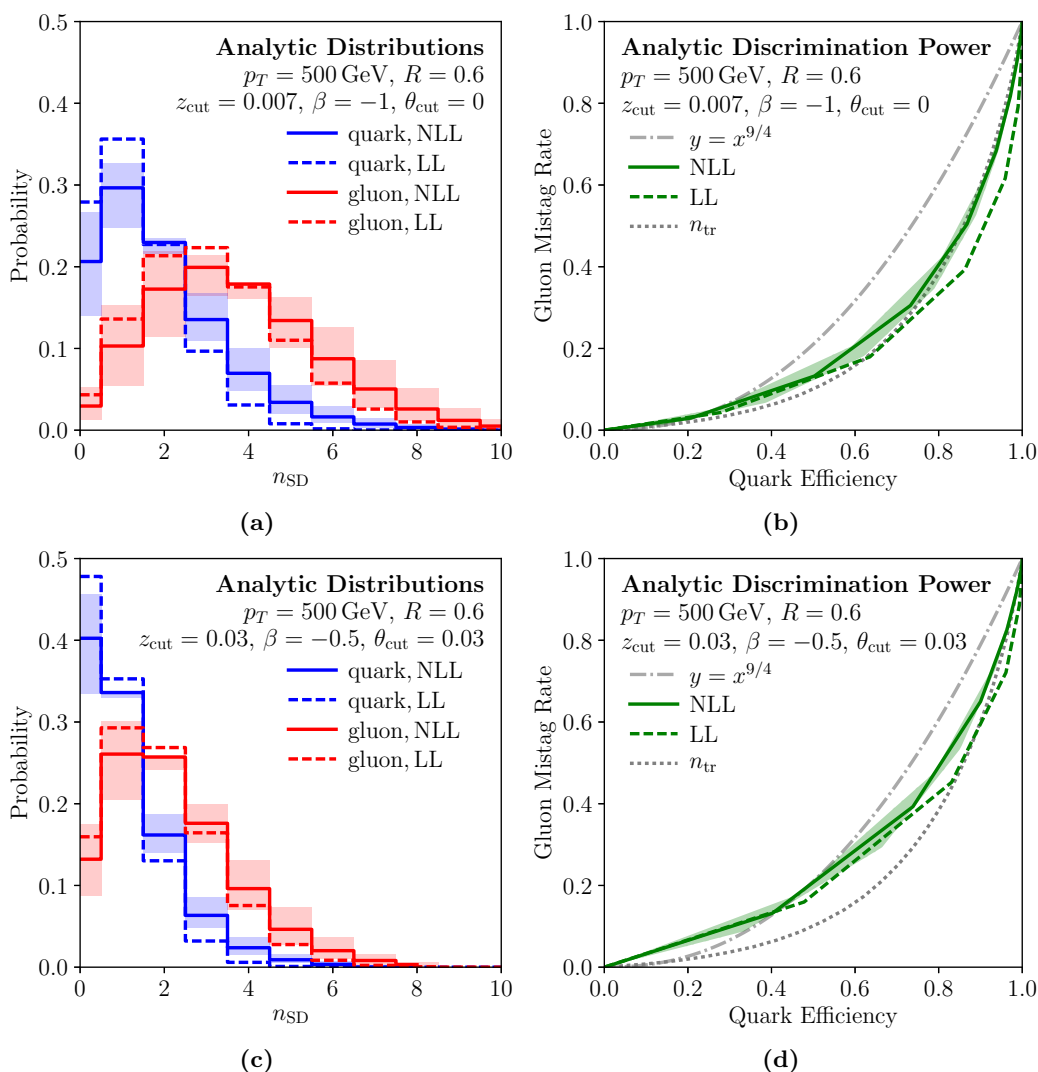
Figure 11. Calculations at LL and NLL accuracy for (left column) n_{SD} distributions and (right column) the corresponding quark/gluon ROC curves. Parameters are chosen according to eqs. (3.17), (3.18) and (3.20) with $\Lambda_{\text{NP}} = 2\,\text{GeV}$ and (top row) $\beta = -1$ and (bottom row) $\beta = -0.5$. The uncertainties in the NLL calculation come from varying the α_s scale by a factor of 2.

effects should generically be minimal provided parameters are chosen at or above the values given in eqs. (3.17), (3.18) and (3.20). To investigate this expectation further, we check the size of nonperturbative corrections in VINCIA by turning hadronization and underlying event off and comparing to results obtained using the default settings. In figure 12, we show n_{SD} with $\beta = -1$ and $\beta = -0.5$, where in each case z_{cut} and θ_{cut} are computed using eqs. (3.17), (3.18) and (3.20) with $\Lambda_{\text{NP}} = 2\,\text{GeV}$. As expected, nonperturbative effects are under control, confirming that our perturbative NLL calculations should indeed be reliable in predicting the n_{SD} distributions. Though not shown, the other three parton shower generators also exhibit comparable nonperturbative shifts.
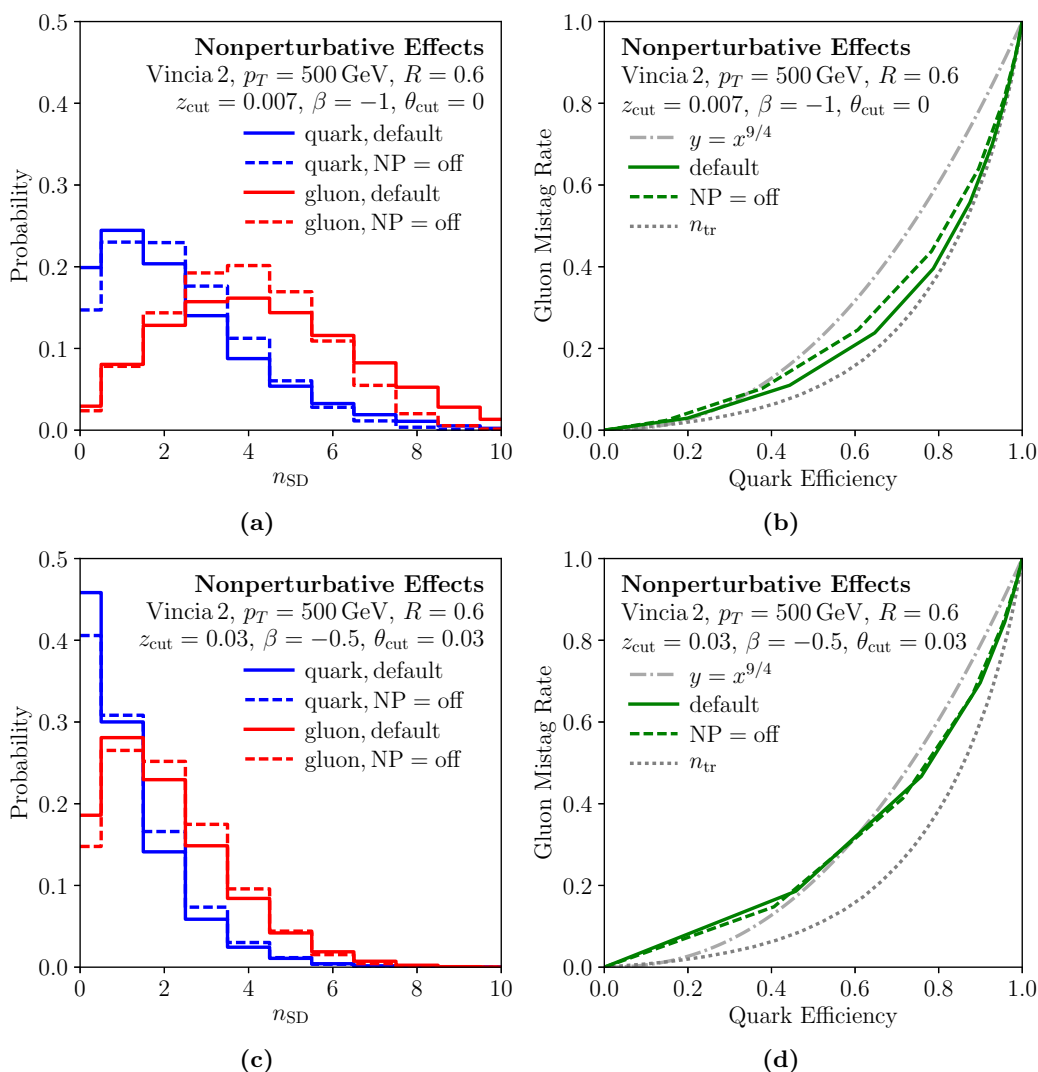
Figure 12. Impact of nonperturbative effects on (left column) n_{SD} distributions and (right column) the corresponding ROC curves. This study employs VINCIA, where parameters are chosen according to eqs. (3.17), (3.18) and (3.20) with $\Lambda_{\mathrm{NP}} = 2\,\mathrm{GeV}$ and (top row) $\beta = -1$ and (bottom row) $\beta = -0.5$.

Next, we show that all parton shower generators predict that soft drop multiplicity is a relatively good quark/gluon discriminant. In figure 13, we compare n_{SD} with $\beta = -1$ and $\beta = -0.5$ to jet mass and n_{tr} for each generator separately. For $\beta = -1$, soft drop multiplicity provides a significant improvement over generic additive observables but does not quite match the performance of track multiplicity. (See, however, figure 10a where nonperturbative parameter values push the performance of n_{SD} to match n_{tr}.) The ordering of the ROC curves is roughly the same between the four generators, though the absolute discrimination power does differ.

Finally, we can directly compare our NLL predictions to the parton shower generators. In figure 14, we show the n_{SD} distributions and ROC curves for both $\beta = -1$ and $\beta = -0.5$.
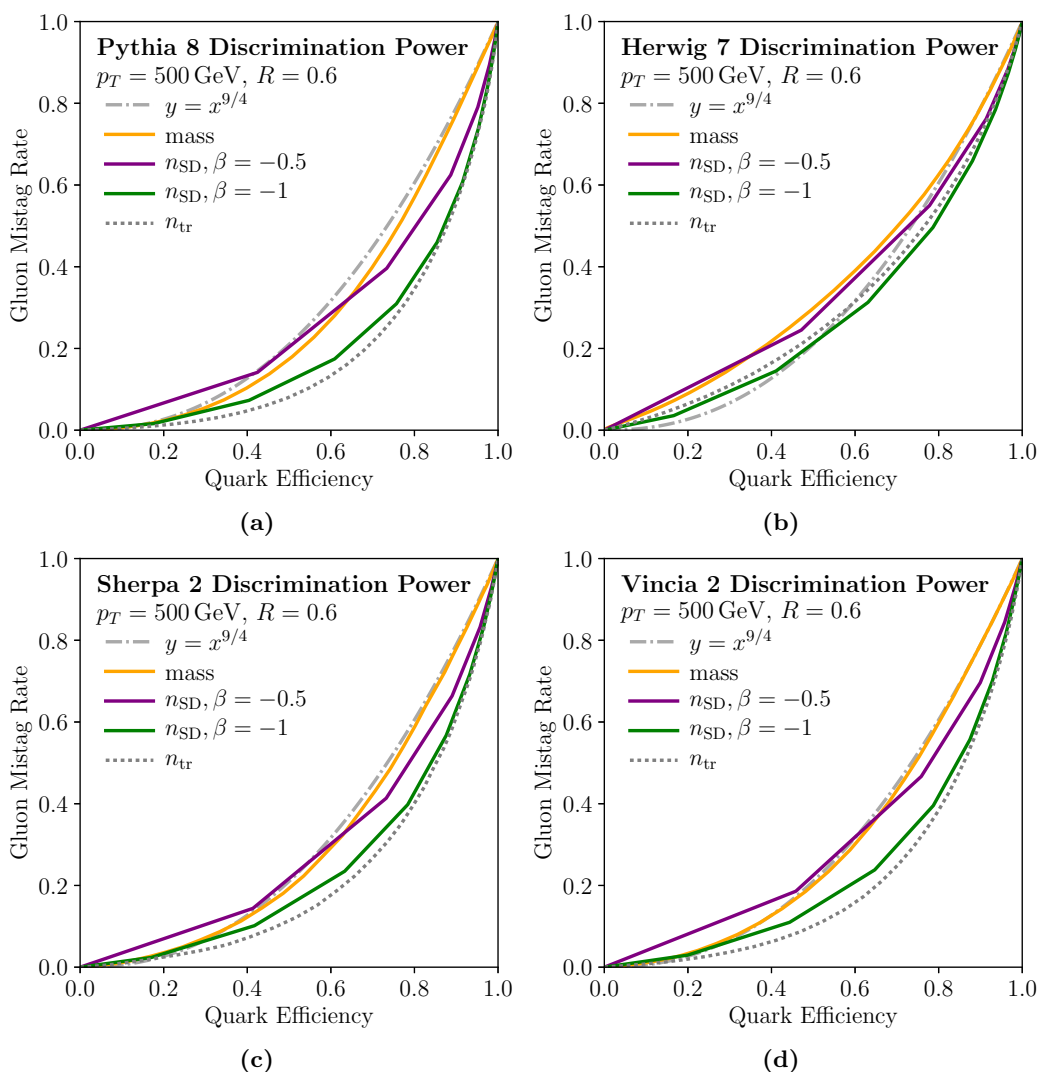
Figure 13. Predicted quark/gluon discrimination power from (a) PYTHIA 8.219, (b) HERWIG 7.0.1, (c) SHERPA 2.2.0, and (d) VINCIA 2.0.01. While the generators disagree about absolute performance, they agree that n_{SD} with $\beta = -1$ outperforms jet mass and approaches the discrimination power of n_{tr}.

When interpreting these curves, one has to remember that the NLL prediction does not include nonperturbative effects. The quark distributions are roughly similar between the various generators, but there is a larger spread in the gluon distributions, a feature also seen in the study of refs. [20, 23]. It is interesting to note that both VINCIA and SHERPA, as well as our NLL calculation, predict rather strong discrimination power, in better agreement with PYTHIA than with HERWIG. This highlights the importance of carrying out these analytic calculations to even higher accuracy, in order to better understand the desired behavior for these parton shower generators.
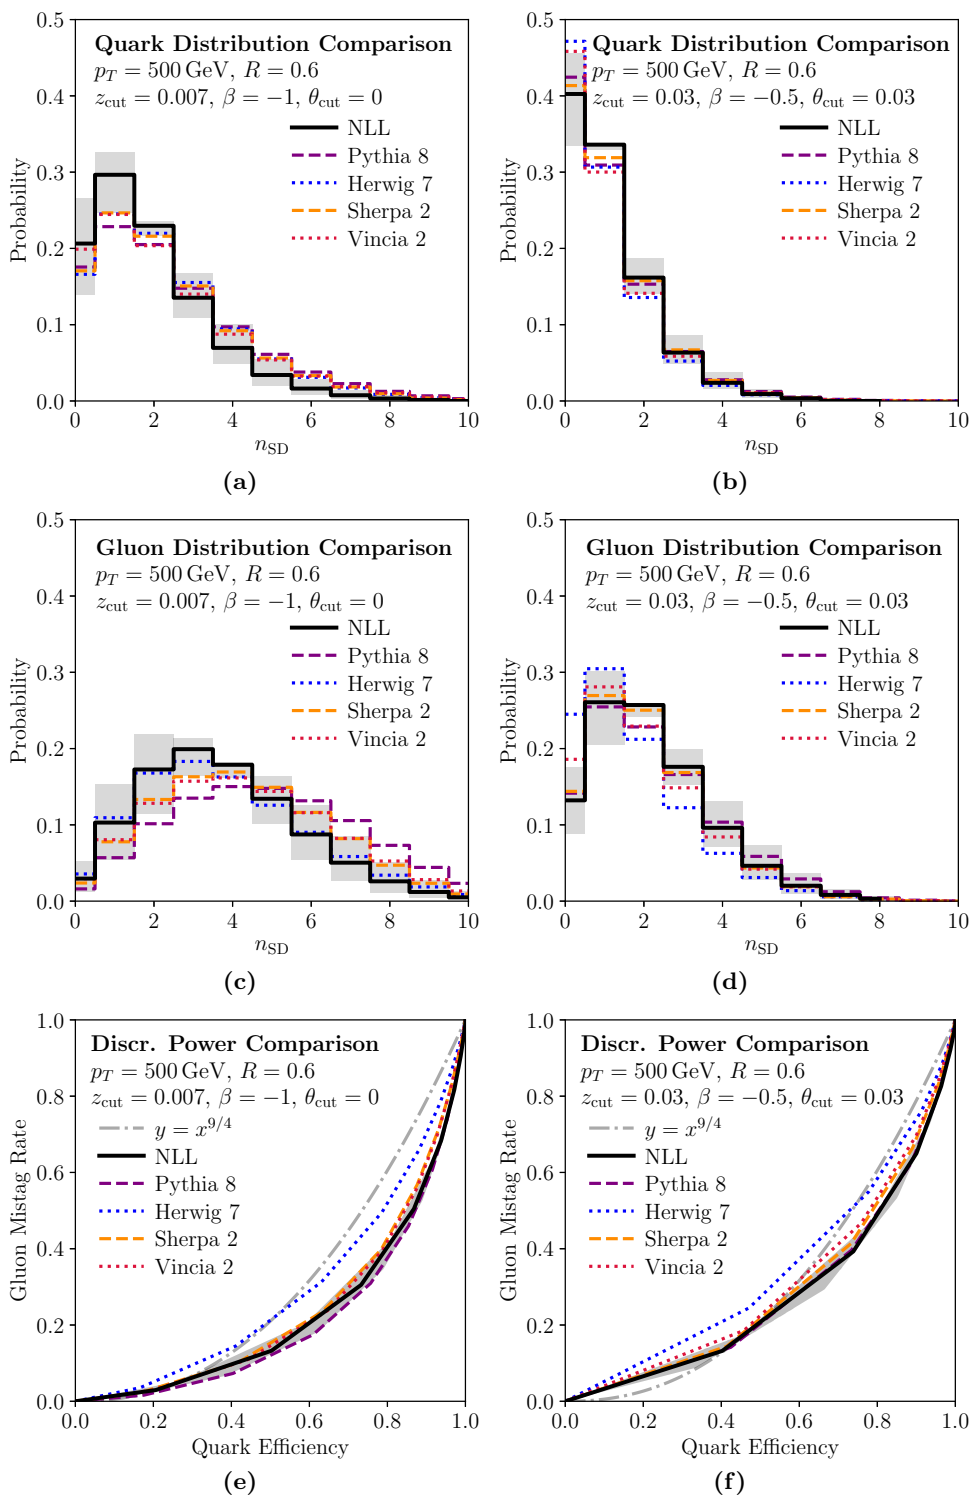
Figure 14. Analytic NLL distributions compared to parton shower generators for (top row) quark jets, (middle row) gluon jets, along with (bottom row) the corresponding ROC curves. Parameters are chosen according to eqs. (3.17), (3.18) and (3.20) with $\Lambda_{\rm NP} = 2\,{\rm GeV}$ and (left column) $\beta = -1$ and (right column) $\beta = -0.5$.

5 Calculations for collinear-unsafe soft drop multiplicity

Thus far, we have focused on choices of ISD parameters where the quark/gluon discrimination power could be predicted using perturbation theory. In the section, we consider the special case of $\theta_{\text{cut}} = 0$ and $\beta = 0$, where the soft drop multiplicity is collinear unsafe but still soft safe, allowing us to calculate its RG evolution.

5.1 Review of generalized fragmentation functions

To study observables with purely collinear final-state divergences, one can use the formalism of GFFs. Ordinary fragmentation functions are well-known objects in QCD which describe the fragmentation of a quark or gluon into a single hadron. GFFs are nonperturbative objects that describe the fragmentation of a quark or gluon into correlated sets of hadrons. The GFF technique has already been applied successfully to weighted jet charge [10, 73], track functions [52, 53], and generalized angularities [18], and a forthcoming paper explores the broader space of observables described by GFFs [45].

Each collinear-unsafe observable x has an associated set of GFFs, $\mathcal{F}_i(x, \mu)$, where i labels each quark flavor, anti-quark flavor, and gluon. They are normalized to have unit integral,

$$\int_{-\infty}^{\infty} \mathrm{d}x \, \mathcal{F}_i(x, \mu) = 1, \tag{5.1}$$

and at leading order, they have the interpretation of the probability of parton i to yield the observable value x. In higher-order partonic calculations, the GFFs absorb collinear divergences and pick up dependence on the RG scale μ. While the GFFs themselves cannot be calculated using perturbation theory, their RG evolution is calculable. Ordinary fragmentation functions exhibit linear DGLAP evolution [65, 74, 75], whereas GFFs in general have non-linear evolution equations which can even involve mixing between different sets of GFFs.

As shown in ref. [45], though, for observables defined on a pairwise clustering tree, the evolution equations for the GFFs greatly simplify. These observables are called fractal jet observables, since their RG evolution is reminiscent of the fractal structure of the parton shower. For $\theta_{\text{cut}} = 0$ and $\beta = 0$, soft drop multiplicity (and its weighted variant) is an example of a fractal jet observable, allowing us to use the GFF formalism.

It is important to emphasize that the GFF formalism only works for purely collinear divergences. For $\theta_{\text{cut}} = 0$ but $\beta > 0$, there are mixed soft-collinear divergences in the simultaneous $z \to 0$ and $\theta \to 0$ limits. These correlated diverges would require additional regulators, similar in spirit to rapidity regularization [76] (see also [77]). The use of fragmentation functions to study the $\beta = 0$ limit was previously considered in ref. [78] to study the soft-dropped z_g distribution (which is the same as z_1 for ISD).

Following ref. [45], consider a fractal observable x defined recursively on an IRC-safe binary clustering tree as follows. Each final-state hadron is assigned a starting weight w_a, which serves as the initial seed for the observable, and the observable x is built recursively according to

$$x = \hat{x}(z, x_1, x_2), \tag{5.2}$$

where $z \in [0, 1]$ is the momentum fraction of the $2 \to 1$ merging, and x_1 and x_2 are the values of the observable (or the starting weight w_a) on the daughter nodes. Note that $\hat{x}$ is independent of the opening angle θ of the merging, and the only angular dependence comes through the choice of clustering tree. The leading-order RG evolution for the GFFs associated with x is

$$\mu \frac{\mathrm{d}}{\mathrm{d}\mu} \mathcal{F}_i(x, \mu) = \frac{1}{2} \sum_{jk} \int \mathrm{d}z \, \mathrm{d}x_1 \, \mathrm{d}x_2 \, \frac{\alpha_s(\mu)}{\pi} P_{i \to jk}(z) \, \mathcal{F}_j(x_1, \mu) \, \mathcal{F}_k(x_2, \mu) \, \delta \left[x - \hat{x}(z, x_1, x_2) \right],$$
(5.3)

where $P_{i \to jk}(z)$ is the splitting function. At this order, the evolution equation (but not the observable itself) is independent of the choice of clustering tree. Note that the evolution equation is also independent of the starting weights w_a, which are effectively encoded in the low-scale initial conditions for $\mathcal{F}_i$. Even though the clustering tree is IRC safe, x is generally collinear unsafe, since eq. (5.2) allows an exactly collinear splitting to change the observable.

The canonical RG scale for a generic GFF is

$$\mu = E_{\mathrm{jet}} R_0,$$
(5.4)

and if we can extract the functional form of $\mathcal{F}_i(x, \mu)$ at a low scale, we can use eq. (5.3) to predict their form at a higher scale. The RG equations have the same recursive structure as a parton shower, and we can use the numerical techniques of ref. [45] to evolve the GFFs in μ. As we will see, our observable of interest actually has a linear evolution equation, which greatly simplifies the numerical treatment.

5.2 Linear evolution for soft drop multiplicity

For $\theta_{\mathrm{cut}} = 0$ and $\beta = 0$, soft drop multiplicity is an example of a fractal observable. More generally, any ISD observable of the form

$$x = \sum_n f(z_n)$$
(5.5)

is a fractal observable. Using C/A for the binary clustering tree with starting weights $w_a = 0$, the recursion relation for this general observable is

$$\hat{x}(z, x_1, x_2) = \begin{cases} x_2 & 0 \leq z < z_{\mathrm{cut}}, \\ x_2 + f(z) & z_{\mathrm{cut}} \leq z \leq 1/2, \\ x_1 + f(1 - z) & 1/2 \leq z \leq 1 - z_{\mathrm{cut}}, \\ x_1 & 1 - z_{\mathrm{cut}} < z \leq 1. \end{cases}$$
(5.6)

The four cases check which subjet is harder and whether the softer subjet passes soft drop. If the softer subjet fails soft drop (i.e. $\min(z, 1 - z) < z_{\mathrm{cut}}$), then the observable value is unchanged. If the softer subjet passes soft drop, then the $f(z)$ (or $f(1 - z)$) value of the splitting enters linearly into the observable.

The recursion relation in eq. (5.6) takes a particularly simple form, since each of the four cases involves either x_1 or x_2, but not both. This allows us to rewrite the RG evolution

from eq. (5.3) in the form

$$
\mu \frac{\mathrm{d}}{\mathrm{d}\mu} \mathcal{F}_i(x,\mu) = \sum_{jk} \frac{\alpha_s(\mu)}{\pi} \left(\int_0^{z_{\mathrm{cut}}} \mathrm{d}z\, P_{i \to jk}(z) \mathcal{F}_k(x,\mu) + \int_{z_{\mathrm{cut}}}^{1/2} \mathrm{d}z\, P_{i \to jk}(z) \mathcal{F}_k(x - f(z), \mu) \right),
$$
(5.7)

where we have simplified using the identity $P_{i \to jk}(z) = P_{i \to kj}(1 - z)$. This evolution equation is linear, and hence is numerically no more difficult to solve than the ordinary DGLAP equations. This form holds both for the ordinary soft drop multiplicity as well as for the weighted variants in appendix A, just with a different choice of $f(z)$.

5.3 Evolution for pure Yang-Mills

Before showing numerical results, it is instructive to consider the case of $n_f = 0$, where there is only a gluon GFF and the evolution can be studied analytically. Of course, this limit cannot teach us anything about quark/gluon discrimination directly, but we will see that the gluon GFF asymptotes to an exact Poisson distribution at sufficiently large μ, such that it behaves like an idealized counting observable.

For pure Yang-Mills, we can drop flavor labels, and write the gluon GFF as $\mathcal{F} \equiv \mathcal{F}_g$ and the relevant splitting function as $P(z) \equiv P_{g \to gg}(z)$. Specializing to soft drop multiplicity (i.e. $f(z) = 1$), the evolution equation in eq. (5.7) becomes

$$
\mu \frac{\mathrm{d}}{\mathrm{d}\mu} \mathcal{F}(x,\mu) = \frac{\alpha_s(\mu)}{\pi} \left(\int_0^{z_{\mathrm{cut}}} \mathrm{d}z\, P(z) \mathcal{F}(x,\mu) + \int_{z_{\mathrm{cut}}}^{1/2} \mathrm{d}z\, P(z) \mathcal{F}(x-1,\mu) \right)
$$
(5.8)

$$
= P_{\mathrm{ave}} \frac{\alpha_s(\mu)}{2\pi} \Big(\mathcal{F}(x-1,\mu) - \mathcal{F}(x,\mu) \Big),
$$
(5.9)

where we have defined

$$
P_{\mathrm{ave}} = \int_{z_{\mathrm{cut}}}^{1-z_{\mathrm{cut}}} \mathrm{d}z\, P(z).
$$
(5.10)

The interpretation of eq. (5.9) is that gluon emissions that pass soft drop are added at a rate of $P_{\mathrm{ave}} \alpha_s(\mu)/2\pi$ in $\log \mu$ evolution. Specifically, in evolving from μ_i to μ_f, the expected number of additional emissions is

$$
\lambda(\mu_i, \mu_f) = \frac{P_{\mathrm{ave}}}{2\pi} \int_{\log \mu_i}^{\log \mu_f} \mathrm{d}(\log \mu)\, \alpha_s(\mu),
$$
(5.11)

so the GFF at μ_f is

$$
\mathcal{F}(x, \mu_f) = \mathcal{F}(x, \mu_i) \otimes \mathrm{Pois}(\lambda(\mu_i, \mu_f))[x],
$$
(5.12)

where the convolution is in x.[12]

As μ_f increases, more emissions are added, so the initial GFF distributions at μ_i becomes less and less important. Substituting in the one-loop running of the strong coupling constant in pure Yang-Mills,

$$
\alpha_s(\mu) = \frac{1}{\beta_0 \log(\mu^2/\Lambda_{\mathrm{QCD}}^2)}, \qquad \beta_0 = \frac{11}{3} C_A,
$$
(5.13)

---

[12]The reader who finds this derivation too slick can explicitly check that eq. (5.12) solves eq. (5.9). It is helpful to note that $\frac{\mathrm{d}}{\mathrm{d}\lambda} \mathrm{Pois}(\lambda)[x] = \mathrm{Pois}(\lambda)[x-1] - \mathrm{Pois}(\lambda)[x]$.

the number of expected emissions is

$$\lambda(\mu_i, \mu_f) = \frac{P_{\text{ave}}}{4\pi\beta_0} \log\left(\frac{\log\frac{\mu_f}{\Lambda_{\text{QCD}}}}{\log\frac{\mu_i}{\Lambda_{\text{QCD}}}}\right). \tag{5.14}$$

Since this quantity continues to grow at high μ_f, the IR boundary condition $\mathcal{F}(x, \mu_i)$ is irrelevant in the $\mu_f \to \infty$ limit, yielding the asymptotic form

$$\mathcal{F}(x, \mu \gg \Lambda_{\text{QCD}}) \approx \text{Pois}(\lambda(\mu))[x], \quad \lambda(\mu) = \frac{P_{\text{ave}}}{4\pi\beta_0} \log\log\frac{\mu}{\Lambda_{\text{QCD}}}. \tag{5.15}$$

Thus, we find a Poisson distribution whose mean scales as $\log\log\mu$, such that the soft drop multiplicity acts like an idealized counting observable.

5.4 Comparison to parton showers

We now compare the results of the GFF approach to parton shower predictions. First, in figure 15, we show the predicted discrimination power for the collinear-unsafe n_{SD} from the same four parton showers studied in section 4.3. We see that for low z_{cut} values, the discrimination power of the collinear-unsafe soft drop multiplicity approaches that of our benchmark IRC-safe soft drop multiplicity, previously shown in figure 13. (It does not, however, reach the power of the nonperturbative soft drop multiplicities shown in figure 10a.) Making z_{cut} any smaller does not significantly improve discrimination power, so we use $z_{\text{cut}} = 0.02$ as our baseline parameter choice.

To make a prediction using the GFF approach, we need to extract the nonperturbative distributions at a low scale and then evolve them to a higher scale. In a full analysis, the low scale distributions would be extracted from data, but here we can use the parton shower generators. For this, we switch to e^+e^- collisions, generating pure quark and gluon samples through the processes $e^+e^- \to \gamma/Z^* \to q\bar{q}$ and $e^+e^- \to H^* \to gg$ in VINCIA 2.0.01. Setting $R_0 = 0.6$ as our baseline, we generate jets with energies in a 10% window of $E_{\text{jet}} = 400\,\text{GeV}$, corresponding to $\mu = E_{\text{jet}}R_0 = 240\,\text{GeV}$. We then extract n_{SD} from the generated events, which at leading order, is a direct measure of the corresponding GFFs.[13]

Using eq. (5.7), we evolve the GFFs to 4 TeV using the energy scale in eq. (5.4) and the two-loop running of α_s.[14] This evolution includes all 10 active quark and antiquark flavors, as $n_f = 5$ in this energy range.[15] There are various sources of theoretical uncertainties in the evolved result, and we highlight two of them in this study. The first contribution is due to the fact that the energy scale eq. (5.4) only depends on the product $E_{\text{jet}}R_0$, though the initial distributions could be extracted with any R_0. To estimate this uncertainty, which serves as a consistency check of the choice of μ scale, we also extract GFFs with $R_0 = 0.3$ and $R_0 = 0.9$, keeping μ fixed. The second contribution is from uncertainty in the absolute

---

[13]At higher orders, one has to perform a matching calculation; see further discussion in ref. [45].

[14]Since we only consider the leading-order evolution of the GFFs, strictly speaking, only leading-order evolution of α_s is needed at this order. Switching to one-loop running has a negligible effect on the results of this section.

[15]For simplicity, we ignore effects due to the $g \to t\bar{t}$ splitting, which would require a matching calculation to the top quark electroweak decay.
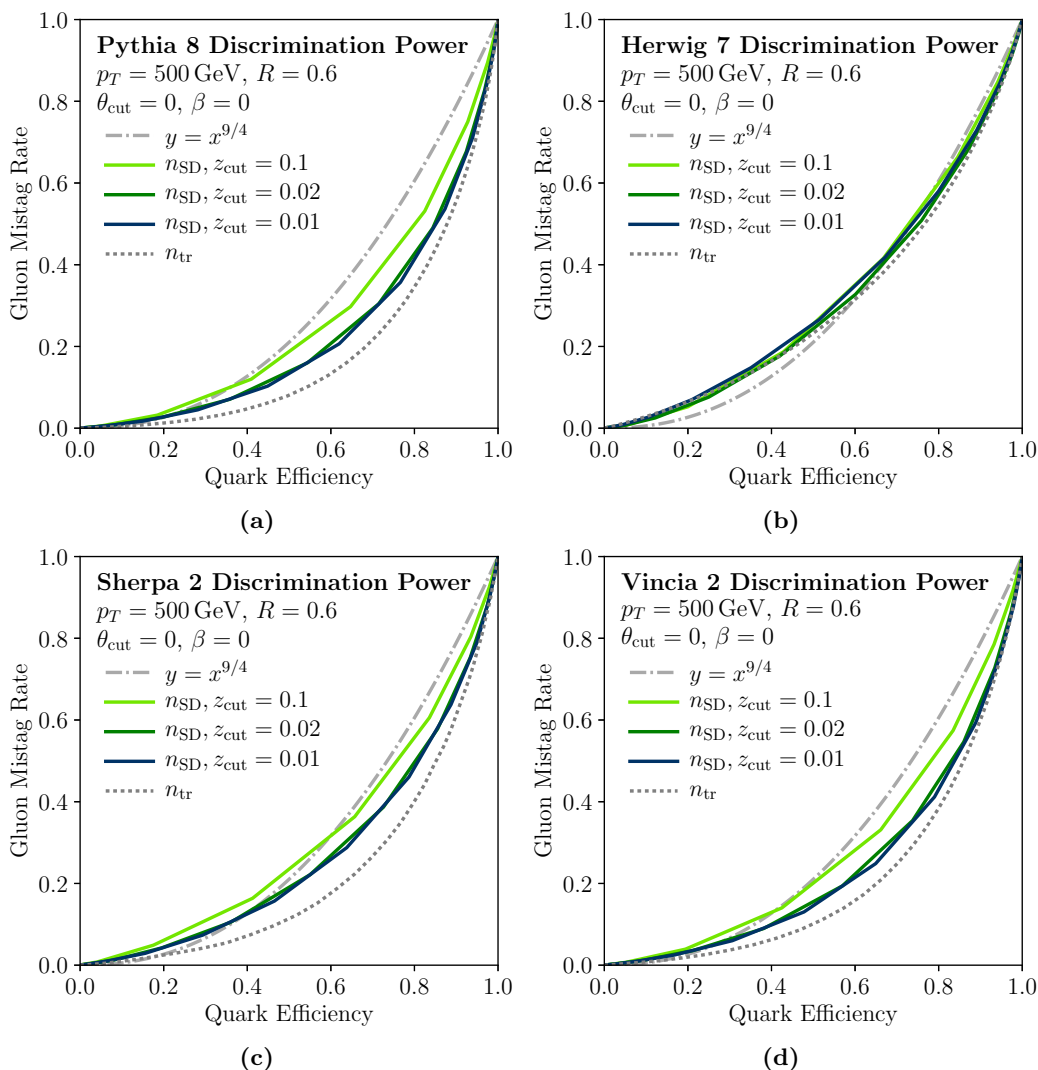
Figure 15. Same as figure 13, but for the collinear-unsafe soft drop multiplicity with $\theta_{\mathrm{cut}} = 0$ and $\beta = 0$.

value of the energy scale itself. To address this, we perform evolution with both half and double the energy scale of eq. (5.4). We plot the envelope of these 9 results in a shaded uncertainty band. Of course, this is only a subset of the possible GFF uncertainties, but a full study is beyond the scope of this work.

The results for $z_{\mathrm{cut}} = 0.02$ and $z_{\mathrm{cut}} = 0.1$ are shown in figure 16, comparing the RG-evolved results to Vincia distributions extracted at the high scale. To show a single curve for quark jets, we plot the quark-singlet distribution

$$\mathcal{Q}(x, \mu) = \frac{1}{2n_f} \sum_{i \in \{u, \bar{u}, ..., b, \bar{b}\}} \mathcal{F}_i(x, \mu) \tag{5.16}$$

as defined in ref. [45] (where it is instead denoted by $\mathcal{S}$). We find reasonable agreement between the RG evolution and Vincia for both z_{cut} values, with a larger range of evolution
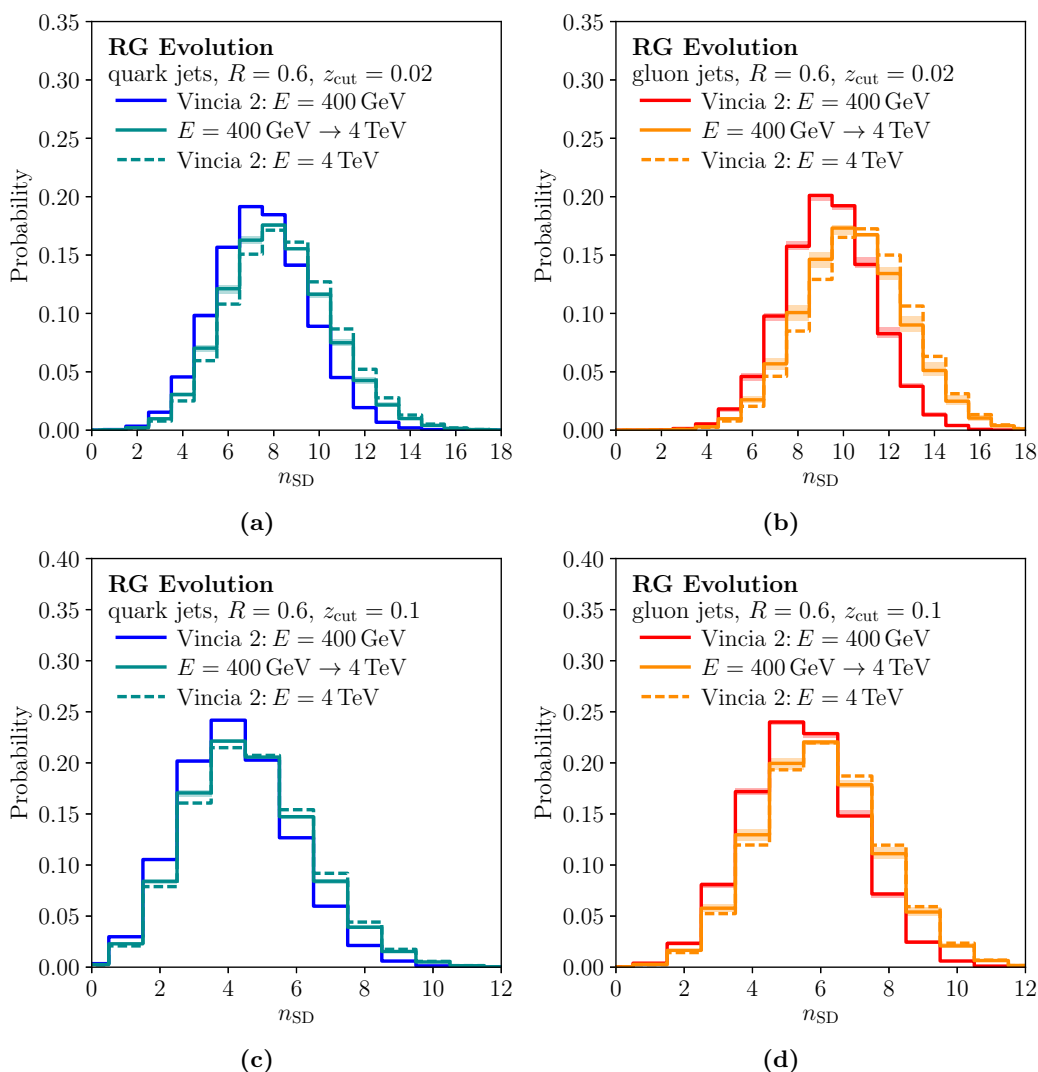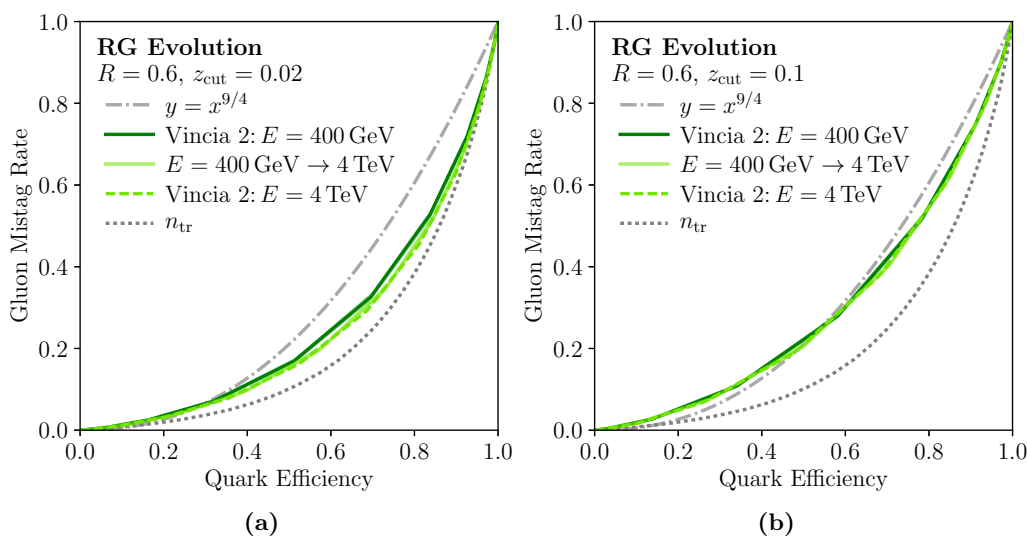
Figure 16. RG evolution of the collinear-unsafe soft drop multiplicity for (left column) the quark singlet GFF and (right column) the gluon GFF. Shown are the results for (top row) $z_{cut} = 0.02$ and (bottom row) $z_{cut} = 0.1$, taking distributions extracted from VINCIA at a low scale and evolving them to a higher scale. The uncertainties in the evolved distributions come from varying the jet radius used for GFF extraction and the μ scale for the RG evolution.

for the case of $z_{cut} = 0.02$. The uncertainties in the RG evolution do not fully cover the high-scale VINCIA distribution, though it is worth emphasizing that we are only using the LO evolution equations.

In figure 17, we show the RG evolution of the quark/gluon ROC curves. Despite the fact that the n_{SD} distributions themselves exhibit significant RG evolution, the corresponding ROC curves do not change significantly with the energy scale μ. This is a key prediction of the GFF approach, and one that we can better understand by studying the moments of the GFF distributions.

Figure 17. RG evolution of ROC curve (quark singlet vs. gluon) for the collinear-unsafe soft drop multiplicity with (a) $z_{\rm cut} = 0.02$ and (b) $z_{\rm cut} = 0.1$. In both cases, there is very little evolution in the discrimination power with energy scale.

5.5 Moment space evolution

To understand the slow evolution of the quark/gluon discrimination power, consider the evolution in moment space. Following ref. [45], the $n^{\rm th}$ moment of a GFF is defined as

$$\overline{\mathcal{F}}_i(n, \mu) = \int \mathrm{d}x \, x^n \mathcal{F}_i(x, \mu). \tag{5.17}$$

In moment space, we denote the gluon GFF by $\overline{\mathcal{G}}(n, \mu)$, and the quark-singlet GFF (as defined in eq. (5.16)) by $\overline{\mathcal{Q}}(n, \mu)$. To derive the moment space evolution equations, we integrate both sides of eq. (5.8) against x^n, shift the final integral by $x \to x + 1$, and then simplify the $n^{\rm th}$ moments with the splitting function identities

$$\int_0^1 \mathrm{d}z \, [P_{g\to gg}(z) + 2n_f P_{g\to q\bar{q}}(z)] = 0, \quad \int_0^1 \mathrm{d}z \, P_{q\to qg}(z) = 0. \tag{5.18}$$

After these manipulations, the moment evolution equation for the $n^{\rm th}$ gluon or quark-singlet GFF can be written solely in terms of the difference $\overline{\mathcal{G}}(n) - \overline{\mathcal{Q}}(n)$, along with lower moments $\overline{\mathcal{G}}(k)$, $\overline{\mathcal{Q}}(k)$ for $k < n$.

For $n = 1$, the evolution equation for the means is

$$\mu \frac{\mathrm{d}}{\mathrm{d}\mu} \begin{pmatrix} \overline{\mathcal{G}}(1) \\ \overline{\mathcal{Q}}(1) \end{pmatrix} = \frac{\alpha_s}{\pi} \left[(\overline{\mathcal{G}}(1) - \overline{\mathcal{Q}}(1)) \begin{pmatrix} \bar{P}_{g\to gg}^{0,1/2} \\ \bar{P}_{q\to qg}^{0,1/2} \end{pmatrix} + \begin{pmatrix} \bar{P}_{g\to gg}^{z_{\rm cut},1/2} + 2n_f \bar{P}_{g\to \bar{q}q}^{z_{\rm cut},1/2} \\ \bar{P}_{q\to qg}^{z_{\rm cut},1/2} + \bar{P}_{q\to gq}^{z_{\rm cut},1/2} \end{pmatrix} \right] \tag{5.19}$$

where we are suppressing the μ arguments and using the abbreviated notation

$$\bar{P}_{i\to jk}^{z_1,z_2} = \int_{z_1}^{z_2} \mathrm{d}z \, P_{i\to jk}(z). \tag{5.20}$$

The appearance of the difference of the moments on the right-hand side has a dramatic effect on the high-energy limit of the evolution. Specifically, the difference in the means evolves as

$$\mu \frac{\mathrm{d}}{\mathrm{d}\mu}\big(\overline{\mathcal{G}}(1) - \overline{\mathcal{Q}}(1)\big) = \frac{\alpha_s}{\pi}\left[c_1 - c_2\big(\overline{\mathcal{G}}(1) - \overline{\mathcal{Q}}(1)\big)\right], \tag{5.21}$$

where c_1 and c_2 are positive constants defined by integrals of the splitting functions. Thus, at high energies, the difference in the means asymptotes to a constant,

$$\overline{\mathcal{G}}(1) - \overline{\mathcal{Q}}(1) \Rightarrow \frac{c_1}{c_2} = \frac{\bar{P}_{g\to gg}^{z_{\mathrm{cut}},1/2} + 2n_f \bar{P}_{g\to\bar{q}q}^{z_{\mathrm{cut}},1/2} - \bar{P}_{q\to qg}^{z_{\mathrm{cut}},1/2} - \bar{P}_{q\to gq}^{z_{\mathrm{cut}},1/2}}{\bar{P}_{q\to qg}^{0,1/2} - \bar{P}_{g\to gg}^{0,1/2}}. \tag{5.22}$$

This asymptotic behavior is strikingly different from the LL analysis of IRC-safe multiplicity in section 3. In the IRC-safe case, the LL prediction is that the gluon and quark means should have a constant ratio determined by C_A/C_F. Here, in the collinear-unsafe case, the gluon and quark means asymptote to having a constant difference. Physically, this occurs because the RG evolution takes flavor mixing effects into account, so that at sufficiently high energies, the n_{SD} distributions for quark and gluon jets become essentially the same. While we have only presented the calculation for the quark-singlet mean, it is straightforward to show that the means for each individual quark flavor behave in the same way, with differences between different quark flavors evolving to zero.

Moving to higher moments, a useful simplification occurs for the variances,

$$\sigma_i^2 = \overline{\mathcal{F}}_i(2) - \overline{\mathcal{F}}_i(1)^2. \tag{5.23}$$

In this case, the evolution of the variances only depends on the difference of the variances and the difference of the means,

$$\mu \frac{\mathrm{d}}{\mathrm{d}\mu}\begin{pmatrix} \sigma_{\overline{\mathcal{G}}}^2 \\ \sigma_{\overline{\mathcal{Q}}}^2 \end{pmatrix} = \frac{\alpha_s}{\pi}\left[\begin{pmatrix} \bar{P}_{g\to gg}^{0,1/2} \\ \bar{P}_{q\to qg}^{0,1/2} \end{pmatrix}\left(\sigma_{\overline{\mathcal{G}}}^2 - \sigma_{\overline{\mathcal{Q}}}^2 - \big(\overline{\mathcal{G}}(1) - \overline{\mathcal{Q}}(1)\big)^2\right) + \begin{pmatrix} \bar{P}_{g\to gg}^{z_{\mathrm{cut}},1/2} + 2n_f \bar{P}_{g\to\bar{q}q}^{z_{\mathrm{cut}},1/2} \\ \bar{P}_{q\to qg}^{z_{\mathrm{cut}},1/2} + \bar{P}_{q\to gq}^{z_{\mathrm{cut}},1/2} \end{pmatrix}\right]. \tag{5.24}$$

At sufficiently high energies, $\overline{\mathcal{G}}(1) - \overline{\mathcal{Q}}(1)$ approaches a constant, so the evolution equation for the variances is of the same form as the evolution equation for the means. We find that, like the means, the difference of variances asymptotes to a constant,

$$\sigma_{\overline{\mathcal{G}}}^2 - \sigma_{\overline{\mathcal{Q}}}^2 \Rightarrow \mathrm{const.} \tag{5.25}$$

Substituting our asymptotic results back into eq. (5.19) and eq. (5.24), we see that both the mean and variance simply grow linearly in $\alpha_s(\mu)\mathrm{d}(\log\mu)$ at high energies, so that they become proportional in the UV limit. Therefore, even with flavor-mixing effects, the soft drop multiplicity maintains a Poisson-like distribution, with $\sigma^2 = O(\mu)$.

We can roughly estimate the discrimination power of the soft drop multiplicity using a relative width, similar to that of eq. (3.13). Since Casimir scaling no longer holds, the distance between the quark-singlet and gluon distributions is no longer characterized by the means, but rather the difference in means. Moreover, in the UV limit, the standard deviations of the quark singlet and gluon distributions approach each other. Thus, the quantity

$$w_{\mathrm{rel}} \equiv \frac{\sqrt{\sigma_{\overline{\mathcal{G}}}^2}}{\overline{\mathcal{G}}(1) - \overline{\mathcal{Q}}(1)} \tag{5.26}$$
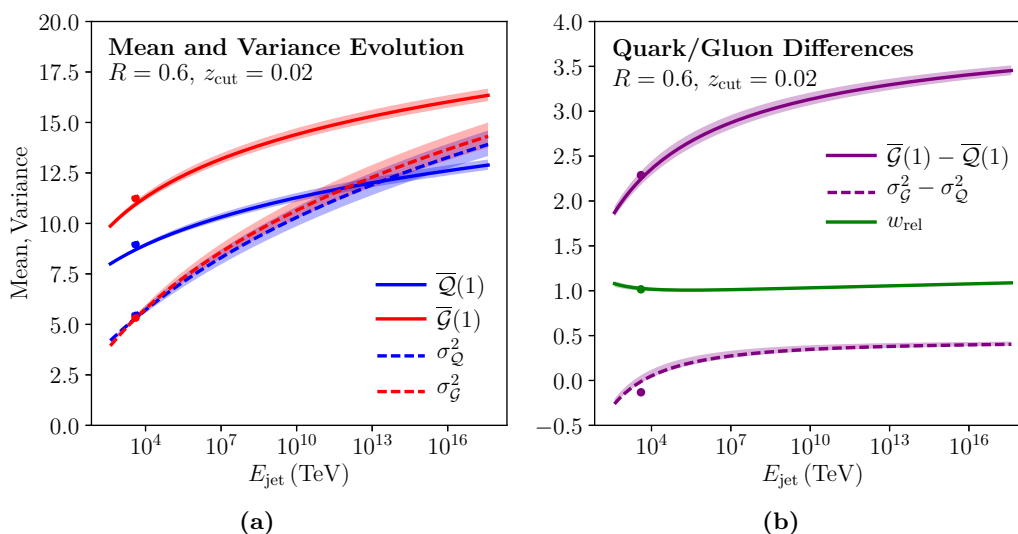
Figure 18. (a) RG evolution of means and variances of the quark-singlet and gluon GFFs for the soft drop multiplicity with $z_{\text{cut}} = 0.02$. (b) RG evolution of the mean/variance differences, which asymptotically approach constants. Also shown is the relative width w_{rel} defined in eq. (5.26), which increases slowly. For comparison, quantities extracted from VINCIA at $E_{\text{jet}} = 4\,\text{TeV}$ are shown as dots.

characterizes the extent to which the distributions overlap, and hence measures the discrimination power of the soft drop multiplicity. We see that, as a result of flavor-mixing effects, the relative width is now expected to increase somewhat as more emissions are counted, roughly as the square root of the mean.

To verify these results, we numerically evolve the GFFs according to eq. (5.7), starting from an initial condition extracted from VINCIA 2.0.01 at $E_{\text{jet}} = 400\,\text{GeV}$ and $R = 0.6$. As in figure 16, we show a theoretical uncertainty band constructed from the envelope of 9 results. In figure 18a, we show the evolution of the mean and variance of the soft drop multiplicity for quark singlets and gluons. As expected from the above analysis, the mean and variance curves become parallel at sufficiently large values of μ. This is confirmed in figure 18b, which shows that the differences do indeed asymptote to constant values.

Crucially, the relative width in figure 18b remains approximately constant over a large energy range, as the increase in the standard deviation is canceled by the increase in the mean difference as it approaches its asymptotic value. This explains the slow evolution of discrimination power seen in figure 17. In this way, even though these collinear-unsafe distributions cannot be predicted directly from first principles, the GFF approach gives us a valuable analytic handle on their RG evolution.

6 Conclusions

Quark/gluon discrimination has a long history, with many proposed discriminants [15, 18, 21, 22, 25, 73, 79–85] though relatively few analytic calculations [17–19]. Because C_A/C_F

is an order 1 number, distinguishing quark- from gluon-initiated jets is an intrinsically hard problem. Moreover, to gain a quantitative understanding of quark/gluon separation power, one has to account for physics effects beyond the LL approximation, including the impact of nonperturbative physics. These physics effects are modeled to differing degrees in parton shower generators, but ultimately one wants quark/gluon studies to be based on systematically-improvable analytic calculations.

In this paper, we introduced an IRC-safe counting observable which approaches the quark/gluon discrimination performance of IRC-unsafe track multiplicity. Through a LL analysis, we demystified the power of multiplicity, showing that Poisson distributions typically yield better quark/gluon separation than Sudakov distributions, even though they are both controlled by the same C_A and C_F Casimir factors. Specifically, we introduced soft drop multiplicity, which depends on multiple soft gluon emissions even at LL accuracy, allowing it to outperform observables like jet mass whose value is dominated by a single gluon emission. Remarkably, there is a choice of ISD parameters where soft drop multiplicity is controlled by perturbative physics, such that its behavior can be reliably studied from first principles.

To gain a more quantitative understanding of $n_{\rm SD}$, we introduced NLL evolution equations, which allowed us to make interesting comparisons to parton shower generators. We also studied a collinear-unsafe (but infrared-safe) version of $n_{\rm SD}$, whose RG evolution could be studied using the formalism of GFFs. In both cases, analytic understanding was aided by the recursive structure of the observable. This motivates further studies into jet measurements performed on (groomed) clustering trees, which can depart significantly from the more commonly studied additive observables.

Ultimately, any single observable will never match the performance of multivariate jet tagging methods. This has been emphasized recently in the context of deep neural networks which exploit subtle correlations to maximize separation power [21, 86–95]. Still, we are encouraged by observables like soft drop multiplicity which offer a balance between discrimination power and analytic tractability. Going beyond LL order where $n_{\rm SD}$ can saturate the discrimination power (see section 3.3), it would be interesting to study correlations between $n_{\rm SD}$ and other IRC-safe observables like jet mass to see if there is additional information in their combination. Because the physics basis for $n_{\rm SD}$ is so transparent, we suspect it will be a useful benchmark for both parton shower tuning and experimental jet analyses. Because the analytic structure of $n_{\rm SD}$ is so unique, we hope it inspires new precision calculations in QCD.

Acknowledgments
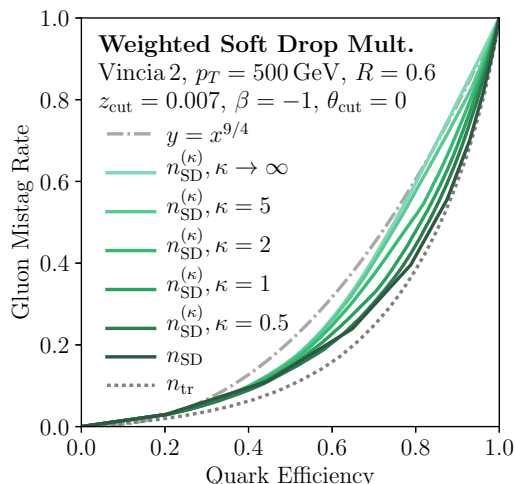
Figure 19. Quark/gluon discrimination power of weighted soft drop multiplicity as a function of κ, at the benchmark parameters from eq. (2.7). We also show the limit $\kappa \to \infty$, which is equivalent to $\max(z_n)$.

A Weighted soft drop multiplicity

At the end of section 3, we used LL reasoning to argue that soft drop multiplicity n_{SD} extracts all of the quark/gluon discriminatory information from the (z_n, θ_n) variables recorded by ISD. In this appendix, we study a variant of n_{SD}, the weighted soft drop multiplicity, defined in eq. (2.8) and repeated for convenience:

$$n_{\mathrm{SD}}^{(\kappa)} = \sum_n z_n^\kappa. \tag{A.1}$$

While quark/gluon performance is not improved by weighting, the purpose of this appendix is to demonstrate that the techniques of this paper are applicable to a variety of observables.

A.1 Discrimination power

For small values of κ, the weighted soft drop multiplicity is still sensitive to all emissions in the region A_{emit}. On the other hand, as $\kappa \to \infty$, only the largest z_n value contributes significantly to the observable. As a result, the weighted multiplicity interpolates between counting and additive behavior, in the limits $\kappa \to 0$ and $\kappa \to \infty$, respectively. The κ dependence of the discrimination power, extracted from VINCIA, is shown in figure 19. One can see that the quark/gluon performance decreases monotonically as κ increases.

The LL distribution of the weighted soft drop multiplicity is analytically complicated. Indeed, any analytic expression for it must contain a sum of distributions, one for each value of the number n of counted emissions. For example, when $\beta \leq 0$, each emission contributes at least z_{cut}^κ, so at most n emissions can contribute to $n_{\mathrm{SD}}^{(\kappa)}$ if its value is below $n \, z_{\mathrm{cut}}^\kappa$. A full analysis along these lines is carried out in appendix A.2 below.

To qualitatively understand the trend in figure 19, consider the limit in which ISD records many emissions. Strictly speaking, this analysis is not quantitatively applicable

in the perturbative regime, where $n \lesssim 10$ emissions are counted. Nor is this reasoning applicable in the collinear-unsafe regime studied in appendix A.3, where solely perturbative reasoning is insufficient. Nonetheless, the many-emission limit serves to build intuition.

In the double-logarithmic approximation, where emissions are soft and collinear and α_s is a fixed coupling, the weighted multiplicity distribution can be found from summing independent identically distributed numbers. By the central limit theorem, this converges to a normal distribution in the limit of many recorded emissions. In this limit, it suffices to compute the mean and variance of $n_{\text{SD}}^{(\kappa)}$ to estimate its discrimination power. These are determined at lowest order from the average values of z^κ and $z^{2\kappa}$ in the allowed emission region as

$$\langle n_{\text{SD}}^{(\kappa)} \rangle_i = \rho_i A_{\text{emit}} \langle z^\kappa \rangle \,, \qquad \text{Var}\left(n_{\text{SD}}^{(\kappa)}\right)_i = \rho_i A_{\text{emit}} \langle z^{2\kappa} \rangle \,, \tag{A.2}$$

where

$$\langle z^\kappa \rangle = \frac{1}{A_{\text{emit}}} \int_{\theta_{\text{cut}}}^{R_0} \frac{\mathrm{d}\theta}{\theta} \int_{z_{\text{cut}}}^{1/2} \frac{\mathrm{d}z}{z} \, z^\kappa \, \Theta\left[z - z_{\text{cut}}\left(\frac{\theta}{R_0}\right)^\beta\right] \,. \tag{A.3}$$

With a fixed coupling, the mean value of z^κ for $\beta > 0$ is

$$A_{\text{emit}} \langle z^\kappa \rangle^{\beta>0} = \frac{1}{2^\kappa \kappa} \log \frac{R_0}{\theta_{\text{cut}}} - \frac{z_{\text{cut}}{}^\kappa}{\beta \kappa^2} \left(1 - \left(\frac{\theta_{\text{cut}}}{R_0}\right)^{\beta\kappa}\right) \,. \tag{A.4}$$

For $\beta < 0$, the mean value is

$$A_{\text{emit}} \langle z^\kappa \rangle^{\beta<0} = \Theta\left[\theta_{\text{cut}} - (2z_{\text{cut}})^{\frac{1}{|\beta|}} R_0\right] \left(\frac{1}{2^\kappa \kappa} \log \frac{R_0}{\theta_{\text{cut}}} - \frac{z_{\text{cut}}{}^\kappa}{\beta \kappa^2} \left[1 - \left(\frac{\theta_{\text{cut}}}{R_0}\right)^{\beta\kappa}\right]\right) \tag{A.5}$$
$$+ \Theta\left[(2z_{\text{cut}})^{\frac{1}{|\beta|}} R_0 - \theta_{\text{cut}}\right] \left(\frac{1}{2^\kappa \beta \kappa} \log(2z_{\text{cut}}) - \frac{z_{\text{cut}}{}^\kappa}{\beta \kappa^2} \left[1 - (2z_{\text{cut}})^{-\kappa}\right]\right) \,.$$

Because of the ρ_i prefactor in eq. (A.2), we see that the mean and variance once again satisfy Casimir scaling as in eq. (3.11). Moreover, both the variance and mean scale with the counted area A_{emit}, establishing that the weighted soft drop multiplicity is Poisson-like distributed as defined in section 3.

The discrimination power is determined by the relative width

$$w_{\text{rel}} \equiv \frac{\sqrt{\text{Var}\left(n_{\text{SD}}^{(\kappa)}\right)_i}}{\left\langle n_{\text{SD}}^{(\kappa)} \right\rangle_i} = \frac{1}{\sqrt{\rho_i A_{\text{emit}}}} \frac{\sqrt{\langle z^{2\kappa} \rangle}}{\langle z^\kappa \rangle} \,. \tag{A.6}$$

We can get a sense for the behavior of w_{rel} by considering two extreme limits. For $\kappa \to 0$ and any choice of β, the mean value $\langle z^\kappa \rangle$ (and hence w_{rel}) approaches a constant, independent of κ. For $\kappa \to \infty$, the mean value scales with κ like

$$A_{\text{emit}} \langle z^\kappa \rangle^{\kappa\to\infty} \sim \frac{1}{2^\kappa \kappa} \,, \tag{A.7}$$

with $z_{\text{cut}} < 1/2$, such that the relative width scales as

$$w_{\text{rel}}^{\kappa\to\infty} \sim \sqrt{\kappa} \,. \tag{A.8}$$
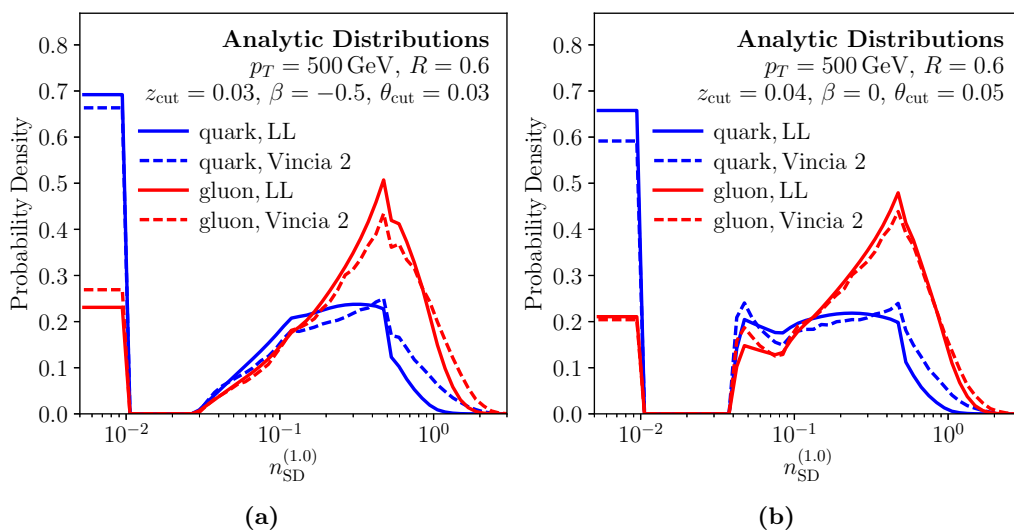
Figure 20. LL calculation of weighted soft drop multiplicity distributions with $\kappa = 1$, compared to VINCIA. The plots have two different sets of ISD parameters which were chosen to display the sharp features characteristic of $n_{\mathrm{SD}}^{(\kappa)}$ in the perturbative regime. The curves shown are the probability distribution functions of $\log n_{\mathrm{SD}}^{(1.0)}$, so that they integrate to one in logarithmic space. The leftmost bin is an underflow bin, showing the probability that no emissions were counted by ISD, such that $n_{\mathrm{SD}}^{(1.0)} = 0$.

Since the relative width increases with increasing κ, this reasoning predicts that the discrimination power decreases as κ increases. This implies the best discrimination power is attained for $\kappa = 0$ (i.e. ordinary soft drop multiplicity) and decreases for higher κ. Physically, the discrimination power of $n_{\mathrm{SD}}^{(\kappa)}$ comes from sensitivity to multiple emissions, and for higher κ, sensitivity to softer emissions is decreased. In the extreme limit of $\kappa \to \infty$, the weighted soft drop multiplicity reduces to the energy fraction of the hardest emission, $n_{\mathrm{SD}}^{\kappa \to \infty} = \max(z_n)$.

This qualitatively explains the trend seen in figure 19, i.e. that the discrimination power monotonically decreasing as κ increases. In the limit $\kappa \to \infty$, the discrimination power reaches the universal result predicted by Casimir scaling (slightly off due to small nonperturbative corrections), as the observable $\max(z_n)$ is determined by a Sudakov form factor.

A.2 Analytic calculation

Using evolution equations similar to those employed in section 4, we can compute the distribution of IRC-safe weighted soft drop multiplicities. We will demonstrate this here at LL for simplicity; by taking into account flavor changes and energy losses, one could obtain NLL evolution equations as in section 4.2. Since $n_{\mathrm{SD}}^{(\kappa)}$ is a continuous observable, however, significantly more computation time would be required to compute its NLL distribution, in comparison to the discrete unweighted case.

Let $p^i(n_{\mathrm{SD}}, \theta_{\mathrm{cut}}) \, dn_{\mathrm{SD}}$ denote the differential probability that, given a flavor i jet, its weighted soft drop multiplicity is measured to be n_{SD}. Here, we leave the z_{cut}, β, and κ
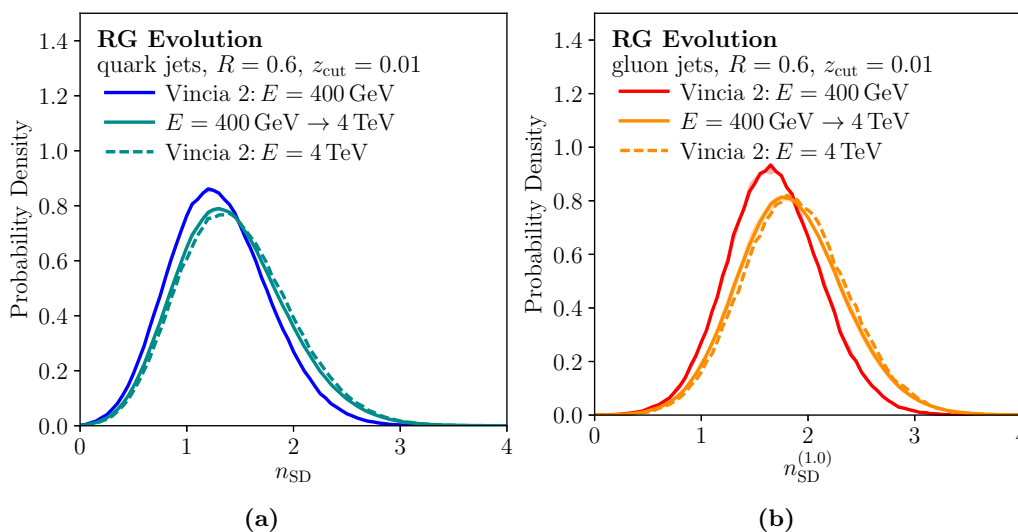
Figure 21. RG evolution of collinear-unsafe weighted soft drop multiplicity with $z_{\text{cut}} = 0.01$ and $\kappa = 1$, for the (a) quark-singlet and (b) gluon cases.

dependence implicit. Though the weighted soft drop multiplicity does not directly count emissions, it is still useful to keep track of the number of contributing emissions, using

$$p^i(n_{\text{SD}}, \theta_{\text{cut}}) = \sum_{n=0}^{\infty} p_n^i(n_{\text{SD}}, \theta_{\text{cut}}), \tag{A.9}$$

where n labels the number of counted emissions as before. If we change the resolution angle from θ_{cut} to $\theta_{\text{cut}} - \delta\theta_{\text{cut}}$, then

$$p_n^i(n_{\text{SD}}, \theta_{\text{cut}} - \delta\theta_{\text{cut}}) = p_n^i(n_{\text{SD}}, \theta_{\text{cut}}) \left(1 - \frac{\delta\theta_{\text{cut}}}{\theta_{\text{cut}}} \int_0^{1/2} dz \, \frac{\alpha_s(z\,\theta\,p_T)}{\pi} P_{i\to i}(z) \, \Theta_{\text{SD}}(z, \theta) \right)$$

$$+ \frac{\delta\theta_{\text{cut}}}{\theta_{\text{cut}}} \int_0^{1/2} dz \, \frac{\alpha_s(z\,\theta\,p_T)}{\pi} P_{i\to i}(z) \, \Theta_{\text{SD}}(z, \theta) \, p_{n-1}^i(n_{\text{SD}} - z^\kappa, \theta_{\text{cut}}). \tag{A.10}$$

This leads to a linear differential equation. Instead of the Poisson distribution found in section 4.1, the solution in this case is differential in $n_{\text{SD}} = \sum_i z_i^\kappa$:

$$p_n^i(n_{\text{SD}}, \theta_{\text{cut}}) \tag{A.11}$$

$$= \frac{e^{-I_{i\to i}(\theta_{\text{cut}}, R_0)}}{n!} \left(\prod_{i=1}^n \int_{\theta_{\text{cut}}}^{R_0} \frac{d\theta_i}{\theta_i} \int_0^{1/2} dz_i \, \frac{\alpha_s(z_i \theta_i p_T)}{\pi} P_{i\to i}(z_i) \Theta_{\text{SD}}(z_i, \theta_i) \right) \delta\left(n_{\text{SD}} - \sum_{i=1}^n z_i^\kappa \right).$$

In the perturbative regime, the behavior of $n_{\text{SD}}^{(\kappa)}$ is most clearly seen on a logarithmic scale. Two example LL distributions are displayed in figure 20 and compared to results from VINCIA. In these examples, soft drop parameters were chosen to demonstrate that the sharp features of the $n_{\text{SD}}^{(\kappa)}$ distributions are indeed captured by the LL evolution equations. These sharp features result from the edges of the $p_n^i(n_{\text{SD}}, \theta_{\text{cut}})$ distributions for different values of n. For example, with $\beta \leq 0$, the $p_n^i(n_{\text{SD}}, \theta_{\text{cut}})$ distribution only has support on the interval $[n\, z_{\text{cut}}^\kappa, \frac{n}{2^\kappa}]$.

A.3 Collinear-unsafe evolution

In the case of a collinear-unsafe weighted soft drop multiplicity with $\beta = 0$ and $\theta_{\text{cut}} = 0$, we can apply the methods of section 5. Specifically, after extracting the GFF at some RG scale μ, we can use eq. (5.7) with the particular choice $f(z) = z^\kappa$ to predict the upwards evolution. In figure 21, we compare the result of the RG evolution for $z_{\text{cut}} = 0.01$ and $\kappa = 1$ to VINCIA, finding overall good agreement. By eye, one can see that these $\kappa = 1$ distributions do not yield as good separation power as the $\kappa = 0$ distributions shown in figure 16, though the degree of RG evolution is similar for both the weighted and unweighted cases.

References

[1] ATLAS collaboration, *Jet energy scale measurements and their systematic uncertainties in proton-proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, [arXiv:1703.09665](https://arxiv.org/abs/1703.09665) [INSPIRE].

[2] CMS collaboration, *Jet energy scale and resolution in the CMS experiment in pp collisions at 8 TeV*, 2017 *JINST* **12** P02014 [arXiv:1607.03663] [INSPIRE].

[3] M. Dasgupta, A. Fregoso, S. Marzani and G.P. Salam, *Towards an understanding of jet substructure*, *JHEP* **09** (2013) 029 [arXiv:1307.0007] [INSPIRE].

[4] 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].

[5] A.J. Larkoski and J. Thaler, *Unsafe but calculable: ratios of angularities in perturbative QCD*, *JHEP* **09** (2013) 137 [arXiv:1307.1699] [INSPIRE].

[6] 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].

[7] 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].

[8] A.J. Larkoski, I. Moult and D. Neill, *Analytic boosted boson discrimination*, *JHEP* **05** (2016) 117 [arXiv:1507.03018] [INSPIRE].

[9] M. Procura, W.J. Waalewijn and L. Zeune, *Resummation of double-differential cross sections and fully-unintegrated parton distribution functions*, *JHEP* **02** (2015) 117 [arXiv:1410.6483] [INSPIRE].

[10] W.J. Waalewijn, *Calculating the charge of a jet*, *Phys. Rev.* **D 86** (2012) 094030 [arXiv:1209.3019] [INSPIRE].

[11] D. Adams et al., *Towards an understanding of the correlations in jet substructure*, *Eur. Phys. J.* **C 75** (2015) 409 [arXiv:1504.00679] [INSPIRE].

[12] A. Altheimer et al., *Boosted objects and jet substructure at the LHC*, report of *BOOST*2012, held at IFIC Valencia, Valencia Spain, 23–27 July 2012, *Eur. Phys. J.* **C 74** (2014) 2792 [arXiv:1311.2708] [INSPIRE].

[13] A. Altheimer et al., *Jet substructure at the Tevatron and LHC: new results, new tools, new benchmarks*, *J. Phys.* **G 39** (2012) 063001 [arXiv:1201.0008] [INSPIRE].

[14] A. Abdesselam et al., *Boosted objects: a probe of beyond the Standard Model physics*, *Eur. Phys. J.* **C 71** (2011) 1661 [arXiv:1012.5412] [INSPIRE].

[15] J. Gallicchio and M.D. Schwartz, *Quark and gluon tagging at the LHC*, *Phys. Rev. Lett.* **107** (2011) 172001 [arXiv:1106.3076] [INSPIRE].

[16] J. Gallicchio and M.D. Schwartz, *Quark and gluon jet substructure*, *JHEP* **04** (2013) 090 [arXiv:1211.7038] [INSPIRE].

[17] A.J. Larkoski, G.P. Salam and J. Thaler, *Energy correlation functions for jet substructure*, *JHEP* **06** (2013) 108 [arXiv:1305.0007] [INSPIRE].

[18] A.J. Larkoski, J. Thaler and W.J. Waalewijn, *Gaining (mutual) information about quark/gluon discrimination*, *JHEP* **11** (2014) 129 [arXiv:1408.3122] [INSPIRE].

[19] B. Bhattacherjee, S. Mukhopadhyay, M.M. Nojiri, Y. Sakaki and B.R. Webber, *Associated jet and subjet rates in light-quark and gluon jet discrimination*, *JHEP* **04** (2015) 131 [arXiv:1501.04794] [INSPIRE].

[20] J.R. Andersen et al., *Les Houches* 2015*: physics at TeV colliders Standard Model working group report*, in 9th *Les Houches Workshop on Physics at TeV Colliders (PhysTeV* 2015*)*, Les Houches France, 1–19 June 2015 [arXiv:1605.04692] [INSPIRE].

[21] 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].

[22] J. Davighi and P. Harris, *Fractal based observables to probe jet substructure of quarks and gluons*, arXiv:1703.00914 [INSPIRE].

[23] P. Gras et al., *Systematics of quark/gluon tagging*, *JHEP* **07** (2017) 091 [arXiv:1704.03878] [INSPIRE].

[24] CMS collaboration, *Performance of quark/gluon discrimination in 8 TeV pp data*, CMS-PAS-JME-13-002, CERN, Geneva Switzerland, (2013).

[25] D. Ferreira de Lima, P. Petrov, D. Soper and M. Spannowsky, *Quark-gluon tagging with shower deconstruction: unearthing dark matter and Higgs couplings*, *Phys. Rev.* **D 95** (2017) 034001 [arXiv:1607.06031] [INSPIRE].

[26] B. Bhattacherjee, S. Mukhopadhyay, M.M. Nojiri, Y. Sakaki and B.R. Webber, *Quark-gluon discrimination in the search for gluino pair production at the LHC*, *JHEP* **01** (2017) 044 [arXiv:1609.08781] [INSPIRE].

[27] ATLAS collaboration, *Light-quark and gluon jet discrimination in pp collisions at $\sqrt{s} = 7$ TeV with the ATLAS detector*, *Eur. Phys. J.* **C 74** (2014) 3023 [arXiv:1405.6583] [INSPIRE].

[28] ATLAS collaboration, *Discrimination of light quark and gluon jets in pp collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector*, ATLAS-CONF-2016-034, CERN, Geneva Switzerland, (2016).

[29] ATLAS collaboration, *Measurement of the charged-particle multiplicity inside jets from $\sqrt{s} = 8$ TeV pp collisions with the ATLAS detector*, *Eur. Phys. J.* **C 76** (2016) 322 [arXiv:1602.00988] [INSPIRE].

[30] P. Bolzoni, B.A. Kniehl and A.V. Kotikov, *Gluon and quark jet multiplicities at* $N^3LO+NNLL$, *Phys. Rev. Lett.* **109** (2012) 242002 [arXiv:1209.5914] [INSPIRE].

[31] P. Bolzoni, B.A. Kniehl and A.V. Kotikov, *Average gluon and quark jet multiplicities at higher orders*, *Nucl. Phys.* **B 875** (2013) 18 [arXiv:1305.6017] [INSPIRE].

[32] R.K. Ellis, W.J. Stirling and B.R. Webber, *QCD and collider physics*, *Camb. Monogr. Part. Phys. Nucl. Phys. Cosmol.* **8** (1996) 1 [INSPIRE].

[33] A. Banfi, G.P. Salam and G. Zanderighi, *Principles of general final-state resummation and automated implementation*, *JHEP* **03** (2005) 073 [hep-ph/0407286] [INSPIRE].

[34] S.D. Ellis, C.K. Vermilion, J.R. Walsh, A. Hornig and C. Lee, *Jet shapes and jet algorithms in SCET*, *JHEP* **11** (2010) 101 [arXiv:1001.0014] [INSPIRE].

[35] A.J. Larkoski, D. Neill and J. Thaler, *Jet shapes with the broadening axis*, *JHEP* **04** (2014) 017 [arXiv:1401.2158] [INSPIRE].

[36] E. Gerwick, T. Plehn, S. Schumann and P. Schichtel, *Scaling patterns for QCD jets*, *JHEP* **10** (2012) 162 [arXiv:1208.3676] [INSPIRE].

[37] Z. Bern et al., *Next-to-leading order* $W + 5$-*jet production at the LHC*, *Phys. Rev.* **D 88** (2013) 014025 [arXiv:1304.1253] [INSPIRE].

[38] S. Badger, B. Biedermann, P. Uwer and V. Yundin, *Next-to-leading order QCD corrections to five jet production at the LHC*, *Phys. Rev.* **D 89** (2014) 034019 [arXiv:1309.6585] [INSPIRE].

[39] ATLAS collaboration, *Measurements of the* W *production cross sections in association with jets with the ATLAS detector*, *Eur. Phys. J.* **C 75** (2015) 82 [arXiv:1409.8639] [INSPIRE].

[40] A.J. Larkoski, S. Marzani, G. Soyez and J. Thaler, *Soft drop*, *JHEP* **05** (2014) 146 [arXiv:1402.2657] [INSPIRE].

[41] Y.L. Dokshitzer, G.D. Leder, S. Moretti and B.R. Webber, *Better jet clustering algorithms*, *JHEP* **08** (1997) 001 [hep-ph/9707323] [INSPIRE].

[42] 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), pg. 270 [hep-ph/9907280] [INSPIRE].

[43] 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].

[44] 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].

[45] B.T. Elder, M. Procura, J. Thaler, W.J. Waalewijn and K. Zhou, *Generalized fragmentation functions for fractal jet observables*, *JHEP* **06** (2017) 085 [arXiv:1704.05456] [INSPIRE].

[46] W.T. Giele, D.A. Kosower and P.Z. Skands, *A simple shower and matching algorithm*, *Phys. Rev.* **D 78** (2008) 014026 [arXiv:0707.3652] [INSPIRE].

[47] W.T. Giele, D.A. Kosower and P.Z. Skands, *Higher-order corrections to timelike jets*, *Phys. Rev.* **D 84** (2011) 054003 [arXiv:1102.2126] [INSPIRE].

[48] T. Sjöstrand, S. Mrenna and P.Z. Skands, *A brief introduction to PYTHIA 8.1*, *Comput. Phys. Commun.* **178** (2008) 852 [arXiv:0710.3820] [INSPIRE].

[49] M. Cacciari, G.P. Salam and G. Soyez, *The anti-k_t jet clustering algorithm*, *JHEP* **04** (2008) 063 [arXiv:0802.1189] [ɪɴSPIRE].

[50] M. Cacciari, G.P. Salam and G. Soyez, *FastJet user manual*, *Eur. Phys. J.* **C 72** (2012) 1896 [arXiv:1111.6097] [ɪɴSPIRE].

[51] *Fastjet contrib webpage*, http://fastjet.hepforge.org/contrib/.

[52] H.-M. Chang, M. Procura, J. Thaler and W.J. Waalewijn, *Calculating track-based observables for the LHC*, *Phys. Rev. Lett.* **111** (2013) 102002 [arXiv:1303.6637] [ɪɴSPIRE].

[53] H.-M. Chang, M. Procura, J. Thaler and W.J. Waalewijn, *Calculating track thrust with track functions*, *Phys. Rev.* **D 88** (2013) 034030 [arXiv:1306.6630] [ɪɴSPIRE].

[54] L. Clavelli, *Jet invariant mass in quantum chromodynamics*, *Phys. Lett.* **B 85** (1979) 111 [ɪɴSPIRE].

[55] S. Catani, G. Turnock and B.R. Webber, *Heavy jet mass distribution in e^+e^- annihilation*, *Phys. Lett.* **B 272** (1991) 368 [ɪɴSPIRE].

[56] S. Catani, L. Trentadue, G. Turnock and B.R. Webber, *Resummation of large logarithms in e^+e^- event shape distributions*, *Nucl. Phys.* **B 407** (1993) 3 [ɪɴSPIRE].

[57] J. Gallicchio, J. Huth, M. Kagan, M.D. Schwartz, K. Black and B. Tweedie, *Multivariate discrimination and the Higgs + W/Z search*, *JHEP* **04** (2011) 069 [arXiv:1010.3698] [ɪɴSPIRE].

[58] C.F. Berger, T. Kucs and G.F. Sterman, *Event shape/energy flow correlations*, *Phys. Rev.* **D 68** (2003) 014012 [hep-ph/0303051] [ɪɴSPIRE].

[59] L.G. Almeida, S.J. Lee, G. Perez, G.F. Sterman, I. Sung and J. Virzi, *Substructure of high-p_T jets at the LHC*, *Phys. Rev.* **D 79** (2009) 074017 [arXiv:0807.0234] [ɪɴSPIRE].

[60] B. Andersson, G. Gustafson, L. Lönnblad and U. Pettersson, *Coherence effects in deep inelastic scattering*, *Z. Phys.* **C 43** (1989) 625 [ɪɴSPIRE].

[61] J.G.M. Gatheral, *Exponentiation of eikonal cross-sections in non-Abelian gauge theories*, *Phys. Lett.* **B 133** (1983) 90 [ɪɴSPIRE].

[62] T. Becher and M. Neubert, *On the structure of infrared singularities of gauge-theory amplitudes*, *JHEP* **06** (2009) 081 [*Erratum ibid.* **11** (2013) 024] [arXiv:0903.1126] [ɪɴSPIRE].

[63] S. Moch, J.A.M. Vermaseren and A. Vogt, *The three loop splitting functions in QCD: the nonsinglet case*, *Nucl. Phys.* **B 688** (2004) 101 [hep-ph/0403192] [ɪɴSPIRE].

[64] J. Frenkel and J.C. Taylor, *Non-Abelian eikonal exponentiation*, *Nucl. Phys.* **B 246** (1984) 231 [ɪɴSPIRE].

[65] G. Altarelli and G. Parisi, *Asymptotic freedom in parton language*, *Nucl. Phys.* **B 126** (1977) 298 [ɪɴSPIRE].

[66] S. Catani, B.R. Webber and G. Marchesini, *QCD coherent branching and semiinclusive processes at large x*, *Nucl. Phys.* **B 349** (1991) 635 [ɪɴSPIRE].

[67] M. Dasgupta and G.P. Salam, *Resummation of nonglobal QCD observables*, *Phys. Lett.* **B 512** (2001) 323 [hep-ph/0104277] [ɪɴSPIRE].

[68] T. Sjöstrand, S. Mrenna and P.Z. Skands, *PYTHIA 6.4 physics and manual*, *JHEP* **05** (2006) 026 [hep-ph/0603175] [ɪɴSPIRE].

[69] T. Sjöstrand et al., *An introduction to PYTHIA* 8.2, *Comput. Phys. Commun.* **191** (2015) 159 [arXiv:1410.3012] [InSPIRE].

[70] M. Bahr et al., *HERWIG++ physics and manual*, *Eur. Phys. J.* **C 58** (2008) 639 [arXiv:0803.0883] [InSPIRE].

[71] J. Bellm et al., *HERWIG 7.0/HERWIG++ 3.0 release note*, *Eur. Phys. J.* **C 76** (2016) 196 [arXiv:1512.01178] [InSPIRE].

[72] T. Gleisberg et al., *Event generation with SHERPA* 1.1, *JHEP* **02** (2009) 007 [arXiv:0811.4622] [InSPIRE].

[73] D. Krohn, M.D. Schwartz, T. Lin and W.J. Waalewijn, *Jet charge at the LHC*, *Phys. Rev. Lett.* **110** (2013) 212001 [arXiv:1209.2421] [InSPIRE].

[74] V.N. Gribov and L.N. Lipatov, *Deep inelastic ep scattering in perturbation theory*, *Sov. J. Nucl. Phys.* **15** (1972) 438 [*Yad. Fiz.* **15** (1972) 781] [InSPIRE].

[75] Y.L. Dokshitzer, *Calculation of the structure functions for deep inelastic scattering and e^+e^- annihilation by perturbation theory in quantum chromodynamics*, *Sov. Phys. JETP* **46** (1977) 641 [*Zh. Eksp. Teor. Fiz.* **73** (1977) 1216] [InSPIRE].

[76] J.-Y. Chiu, A. Jain, D. Neill and I.Z. Rothstein, *A formalism for the systematic treatment of rapidity logarithms in quantum field theory*, *JHEP* **05** (2012) 084 [arXiv:1202.0814] [InSPIRE].

[77] D. Bertolini, J. Thaler and J.R. Walsh, *The first calculation of fractional jets*, *JHEP* **05** (2015) 008 [arXiv:1501.01965] [InSPIRE].

[78] A.J. Larkoski, S. Marzani and J. Thaler, *Sudakov safety in perturbative QCD*, *Phys. Rev.* **D 91** (2015) 111501 [arXiv:1502.01719] [InSPIRE].

[79] H.P. Nilles and K.H. Streng, *Quark-gluon separation in three jet events*, *Phys. Rev.* **D 23** (1981) 1944 [InSPIRE].

[80] L.M. Jones, *Tests for determining the parton ancestor of a hadron jet*, *Phys. Rev.* **D 39** (1989) 2550 [InSPIRE].

[81] Z. Fodor, *How to see the differences between quark and gluon jets*, *Phys. Rev.* **D 41** (1990) 1726 [InSPIRE].

[82] L. Jones, *Towards a systematic jet classification*, *Phys. Rev.* **D 42** (1990) 811 [InSPIRE].

[83] L. Lönnblad, C. Peterson and T. Rognvaldsson, *Finding gluon jets with a neural trigger*, *Phys. Rev. Lett.* **65** (1990) 1321 [InSPIRE].

[84] J. Pumplin, *How to tell quark jets from gluon jets*, *Phys. Rev.* **D 44** (1991) 2025 [InSPIRE].

[85] CMS collaboration, *Search for a Higgs boson in the decay channel $H \to ZZ^{(*)} \to q\bar{q}\ell^-\ell^+$ in pp collisions at $\sqrt{s} = 7$ TeV*, *JHEP* **04** (2012) 036 [arXiv:1202.1416] [InSPIRE].

[86] J. Cogan, M. Kagan, E. Strauss and A. Schwarztman, *Jet-images: computer vision inspired techniques for jet tagging*, *JHEP* **02** (2015) 118 [arXiv:1407.5675] [InSPIRE].

[87] L.G. Almeida, M. Backović, M. Cliche, S.J. Lee and M. Perelstein, *Playing tag with ANN: boosted top identification with pattern recognition*, *JHEP* **07** (2015) 086 [arXiv:1501.05968] [InSPIRE].

[88] 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].

[89] P. Baldi, K. Bauer, C. Eng, P. Sadowski and D. Whiteson, *Jet substructure classification in high-energy physics with deep neural networks*, *Phys. Rev.* **D 93** (2016) 094034 [arXiv:1603.09349] [INSPIRE].

[90] J.S. Conway, R. Bhaskar, R.D. Erbacher and J. Pilot, *Identification of high-momentum top quarks, Higgs bosons and W and Z bosons using boosted event shapes*, *Phys. Rev.* **D 94** (2016) 094027 [arXiv:1606.06859] [INSPIRE].

[91] D. Guest, J. Collado, P. Baldi, S.-C. Hsu, G. Urban and D. Whiteson, *Jet flavor classification in high-energy physics with deep neural networks*, *Phys. Rev.* **D 94** (2016) 112002 [arXiv:1607.08633] [INSPIRE].

[92] J. Barnard, E.N. Dawe, M.J. Dolan and N. Rajcic, *Parton shower uncertainties in jet substructure analyses with deep neural networks*, *Phys. Rev.* **D 95** (2017) 014018 [arXiv:1609.00607] [INSPIRE].

[93] L. de Oliveira, M. Paganini and B. Nachman, *Learning particle physics by example: location-aware generative adversarial networks for physics synthesis*, arXiv:1701.05927 [INSPIRE].

[94] 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].

[95] G. Louppe, K. Cho, C. Becot and K. Cranmer, *QCD-aware recursive neural networks for jet physics*, arXiv:1702.00748 [INSPIRE].