Pileup and Infrared Radiation Annihilation (PIRANHA): a paradigm for continuous jet grooming

Samuel Alipour-fard,[a,b] **Patrick T. Komiske,**[a] **Eric M. Metodiev**[a] **and Jesse Thaler**[a,b]

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

[b]*The NSF AI Institute for Artificial Intelligence and Fundamental Interactions*

E-mail: samuelaf@mit.edu, pkomiske@mit.edu, metodiev@mit.edu, jthaler@mit.edu

ABSTRACT: Jet grooming is an important strategy for analyzing relativistic particle collisions in the presence of contaminating radiation. Most jet grooming techniques introduce hard cutoffs to remove soft radiation, leading to discontinuous behavior and associated experimental and theoretical challenges. In this paper, we introduce Pileup and Infrared Radiation Annihilation (PIRANHA), a paradigm for continuous jet grooming that overcomes the discontinuity and infrared sensitivity of hard-cutoff grooming procedures. We motivate PIRANHA from the perspective of optimal transport and the Energy Mover's Distance and review Apollonius Subtraction and Iterated Voronoi Subtraction as examples of PIRANHA-style grooming. We then introduce a new tree-based implementation of PIRANHA, Recursive Subtraction, with reduced computational costs. Finally, we demonstrate the performance of Recursive Subtraction in mitigating sensitivity to soft distortions from hadronization and detector effects, and additive contamination from pileup and the underlying event.

Contents

1 Introduction

Jets of hadrons produced in high-energy particle collisions are an important guide in our understanding of the subatomic universe and quantum chromodynamics (QCD), our fundamental theory of the strong nuclear force. Quantifying the internal structure of hadronic jets plays an important role in translating between experimental results and theoretical predictions. The toolkit of jet substructure, including techniques from jet grooming [1–7] and pileup mitigation [8–24], helps characterize the radiation patterns produced by jets in a wide variety of physical processes at colliders. Jet substructure can distinguish jets produced by light quarks, gluons, and the decays of electroweak bosons and top quarks [25–53], provide accurate predictions of the behavior of boosted objects in many contexts [2, 14, 54–93], and even reveal non-perturbative features of the infrared structure of QCD [94–98].

Jet grooming facilitates our understanding of jet substructure by removing low-energy pollution from jets. The major goal of jet grooming is to reduce the sensitivity of jet observables to two conceptually distinct categories of pollution:

- **Soft Distortions** are low-energy effects that distort the radiation coming directly from hard jet production, such as hadronization and detector effects;

- **Additive Contamination** denotes a backdrop of low-energy radiation that pollutes a jet but is not directly connected to jet production, including the pileup (PU) of radiation from overlapping particle collisions and an underlying event (UE) of secondary parton interactions.

The low-energy phenomena of QCD that produce both soft contamination and additive contamination are notoriously difficult to model, and jet grooming has become an important toolkit used to empower both experimental and theoretical analyses of jet substructure [1–7, 14, 25–93, 99–101].

A complication of many popular jet grooming procedures is that they are *discontinuous* in event space: they may map similar ungroomed jets, which differ only by small changes in their distributions of energy, into extensively distinct groomed jets [3, 4]. Discontinuous behavior in jet grooming algorithms leads in turn to undesirably sensitive responses to low-energy pollution. Experimentally, discontinuities in grooming lead to large responses of groomed substructure to the interactions of jets with experimental detectors, and resulting difficulties in unfolding [99–101]. Theoretically, discontinuities in grooming lead to uncertainties in fixed-order calculations [3] and complications in describing the effects of hadronization on groomed jet substructure [96].

In this paper, we introduce **P**ileup and **I**nfrared **R**adiation **A**nnihilation (PIRANHA), a paradigm for continuous jet grooming, using recent geometric perspectives in collider

physics. Piranha applies intuition from optimal transport theory and the Energy Mover's Distance (EMD) between collider events [95, 102] to overcome the discontinuous behavior of existing jet grooming algorithms. We introduce the new terminology of *soft* and *angular discontinuities* to categorize discontinuous behavior under infinitesimal changes to the energies and angles of jet constituents, respectively. We explore the improved continuity of Piranha conceptually through several toy examples, and more practically by studying the responses of Piranha grooming to several sources of soft distortions and additive contamination.

At present, Piranha has three concrete implementations:

- **Apollonius Subtraction** (P-AS), introduced in ref. [102], is a continuous, conceptually simple application of the Piranha paradigm with a computationally expensive implementation;

- **Iterated Voronoi Subtraction** (P-IVS), also introduced in ref. [102], is a continuous, computationally efficient modification of P-AS with a simple geometric implementation;

- **Recursive Subtraction** (P-RS) is a family of nearly-continuous, computationally efficient, tree-based algorithms motivated by P-AS and P-IVS. We develop the concrete implementation of **Recursive Subtraction with a Fraction f** (P-RSF$_f$). We discuss how P-RSF$_f$ is discontinuous in suppressed regions of parameter space, and focus on P-RSF$_{f=1/2}$, which is the only fully soft-continuous version of P-RSF$_f$.

Table 1 expresses the infrared/collinear safety (using the definition of ref. [102]) and soft/angular continuity properties of the traditional and Piranha groomers we discuss in this paper. The soft continuity properties expressed in table 1 are simple to understand by comparing the hard-cutoff framework of popular groomers to the subtractive framework of Piranha. The angular discontinuities expressed in table 1 are more subtle and are inherited from angular-ordered jet clustering. We note that the Recursive Subtraction algorithms we introduce are discontinuous only in a highly suppressed region of phase space, and we still find it suitable to call P-RS a family of Piranha groomers.

Popular jet grooming algorithms are soft discontinuous because they introduce hard cutoffs on the energy of jet constituents, removing particles that are too low-energy to be classified as important probes of hard physics; we call such groomers *hard-cutoff groomers*. The modified Mass Drop Tagger (mMDT) [4], for example, requires that certain sub-jets within a jet carry a sufficient energy fraction, $z > z_{\text{cut}}$, to survive the grooming procedure. If a sub-jet has an energy fraction near z_{cut}, low-energy pollution may push its energy to the other side of the cutoff, leading to soft discontinuous behavior and associated theoretical and experimental complications. Figure 1 demonstrates this soft discontinuous behavior, the corresponding soft continuity of P-RSF$_{f=1/2}$, by comparing their action on two sets of similar events: two toy jet events $\mathcal{E}^{\pm}$ consisting of two particles each, and a QCD jet produced in Pythia 8.244 along with the same QCD event with a small amount of gaussian noise applied to the p_T, rapidity, and azimuthal angle of each particle. We revisit the events $\mathcal{E}^{\pm}$ as tools to characterize soft discontinuity more precisely in section 3.2. The

Groomer	Acronym	Safety		Continuity	
		Infrared	Collinear	Soft	Angular
Apollonius Subtraction	P-AS	✓	✓	✓	✓
Iterated Voronoi Subtraction	P-IVS	✓	✓	✓	✓
Recursive Sub. w/Fraction $f=1/2$	P-RSF$_{f=1/2}$	✓	✓	✓	⚠
Recursive Sub. w/Fraction $f\neq1/2$	P-RSF$_{f\neq1/2}$	✓	✓	⚠	⚠
Soft Drop (SD), $\beta_{\rm SD}>0$	SD$_{\beta>0}$	✓	✓	✗	⚠
SD (grooming mode), $\beta_{\rm SD}\leq0$	SD$_{\beta\leq0}$	✓	✗	✗	⚠

Table 1. Safety and continuity properties of the Piranha and Soft Drop grooming algorithms studied in this paper. The mark ✓ indicates that a groomer satisfies the associated condition, while the mark ✗ indicates that it does not. The mark ⚠ indicates that a groomer does not satisfy the condition, but only in a boundary region of phase space that is exponentially suppressed at leading order in perturbative QCD. The mark ⚠ indicates that a groomer does not satisfy the condition, but only in a boundary region of phase space that is exponentially suppressed at leading order in perturbative QCD.[1]

mMDT grooming procedure maps $\mathcal{E}^+$ and $\mathcal{E}^-$ to very distinct events, and similarly maps the noiseless and noisy QCD jets to very distinct groomed jets. We chose an unusually high value of $z_{\rm cut}=0.25$ when grooming the QCD events to make the discontinuity of mMDT more apparent. We will also use Soft Drop [3], a generalization of mMDT reviewed in section 3.1 which reproduces mMDT when a parameter $\beta_{\rm SD}$ is set to zero, as another example of soft discontinuous hard-cutoff grooming.

Piranha borrows technology from optimal transport to produce soft continuous, subtractive jet grooming algorithms without hard cutoffs. Piranha groomers may be roughly imagined as optimally transporting hungry piranhas to eat up delicious low-energy pollution. Given an ungroomed event with energy density $\mathcal{E}_0$ in the pseudorapidity-azimuth angle plane and any ansatz $\rho\,\mathcal{C}$ for the energy density of low-energy pollution, one particularly succinct Piranha algorithm is

$$\mathcal{E}_g\left[\mathcal{E}_0\right]=\arg\min_{\mathcal{E}'}\ {\rm EMD}_{\beta,R}\left(\mathcal{E}_0,\mathcal{E}'+\rho\,\mathcal{C}\right),\tag{1.1}$$

where $\mathcal{E}_g$ is a groomed energy distribution, $\rho\,\mathcal{C}$ is a free input to the algorithm indicating a model of the energy flow of the contaminating radiation (see section 2.1), and explicit dependence on the EMD metric makes the continuity of $\mathcal{E}_g\left[\mathcal{E}_0\right]$ manifest.

P-AS is a direct, continuous application of eq. (1.1) with $\beta=1$, and for which $\rho\,\mathcal{C}$ is uniform in the pseudorapidity-azimuth plane to approximate common sources of additive contamination such as PU and UE [23, 103–112]. P-IVS approximates P-AS and preserves

---

[1]Soft Drop with $\beta_{\rm SD}\leq0$ is not collinear safe in grooming mode, used in this work, where jets consisting of a single particle are retained as groomed jets. Grooming mode is a natural choice in the context of pileup mitigation (see section 3.2). Note that Soft Drop with $\beta_{\rm SD}\leq0$ is collinear safe in tagging mode, where jets consisting of a single particle are groomed away completely.
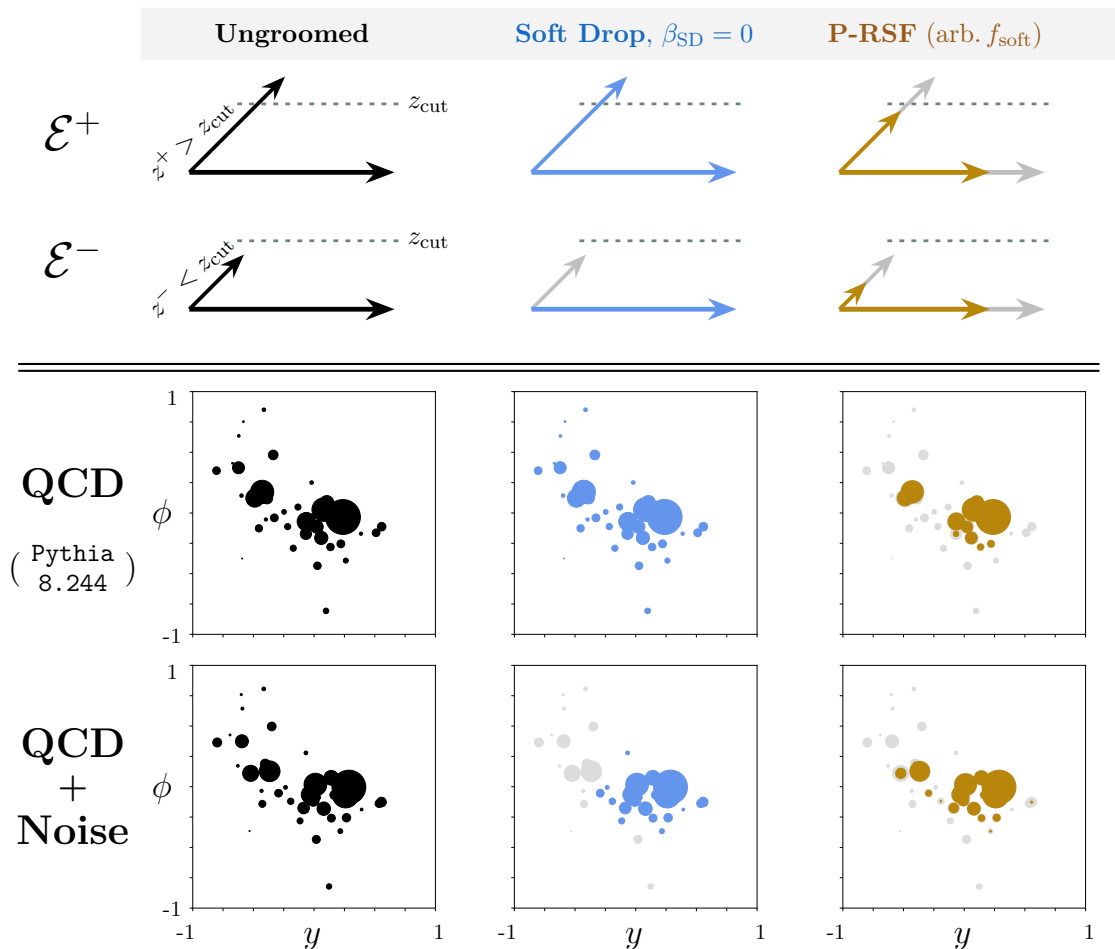
Figure 1. A visualization of how similar ungroomed events (first column, in black) may be mapped discontinuously to distinct groomed jets by the hard-cutoff mMDT algorithm (Soft Drop with $\beta_{\rm SD} = 0$, or $\mathrm{SD}_{\beta=0}$; middle column, in blue), while the continuous P-RSF$_{f=1/2}$ algorithm will map them to similar groomed results (final column, in orange). In the first two rows, we compare the action of mMDT and P-RSF$_{f=1/2}$ on simple toy events $\mathcal{E}^{\pm}$, with particles visualized as arrows whose size corresponds to their energy. In the final two rows, we compare the groomers acting on a more realistic QCD event produced with `Pythia 8.244`, and on the same QCD event with additional Gaussian noise, respectively.

its continuity. Finally, the Recursive Subtraction with a Fraction $1/2$ (P-RSF$_{f=1/2}$) algorithm we introduce in section 2.5 furnishes, to our knowledge, the first example of a soft continuous tree-based jet grooming algorithm; the soft discontinuities of other P-RS algorithms occur only in highly suppressed regions of phase space.

The rest of the paper proceeds as follows: in section 2, we review the Energy Mover's Distance (EMD) of ref. [95], propose the PIRANHA paradigm for continuous jet grooming, review the Apollonius Subtraction (P-AS) and Iterated Voronoi Subtraction (P-IVS) algorithms of ref. [102], and introduce Recursive Subtraction (P-RS) with a Fraction (P-RSF) as a new, tree-based implementation of PIRANHA. In section 3, we discuss the soft discontinuities present in hard-cutoff groomers and P-RSF, demonstrate the soft continuity of

P-RSF$_{f=1/2}$, and examine angular discontinuity in tree-based groomers. In the remainder of the paper, we demonstrate the soft insensitivity enjoyed by PIRANHA groomers over hard-cutoff groomers through several phenomenological studies. In section 4, we compare the responses of the hard-cutoff Soft Drop algorithm and P-RSF$_{f=1/2}$ to soft distortions from hadronization and to the removal of neutral particles, as a proxy for detector effects. In section 5, we compare the ability of Soft Drop and P-RSF$_{f=1/2}$ to effectively remove additive contamination from pileup and the underlying event. We give concluding thoughts and avenues for future exploration in section 6, delegating discussions of PIRANHA grooming in EMD mode, comparisons of other hard-cutoff and PIRANHA groomers, and a first look at the perturbative structure of PIRANHA grooming to appendices A, B and C, respectively. The PIRANHA algorithms used in our analysis are available on GitHub [113].

2 The PIRANHA paradigm for continuous grooming

This section introduces the PIRANHA paradigm for jet grooming, which uses optimal transport to continuously remove low-energy pollution from jets. We first review the Energy Mover's Distance (EMD) of ref. [95], which provides a framework to introduce the principles of PIRANHA groomers and to enforce their continuity. We then discuss three implementations of PIRANHA: Apollonius Subtraction (P-AS) and Iterated Voronoi Subtraction (P-IVS) — both introduced in ref. [102] — and the new tree-based Recursive Subtraction (P-RS) family of PIRANHA algorithms. We describe a variant of PIRANHA — grooming in EMD mode — in appendix A. We end this section with section 2.7, which may be skipped on a first reading, where we present definitions of stronger notions of continuity that could potentially pave the way for more constrained and well-behaved strategies for continuous grooming.

2.1 Review of the Energy Mover's Distance

The Energy Mover's Distance (EMD) of ref. [95] provides a quantitative measure of the similarity between two jets. We introduce the EMD here both to facilitate our definition of continuity in jet grooming and as a useful tool for quantifying the responses of jet groomers to low-energy pollution. Readers familiar with the EMD should feel free to skip to section 2.2.

The EMD is a metric on the space of collider events. Here and in the remainder of the paper, we borrow the terminology of ref. [102] and use the terms "event" and "collider event" to refer to the *energy flow* of the event, or the angular distribution of energy of its outgoing radiation.[2] The EMD may be thought of as the amount of "work" required to rearrange one event into another. For events that consist of a finite number of outgoing

---

[2]For example, the energy flow of an event $\mathcal{E}$ with M outgoing particles is the angular distribution

$$\mathcal{E}(\hat{n}) = \sum_{i=1}^{M} E_i \, \delta(\hat{n} - \hat{n}_i),$$

where $E_i > 0$ denotes the energy carried by particle i, and $\hat{n}_i$ denotes its outgoing angular direction.

particles, the EMD is defined as the solution to the optimal transport problem

$$\mathrm{EMD}_{\beta,R}\left(\mathcal{E},\mathcal{E}'\right) = \min_{f_{ij}>0} \sum_{i=1}^{M} \sum_{j=1}^{M'} f_{ij} \left(\frac{\theta_{ij}}{R}\right)^{\beta} + \left| \sum_{i=1}^{M} E_i - \sum_{j=1}^{M'} E_j' \right|, \tag{2.1}$$

$$\sum_{i=1}^{M} f_{ij} \le E_j', \quad \sum_{j=1}^{M'} f_{ij} \le E_i, \quad \sum_{i=1}^{M} \sum_{j=1}^{M'} f_{ij} = \min\left(\sum_{i=1}^{M} E_i, \sum_{j=1}^{M'} E_j'\right), \tag{2.2}$$

where $i \in \{1, \dots, M\}$ indicates a final-state particle of $\mathcal{E}$ with energy E_i, $j \in \{1, \dots, M'\}$ indicates a final-state particle of $\mathcal{E}'$ with energy E_j', and θ_{ij} denotes an angular distance between particles i and j. $\beta > 0$ and $R > 0$ are free parameters that control the behavior and relative importance of the first term on the right-hand side of eq. (2.1).[3]

In this work, we adopt the hadronic angular measure of ref. [102], given by

$$\theta_{ij} = \sqrt{-(n_i - n_j)^2}, \qquad n_i^{\mu} = \frac{p_i^{\mu}}{E_{Ti}} = (\cosh(y_i), \mathbf{v}_{Ti}, \sinh(y_i))^{\mu}. \tag{2.3}$$

Here, n_i^{μ} is a vector describing the motion of particle i, parameterized by its rapidity y_i and its velocity $\mathbf{v}_{Ti}$ in the directions transverse to the beampipe. This angular metric reproduces the rapidity-azimuth distance between particles i and j in the small rapidity-azimuth distance limit.[4]

The EMD between a groomed jet in the presence and absence of low-energy pollution also provides an observable-independent tool for quantifying the sensitivity of jet grooming to pollution from soft distortions and additive contamination. Since the EMD furnishes a metric on the space of events, two events are separated by zero EMD iff they have identical energy flow, and two jets separated by zero EMD will yield the same value for every infrared/collinear (IRC) safe observable (see Lemma 1 of ref. [102]). Further, the EMD between two events also *directly* bounds the differences in a large class of IRC-safe observable quantities (Lipschitz continuous observables) between the events [95]. For example, the EMD between a parton-level jet and the same jet after hadronization (the parton-hadron EMD) directly bounds the difference between a large class of parton- and hadron-level IRC-safe observables. The parton-hadron EMD is therefore a powerful, observable independent tool for characterizing the response of a jet to soft distortions from hadronization. We use the EMD in our analysis of grooming sensitivity to several sources of low-energy pollution in sections 4 and 5.

---

[3]Strictly speaking, the EMD defines a metric only if $2R$ is greater than the maximum possible angular distance θ_{ij}. Furthermore, for $\beta > 1$, one must raise the first term of eq. (2.1) to the $1/\beta$ power. We use the EMD without a subscript, $\mathrm{EMD}\left(\mathcal{E}, \mathcal{E}'\right)$ to denote the EMD with $\beta = 1$ and with large enough R that it furnishes a metric.

[4]Another popular choice for the form of the n_i^{μ} is captured by

$$n_{i,\ e^+e^-}^{\mu} = \frac{p_i^{\mu}}{E_i} = (1, \mathbf{v}_i)^{\mu}, \tag{2.4}$$

which is a natural choice in the study of electron-positron collisions [102]. This choice reproduces the real-space opening angle between particles i and j in the small angle limit.

2.2 PIRANHA and continuity

PIRANHA grooming may be intuitively imagined as the optimal transport of a group of hungry piranhas towards the low-energy pollution within a jet. In this analogy, henceforth the piranha analogy, a PIRANHA groomer populates the η-ϕ plane with a school of piranhas, distributed according to a particular model of the polluting energy density. When they are given an event to groom, the piranhas discuss which piranhas will get to eat which jet constituent; in the language of optimal transport, we say that the piranhas decide on a set of transport plans. The piranhas then implement the plan and feed on the constituents of the jet, or subtract the p_T of each jet constituent.

PIRANHA grooming uses optimal transport to design *continuous* grooming procedures that have reduced sensitivity to low-energy pollution. In particular, continuity requires that minuscule changes to the jet do not dramatically change the resulting groomed jet. In the piranha analogy, continuity emerges when the piranhas' transport plans do not respond dramatically to small changes in the constituents of an ungroomed jet. Since both soft distortions and fluctuations in additive contamination may lead to small changes in an event, continuity helps ensure that neither dramatically changes the information carried by groomed jets.

We may use the EMD to define continuity more precisely for maps on energy flows, in the spirit of Definition 1 of ref. [102]:

Definition 1 *A map M from energy flows to energy flows is* continuous *at an event $\mathcal{E}$ if, for any $\varepsilon > 0$, there exists a $\delta > 0$ such that for all $\mathcal{E}'$,*

$$\mathrm{EMD}(\mathcal{E}, \mathcal{E}') < \delta \quad \implies \quad \mathrm{EMD}(M(\mathcal{E}), M(\mathcal{E}')) < \varepsilon.$$

*We say that M is **continuous** if it is continuous at all events $\mathcal{E}$.*

Definition 1 encodes the requirement that a continuous grooming algorithm must map jets that are infinitesimally nearby in event space to groomed results that are also infinitesimally close.[5]

The EMD itself also furnishes an example of an optimal transport problem, allowing us to immediately introduce a simple PIRANHA groomer. To begin, let us model the radiation that we would like to remove from our jet as a distribution of energy $\rho\mathcal{C}$ in a portion of the rapidity-azimuth (η-ϕ) plane with area A. We use ρ to indicate the mean energy density and $\mathcal{C}$ to indicate a distribution modeling the shape of the contaminating radiation, normalized such that

$$\rho \int \mathrm{d}y \, \mathrm{d}\phi \, \mathcal{C}(y, \phi) = \rho A = E \tag{2.5}$$

gives the total energy of the additive contamination we would like to remove. Again echoing ref. [102], we may phrase the subtraction of the energy distribution $\rho\mathcal{C}$ from an event $\mathcal{E}$ as

---

[5]Equivalently, we may say that a grooming procedure is continuous if it maps nearby jets to nearby groomed jets. An event $\mathcal{E}$ is **near** an open ball $\mathcal{B}$ if any neighborhood of $\mathcal{E}$ contains events in $\mathcal{B}$. On the real line, for example, the point $E = 1$ is near the open ball $B = (-1, 1)$. A grooming procedure G is then **continuous** if for any event $\mathcal{E}$ and open ball $\mathcal{B}$, $\mathcal{E}$ *is near* $\mathcal{B} \implies G(\mathcal{E})$ *is near* $G(\mathcal{B})$.

the solution to the optimal transport problem

$$\mathcal{E}_g[\mathcal{E}_0] = \arg\min_{\mathcal{E}'} \text{EMD}_{\beta,R}\left(\mathcal{E}_0, \mathcal{E}' + \rho\mathcal{C}\right). \tag{2.6}$$

It was shown in Lemma 3 of ref. [102] that eq. (2.6) defines a subtractive algorithm that merely decreases the energy of jet constituents, and does not add new particles to an event. In the piranha analogy, $\rho\mathcal{C}$ roughly describes the distribution of our piranhas, with a "piranha energy density" ρ, while the minimization problem of eq. (2.6) provides them with the transport plans that guide them to subtract from the p_T of jet constituents. Furthermore, the form of $\mathcal{E}_g[\mathcal{E}_0]$ in eq. (2.6) is manifestly continuous due to its dependence on the EMD metric.[6]

For the more complicated PIRANHA groomers we introduce in this paper, it will be helpful to define two additional properties that are required by full continuity:

- **Soft Continuity** is the requirement that a groomer is invariant under infinitesimally soft perturbations of the energies of jet constituents;

- **Angular Continuity** is the requirement that a groomer is invariant under infinitesimally small angular changes in the directions of jet constituents.

We will see that hard-cutoff groomers and even many recursive subtraction algorithms suffer from soft discontinuity in section 3.1 and 3.4 respectively. In section 3.5, we will explore how tree-based grooming algorithms — hard-cutoff and PIRANHA alike — may inherit angularly discontinuous behavior from angular-ordered jet clustering.[7]

2.3 Apollonius Subtraction (P-AS)

Apollonius Subtraction (P-AS) is a direct application of eq. (2.6), in which we take contaminating radiation to be uniformly distributed in a region of the η-ϕ plane bounded by a maximum pseudorapidity, $|\eta| < \eta_{\max}$. We denote the uniform event with energy density ρ by $\rho\mathcal{U}$. This has proven to be an effective model for pileup, the underlying event, and initial state radiation [23, 103–112], and motivates our approach to jet grooming for the remainder of the paper.

More precisely, using $\beta = 1$, $R = 1$, and replacing $\rho\mathcal{C}$ with $\rho\mathcal{U}$ in eq. (2.6) yields an optimal transport problem known as an *Apollonius problem*. The Apollonius problem is well studied in the optimal transport literature and can be solved with an *Apollonius diagram*, or additively weighted Voronoi diagram, which assigns an *Apollonius region* to

---

[6]While $\mathcal{E}_g[\mathcal{E}_0]$ in eq. (2.6) is continuous, manifest continuity due to dependence on the EMD should be treated skeptically. For example, the "grooming procedure" associated with eq. (2.6) may not be continuous if the arg min operation is restricted to a non-convex sub-space of the full space of events. Analogously, $f(x) = \arg\min_{y} ||x, y+1|| = x-1$ is continuous as a map from $\mathbb{R}$ to $\mathbb{R}$ with the standard topology. However, $g(x) = \arg\min_{y \in \{0,1\}} ||x, y+1||$, which differs only by the allowed values of y in the arg min, is not.

[7]Collinear safety, or invariance under exact collinear splittings which replace a final-state particle of momentum p^μ by two particles with momenta λp^μ and $(1-\lambda)p^\mu$, is a weaker condition. Exact collinear splittings do not change the energy distribution associated with an event; therefore, any well-defined map on the space of particle events must be collinear safe [102].

Cartoon of a PIRANHA groomer
acting on the space of energy flows

Visualization of P-AS acting on
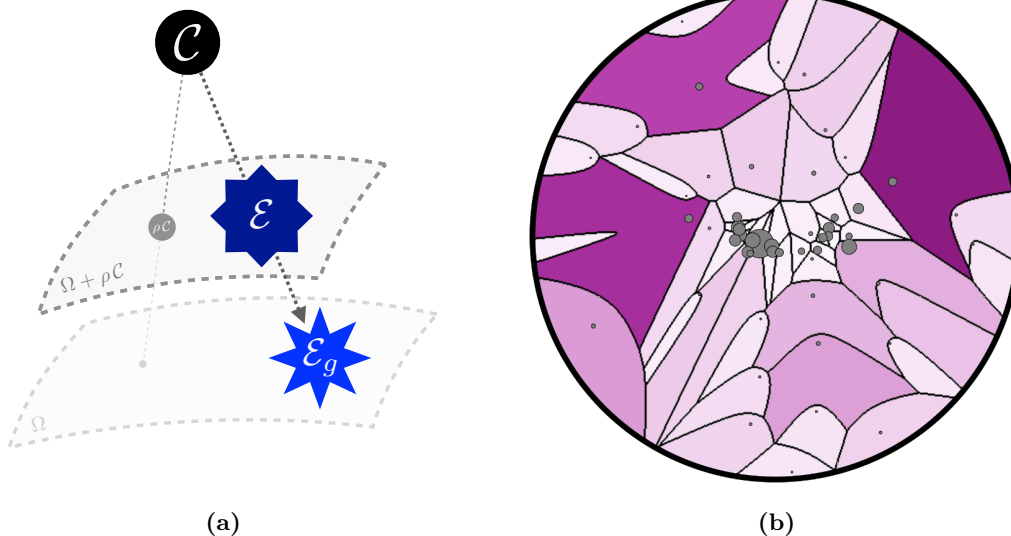an event generated in `Pythia 8.244`

(a)

(b)

Figure 2. (a) A cartoon of a generic implementation of the PIRANHA paradigm described by eq. (2.6), based on a figure from ref. [102]. $\mathcal{E} \in \Omega$ indicates a particular event $\mathcal{E}$ in the space Ω of events to be groomed, $\mathcal{C}$ indicates a model energy distribution for contaminating radiation, and $\rho \in \mathbb{R}$ parameterizes the strength of the grooming. Eq. (2.6) uses these ingredients to produce the groomed event $\mathcal{E}_g$. (b) An Apollonius diagram in the rapidity-azimuth plane, also from ref. [102], used by the Apollonius Subtraction (P-AS) algorithm. P-AS, an implementation of the PIRANHA paradigm described in section 2.3, models contaminating radiation as uniform in the rapidity-azimuth plane $\mathcal{C} = \mathcal{U}$. The color intensity of each Apollonius region is proportional to the amount by which the corresponding particle i is groomed, and thus to the area $A_i^{\text{Apoll.}}$ of the region.

each particle in the event. The precise structure of the Apollonius diagram corresponding to an event is described in a physics context in section 5.4 of ref. [102], and in the context of optimal transport in refs. [114–116].[8]

P-AS may then be phrased as a constituent-level area subtraction procedure that solves the Apollonius problem for a given event. To solve the Apollonius problem associated with a particular Apollonius diagram, we subtract from the p_T of each particle an amount proportional to the area of its Apollonius region, $p_{T\,i}^{\text{AS}} = p_{T\,i} - \rho A_i^{\text{Apoll.}}$. Letting $z_{\text{cut}} = \rho\, A_{\text{tot}}/p_{T\,\text{tot}}$, in analogy to the z_{cut} of traditional groomers such as Soft Drop, we may equivalently write $p_{T\,i}^{\text{AS}} = p_{T\,i} - z_{\text{cut}}\, p_{T\,\text{tot}}\, A_i^{\text{Apoll.}}/A_{\text{tot}}$. In the piranha analogy, the transport plans for an event are determined by its Apollonius diagram, and all of the piranhas within a given Apollonius region feed on the associated jet constituent. A rigorous proof of the continuity of P-AS is given in Lemma 3.3 of ref. [114], and a visual representation of P-AS, taken from ref. [102], is depicted in figure 2.

---

[8]Eq. (2.6) with $\rho \mathcal{C} = \rho \mathcal{U}$ and arbitrary positive β also describes a valid optimal transport problem and an associated PIRANHA groomer. The solution to the associated optimal transport problem is described by a *generalized Laguerre diagram* [102, 116]. We leave the study of PIRANHA grooming motivated by generalized Laguerre diagrams for future work.

P-AS is closely related to constituent-level area subtraction techniques for pileup mitigation. Indeed, P-AS was initially introduced in the context of pileup mitigation in ref. [102], where it was shown that in certain limits, Apollonius Subtraction aligns with the discontinuous Voronoi Area Subtraction (VAS) procedure for pileup mitigation [8, 12, 117]. Thus, P-AS may be thought of as a continuous analog to existing techniques for constituent-level area subtraction. P-AS also aligns in certain limits with continuum Constituent Subtraction (CS), a continuous but computationally expensive algorithm for the removal of pileup [102, 118].

Unfortunately, the computational cost of finding the Apollonius diagram for a given particle event is relatively high because we rely on directly solving the Apollonius problem using numerical ghosts [102]: we use a uniform grid of "ghost" particles in the η-ϕ plane to obtain the Apollonius diagram associated with a particular event computationally. Data collection and analysis for theoretical and experimental studies at colliders motivate the computationally efficient analogs of P-AS that we describe in the remainder of this section.[9] A comparison of the runtime for our different PIRANHA algorithms is given in figure 17 of appendix B.

2.4 Iterated Voronoi Subtraction (P-IVS)

Iterated Voronoi Subtraction (P-IVS) is a PIRANHA groomer similar to both P-AS and VAS that overcomes the computational inefficiency of P-AS. P-IVS also uses the uniform event $\rho \mathcal{U}$ as a model for additive contamination, but crucially uses Voronoi diagrams rather than Apollonius diagrams to subtract this contamination away.

Like P-AS, the amount of grooming performed by the P-IVS grooming procedure is encoded in a parameter ρ. We may describe P-IVS quite succinctly as the solution to the iterated series of optimal transport problems [102]:

$$\mathcal{E}_{\mathrm{IVS}}^{(n+1)} = \arg\min_{\mathcal{E}'} \mathrm{EMD}(\mathcal{E}^{(n)}, \mathcal{E}' + \rho^{(n)}\mathcal{U}), \tag{2.7}$$

with

$$\rho^{(n)} = \min\left\{\rho - \sum_{i=0}^{n-1} \rho^{(i)}, \ \min_i \frac{p_{Ti}^{(n)}}{A_i^{(n)}}\right\}. \tag{2.8}$$

The final groomed event is given by the limit of the recursive procedure in eq. (2.7), $\mathcal{E}_{\mathrm{IVS}} = \lim_{n\to\infty} \mathcal{E}_{\mathrm{IVS}}^{(n)}$. Here, $\rho^{(0)} = 0$, $p_{Ti}^{(n)}$ is the transverse momentum of particle i after n applications of the recursive algorithm in eq. (2.7), and $A_i^{(n)}$ is the Voronoi region for the event after n applications of eq. (2.7). For example, $p_{Ti}^{(0)}$ and $A_i^{(0)}$ describe the transverse momenta and Voronoi areas of the ungroomed event.

More precisely, we may describe our ungroomed event as a collection of points in the η-ϕ plane describing the directions of outgoing momenta, each weighted by its transverse momentum. We may then enumerate the steps of the IVS procedure as follows:

---

[9]Another approach to developing analogs of P-AS that we do not pursue in this work is to use the formalism of *linearized optimal transport*, explored mathematically in refs. [119, 120] and applied to particle collisions in refs. [121, 122], which preserves the strengths of the EMD and significantly reduces the computational costs associated with optimal transport problems.

1. P-IVS constructs a Voronoi diagram in the η-ϕ plane for the particles in the event. We will label the stage of the P-IVS procedure by an integer n, starting at $n = 1$ and going up to at most the number of particles in the event.

2. P-IVS modifies the p_T of every particle in the event, subtracting p_T proportional to the area of the particle's Voronoi region until a particle would be removed or the grooming is complete:

$$p_{T\,i}^{(n)} = p_{T\,i}^{(n-1)} - \rho^{(n)} A_i^{(n-1)}, \tag{2.9}$$

with $\rho^{(n)}$ given by eq. (2.8).

3. If $\rho - \sum_{i=1}^{n} \rho^{(i)} > 0$, P-IVS increments n by 1 and continues from Step 1 by drawing a new Voronoi diagram for the modified event with a particle removed. Otherwise, the grooming is complete.

As with P-AS, the simple expression of P-IVS in terms of the EMD already showcases its continuity and connection to optimal transport. One may worry that the presence of an infinitesimally soft particle in an event may dramatically change the Voronoi diagram for an event, and therefore dramatically change the result of grooming using Voronoi areas. P-IVS overcomes this challenge by only using Voronoi areas until a particle is removed; when a particle is removed from the event, P-IVS computes an updated set of Voronoi areas that do not rely on the removed particle, continuing recursively until it has removed the correct amount of energy. A rigorous proof of the continuity of P-IVS is given in Lemma 3.3 of ref. [114].

Unweighted Voronoi diagrams for the original event $\mathcal{E}^{(0)}$ can be found efficiently, unlike the weighted Voronoi/Apollonius diagrams needed for P-AS. Furthermore, the Voronoi diagrams for the subtracted events $\mathcal{E}_{\mathrm{IVS}}^{(n)}$ used throughout the stages of the P-IVS algorithm do not need to be computed from scratch, and can be found in constant (amortized) time [102]. P-IVS thus retains the continuous grooming properties of P-AS while remaining amenable to numerical computation and data collection.

2.5 Tree-based PIRANHA: Recursive Subtraction (P-RS) with a Fraction (F)

Recursive Subtraction (P-RS) is an extension of the PIRANHA paradigm into the space of tree-based grooming algorithms and draws on the strengths of both hard-cutoff and PIRANHA grooming. We use "Recursive Subtraction" to denote an algorithm that takes in a binary tree of emissions describing a jet and recursively subtracts from the momenta of its jet constituents. The dependence of P-RS on binary trees is reminiscent of many hard-cutoff grooming strategies, which are computationally efficient and bear a close connection to the physics of perturbative QCD [3, 123]. The subtractive nature of P-RS echoes the area subtraction techniques of P-AS and P-IVS, which have a close connection to optimal transport and enjoy the advantage of continuity. P-RS combines these strengths and paves the way for PIRANHA groomers that are experimentally useful and theoretically tractable.

P-RS acts subtractively on jets associated with a *binary tree* or *clustering history*: a tree structure that emerges from recursively combining jet constituents, two at a time, until only the full jet remains. In particular, P-RS is a *de-clustering algorithm* on the jet tree: it splits the full jet tree into its two sub-jets, then splits each sub-jet into two sub-jets as

dictated by the clustering history, continuing until only final-state particles remain. In the discussion that follows, we also refer to sub-jets in the binary tree as *branches* or *emissions*.

Following the procedures of traditional grooming algorithms such as Soft Drop [3], the P-RS grooming algorithms we consider in this paper are *angular-ordered*: they act on jets associated with a binary tree whose splittings are ordered by the angular separation of their sub-jets. In particular, we consider jets are first clustered using the anti-k_t algorithm [124] and then *re-clustered* using the pair-wise angular-ordered Cambridge-Aachen (C/A) clustering algorithm [125].[10] C/A recursively clusters the jet constituents that are closest in angle, so that as we de-cluster each branch of the C/A tree we recover narrower and narrower emissions. Since the branching structure of an angular-ordered tree of emissions is comparable to the history of emissions within the parton model, a phenomenon known as coherent parton branching [123], angular-ordered grooming procedures have a closer connection to the physics of perturbative QCD.

P-RS also takes in a total amount of transverse momenta to subtract from the clustered jet, analogous to the choice of ρA_{jet} in the P-AS and P-IVS algorithms. At each splitting in the tree, P-RS assigns some amount of the grooming associated with a given branch to its two emissions. When the de-clustering finally reaches the final-state particles of a jet, P-RS subtracts from their transverse momenta, much like P-AS and P-IVS, by an amount proportional to the assigned grooming.

Recursive Subtraction with a Fraction f_{soft}. (P-RSF$_f$, or P-RSF) is a simple implementation of P-RS that does not depend on the kinematic information of each splitting. P-RSF depends on two real parameters, each between 0 and 1:

- z_{cut} is the fraction of transverse momentum the grooming will remove from the entire jet, $p_{T,\text{P-RSF}} = (1 - z_{\text{cut}})p_{T,\text{tot}}$. We may write $z_{\text{cut}} = z_{\text{cut}}^{(0)} = \rho A_{\text{jet}}/p_{T,\text{jet}}$, in analogy to the "piranha energy density" ρ of P-AS and P-IVS. We denote the total transverse momentum removed from the jet as $\Delta^{(0)} = z_{\text{cut}} p_{T,\text{tot}}$, and the transverse momentum removed from a branch i of the jet is $\Delta^{(i)} = z_{\text{cut}}^{(i)} p_{T,\text{tot}}$;

- f_{soft} is the fraction of the grooming assigned to the softer sub-jet at each stage of the de-clustering, as detailed below. P-RSF prefers to groom softer radiation more and more as f_{soft} increases.

In the steps of the algorithm below, we index the current branch of the jet on which P-RSF is acting by i. The softer sub-jet of branch i is denoted "i, soft", and the harder branch is denoted "i, hard". If the algorithm is at its starting point and considering the entire jet, we write $i = 0$.

The concrete prescription of P-RSF is as follows:

1. P-RSF attempts to de-cluster sub-jet i into two sub-jets. If this is impossible, the branch corresponds to a final-state particle and P-RSF proceeds to Step 5.

---

[10]Starting with the anti-k_t algorithm both mitigates fluctuations due to contamination in jet areas [124] and reduces the effects of clustering logarithms in analytic calculations [3].

2. P-RSF divides the grooming assigned to branch i between the softer sub-jet (i, soft) and the harder sub-jet (i, hard):

$$\Delta^{(i, \text{soft})} = f_{\text{soft}} \, \Delta^{(i)} \triangleq z_{\text{cut}}^{(i, \text{soft})} p_{T, \text{tot}} \tag{2.10a}$$

$$\Delta^{(i, \text{hard})} = (1 - f_{\text{soft}}) \, \Delta^{(i)} \triangleq z_{\text{cut}}^{(i, \text{hard})} p_{T, \text{tot}}, \tag{2.10b}$$

so that $\Delta^{(i, \text{soft})} + \Delta^{(i, \text{hard})} = \Delta^{(i)}$. In words, P-RSF assigns a fraction f_{soft} of the grooming to the softer sub-jet, and the remaining fraction $1 - f_{\text{soft}}$ to the harder sub-jet. Equivalently,

$$z_{\text{cut}}^{(i, \text{soft})} = f_{\text{soft}} z_{\text{cut}}^{(i)} \tag{2.11a}$$

$$z_{\text{cut}}^{(i, \text{hard})} = (1 - f_{\text{soft}}) z_{\text{cut}}^{(i)}, \tag{2.11b}$$

so that $z_{\text{cut}}^{(i, \text{soft})} + z_{\text{cut}}^{(i, \text{hard})} = z_{\text{cut}}^{(i)}$.

3. P-RSF checks if one of the sub-jets of the current branch is removed by the grooming:

   (a) If the grooming assigned to the softer sub-jet is greater than its actual transverse momentum, $\Delta^{(i, \text{soft})} > p_T^{(i, \text{soft})}$, P-RSF removes the softer sub-jet from the event. P-RSF then assigns the remaining grooming to the harder sub-jet, $\Delta^{(i, \text{hard})} = \Delta^{(i)} - p_T^{(i, \text{soft})}$.

   (b) Similarly, if the grooming assigned to the harder sub-jet is greater than its transverse momentum, $\Delta^{(i, \text{hard})} > p_T^{(i, \text{hard})}$, P-RSF removes the harder sub-jet from the event. It then assigns the remaining grooming to the softer sub-jet, $\Delta^{(i, \text{soft})} = \Delta^{(i)} - p_T^{(i, \text{hard})}$. This is possible only if $f_{\text{soft}} < 1/2$.

4. P-RSF implements Step 1 on the surviving sub-jets of i using the values of $\Delta^{(i, \text{soft})}$ and $\Delta^{(i, \text{hard})}$ calculated in the previous two steps.

5. When P-RSF reaches a final-state particle (a branch that cannot be further divided), it subtracts the grooming from its transverse momentum, effectively adjusting its energy fraction. Explicitly,

$$p_{T, \text{ P-RSF}}^{(i)} = p_T^{(i)} - \Delta^{(i)}. \tag{2.12}$$

Here, $p_{T, \text{ P-RSF}}^{(i)}$ denotes the groomed p_T of the final-state particle associated with the current branch i. This may also be described as a shift in the fraction of the total transverse momentum carried by the particle i,

$$z_{\text{P-RSF}}^{(i)} = (z^{(i)} - z_{\text{cut}}^{(i)})/(1 - z_{\text{cut}}), \tag{2.13}$$

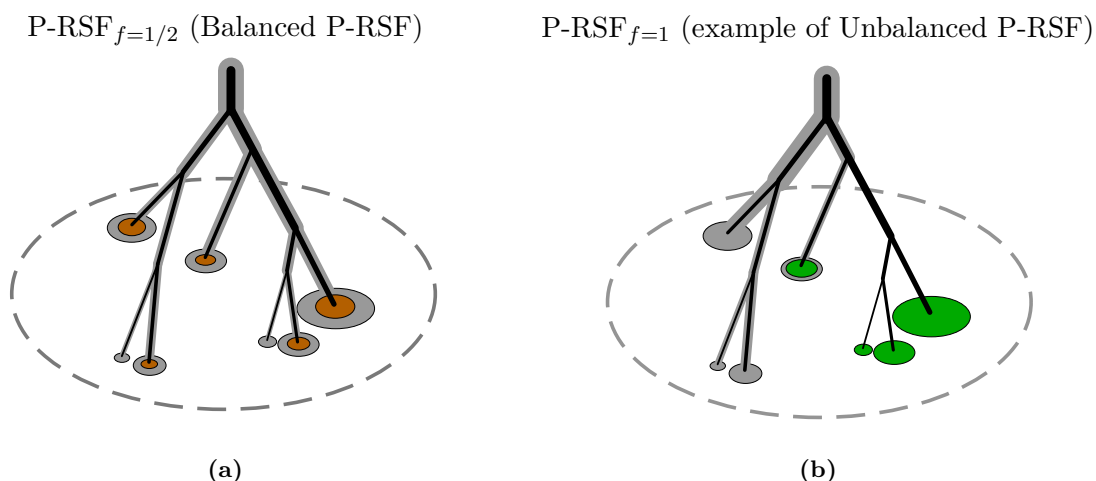where the additional factor of $1 - z_{\text{cut}}$ is due to normalization to the groomed transverse momentum.

Figure 3. Visualization of (a) P-RSF$_{f=1/2}$ and (b) P-RSF$_{f=1}$ acting on a toy jet depicted as an angular-ordered tree of emissions. The jet cone is indicated by the dashed ellipse, and the groomed final-state particles are indicated by colored ellipses, in dark orange (P-RSF$_{f=1/2}$) and green (P-RSF$_{f=1}$); the size of each circle is proportional to the groomed p_T of the corresponding particle. The grey highlights behind each emission in the tree indicate the transverse momentum subtracted from each branch, and eventually from the final-state particles of that branch, as described in section 2.5. The size of the grey highlights behind each final-state particle indicates the amount of transverse momentum subtracted from that particle.

P-RSF builds on the strengths of the P-RS framework, providing a simple implementation that does not rely on detailed kinematic information. By combining the advantages of tree-based hard-cutoff and continuous PIRANHA grooming paradigms, P-RS and P-RSF offer a promising path for experimentally useful and theoretically tractable continuous groomers. A visualization of P-RSF acting on a tree of emissions for $f_{\rm soft} = 1/2$ (which we also call **Balanced P-RSF**) and $f_{\rm soft} = 1$ (an example of P-RSF$_{f_{\rm soft}\neq 1/2}$ or **Unbalanced P-RSF**) is shown in figure 3.

2.6 P-RSF is (almost) continuous

P-RSF is well-defined on the space of energy distributions of particle events because it is invariant under exactly collinear splittings. To see this, we first notice that exact collinear splittings are infinitely narrow, and always correspond to branchings at the final layer of an angular-ordered tree of emissions. With this in mind, let us implement an exact collinear splitting, replacing a final-state particle i within a jet by two exactly collinear final-state particles, "i, soft" and "i, hard". Following the presentation of section 2.5, the transverse momentum subtracted from the new final state particles adds up to the transverse momentum subtracted from i, $\Delta^{(i,{\rm soft})} + \Delta^{(i,{\rm hard})} = \Delta^{(i)}$. Since i, soft and i, hard are exactly collinear, the energy distribution of i after grooming is equal to the sum of the energy distributions for i, soft and i, hard after grooming. The result of the P-RSF grooming procedure is therefore robust against exactly collinear splittings for any value of $f_{\rm soft}$.

While the subtractive algorithm of P-RSF cannot be expressed simply in terms of geometry, P-RSF echoes the features of P-AS and P-IVS that grant their continuity. For example, Steps 3(a) and 3(b) of the P-RSF algorithm ensure that P-RSF does not assign an area to any final-state particle that is larger than its transverse momentum.

However, P-RSF still suffers from discontinuities in suppressed regions of parameter space. We show in section 3.4 that Unbalanced P-RSF, or P-RSF with $f_{\text{soft}} \neq 1/2$, suffers from soft discontinuities: small changes to the energy of the jet have the potential to change which emissions of the jet are softer. The balanced recursive subtraction procedure, P-RSF$_{f=1/2}$, overcomes this weakness and is soft-continuous. Furthermore, section 3.5 discusses how recursive tree-based grooming algorithms suffer from discontinuities inherited from pairwise clustering, leading to angular discontinuities in both Balanced and Unbalanced P-RSF as well as traditional grooming algorithms such as Soft Drop.

2.7 Stronger notions of continuity

As a final note for this section, we point out that there are analogs of continuity that provide stronger constraints. One example of a stronger form of continuity is *uniform continuity*:

Definition 2 *A map M from energy flows to energy flows is* uniformly continuous *if, for any $\varepsilon > 0$, there exists a $\delta > 0$ such that for all $\mathcal{E}$ and $\mathcal{E}'$,*

$$\text{EMD}(\mathcal{E}, \mathcal{E}') < \delta \quad \implies \quad \text{EMD}(M(\mathcal{E}), M(\mathcal{E}')) < \varepsilon.$$

Roughly, we might say that uniform continuity requires that when we pick ε, M must be "continuous with the same δ for every event".

An even stronger condition is that of *Hölder continuity*:

Definition 3 *A map M from energy flows to energy flows is* Hölder continuous *with exponent $\alpha \in \mathbb{R}$ if, for all $\mathcal{E}$ and $\mathcal{E}'$,*

$$\text{EMD}(M(\mathcal{E}), M(\mathcal{E}')) < K \cdot \left(\text{EMD}(\mathcal{E}, \mathcal{E}')\right)^{\alpha},$$

with $K \in \mathbb{R}$ a constant. The special case $\alpha = 1$ is called Lipschitz continuity*.*

Placing stronger constraints on grooming, such as uniform continuity and Hölder continuity, has the potential to further constrain the effects of soft contamination and fluctuations in additive radiation on groomed results. For example, in which regions of parameter space are the hard-cutoff or Piranha groomers presented in this paper uniformly or Hölder continuous? Are other methods for continuous grooming, such as EMD-mode Piranha introduced in appendix A or Piranha groomers designed using linearized optimal transport [119–122], more amenable to strongly continuous generalizations? Can this knowledge help us identify obstacles to achieving these stronger types of continuity and guide us in designing more powerful, robust observables and groomers?

Though we do not investigate grooming methods that utilize these more restrictive analogs of continuity in this work, the invention and understanding of such groomers offer

an intriguing and promising avenue for future research. Instead, the remainder of this work is devoted to introducing hard-cutoff grooming, examining the types of discontinuities that may emerge in jet grooming, and studying the sensitivity of the continuous groomers we introduced in this section to important examples of low-energy pollution that appear in the study of particle collision data.

3 Discontinuities in grooming

In this section, we begin by reviewing the hard-cutoff Soft Drop grooming algorithm (section 3.1). We then compare the leading order (LO) soft continuity properties of Soft Drop and Recursive Subtraction both conceptually (section 3.2) and through a perturbative calculation (section 3.3).

In section 3.4 and 3.5, we explore potential challenges to continuity that can arise in more intricate scenarios. Specifically, we discuss suppressed soft discontinuities that emerge in Unbalanced Recursive Subtraction (section 3.4) and angular discontinuities induced by angular-ordered jet clustering that emerge beyond LO in tree-based grooming algorithms (section 3.5). We do not delve into the mathematical intricacies or perturbation theory of these examples, and they may be skipped on an initial reading.

3.1 Review of hard-cutoff grooming: soft drop

A common and elegant approach to jet grooming is to provide a hard cutoff on the energy of the radiation within a jet; this approach retains high-energy information, removes radiation that may be sensitive to soft distortions or due to additive contamination, and even facilitates perturbative calculations of jet substructure observables. In this section, we review a commonly used hard-cutoff grooming algorithm: the Soft Drop de-clustering algorithm [3]. Readers familiar with Soft Drop may skip to section 3.2, where we discuss the soft discontinuities of the Soft Drop algorithm.

The Soft Drop de-clustering algorithm is tree-based, computationally efficient, and serves as a representative example of a hard-cutoff grooming algorithm. Much like P-RSF, Soft Drop begins with an angular ordered tree of emissions and grooms at each branch of the tree. For Soft Drop, however, the parameter z_{cut} provides a hard cutoff, and Soft Drop sometimes maps similar events into groomed events that are very distinct. The resulting discontinuity in the grooming procedure leads to theoretical challenges and uncertainties, unpredictable and non-linear responses to non-perturbative physics, and difficulties in interpreting experimental data [99–101].

The Soft Drop algorithm depends on two parameters. The first, z_{cut}, is a hard cutoff that parameterizes the strength of the grooming. The second, β_{SD}, is a parameter that controls the angular dependence of the grooming procedure. Soft Drop preferentially removes soft radiation at wider and wider angles as β_{SD} is increased.

Soft Drop acts on a jet with characteristic radius R_0 that is first clustered using an algorithm such as the anti-k_t algorithm, and then re-clustered with the Cambridge-Aachen (C/A) algorithm. As in section 2.5, this produces a binary tree of clustered particles within

the jet, such that particles deeper in the tree tend to be closer in angle. Soft Drop then implements the following algorithm:

1. The most recent stage of the C/A clustering is undone, breaking the jet into two sub-jets.

2. The softer sub-jet is removed from the jet if it is not energetic enough to be considered relevant to the hard physics under study; in particular, let $z_{\text{soft}} = \min(p_{T,1}, p_{T,2})/(p_{T,1} + p_{T,2})$ denote the fraction of the p_T carried by the softer sub-jet, and let θ denote the angular distance between the two sub-jets. If the two sub-jets pass the Soft Drop grooming criterion:

$$z_{\text{soft}} > z_{\text{cut}} \left(\frac{\theta}{R_0} \right)^{\beta_{\text{SD}}} , \tag{3.1}$$

the grooming procedure is stopped. The two sub-jets are then re-merged and labeled as the groomed jet without further modification.

3. If the inequality of eq. (3.1) is not satisfied, the softer sub-jet is not energetic enough to be considered relevant to the hard physics of jet production. In this case, the softer of the two sub-jets is removed from the event and the procedure is repeated recursively by de-clustering the harder of the two sub-jets.

In this work, we use Soft Drop in grooming mode: if Soft Drop reaches a final-state particle that cannot be de-clustered as required by Step 1, we keep it as the groomed jet. Tagging mode would instead remove it from the jet entirely and leave a groomed event with no particles [3].

Grooming mode is a natural choice in the context of pileup mitigation because of the extreme behavior of tagging mode, which completely eliminates a single particle jet even when $z_{\text{cut}} = 0$. In particular, if we have an event in which we estimate there is *no* contamination from pileup (see section 5.1), we want to avoid grooming the jet at all and correspondingly set $z_{\text{cut}} = 0$. For a tree-level jet, consisting of a single quark, grooming mode respects our wish to leave the jet untouched, while tagging mode would nonetheless remove it entirely.

Soft Drop de-clustering is a generalization of the modified Mass Drop Tagger (mMDT) algorithm [4], and reproduces mMDT when $\beta_{\text{SD}} = 0$. Soft Drop has seen success in a wide variety of phenomenological applications, such as the characterization of boosted objects [2, 14, 25–93] and the extraction of parameters of the Standard Model, such as the top quark mass [126–128] and the strong coupling constant [129], from particle collision data.

Soft Drop with $\beta_{\text{SD}} > 0$ is collinear safe (invariant under exactly collinear splittings), and is therefore well-defined on the space of energy flows. However, when $\beta_{\text{SD}} \leq 0$, Soft Drop in grooming mode is not collinear safe. The lack of collinear safety in grooming mode can be quickly derived when considering a collinear splitting with $z < z_{\text{cut}}$ which fails the Soft Drop criterion. A similar problem emerges if one makes appropriate definitions for the algorithm in the limit $\theta \to 0$, and Soft Drop in grooming mode is also not collinear safe when $\beta_{\text{SD}} < 0$. Therefore, Soft Drop in grooming mode with $\beta_{\text{SD}} \leq 0$ is not well-defined as a map on the space of energy distributions, and we do not consider it in the remainder of the paper.
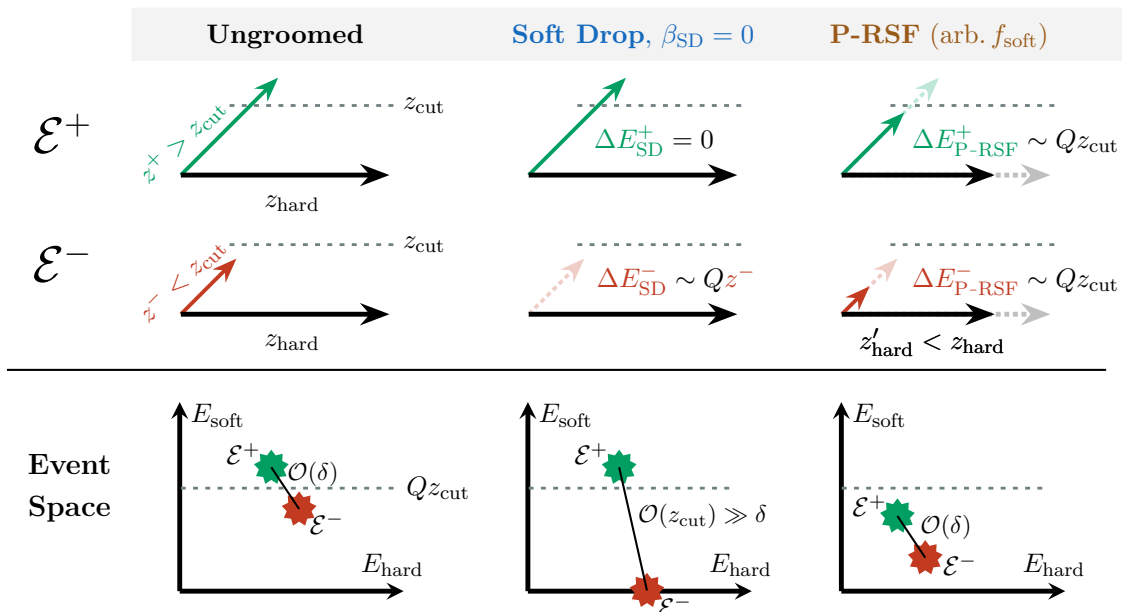
Figure 4. A cartoon comparing the discontinuous action of hard-cutoff grooming to the continuous action of PIRANHA on the nearly identical two-particle events with energy Q, $\mathcal{E}^+$ and $\mathcal{E}^-$. We use Soft Drop (with $\beta_{\mathrm{SD}} = 0$) and P-RSF (with arbitrary f_{soft}) with a grooming parameter z_{cut} as representative examples of the hard-cutoff and PIRANHA paradigms, respectively. The events $\mathcal{E}^{\pm}$ differ only by the energy fraction of the softer particle, $z^{\pm} = z_{\mathrm{cut}} \pm \delta/2$, and are separated by an infinitesimal EMD, $\mathrm{EMD}(\mathcal{E}^+, \mathcal{E}^-) \sim \delta \ll z_{\mathrm{cut}}$. Despite their similarity, Soft Drop maps the events to distinct groomed results separated by a large EMD: $\mathrm{EMD}(\mathcal{E}^+_{\mathrm{SD}}, \mathcal{E}^-_{\mathrm{SD}}) \sim z_{\mathrm{cut}} \gg \delta$. P-RSF instead maps the events continuously to infinitesimally similar groomed events: $\mathrm{EMD}(\mathcal{E}^+_{\mathrm{P\text{-}RSF}}, \mathcal{E}^-_{\mathrm{P\text{-}RSF}}) \sim \delta$. $\Delta E^{\pm}_G$ indicates the amount of energy removed from the event $\mathcal{E}^{\pm}$ by the groomer G.

3.2 An invitation to soft discontinuities: soft drop versus P-RS

Traditional groomers are discontinuous in the regions of parameter space near a hard cutoff. For example, Soft Drop is discontinuous in the region of parameter space where $z_{\mathrm{soft}} = z_{\mathrm{cut}} \theta^{\beta_{\mathrm{SD}}} / R_0^{\beta_{\mathrm{SD}}}$ due to the Soft Drop criterion of eq. (3.1). In this subsection, we discuss the discontinuity of hard-cutoff grooming and the relative continuity of PIRANHA.

The following discussion centers around the illustrative example shown in figure 4, which shows the simplest instance of discontinuous behavior near a hard cutoff as well as the resolution offered by PIRANHA. In this example, we consider two nearly identical events, $\mathcal{E}^+$ and $\mathcal{E}^-$, that each contain two particles. The softer constituent of $\mathcal{E}^+$ has an energy fraction $z^+ = z_{\mathrm{cut}} + \delta$ slightly above the grooming parameter z_{cut}, while the softer constituent of $\mathcal{E}^-$ has an energy fraction $z^- = z_{\mathrm{cut}} - \delta$ slightly below z_{cut}. The energy flows of the two events are nearly indistinguishable, in the sense that $\mathrm{EMD}(\mathcal{E}^+, \mathcal{E}^-) \sim \delta$, and we choose δ for consistency with Definition 1. We use Soft Drop with $\beta_{\mathrm{SD}} = 0$, or mMDT, as a representative example of hard-cutoff grooming, and P-RSF with arbitrary f_{soft} as a representative example of the PIRANHA paradigm; for each, we use the parameter z_{cut}.

Soft Drop treats $\mathcal{E}^+$ and $\mathcal{E}^-$ very differently: it does not modify $\mathcal{E}^+$, but completely removes the softer particle of $\mathcal{E}^-$. The energy flow of $\mathcal{E}^+$ is unchanged, while that of

$\mathcal{E}^-$ is changed dramatically, and the EMD between the groomed jets is relatively large: $\mathrm{EMD}(\mathcal{E}_{\mathrm{SD}}^+, \mathcal{E}_{\mathrm{SD}}^-) \sim z_{\mathrm{cut}} \gg \delta$. Direct application of Definition 1 to our example shows that Soft Drop is discontinuous.[11] Since the events $\mathcal{E}^+$ and $\mathcal{E}^-$ differ by small changes to the energy of jet constituents, we see that Soft Drop is soft discontinuous in the region of parameter space near the cutoff z_{cut}.

The procedure of P-RSF is quite different. When grooming both events, P-RSF$_{f=1/2}$ grooms the energy of both emissions by the same amount, leading to no large discrepancy in the two groomed results. Indeed, after the P-RSF$_{f=1/2}$ grooming procedure, these two events are still separated by an infinitesimal EMD: $\mathrm{EMD}(\mathcal{E}_{\mathrm{RSF}}^+, \mathcal{E}_{\mathrm{RSF}}^-) \sim \delta$. This simple example demonstrates how hard-cutoff grooming methods may discontinuously map two nearly identical events into vastly different groomed results, and showcases the soft continuous resolution of the P-RSF$_{f=1/2}$ grooming procedure and other PIRANHA groomers.

3.3 Soft discontinuities in perturbation theory: soft drop versus P-RS

Next, we examine the manifestations of the soft discontinuous behavior of hard-cutoff grooming more quantitatively by comparing Soft Drop and P-RSF at leading order in perturbation theory (LO). In particular, we examine the LO effects of discontinuity on the distributions of the two-prong generalized energy correlation functions (ECFs) of ref. [39]. For jets that are central ($y = 0$) and narrow ($R_0 \ll 1$), we may write the ECFs in the form

$$C_1^{(\varsigma)} \simeq \frac{1}{2} \sum_{i=1}^{M} \sum_{j=1}^{M} z_i z_j \left(\frac{\theta_{ij}}{R_0} \right)^{\varsigma}, \tag{3.2}$$

up to non-singular corrections in powers of the jet radius, which we neglect [3]. In eq. (3.2), z_i represents the fraction of the jet energy carried by particle i, θ_{ij} indicates the angle between particles i and j, R_0 indicates the jet radius, and the sum is over all particles in the groomed jet.

In our discussion below, we focus on the calculation of LO substructure because it demonstrates the effects of the discontinuity of Soft Drop and the continuity of PIRANHA in a simple context. We provide a more detailed description of our LO calculation of PIRANHA-groomed substructure in appendix C.1, and extend our substructure analysis to leading logarithmic (LL) accuracy in appendix C.3. The resummed analysis of appendix C.3 does not change the qualitative conclusions of our discussion below. However, it highlights subtleties in systematically improving our resummed substructure calculations due to the global, subtractive nature of PIRANHA grooming. In appendix C.2, we also discuss the calculation of the groomed energy fraction z_g — another common observable used in the study of jet grooming — which exhibits similar subtleties. Therefore, while our resummation does not change the qualitative conclusions of the following discussion, it indicates

---

[11]We note that Soft Drop with $\beta \neq 0$ is also soft discontinuous. Indeed, the above arguments still hold for the case $\beta \neq 0$, verbatim, when the angle between the particles of $\mathcal{E}^\pm$ is fixed to $\theta = R_0$. As expressed in eq. (3.1), the Soft Drop criterion at this fixed angle is still $z_{\mathrm{soft}} > z_{\mathrm{cut}}$. However, since Soft Drop with $\beta \neq 0$ drops the softer particle if it does not satisfy $z/\theta^\beta > z_{\mathrm{cut}}/R_0^\beta$, we may also say that Soft Drop with $\beta \neq 0$ is "z/θ^β-discontinuous", where z-discontinuity indicates the soft discontinuity for the case of $\beta = 0$ that we discuss above.

that the computation of resummed Piranha observables at higher accuracy may require new calculational tools.

At LO, the emission of a single parton with energy fraction z and angle θ is described approximately by the pseudo-probability distribution

$$\frac{\alpha_s}{\pi} p_i(z) \frac{1}{\theta} \approx \frac{2C_{R_i}\alpha_s}{\pi} \frac{1}{z} \frac{1}{\theta}, \tag{3.3}$$

where $p_i(z)$ is a Dokshitzer-Gribov-Lipatov-Altarelli-Parisi (DGLAP) splitting function [131–133] describing the splitting of a mother parton i, and C_{R_i} is the quadratic Casimir for the SU(3) color representation R_i of the mother parton. $C_F = (N_C^2 - 1)/(2N_C) = 4/3$ for quarks and $C_A = N_C = 3$ for gluons. We are interested in particular in the *reduced* splitting functions,

$$\overline{p}_i(z) = p_i(z) + p_i(1-z), \tag{3.4}$$

where $z \in (0, 1/2)$. Reduced splitting functions are relevant for our groomed calculations because they describe the energy distribution of the softer of the two partons after the splitting.

The probability distribution of $C_1^{(\varsigma)}$ for Soft Drop was studied in ref. [3], which found, for $\beta_{\text{SD}} > 0$, the LO result

$$\rho_{i,\,\text{SD}}(C_1^{(\varsigma)}) \approx \frac{2\alpha_s C_{R_i}}{\pi\varsigma} \frac{1}{C_1^{(\varsigma)}} \times \begin{cases} -\log C_1^{(\varsigma)} + B_i, & C_1^{(\varsigma)} > z_{\text{cut}}; \\ -\frac{\beta_{\text{SD}}}{\varsigma+\beta_{\text{SD}}} \log C_1^{(\varsigma)} - \frac{\varsigma}{\varsigma+\beta_{\text{SD}}} \log z_{\text{cut}} + B_i, & C_1^{(\varsigma)} < z_{\text{cut}}, \end{cases} \tag{3.5}$$

away from $C_1^{(\varsigma)} = 0$, up to $\mathcal{O}(\alpha_s^2)$ and terms that are power suppressed in $C_1^{(\varsigma)}$, z_{cut}, or both. B_i is a factor due to hard-collinear pieces of the splitting function that are not singular as z approaches 0: $B_q = -3/4$ for quarks and $B_g = -11/12 + n_f/(6C_A)$ for gluons, where n_f is the number of active quark flavors. The case $\beta_{\text{SD}} < 0$ leads to an identical result up to another constraint on $C_1^{(\varsigma)}$, which simply multiplies the expression above by $\Theta(C_1^{(\varsigma)} > z_{\text{cut}}^{\varsigma/|\beta_{\text{SD}}|})$.

The piece-wise behavior of eq. (3.5) and the associated kink in the Soft Drop $C_1^{(\varsigma)}$ distribution are due to the discontinuous behavior of Soft Drop. As noted by ref. [3], when $\beta_{\text{SD}} \geq 0$ any two-parton configurations with $C_1^{(\varsigma)} = z\,(\theta/R_0)^\varsigma > z_{\text{cut}}$ must have $z\,(\theta/R_0)^{\beta_{\text{SD}}} > z_{\text{cut}}$, and are therefore not affected by Soft Drop. On the other hand, configurations with $C_1^{(\varsigma)} < z_{\text{cut}}$ may be Soft Dropped, leading to the piece-wise change in groomed substructure in this region of phase space.

We discuss the calculation of the $C_1^{(\varsigma)}$ distribution for P-RSF groomed jets in appendix C.1, where we find the LO result

$$\rho_{i,\,\text{P-RSF}}(C_1^{(\varsigma)}; f) \approx \frac{2\alpha_s C_{R_i}}{\varsigma\,\pi} \frac{1}{C_1^{(\varsigma)}} \left(-\log\left(C_1^{(\varsigma)} + f\,z_{\text{cut}}\right) + B_i \right), \tag{3.6}$$

away from $C_1^{(\varsigma)} = 0$, using the same factors of B_i as for Soft Drop, and up to terms that are power-suppressed in $C_1^{(\varsigma)}$, z_{cut}, or both.[12]

---

[12]We include some additional power-suppressed terms and additional contributions when $C_1^{(\varsigma)} = 0$ in the LO discussion of appendix C.1.

Notably, the P-RSF distribution for $C_1^{(\varsigma)}$ is smooth and does not need to be defined in a piece-wise fashion. Even at LO, the subtractive nature of P-RSF leads to a smooth interpolation between double-logarithmic behavior at large $C_1^{(\varsigma)}$ and single-logarithmic behavior at small $C_1^{(\varsigma)}$.[13] Much as the kink in the Soft Drop distribution can be attributed to the sudden activation of the grooming procedure in certain regions of phase space, the smoothness of the P-RSF distribution reflects in part that the grooming procedure is always active.[14] While the piece-wise behavior of Soft Drop observable distributions is smoothed out by all-orders effects [134], the smoothness of P-RSF observable distributions at LO is a manifestation of the subtractive behavior of P-RSF that ensures its continuity on two-parton jets.

We also note that the substructure distributions of other PIRANHA grooming procedures can be calculated at LO with a similar method by upgrading f_{soft} to a function of z and θ. Further, we expect that such an $f(z,\theta)$ may be well approximated by $f(0,0)$, with sub-leading corrections proportional to $\mathcal{O}(z,\theta,z_{\mathrm{cut}})$.

In our discussion above, we emphasized that the kinks in Soft Drop groomed substructure distributions are manifestations of the discontinuity of Soft Drop itself. Generalizations of these arguments, however, must be verified carefully. In particular, there is not generically a one-to-one correspondence between kinks in distributions and discontinuities in the associated observable. As a toy one-dimensional example, consider a random variable X that is uniformly distributed in $(-1/2, 1/2)$. The discontinuous, "zig-zagging" function $f(X) = X \pm 1/2$, choosing the upper sign for $X < 0$ and the lower sign for $X > 0$, is uniformly distributed. Therefore, discontinuous functions on our toy "phase space" — the domain of X — do not necessarily have kinks in their distributions. Conversely, there is a kink in the distribution of $f(X) = \mp(1 - \sqrt{1 \pm 2X})/2$, where we take the upper sign when $X < 0$ and the lower sign for $X \geq 0$, even though $f(x)$ is a continuous function of x and has a continuous first derivative.

3.4 Suppressed soft discontinuities in unbalanced P-RS

Unbalanced P-RSF algorithms are also discontinuous in the suppressed region of parameter space where $z = 1/2$.[15] The discontinuous behavior emerges because Unbalanced P-RSF algorithms treat harder and softer sub-jets differently; since an infinitesimally soft perturbation can lead to either sub-jet having more energy when $z = 1/2$, an infinitesimal change in an ungroomed energy flow can lead to macroscopic differences in groomed results. In the language of section 2.2, we say that Unbalanced P-RSF is soft discontinuous in the suppressed, measure zero region of parameter space where $z = 1/2$, and that Balanced P-RSF is the only soft continuous P-RSF algorithm.

---

[13]The terms "double logarithmic" and "single logarithmic" here refer to the behavior of the LO cumulative distribution function. In particular, we note that when $C_1^{(\varsigma)} \gg f z_{\mathrm{cut}}$, the pseudo-probability distribution $\rho \propto \log{(C)}/C$ demonstrates double-logarithmic behavior. On the other hand, when $0 < C_1^{(\varsigma)} \ll f z_{\mathrm{cut}}$, the distribution $\rho \propto \log{(f z_{\mathrm{cut}})}/C$ demonstrates single-logarithmic behavior.

[14]We say that the smoothness of P-RSF only reflects the global activity of the grooming *in part* because we expect that PIRANHA groomers that turn on gradually will also have smooth substructure distributions.

[15]The probability density of seeing a softer sub-jet with energy fraction z scales as $1/z$ at leading logarithmic accuracy in perturbative QCD. Therefore, the region with $z \sim 1/2$ at the boundary of the jet phase space is far less populated than the regions of phase space near a hard-cutoff with $z \sim z_{\mathrm{cut}} < 1/2$.
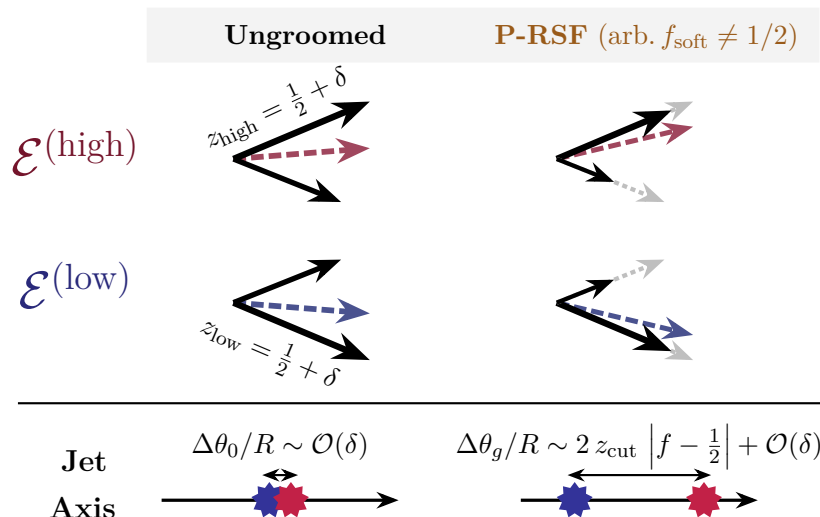
Figure 5. A cartoon demonstrating the soft discontinuity of P-RSF for $f_{\text{soft}} \neq 1/2$ acting on the nearly identical events $\mathcal{E}^{(\text{high},\text{low})}$, each with two particles with nearly identical energy fractions, $z \sim 1/2$, and with opening angle R. Red denotes the event $\mathcal{E}^{(\text{high})}$, blue denotes $\mathcal{E}^{(\text{low})}$, and a thick dotted arrow indicates the jet axis for each event. The hardest particle of $\mathcal{E}^{(\text{high})}$ is its upper particle, while the hardest particle of $\mathcal{E}^{(\text{low})}$ is its lower particle. Since P-RSF with $f_{\text{soft}} \neq 1/2$ preferentially subtracts radiation from the softer particle of each event, the events are mapped discontinuously to distinct groomed results. The discontinuous behavior of the groomed jet axis is one manifestation of this soft discontinuous behavior: while the ungroomed energy-weighted jet axes of $\mathcal{E}^{(\text{high})}$ and $\mathcal{E}^{(\text{low})}$ differ by an infinitesimal amount ($\Delta\theta_0 \sim R\epsilon$), the groomed jet axes are widely separated for any $f_{\text{soft}} \neq 1/2$.

In the following discussion, we explore the soft discontinuity of Unbalanced P-RS in the context of a jet containing only two partons. We find that the discontinuity at $z = 1/2$ can lead to macroscopic changes to the groomed jet axis. However, only the orientation of the jet is affected by the soft discontinuity in our two-parton example, and the soft discontinuity of Unbalanced P-RS is not reflected at LO by traditional jet substructure variables. Nonetheless, important features of a jet, such as its thrust axis [136, 137], are susceptible to discontinuities at LO. Furthermore, even traditional jet substructure observables are impacted by the soft discontinuity of Unbalanced P-RS in three-parton jet configurations that emerge beyond LO.

In the example of figure 5, we explore the effect of the soft discontinuity of Unbalanced P-RSF on the axis of a jet. In particular, we examine nearly identical two-parton jets in the region of phase space near $z \sim 1/2$ that differ only by an infinitesimally soft perturbation. We denote these events by $\mathcal{E}^{(i)}$, where $i \in \{\text{high}, \text{low}\}$. The $\mathcal{E}^{(i)}$ each consist of two particles: a "high" and "low" particle separated by an angular distance R. In $\mathcal{E}^{(i)}$, particle i has just slightly more energy, $z_i = 1/2 + \delta$.

For simplicity, we focus on the discontinuous behavior of the energy-weighted jet axis,

$$\hat{n}_{\text{jet}} \propto z_{\text{high}} \hat{n}_{\text{high}} + z_{\text{low}} \hat{n}_{\text{low}}, \tag{3.7}$$

for the groomed events $\mathcal{E}^{(i)}$, and the proportionality holds up to normalization. Since

the ungroomed $\mathcal{E}^{(i)}$ are nearly identical, the angle between the ungroomed jet axes is infinitesimally small:

$$\Delta\theta_0 = 2\,R\,\delta + \mathcal{O}(R^3,\,\delta^3), \tag{3.8a}$$

$$\lim_{\delta\to 0}\Delta\theta_0 = 0, \tag{3.8b}$$

where $\Delta\theta_0$ indicates the angle between the jet axes of $\mathcal{E}^{(\text{high})}$ and $\mathcal{E}^{(\text{low})}$ before grooming. After the application of P-RSF, however, the softer branch of each event will be dramatically diminished, and there is a large angle between the two groomed jet axes:

$$\lim_{\delta\to 0}\Delta\theta_g = 2\,R\,\frac{z_{\text{cut}}}{1-z_{\text{cut}}}\,|f_{\text{soft}} - 1/2| + \mathcal{O}(R^3) \neq 0, \tag{3.9}$$

where $\Delta\theta_g$ indicates the angle between the jet axes of $\mathcal{E}^{(\text{high})}$ and $\mathcal{E}^{(\text{low})}$ after grooming with P-RSF. The continuous behavior indicated by $\lim_{\delta\to 0}\Delta\theta_0 = 0$ before grooming is in sharp contrast to the discontinuous behavior after grooming, $\lim_{\delta\to 0}\Delta\theta_g \neq 0$. In eq. (3.9), we have assumed that $\max(f, 1-f)\,z_{\text{cut}} < 1/2-\delta$, so that neither particle is fully removed from the event. Of course, when $f_{\text{soft}} = 1/2$, the grooming algorithm treats the harder and softer sub-jets equally, so that we once again have $\Delta\theta_g \to \mathcal{O}(\delta)$ and the soft discontinuous behavior fades away.

This simple example demonstrates a general principle for designing Recursive Subtraction algorithms: if they are to be soft-continuous, they must treat the softer and harder sub-jets identically in the limit $z \to 1/2$. For example, let us imagine a P-RSF-type grooming algorithm with f_{soft} upgraded to a function of the energy fraction z and angle θ of a branch, $f_{\text{soft}} \to f(z, \theta)$. Any P-RSF-type groomer with $f(1/2, \theta) = 1/2$ overcomes the soft discontinuity discussed of the discussion above. We call such P-RS groomers that treat emissions in the region $z = 1/2$ identically "Hard-Balanced" Recursive Subtractors, as they are balanced in the region of phase space where two sub-jets are equally hard. A P-RSF algorithm with $f(z, \theta) = 1-z$, for example, is Hard-Balanced, and still preferentially grooms infinitesimally soft radiation. We leave the study of Hard-Balanced Recursive Subtractors to future work.

Finally, we note that while the soft discontinuity we detailed above is an important formal feature of P-RSF, the regions of phase space where $z \sim 1/2$ contribute far less to observable distributions than those for which $z \sim z_{\text{cut}}$. In particular, the pseudo-probability distribution for z scales as $1/z$ in perturbative QCD, implying that the soft discontinuity of Unbalanced P-RSF that we explore above is suppressed relative to the soft discontinuities of hard-cutoff groomers such as Soft Drop. Indeed, the soft discontinuities of P-RSF do not manifest in leading logarithmic substructure distributions, and we may even use P-RSF with $f_{\text{soft}} \neq 1/2$ to gain perturbative insight into the soft continuous analog, P-RSF$_{f=1/2}$, and into PIRANHA grooming in general; we explain this reasoning in more depth and pursue this goal in appendix C.

On the other hand, soft discontinuities associated with hard cutoffs in traditional grooming procedures lie on the interior of the jet phase space, straddled by the regions $0 < z < z_{\text{cut}}$ and $z_{\text{cut}} < z < 1/2$. As we saw in section 3.3, the associated soft discontinuities lead to piece-wise definitions of groomed observable distributions even at LO to isolate the behavior of radiation above and below the associated hard cutoff.
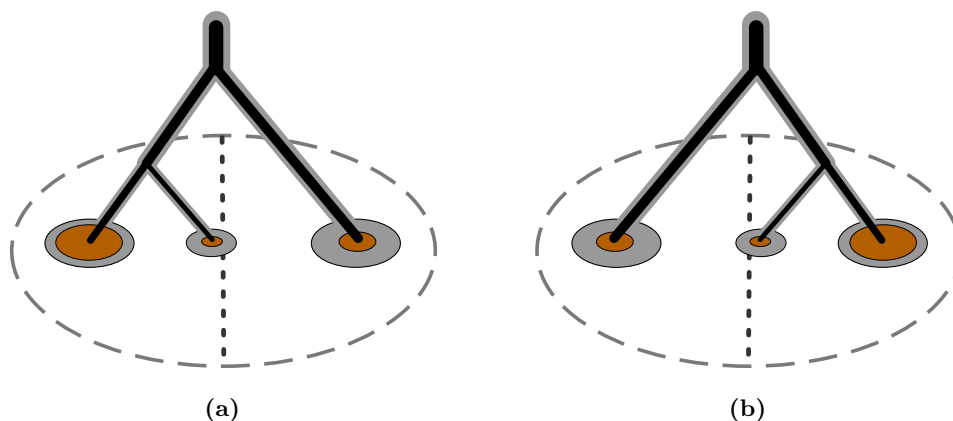
Figure 6. A visualization of angular discontinuity in tree-based grooming in the style of figure 3, with the angular-ordered P-RSF$_{f=1/2}$ algorithm as a representative example. The thick dashed line in the center of the jet is equidistant in angle from the left and right sub-jets of the event, and the middle emission lies nearly on this line. In (a), the middle emission lies closer to the left sub-jet, and the right sub-jet is therefore groomed more. In (b), the middle emission is closer to the right sub-jet, and the left sub-jet is groomed more. We see that a small change in the angle of the middle emission may drastically change the clustering history of the jet, and the result of a tree-based grooming procedure.

3.5 Clustering discontinuities beyond leading order in tree-based grooming

Recursive tree-based grooming algorithms also suffer from *clustering discontinuities*, or discontinuities inherited from the procedure of clustering the jet into a tree of emissions (as discussed in section 2.5). Clustering discontinuities are inevitable in tree-based grooming because any clustering algorithm relies on the discontinuous notion of which particles are "closest" to one another. However, since there is only one way to cluster a two-parton jet, clustering discontinuities only emerge beyond LO and when a jet contains three or more partons.

Soft Drop and P-RSF in particular suffer from an *angular clustering discontinuity*. Both rely on an angular-ordered jet clustering history obtained from the angular-ordered C/A clustering algorithm; the notion of the C/A clustering history is discontinuous because small changes to the angles of jet constituents may lead to distinct clustering histories. In the language of section 2.2, we say angular-ordered clustering algorithms and angular-ordered, tree-based grooming algorithms are angularly discontinuous in a measure zero region of phase space.

In our discussion, we focus on a simple manifestation of angular clustering discontinuities in the three-parton jet shown in figure 6. In particular, we consider the region of phase space where a particular sub-jet is nearly equidistant from a sub-jet on the left and another on the right. We call the left and right sub-jets L and R, and we denote the equidistant emission as M, for "middle".

A small change in the location of emission M may lead to distinct C/A clustering histories. If M is slightly closer to the left emission L, then L will receive a quarter of the

total grooming while R will receive half (as in figure 6a); the result is that L is groomed less than R, and much more of the structure of the left side of the jet is preserved after the grooming. Alternatively, if M is slightly closer to the right, then R will be groomed less than L; the right side of the jet will be preserved while much of the left side will be groomed away (as in figure 6b). Despite the similarity of the ungroomed events, the resulting groomed events will have macroscopic differences and be separated by a relatively large EMD.

To emphasize this point quantitatively, let us consider P-RSF acting on the simple example where all three emissions lie in a plane, the left and right emissions each have an energy fraction $z_\ell = z_r = z_0$, and the middle emission — nearly equidistant from the left and right emissions — has $z_m = 1 - 2z_0$.[16] We use R to denote the angle between the left and right emissions, and δ to denote a small perturbation in the angle of the middle emission: $\theta_{\ell m} = R/2 + \delta$ in one event, while $\theta_{\ell m} = R/2 - \delta$ in the other. We again examine the discontinuity in the energy-weighted jet axis, which takes the more general form

$$\hat{n}_{\text{jet}} \propto \sum_i z_i \hat{n}_i, \tag{3.10}$$

where the proportionality holds up to normalization. A simple computation then yields

$$\Delta\theta_0 = 2\, z_m\, \delta + \mathcal{O}\left(R^2, \delta^3\right), \tag{3.11a}$$

$$\lim_{\delta \to 0} \Delta\theta_0 = 0, \tag{3.11b}$$

for the angle $\Delta\theta_0$ between the two jet axes before grooming, while

$$\lim_{\delta \to 0} \Delta\theta_g = R\, \frac{z_{\text{cut}}}{1 - z_{\text{cut}}}\, |f(f-3)+1| + \mathcal{O}(R^3) \neq 0, \tag{3.12}$$

for the angle between the two jet axes after the application of P-RSF, where we have assumed that z_{cut} is sufficiently small that all particles survive the grooming. We again see a discontinuous behavior in the energy-weighted axis of the groomed jets, $\lim_{\delta \to 0} \Delta\theta_g \neq 0$, despite the fact that the ungroomed events were nearly identical and $\lim_{\delta \to 0} \Delta\theta_0 = 0$.

Fortunately, the angular discontinuities associated with tree-based clustering algorithms are also suppressed in the phase space of perturbative QCD. First, we note that the angular discontinuities in angular-ordered clustering algorithms emerge only at higher orders of perturbative accuracy for which a jet has three or more constituents. These regions of phase space are generically suppressed in high-energy QCD calculations, where the emission of extra final-state partons is suppressed by the strong coupling constant α_s. Furthermore, perturbative QCD predicts a strong hierarchy in the angles between different emissions in a jet, $\theta_1 \gg \theta_2 \gg \cdots$ [123]. Therefore, even among phase space configurations with three or more constituents, it is very unlikely in perturbative QCD to find a configuration in which a particular final-state particle lies nearly equidistant in angle between two others.

---

[16]We focus on P-RSF in this example to avoid the more complex analysis required for Soft Drop that may obfuscate the physics at hand. Soft Drop also exhibits an angular clustering discontinuity that can be seen in the context of this simple example, but that is complicated by the possible behaviors of the grooming: z_{cut} and z_0 may conspire such that zero, one, or two emissions are groomed from the event. The effects of the clustering discontinuity on the energy-weighted jet axis must therefore be considered piece-wise in the (z_{cut}, z_0) parameter space.

We also emphasize that the angular discontinuity of the C/A clustering algorithm is representative of more general discontinuities in tree-based clustering and in jet algorithms.[17] For example, the k_t and anti-k_t algorithms exhibit a similar k_t discontinuity by identical arguments. Discontinuities of different clustering algorithms may be more or less suppressed in the phase space of perturbative QCD, or more or less sensitive to particular sources of low-energy pollution. Therefore, a more detailed discussion of clustering discontinuities and re-clustering schemes that minimize their effects is a potential direction for future research.

4 Responses of grooming to soft distortions

We now compare the responses of PIRANHA and traditionally groomed jets to soft distortions in more detail, focusing first on hadronization and then on the exclusion of neutral particles as a simplified probe of smearing effects. In particular,

- We use P-RSF$_{f=1/2}$ as a representative of PIRANHA grooming procedures;

- We use Soft Drop with $\beta_{SD} = 2$ as a representative of traditional grooming procedures.

For details regarding other PIRANHA grooming procedures and Soft Drop with different β_{SD}, see appendix B.

4.1 Samples and observables

In the studies of this section, we use `Pythia 8.244` [105] with the default 4C tune [140], and work with samples of QCD quark and gluon jets without multiple parton interactions. We consider jets with transverse momentum $p_T > 500$ GeV, maximum pseudorapidity $|\eta| < 4$, and clustered with the anti-k_T algorithm [124] with $R = 1$; we then re-cluster using the Cambridge-Aachen algorithm [125]. Our analyses focus on Soft Drop with $\beta_{SD} = 2$ as a representative for traditional grooming procedures and on Balanced Recursive Subtraction (P-RSF$_{f=1/2}$) as a representative for PIRANHA grooming; the Balanced Recursive Subtraction algorithm is available on GitHub [113].

We study the responses of groomed jets to soft distortion through three qualitatively different lenses:

- **Energy Flow: EMD.**
  To understand the overall response of groomed jet energy flow to soft radiation, we use the EMD of section 2.1. The EMD between two jets bounds the difference of

---

[17]To mitigate the discontinuities of tree-based clustering, one can introduce a continuous weighting procedure for clustering histories. One way to assign continuous weights to different clustering histories is the Q-JET scheme [138, 139], which notably leads to reduced statistical fluctuations in observable distributions compared to the standard scheme of using a single clustering history for jet substructure calculations. While the Q-JET scheme has not been applied directly to grooming in the context of energy flows, to our knowledge, it provides a particularly simple way to overcome the clustering discontinuities we discuss here: we may simply define an overall groomed energy flow — devoid of clustering discontinuities — as the weighted average of the groomed energy flows associated with each possible clustering history.

a large class of IRC-safe observables between the jets and provides an observable-independent tool to study the robustness of different grooming procedures to the presence of soft contamination.[18]

- **Extensive Properties: p_T.**
  To understand how the extensive properties of groomed jets are modified by soft contamination, we examine the responses of the transverse momentum of groomed jets. We are concerned in particular with the amount by which transverse momentum in the presence of contamination, $\widetilde{p_T}$, is shifted relative to the associated uncontaminated value, p_T:

$$\Delta p_T = \widetilde{p_T} - p_T. \tag{4.1}$$

p_T indicates transverse momentum in the absence of distortion or contamination and $\widetilde{p_T}$ captures the response of transverse momentum to different sources of low-energy physics. This is specified more precisely in each of the subsequent sections.

- **Substructure: $C_1^{(\varsigma)}$.**
  We use jet substructure to capture the response of the intensive properties of jets to contaminating radiation. In particular, we study the response of the two-prong energy correlation functions (ECFs) of ref. [39]:

$$C_1^{(\varsigma)} = \frac{1}{2} \frac{\sum_{i=1}^{M} \sum_{j=1}^{M} p_{T,i} \; p_{T,j} \; (R_{ij}/R_0)^{\varsigma}}{p_{T,\,\text{tot.}}^2}. \tag{4.2}$$

$p_{T,i}$ represents the fraction of the jet transverse momentum carried by particle i, R_{ij} indicates the rapidity-azimuth distance between particles i and j, and the sum is over all particles in the jet. In the limit of jets that are central ($y = 0$) and narrow (have jet radius $R_0 \ll 1$), we can write the ECFs in the approximate form of eq. (3.2). As a concrete example of the response of substructure to soft radiation, we focus on $\Delta C_1^{(2)} = \widetilde{C_1^{(2)}} - C_1^{(2)} \approx \Delta \left(m^2/p_T^2 \right)$, where a tilde again indicates the associated observable in the presence of a particular source of low-energy physics, while the absence of a tilde indicates the absence of low-energy effects.

We emphasize that in all of our phenomenological studies (except our studies of the underlying event in section 5.2 and 5.3), we examine how each source of contamination leads to changes on a "per-jet" level, rather than at the level of distributions of observables. In particular, we examine the EMD between each jet before and after contamination, and the changes in p_T and $C_1^{(2)}$ induced by contamination for individual jets. To the best of our knowledge, the effects of hadronization on a per-jet basis have not been studied in this way before, as the sensitivity of groomed jets to hadronization is often studied instead at a statistical level by studying the distributions of groomed jet observables. Therefore, the results we present offer a more detailed examination of the effects of the hadronization model of

---

[18]In particular, the EMD of section 2.1 tends to reflect the response of extensive properties of jets, as we see concretely in our examples; a less concrete way to see this is by directly looking at eq. (2.1), and noticing that the EMD has units of energy.

`Pythia 8.244`; however, these results may not be precisely reflected solely by the statistical changes in distributions of groomed jet observables that are induced by hadronization.

4.2 Hadronization

We first provide phenomenological evidence that continuous grooming is less sensitive than hard-cutoff grooming to the physics of hadronization. A deep understanding of the physics of traditional jets requires careful consideration of hadronization corrections, such as in jet masses [96, 112, 134] and hadronic event shapes [111, 141]. Hard-cutoff groomed jet observables in particular can undergo large corrections from the hadronization process due to the discontinuity of the hard-cutoff paradigm: hadronization can transform partons with energy below a hard cutoff into hadrons with energy above the hard cutoff, and vice versa. These subtleties lead to additional complications in the theoretical calculation of non-perturbative effects on traditionally groomed jet observables [96].

In figure 7, we compare how ungroomed jets, traditionally groomed jets, and PIRANHA-groomed jets respond to hadronization in the context of energy flow, extensive properties, and substructure. Our studies focus on the parton-hadron EMD (the EMD between parton-level jets and their hadron-level counterparts), the parton-hadron Δp_T, and the parton-hadron $\Delta C_1^{(2)}$. PIRANHA-groomed jets typically exhibit smaller and more predictable responses to hadronization, evinced by the reduced variance in their hadronization response. One key factor in the improved responses of PIRANHA-groomed jets is the scaling-down of distortions due to hadronization in PIRANHA-groomed jets: the subtractive approach of PIRANHA (and in particular P-RSF when $f_{\text{soft}} \neq 0, 1$) removes energy from every particle in the event, resulting in a controlled and uniform hadronization response. On the other hand, Soft Drop grooming does not implement continuous, event-wide subtraction. Jets groomed with Soft Drop either retain or remove particles as a result of distortions due to hadronization, resulting in a larger variance in the properties of jets groomed with Soft Drop.

We plot distributions of the parton-hadron EMD for groomed QCD jets for the benchmark value of $z_{\text{cut}} = 0.1$ in figure 7a. We see a sharper peak at a smaller EMD in the parton-hadron EMD distributions for P-RSF$_{f=1/2}$ and longer tails for Soft Drop, already indicating that hadronization is more likely to dramatically change the result of Soft Drop. Figures 7b and 7c show that the average parton-hadron EMD and the variance in the parton-hadron EMD are both smaller for P-RSF$_{f=1/2}$ groomed jets than for Soft Drop groomed jets for a wide range of z_{cut} values, again evincing that PIRANHA-groomed jets have less sensitive responses to hadronization.

The extensive observable p_T and the substructure observable $C_1^{(2)}$ also reflect the increased robustness of PIRANHA-groomed jets to hadronization. In our discussion of hadronization, we characterize the response of p_T by using eq. (4.1) with $p_T = p_T^{(\text{parton})}$ indicating the parton-level transverse momentum, before the addition of model-dependent hadronization effects, and $\widetilde{p_T} = p_T^{(\text{hadron})}$. In figure 7d, we see that for $z_{\text{cut}} = 0.1$, PIRANHA-groomed p_T again tends to have a sharper response to hadronization, while figures 7e and 7f indicates that this sharper response holds for a wide range of z_{cut} values. The behavior of the p_T closely mimics that of the parton-hadron EMD, evincing our arguments that the parton-hadron EMD provides a probe of generic jet observables to hadronization. In
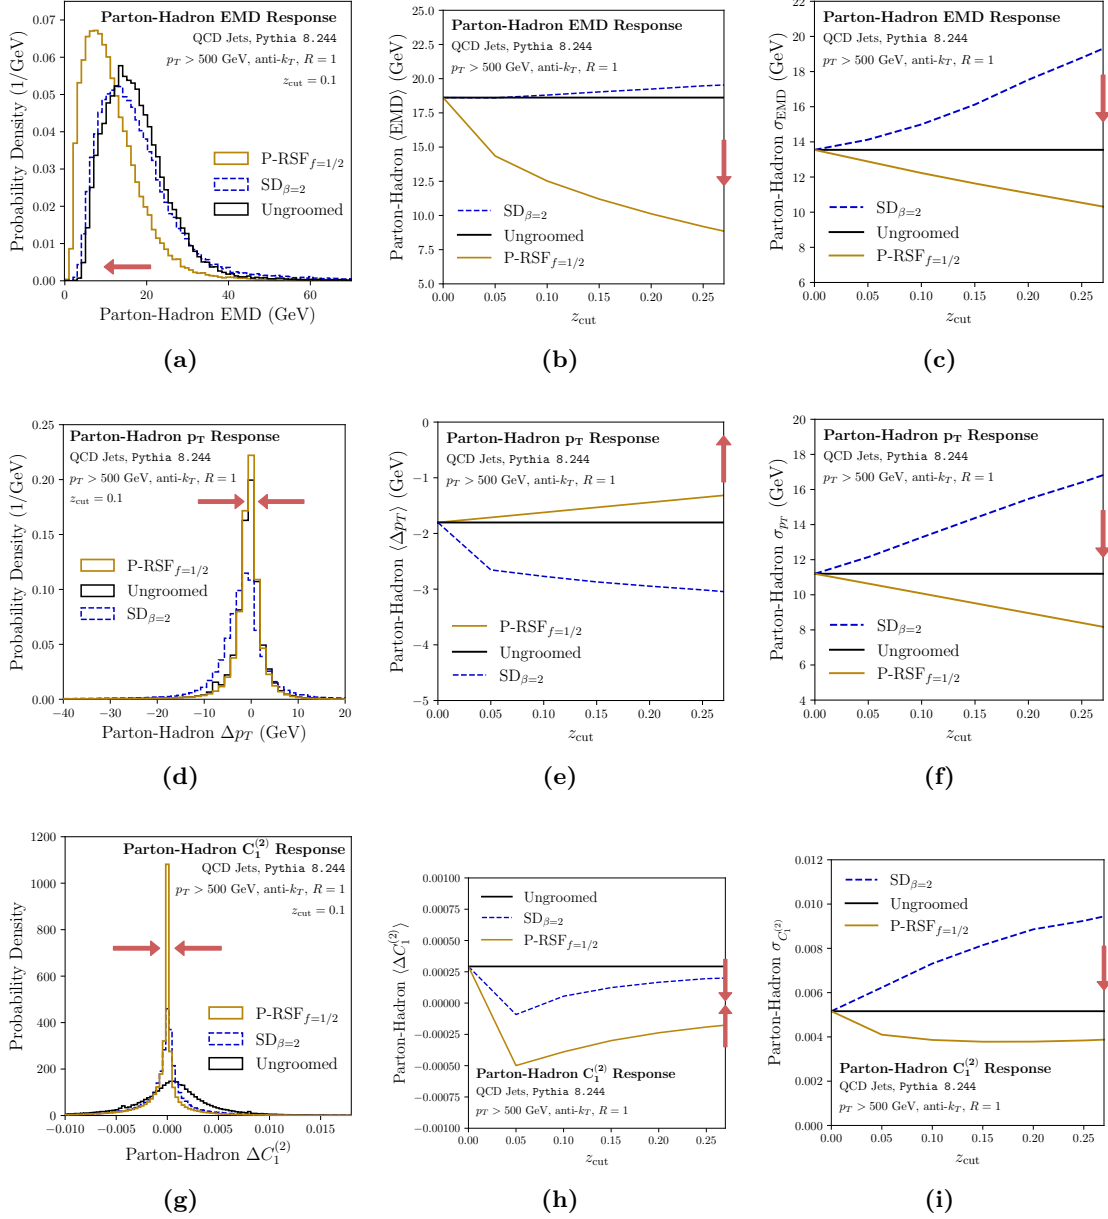
Figure 7. Per-jet hadronization responses of (top row) EMD, (middle row) Δp_T, and (bottom row) $\Delta C_1^{(2)}$ using Balanced Recursive Subtraction (P-RSF$_{f=1/2}$, orange) and Soft Drop with $\beta_{\mathrm{SD}} = 2$ (SD$_{\beta=2}$, blue). We display (left column) the distribution of the response for z_{cut}=0.1, (middle column) the mean response as a function of z_{cut}, and (right column) the standard deviation of the response as a function of z_{cut}. The red arrows indicate the direction corresponding to better performance.

figures [7g], [7h] and [7i], we make similar conclusions for $C_1^{(2)} \approx m^2/p_T^2$. We first note that the parton-hadron shift in the $C_1^{(2)}$ substructure observable is larger for P-RSF$_{f=1/2}$, at least for fixed z_{cut}, and the substructure of P-RSF$_{f=1/2}$ groomed jets may undergo larger changes due to hadronization than that of Soft Drop groomed jets. However, the variance in the shift in P-RSF$_{f=1/2}$ groomed substructure is smaller than that of Soft Drop, indicating that the response of PIRANHA-groomed substructure to hadronization is more predictable than that of Soft Drop.

Overall, we find that PIRANHA-groomed jets have more predictable responses to hadronization than traditionally groomed jets, supported by the observable independent EMD, as well as by extensive observables and jet substructure. We have provided evidence that IRC-safe observables of jets groomed with Soft Drop have generically larger responses to hadronization, while PIRANHA-groomed observables are generally less affected by the physics of hadronization. We hope that the stability of PIRANHA-groomed jets to the non-perturbative physics of hadronization may facilitate even further communication between theoretical predictions and experimental results for groomed jet observables.

4.3 All versus charged

We next compare the response of hard-cutoff and PIRANHA grooming to the exclusion of neutral particles, as a rough analog of the smearing of neutral particles due to detector effects. Small tweaks in experimental signatures due to detector effects may produce large changes in hard-cutoff groomed observables, providing another obstacle in extracting fundamental physics from traditionally groomed jets. The continuity of PIRANHA again leads us to expect that PIRANHA-groomed observables are more robust to detector effects than traditionally groomed jet observables.

To counteract the leading bias due to the loss of neutral energy in the jet, we apply a rescaling to the p_T of the charged-only jet constituents:

$$
\tilde{p}_T^i = \begin{cases} 0, & \text{particle } i \text{ is neutral} \\ p_T^i \dfrac{\left\langle p_T^{(\text{all})} \right\rangle}{\left\langle p_T^{(\text{charged})} \right\rangle}, & \text{particle } i \text{ is charged} \end{cases}, \tag{4.3}
$$

where p_T^i is the transverse momentum of particle i in the original jet, $\mathcal{J}$, containing both charged and neutral particles, and $\tilde{p}_T^i$ is the transverse momentum we assign to particle i in the charged-only, rescaled jet, $\tilde{\mathcal{J}}$, with neutral particles excluded. Here, $p_T^{(\text{all})} = \sum_i p_T^i$ is the total transverse momentum including both charged and neutral particles, while $p_T^{(\text{charged})} = \sum_{i \text{ charged}} p_T^i$ includes only charged constituents. The rescaling factor $\left\langle p_T^{(\text{all})} \right\rangle / \left\langle p_T^{(\text{charged})} \right\rangle \approx 1.47$ is defined such that the average p_T is preserved after removing neutral particles.[19]

---

[19]We can gain intuition for this rescaling factor by considering the isospin-preserving limit, where we expect similar numbers of π^+, π^-, and π^0 mesons to be produced. Events for which we discard the π^0 particles should have roughly 2/3 the total transverse momentum, ignoring subtleties associated with kaons and other heavier states. This leads to an estimate of $\left\langle p_T^{(\text{all})} \right\rangle / \left\langle p_T^{(\text{charged})} \right\rangle \approx 3/2$ — remarkably close to the numerical value of 1.47 found in Pythia 8.244.
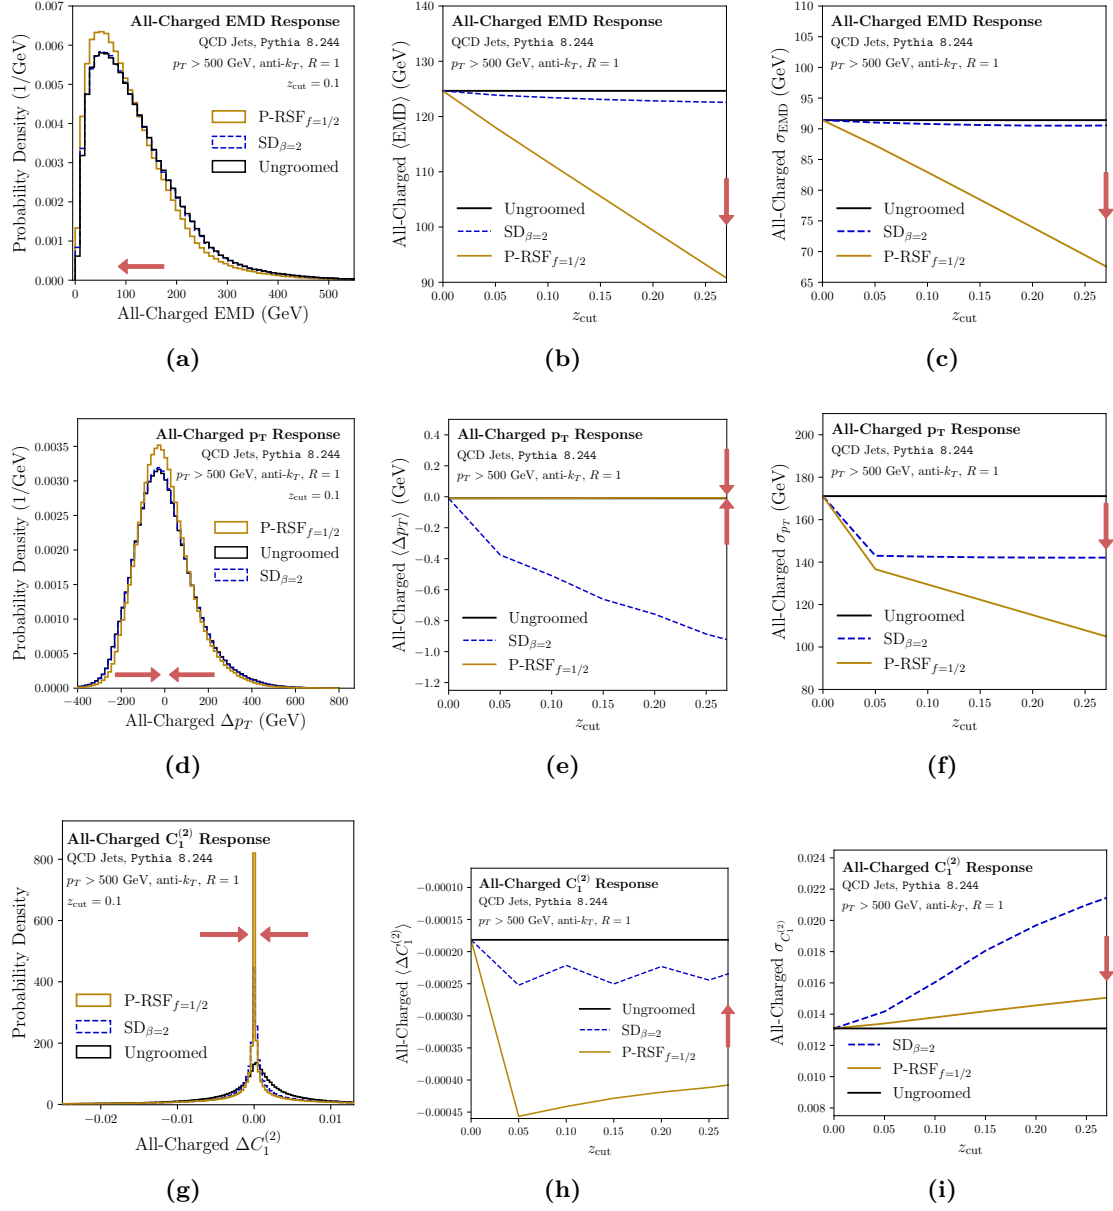
Figure 8. Same as figure 7, but for per-jet responses to the exclusion of neutral particles. We rescale the p_T of the charged particles in order to eliminate the leading p_T bias due to the loss of neutral energy.

In our discussion of EMD, we compute the response $\text{EMD}\left(G(\mathcal{J}), G(\tilde{\mathcal{J}})\right)$, where G indicates the grooming algorithm under consideration; we call this the *all-charged EMD*. Similarly, in our discussion of eq. (4.1), we take $p_T = p_T^{(\text{all})}$ to be the transverse momentum including both charged and neutral particles and $\widetilde{p_T} = \sum_i \tilde{p}_T^i = p_T^{(\text{charged, rescaled})}$ to be the rescaled contribution from charged particles only. We refer to the difference as the all-charged p_T shift, or the all-charged Δp_T. Note that the rescaling procedure does not impact dimensionless charged-only substructure observables like $C_1^{(2)}$, since they are normalized by the jet momentum.

Our comparison of the all-charged response of PIRANHA grooming to that of traditional grooming procedures is shown in figure 8. As in our study of hadronization, we begin our study with a discussion of the EMD. The all-charged EMD bounds changes in IRC-safe observables due to the exclusion of neutral particles and corresponding rescaling of charged particles. Since neutral particles are the most susceptible to smearing effects due to detector responses, the all-charged EMD furnishes an observable-independent probe for the effects of experimental detectors on groomed jets. We demonstrate distributions of the all-charged EMD for groomed QCD jets with several choices of grooming in figure 8a. The contrast between PIRANHA and traditionally groomed all-charged EMD is not as sharp as for the parton-hadron EMD. However, PIRANHA grooming again enjoys smaller and more sharply peaked EMD responses to the exclusion of neutral particles. Figure 8b shows that the all-charged EMD is smaller for P-RSF$_{1/2}$ groomed jets than for Soft Drop groomed jets for a wide range of z_{cut} values. Figure 8c shows that the fluctuations in the all-charged EMD, and therefore the fluctuations in the response of grooming to our naive model of smearing, are noticeably smaller for PIRANHA groomers than for traditional groomers.

Our results for the all-charged p_T shifts are similar to our results for the all-charged EMD. Figure 8d demonstrates that the p_T response of P-RSF$_{f=1/2}$ is smaller and more sharply peaked for $z_{\text{cut}} = 0.1$, while figures 8e and 8f respectively show that the P-RSF$_{f=1/2}$ all-charged response remains small and sharply peaked for a wide range of z_{cut} values. Indeed, the fluctuations of the all-charged p_T response are nearly identical for each PIRANHA groomer.

Figures 8g, 8h and 8i compare the robustness of P-RSF$_{f=1/2}$ and Soft Drop for the substructure observable $C_1^{(2)}$. Figure 8g again shows that the all-charged response of P-RSF$_{f=1/2}$ is more sharply peaked than that of Soft Drop. Similarly, figures 8h and 8i demonstrate that the all-charged shifts to PIRANHA-groomed jet substructure are generally larger but more predictable than those of Soft Drop for a wide range of z_{cut} values. The increased robustness of PIRANHA grooming procedures compared to Soft Drop is less sharp than in the case of hadronization, but the PIRANHA-groomed observables nonetheless have less spread and a stronger linear correlation in response to our naive model of smearing.

5 Responses of grooming to additive contamination

We examine the ability of PIRANHA grooming to mitigate the effects of additive contamination from pileup (PU) and the underlying event (UE), which both consist of *extra* soft radiation that sits on top of a hard process.

- We use P-RSF$_{f=1/2}$ as a representative of PIRANHA grooming procedures;

- We use SD$_{\beta=2}$ as a representative of traditional grooming procedures;

- In our PU studies, we use Constituent Subtraction (CS), a constituent-level area subtraction method for pileup mitigation [118], as a representative of PU mitigation procedures.

The case of additive contamination is qualitatively different than that of soft distortions; we do not want to examine the robustness of the grooming procedure to additive contamination, but rather the ability of the grooming to *remove* the contamination. In our PU mitigation studies, we find that P-RSF$_{f=1/2}$ — which behaves comparably to CS — is a more effective tool than SD$_{\beta=2}$. In our UE studies, we find that P-RSF$_{f=1/2}$ and SD$_{\beta=2}$ are comparable tools for the subtraction of UE effects from jet substructure.

We simulate PU using the dijet and minimum bias samples from ref. [103], produced using `Pythia 8.185` with tune 4C for proton-proton collisions at $\sqrt{s} = 14\,\text{TeV}$. We layer minimum bias events on top of the hard dijet events, taking the number of pileup events to be Poisson distributed with a mean of $\langle n_{\text{PU}} \rangle = 50$ events.

We simulate UE by turning on multiple parton interactions in `Pythia 8.244` [105] with the default 4C tune [140]. We begin with a study of QCD jets and, to roughly echo similar studies in the original Soft Drop paper [3], also explore the groomed mass resolution of jets produced by boosted W bosons and top quarks in the presence of UE.

5.1 Minimum bias events and pileup

We first examine the effectiveness of different groomers in mitigating the effects of PU on the dijet events of ref. [103]. To quantify the ability of different groomers to remove PU, we perform an event-by-event comparison of *hard events* to the associated *PU-subtracted events*:

- A **hard event** is a dijet event without any PU, representing the physics of a hard process;

- A **PU-subtracted event** is produced by simulating the additive contamination due to PU on top of a hard event and then attempting to groom the PU away.

To produce a PU-subtracted event, we first simulate PU by adding the energy distributions of a Poisson-distributed number of minimum bias events to the energy distribution of a given hard event. We then estimate the contaminating energy density due to PU using the `GridMedianBackgroundEstimator` (GMBE) method of `FastJet` [117]. We tune an additional correction factor that scales our estimate $\rho_{\text{est.}}$ relative to the GMBE value ρ_{GMBE} for each grooming algorithm to subtract pileup more effectively, as discussed in more detail in appendix B. Finally, we produce the PU-subtracted event by removing a fraction z_{cut} of the jet p_T consistent with the estimated pileup density:

$$z_{\text{cut}}\, p_T^{(\text{jet})} = \rho_{\text{est.}}\, A_{\text{jet}} \approx p_T^{(\text{PU in jet})}. \tag{5.1}$$
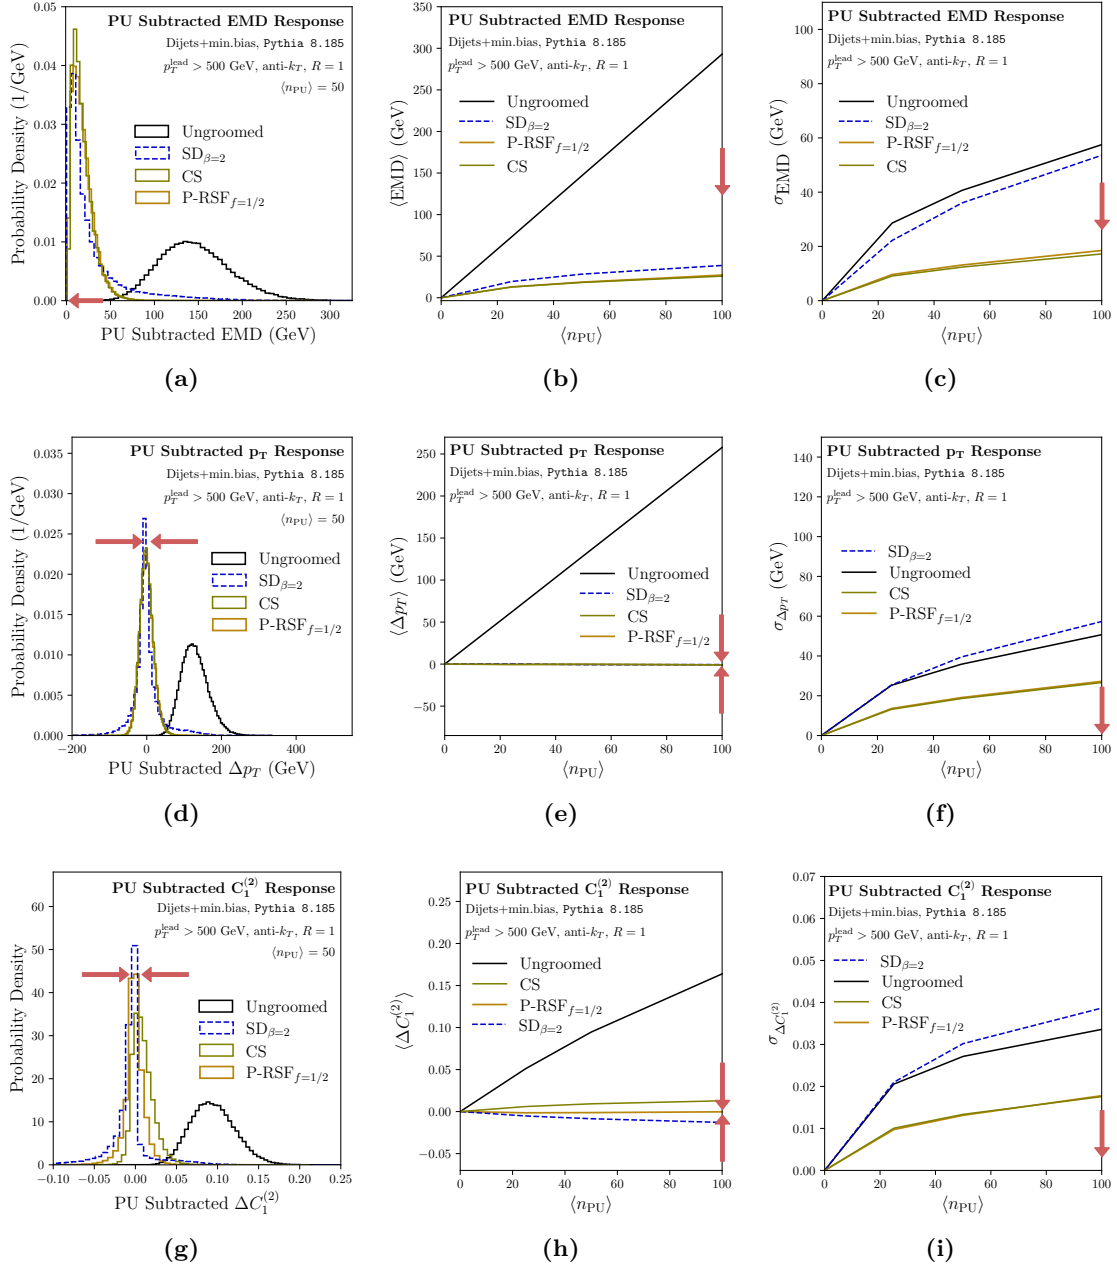
Figure 9. Properties of per-jet PU-subtracted (top row) EMD, (middle row) Δp_T, and (bottom row) $\Delta C_1^{(2)}$, discussed in section 5.1 for Balanced Recursive Subtraction (P-RSF$_{f=1/2}$, orange), Soft Drop with $\beta_{\mathrm{SD}} = 2$ (SD$_{\beta=2}$, blue), and Constituent Area Subtraction (CS, green). We display (left column) the distribution of the each PU-subtracted observable for $\langle n_{\mathrm{PU}} \rangle = 50$, (middle column) the mean PU-induced shift in each observable as a function of $\langle n_{\mathrm{PU}} \rangle$, and (right column) the standard deviation of the PU-induced shifts in each observable as a function of $\langle n_{\mathrm{PU}} \rangle$. The red arrows indicate the direction corresponding to better performance.

We examine the EMD, Δp_T, and $\Delta C_1^{(2)}$ between the *leading jets* of the hard and pileup-events in figure 9.[20] In figures 9a, 9b and 9c, we examine the EMD between the leading jet in the hard dijet event and the leading jet in the subtracted event with pileup. Figure 9a demonstrates that when $\langle n_{\mathrm{PU}} \rangle = 50$, P-RSF$_{f=1/2}$ tends to subtract pileup more accurately and predictably, evinced by a more sharply peaked EMD distribution and a smaller average PU-subtracted EMD than found for CS and Soft Drop. This strength of P-RSF$_{f=1/2}$ persists as the amount of pileup is increased, as suggested by figures 9b and 9c.

Figures 9d, 9e and 9f offer similar conclusions for the jet p_T. In our study of pileup responses, $p_T = p_T^{(\mathrm{hard})}$ indicates the transverse momentum of the hard event, before the layering of minimum bias/pileup events, and $\widetilde{p_T} = p_T^{(\mathrm{groomed\ PU})}$ indicates the transverse momentum after pileup has been added and then groomed away. RSF$_{1/2}$ again reproduces the extensive quantity p_T of the jet in the absence of PU for a wide range of $\langle n_{\mathrm{PU}} \rangle$ values. Identical conclusions hold for the substructure of the jet, as suggested by the behavior of $C_1^{(2)}$ shown in figures 9g, 9h and 9i.

Figures 15 and 16 of appendix B evince further that with our tuned procedure for PU mitigation, PIRANHA groomers and CS perform more reliably than traditional groomers in the removal of pileup. Furthermore, P-RSF algorithms are tree-based and offer orders-of-magnitude faster PU mitigation over CS, P-AS, and P-IVS (see figure 17 of appendix B).

5.2 Underlying event

We next examine the ability of grooming to eliminate additive contamination from UE. Since we model UE by generating events in the presence of multiple parton interactions in `Pythia 8.244`, it is less straightforward to study the effects of UE on an event-by-event basis. We instead work at the level of *distributions*: we tune both P-RSF$_{f=1/2}$ and Soft Drop so that, when acting on events with UE, they most accurately reproduce the substructure distribution of events without UE. We focus on substructure alone in our discussion of UE correction because we cut on jets with $p_T > 500$ GeV both for events with and without UE, and it is less meaningful to focus on the associated p_T distributions; we leave a more detailed study of the ability of PIRANHA groomers to remove UE at an event-by-event level for future work.

We characterize the ability of groomers to correct for the presence of UE by comparing *base distributions* to *UE-corrected* distributions:

- A **base distribution** is a substructure distribution with a fixed grooming parameter, $z_{\mathrm{cut}}^{(\mathrm{base})}$;

- A **UE-corrected distribution** is a substructure distribution in the presence of UE and with slightly more grooming tuned to soak up the additional energy contribution from UE: $z_{\mathrm{cut}}^{(\mathrm{corr.})} = z_{\mathrm{cut}}^{(\mathrm{base})} + \delta z_{\mathrm{cut}}$.

---

[20]In some situations, the pileup contributes enough energy that the leading jet after the addition of pileup corresponds to the direction of the sub-leading jet before the addition of pileup. This phenomenon occurs only for a small fraction of events, and we do not include these events in the following discussion.

We tune δz_{cut} by hand, for every $z_{\text{cut}}^{(\text{base})}$, to optimally correct for the presence of UE[21] by minimizing the *Wasserstein distance* between the two distributions. The Wasserstein distance is a metric on the space of distribution, much in the same way that the EMD is a metric on the space of energy flows, and provides a useful and quantitative measure of the ability of the grooming to correct for UE.

Figure 10 compares base distributions to UE-corrected distributions for P-RSF$_{f=1/2}$ and Soft Drop. Figures 10a and 10c examine the ability of the groomers to reproduce ungroomed base distributions, with $z_{\text{cut}}^{(\text{base})} = 0$. Figures 10b and 10d examine instead groomed base distributions with $z_{\text{cut}}^{(\text{base})} = 0.025$.

Figures 11a and 11b quantify the effectiveness and robustness of P-RSF$_{f=1/2}$ and Soft Drop as a function of $z_{\text{cut}}^{(\text{base})}$ by studying the un-corrected Wasserstein distance and the UE-corrected Wasserstein distance, respectively. Figure 11b shows the UE-corrected Wasserstein distance as a function of z_{cut}, and suggests that P-RSF$_{f=1/2}$ is better at correcting for UE when $z_{\text{cut}}^{(\text{base})} \lesssim .05$, and that SD$_{\beta=0}$ and P-RSF$_{f=1/2}$ have similar UE correction ability for values of $z_{\text{cut}}^{(\text{base})}$ between 0.05 and .1.

We note that we expect hard-cutoff groomers to be generically more *robust* to the presence of UE because even at fixed z_{cut}, they can remove an arbitrary amount of energy due to soft, wide-angle radiation in an event. We correspondingly expect that hard-cutoff groomers require less additional grooming to remove the effects of UE. PIRANHA groomers, on the other hand, cannot remove arbitrary amounts of energy from an event; we expect that we must always slightly increase the strength of PIRANHA grooming to soak up additive contamination from UE. We demonstrate that hard-cutoff groomers indeed require less additional δz_{cut} for optimal UE correction in figure 18 of appendix B.3.

We also point out that Recursive Subtraction techniques that are more geared towards the removal of the soft and wide-angle radiation which is characteristic of UE may perform even better in UE correction. In hard-cutoff grooming, for example, Soft Drop with $\beta_{\text{SD}} > 0$ preferentially grooms wide-angle radiation, Soft Drop with positive β_{SD} is particularly well suited to remove UE (see, for example, figure 18 of appendix B.3). Perhaps P-RS grooming procedures that upgrade the parameter f_{soft} of P-RSF into a function of the soft energy fraction z and the splitting angle θ of a branch (such as Hard-Balanced P-RSF, as in section 3.4) may be more suited to the removal of soft wide-angle radiation from UE.

We have argued that P-RSF$_{f=1/2}$ and SD$_{\beta=2}$ are comparable tools for the subtraction of UE effects in jet substructure, though P-RSF$_{f=1/2}$ requires more tuning to achieve the same level of UE correction. We emphasize that further investigation of such P-RS algorithms might improve the ability of PIRANHA in the tagging of boosted objects and beyond.

---

[21]We note that the optimized δz_{cut} is energy-dependent and process-dependent. For example, we expect that δz_{cut} should be parametrically smaller for processes at higher energies since the energy removed from a jet scales as $Q z_{\text{cut}}$, where Q is the energy scale of the hard process, while the energy scales associated with UE are only weakly correlated with the energy of the hard process (see, for example, figure 3 of ref. [142] or section 7.2 of ref. [143]).
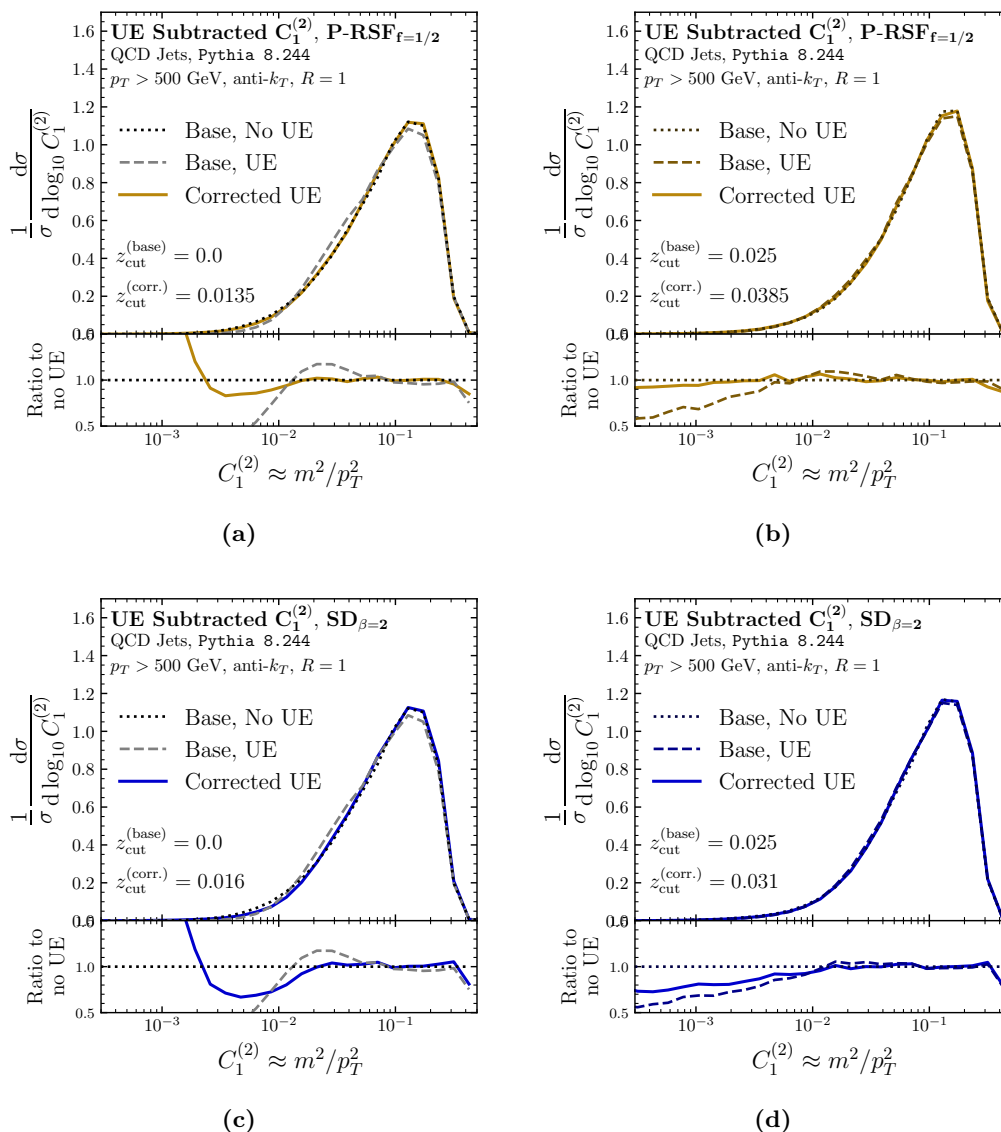
(a)

(b)

(c)

(d)

Figure 10. The response of groomed substructure distributions of QCD jets to the underlying event (UE), using P-RSF$_{f=1/2}$ (top row) and Soft Drop with $\beta_{\rm SD} = 2$ (bottom row). We examine the robustness of groomed substructure to the presence of UE by comparing "Base" substructure distributions with and without UE for fixed $z_{\rm cut} = z_{\rm cut}^{\rm (base)}$. We also examine the possibility of adding *additional* grooming to remove the effects of UE, picking a $z_{\rm cut}^{\rm (corr.)}$ that removes UE to most accurately reproduce substructure distributions using $z_{\rm cut} = z_{\rm cut}^{\rm (base)}$ in the absence of UE. The left column shows the substructure distributions for $z_{\rm cut}^{\rm (base)} = 0$, while the right column shows the substructure distributions for $z_{\rm cut}^{\rm (base)} = 0.025$.
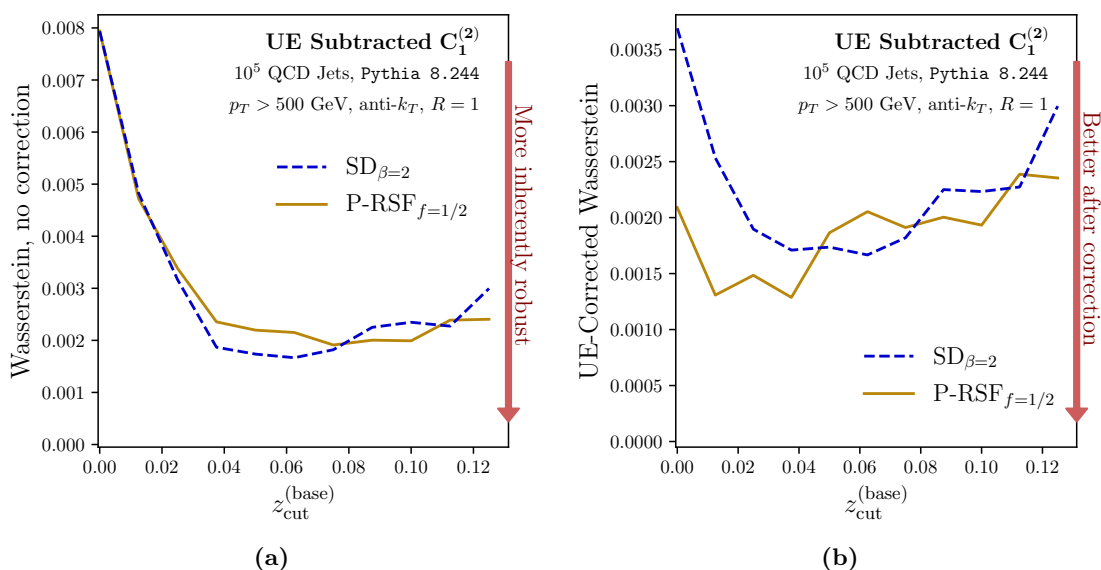
Figure 11. The behavior of the Wasserstein distance, in arbitrary units, between the $C_1^{(2)}$ distribution: (a) before UE correction, and (b) after UE correction, for $\mathrm{SD}_{\beta=2}$ and P-RSF$_{f=1/2}$ as a function of the base amount of grooming $z_{\mathrm{cut}}^{(\mathrm{base})}$. In (b), a lower UE-corrected Wasserstein distance quantitatively indicates better correction against the effects of UE: the substructure distribution of UE-corrected jets more closely resembles the substructure distribution of jets without UE.

5.3   Case study: mass resolution for *W* and top jets

As a final case study, we explore the use of PIRANHA grooming strategies in the tagging of boosted objects, another useful application of traditional grooming strategies such as Soft Drop. We must once again overcome the subtlety that PIRANHA grooming strategies remove energy proportional to the grooming parameter (such as z_{cut} in Recursive Safe Subtraction). To successfully tag boosted objects decaying into jets, we must tune the grooming parameter for PIRANHA strategies more than we would for Soft Drop. Nonetheless, we examine the mass resolution of PIRANHA-groomed jets from the decays of boosted *W* bosons and top quarks and discover that PIRANHA may offer some advantages that are complementary to those of traditional grooming strategies. The tagging procedure we implement in the following discussion is quite simple, however, and the applications of PIRANHA in boosted object tagging should be subjected to a more complete analysis.

In figure 12, we assess the behavior of PIRANHA in tagging boosted objects. We focus on jets of $p_T \geq 500$ GeV produced from the decays of *W* bosons and top quarks in figures 12a and 12c, respectively. We then plot the associated background rejection as a function of signal efficiency in figures 12b and 12d.

To produce the groomed mass distributions of figures 12a and 12c, we groom jets in the presence of UE. Afterward, we rescale the mass of the jet to correct for the shift of the mass peak due to the removal of energy from the jet during grooming. This rescaling procedure is necessary to reproduce the correct mass peak for P-RSF$_{f=1/2}$; however, rescaling Soft
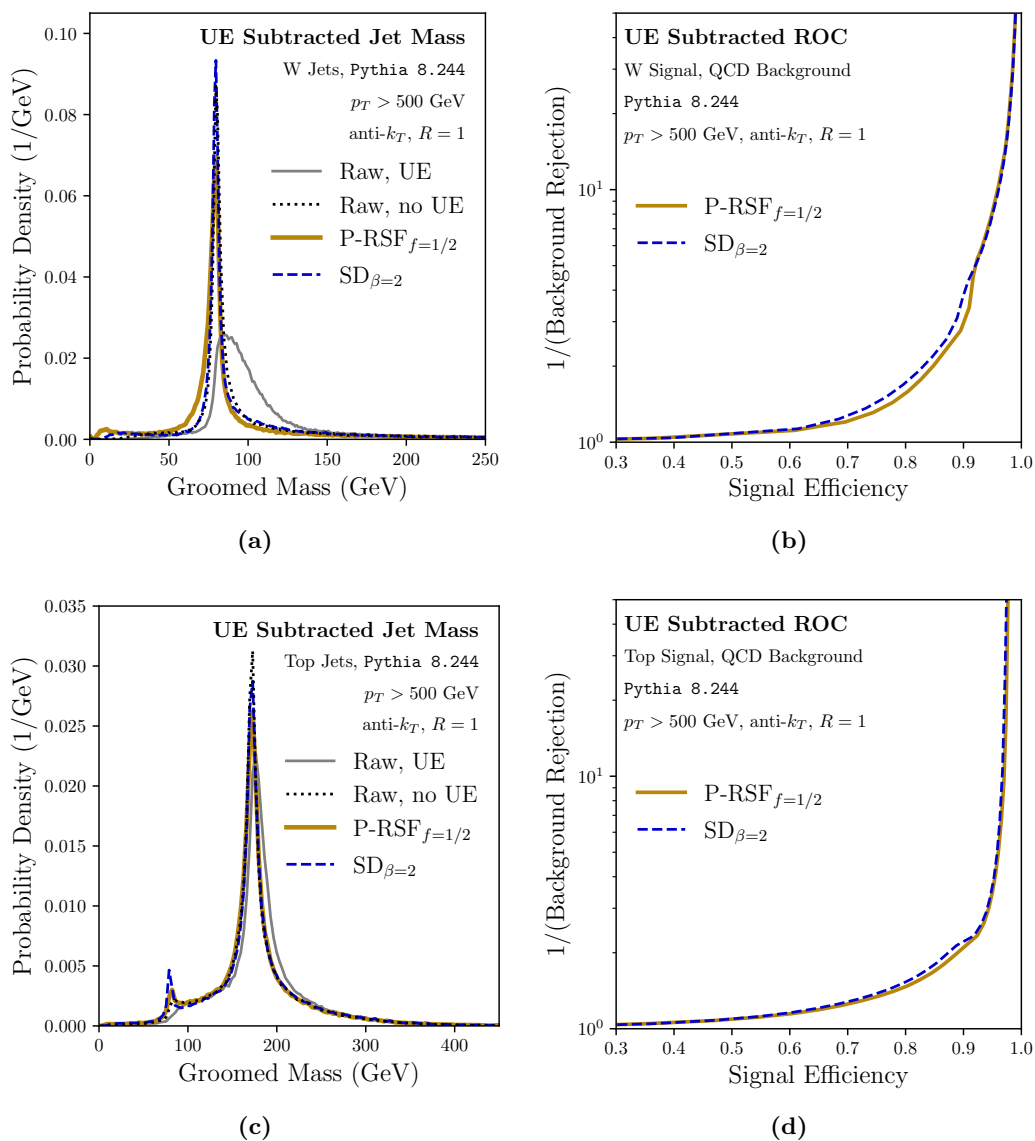
Figure 12. A comparison of traditional and PIRANHA grooming procedures in the tagging of (top row) boosted W bosons and (bottom row) top quarks. The data in the plots correspond to 100,000 jet events with $p_T > 500$ GeV in `Pythia 8.244`. The left column displays the mass distributions of W and top jets with and without UE, and with UE groomed away using either Soft Drop with $\beta_{\mathrm{SD}} = 2$ or P-RSF$_{f=1/2}$ with additional post-processing discussed in the text. The right column displays the relationship between background rejection and signal efficiency associated with W and top tagging, respectively, produced by considering symmetric windows around the mass of the relevant boosted object and comparing with a QCD background.

Drop groomed mass distributions has a minimal effect since Soft Drop tends to remove wide-angle soft radiation without severely affecting radiation from the hard event.

Figures 12a and 12c show rescaled mass distributions for each grooming algorithm for the value of z_{cut} that most precisely removes UE, in the sense that it most closely reproduces the mass distribution of jets without UE. We find that the best performance is achieved by $z_{\text{cut}}^{(\text{P-RSF}_{1/2})} = 0.034$ and $z_{\text{cut}}^{(\text{SD}_2)} = 0.025$ when grooming W jets, and $z_{\text{cut}}^{(\text{P-RSF}_{1/2})} = 0.023$ and $z_{\text{cut}}^{(\text{SD}_2)} = 0.035$ when grooming top jets. In each case, Soft Drop does not require any additional rescaling, while we find the best performance if we rescale the energy of P-RSF$_{f=1/2}$ groomed jet constituents by about $+4\%$.

In this simple context, P-RSF$_{f=1/2}$ offers slightly better background rejection for reasonable signal efficiencies. We evince this in figures 12b and 12d, where we show the signal efficiency versus background rejection curves associated with comparing signal W or top jets to QCD background, again with $p_T \geq 500$ GeV. Our tagging procedure is minimal: we accept events within a symmetric window of width δm around the mass of the W boson or top quark, respectively, and figures 12b and 12d are produced by varying the window size δm. For reasonable values of signal efficiency, between around 60% and 90%, the P-RSF$_{f=1/2}$ algorithm gives greater background rejection than Soft Drop.

Again, the results in figures 12a and 12c and figures 12b and 12d suggest that P-RSF$_{f=1/2}$ and SD$_{\beta=2}$ behave comparably as tools to remove the effects of UE in jet mass distributions, and to tag underlying boosted object. These results are preliminary, and to state these conclusions with more confidence will require more detailed tagging procedures than the simple algorithm described above. We leave a more detailed study of the tagging potential of P-RSF$_{f=1/2}$ and PIRANHA to future work.

6   Discussion and conclusions

In this work, we proposed the paradigm of Pileup and Infrared Radiation Annihilation (PIRANHA) for continuous jet grooming. Motivated by optimal transport theory and the Energy Mover's Distance (EMD) of ref. [95], we re-framed the Apollonius Subtraction (P-AS) and Iterated Voronoi Subtraction (P-IVS) procedures of ref. [102] as implementations of the PIRANHA paradigm. We additionally introduced Recursive Subtraction with a Fraction (P-RSF) as a family of new groomers motivated by the PIRANHA paradigm. We showed that a particular Recursive Subtractor, P-RSF$_{f=1/2}$, overcomes the soft discontinuities of traditional hard-cutoff grooming procedures, and that general P-RSF algorithms only have soft discontinuities in suppressed regions of phase space. We highlighted the unprecedented robustness of P-RSF$_{f=1/2}$ to hadronization, detector effects, and pileup. We showed also that P-RSF$_{f=1/2}$ may be able to correct for the presence of the underlying event. Though hard-cutoff groomers may be more robust against effects from UE without additional tuning, PIRANHA groomers can be tuned to remove additive contamination from the underlying event. We used the example of additive contamination from UE to argue that PIRANHA may also have applications in the tagging of boosted objects.

There are several immediately evident avenues for future phenomenological and theoretical exploration. While we argued that P-RSF had more robust responses to soft

distortions, it will be interesting to quantify this including theoretical, model-dependent uncertainties. For example, it will be interesting to study the robustness of P-RSF$_{f=1/2}$ to hadronization when using different models of color recombination, with different `Pythia` tunes, and from the perspective of effective field theory as in ref. [96]. Similarly, the effects of experimental detectors on P-RSF$_{f=1/2}$ groomed quantities must be explored with more realistic models of detector responses and compared in more detail with traditional grooming techniques. More detailed studies of these model-dependent uncertainties may facilitate a more precise, and even process-dependent method to tune the parameters of PIRANHA groomers such as P-RSF to the removal of PU and UE in more realistic scenarios.

Designing and studying variants of P-RSF in which the amount of grooming for a particular emission depends on the angle of the emission, as in Soft Drop with $\beta_{\mathrm{SD}} \neq 0$, and even the energy fraction, may be another easy way to improve the robustness and precision of P-RS algorithms in removing specific models of soft contamination. Indeed, any method that varies f_{soft}, z_{cut}, ρ, or other parameters of the jet grooming procedure in a way that depends on more detailed information within the jet could be useful in more precise removal of contaminating radiation, such as the removal of PU, UE, or the thermal noise that contaminates jets originating in heavy ion collisions.

A first-principles calculation of Recursive Subtraction groomed jet observables will rely on analytic exploration of the elaborately correlated emissions of P-RS groomed jets, as well as contributions from non-global configurations. We begin this journey in appendix C, but further exploration may depend on a more precise application of existing tools for jet substructure or even on newer tools in perturbative QCD, such as the techniques of ref. [144].

The study of EMD-mode PIRANHA grooming, introduced in appendix A, is another interesting avenue for the development of continuous grooming techniques. Unlike the PIRANHA groomers explored in the main text, EMD-mode PIRANHA groomers have the potential to remove an arbitrary amount of p_T from an event. We look forward to future studies of EMD-mode grooming for its potential to address stochastic fluctuations in the levels of jet contamination.

Finally, the use of alternative methods for optimal transport, such as the EMD without the restriction of $\beta = 1$ that we chose in this work, or the flexible and computationally efficient formalism of linearized optimal transport [119–122], may offer new tools and additional insights into continuous grooming.

The PIRANHA grooming strategy has intricate geometric origins, but its goal is simple: the optimal and continuous removal of contaminating low-energy radiation. We hope that the simple strategies and examples of continuous grooming discussed in this work may provide and inspire new tools for clear communication between experimental results and theoretical predictions regarding our microscopic universe.

Acknowledgments

A Grooming in EMD mode

In this appendix, we describe another route for developing Piranha groomers: *EMD mode.* EMD-mode grooming produces a groomed event with a fixed EMD relative to its ungroomed counterpart, and furnishes a complementary approach to "p_T-mode" groomers. We provide a conceptual introduction to EMD-mode grooming below, leaving an exploration of the phenomenological implications of EMD mode to future work.

The Piranha groomers described above all subtract a fixed amount of transverse momentum from the energy flow of an event. Let us call them p_T-*mode groomers.* p_T-mode grooming gives us a great deal of control over the amount of energy we remove from an event. When using Piranha groomers in p_T mode to remove additive contamination, however, this precise control over Δp_T leads to additional complications; since additive contamination may add arbitrary amounts of contaminating soft radiation to an event, removing a fixed amount of p_T from the contaminated event may not be the best strategy for reproducing the un-contaminated event. These complications lead to important considerations when using Piranha for pileup mitigation, as in section 5.1 and appendix B.2, and they are especially disadvantageous when using Piranha to correct for the presence of the underlying event, as in section 5.2 and appendix B.3.

Grooming in EMD mode is a complementary approach that allows Piranha groomers to remove arbitrary amounts of soft radiation from an event. Piranha in EMD mode therefore furnishes a conceptually interesting alternative for continuous grooming.

First, let us briefly review how the p_T-mode algorithms we have presented for P-AS, P-IVS, and P-RS all subtract a fixed amount of transverse momentum, $\Delta p_T = \rho A_{\mathrm{tot}}$, or $\Delta p_T = z_{\mathrm{cut}} p_T$, from the event under consideration.

- P-AS subtracts the pre-specified p_T all at once by finding the Apollonius regions associated with the event, as in section 2.3.

- P-IVS subtracts the pre-specified p_T by subtracting p_T from each particle in the event until it removes a particle, and then continues recursively until the specified p_T has been removed, as in section 2.4.

- P-RS subtracts the pre-specified p_T by recursively assigning fractions of the total p_T to be removed to each branch of the jet tree until it reaches the final-state particles of the jet, as in section 2.5.

The approach for EMD-mode grooming is quite similar. In EMD mode, however, each algorithm continues until a specific amount of EMD has been subtracted from the event, often in an iterative procedure. As a concrete example, let us construct the algorithm for P-IVS in EMD mode, or P-IVS$^{(\text{EMD})}$:

1. P-IVS$^{(\text{EMD})}$ finds the Voronoi diagram for the event, as discussed in section 2.4, and indexes the steps of the algorithm by n, starting at $n = 1$. We define the energy flow of the algorithm after step n is completed by $\mathcal{E}_n$, where $\mathcal{E}_0$ is the energy flow of the original event. We also define the EMD between $\mathcal{E}_{n-1}$ and $\mathcal{E}_n$ as EMD_n.

2. P-IVS$^{(\text{EMD})}$ attempts to modify the p_T of every particle in the event in the same way as P-IVS in p_T-mode:
$$p_{Ti}^{(n)} = p_{Ti}^{(n-1)} - \rho^{(n)} A_i^{(n-1)}, \tag{A.1}$$
where the $p_{Ti}^{(n-1)}$ and $A^{(n-1)}$ are the transverse momenta/Voronoi areas after step $n - 1$, and
$$\rho^{(n)} = \min_i p_{Ti}^{(n)}/A_i^{(n)}, \tag{A.2}$$
in analogy with eq. (2.8). As in section 2.4, $\rho^{(n)}$ describes the maximum p_T that may be subtracted from an event before one of the particles in the event is groomed away entirely.

3. P-IVS$^{(\text{EMD})}$ next calculates EMD_n for this proposed groomed event, and asks whether it has subtracted a total EMD of EMD_{cut} from the event. If $\sum_{k=1}^n \text{EMD}_n < \text{EMD}_{\text{cut}}$, P-IVS$^{(\text{EMD})}$ has not yet subtracted the full EMD_{cut} from the event. In this case, we have more grooming to do: P-IVS$^{(\text{EMD})}$ continues recursively, going back to the first step of the algorithm.

4. If $\sum_{k=1}^n \text{EMD}_k > \text{EMD}_{\text{cut}}$, P-IVS$^{(\text{EMD})}$ has subtracted too much EMD from the event. In this case, P-IVS$^{(\text{EMD})}$ revises its proposed groomed event at this step in the algorithm by evaluating the value of $\rho^{(n)}$ that would give $\sum_{k=1}^n \text{EMD}_k > \text{EMD}_{\text{cut}}$, and finding the associated groomed event by using eq. (A.1). It returns the resulting event as the final groomed event.

As in P-IVS, P-IVS$^{(\text{EMD})}$ does not need to re-compute the Voronoi diagram from scratch at each step of the algorithm. Similarly, our procedure for finding the value of $\rho^{(n)}$ that subtracts the correct amount of total EMD in the final step of the algorithm is computationally efficient, and does not rely on a complete re-evaluation of the EMD at each step of the algorithm.

At present, we do not have an implementation of P-RS in EMD mode. That said, we expect that the generalization of EMD-mode grooming to P-RS and P-RSF will be similar: at step n of the EMD mode algorithm, we perform grooming with the maximum value of z_{cut} until a particle would be removed, and define EMD_n as the EMD between the energy flows at steps $n - 1$ and n of the grooming. If the sum of the subtracted EMDs is greater than EMD_{cut}, $\sum_{k=1}^n \text{EMD}_k > \text{EMD}_{\text{cut}}$, we must compute the value of z_{cut} that fixes $\text{EMD}_n > \text{EMD}_{\text{cut}} - \sum_{k=1}^{n-1} \text{EMD}_k$, groom with this value of z_{cut}, and return

the resulting groomed event. Otherwise, if $\sum_{k=1}^{n} \text{EMD}_k < \text{EMD}_{\text{cut}}$, we would continue grooming recursively.

We hope that future exploration of EMD-mode grooming may offer additional utility in the use of continuous grooming in the subtraction of additive contamination, and in the mitigation of obfuscating radiation in general.

B Feeding frenzy: comparing a plethora of grooming options

In this appendix, we present a brief collection of additional results regarding the responses of PIRANHA and Soft Drop groomers to soft distortions and additive contamination. We extend the results in the main text by comparing P-RSF$_{f=1/2}$, the focus of our phenomenological studies in the main text, to three groups of grooming algorithms:

- **Hard-Cutoff Groomers:** SD$_{\beta=0}$, SD$_{\beta=1}$, and SD$_{\beta=2}$;

- **Fully Continuous Groomers:** P-AS and P-IVS;

- **Recursive Subtractors:** P-RSF$_{f=0}$, P-RSF$_{f=1/2}$, P-RSF$_{f=3/4}$, and P-RSF$_{f=1}$.

In our pileup studies, we also count Constituent Subtraction (CS) [118] among the continuous groomers due to its continuity properties in the continuum limit. The comparisons of this appendix demonstrate that the conclusions we drew in the main text by comparing P-RSF$_{f=1/2}$ to SD$_{\beta=2}$ seem to hold in greater generality: PIRANHA groomers tend to have more robust responses to soft distortions and additive contamination than traditional grooming methods. Furthermore, the responses of P-RSF$_{f=1/2}$ are similar to the responses of the fully continuous groomers P-AS and P-IVS (and CS, in the case of pileup mitigation); each of the PIRANHA algorithms we explore in this appendix are available on GitHub [113].

As in the main text, we first extend our results involving hadronization corrections (section 4.2) and all-charged corrections (section 4.3) to explore how several choices of grooming and grooming parameters may respond to soft distortions. We then begin our discussion of additive contamination by providing more detail regarding our pileup studies: we explain in greater detail our procedure for pileup mitigation used in section 5.1 and extend our results to several groomers and values of $\langle n_{\text{PU}} \rangle$. Finally, we similarly explore how different groomers respond to the underlying event, extending the results of section 5.2.

B.1 Soft distortions

In the main text, we made the argument that continuity provides more predictable responses to soft distortions; we therefore focus on the fluctuations induced by soft distortions by studying the standard deviations of the groomed EMD and $C_1^{(2)}$ as a function of z_{cut}.

Figure 13 shows the fluctuations in the parton-hadron groomed EMD and the parton-hadron $\Delta C_1^{(2)}$, as discussed in section 4.2, for a variety of grooming options. The first column of figure 13 demonstrates that P-RSF$_{f=1/2}$ groomed jets tend to have significantly smaller fluctuations due to hadronization, echoing the conclusions of section 4.2. We note that SD$_{\beta=2}$ is significantly more robust to hadronization than SD$_{\beta=0}$ when considering extensive observables (EMD and p_T). The second two columns of figure 13 show that the
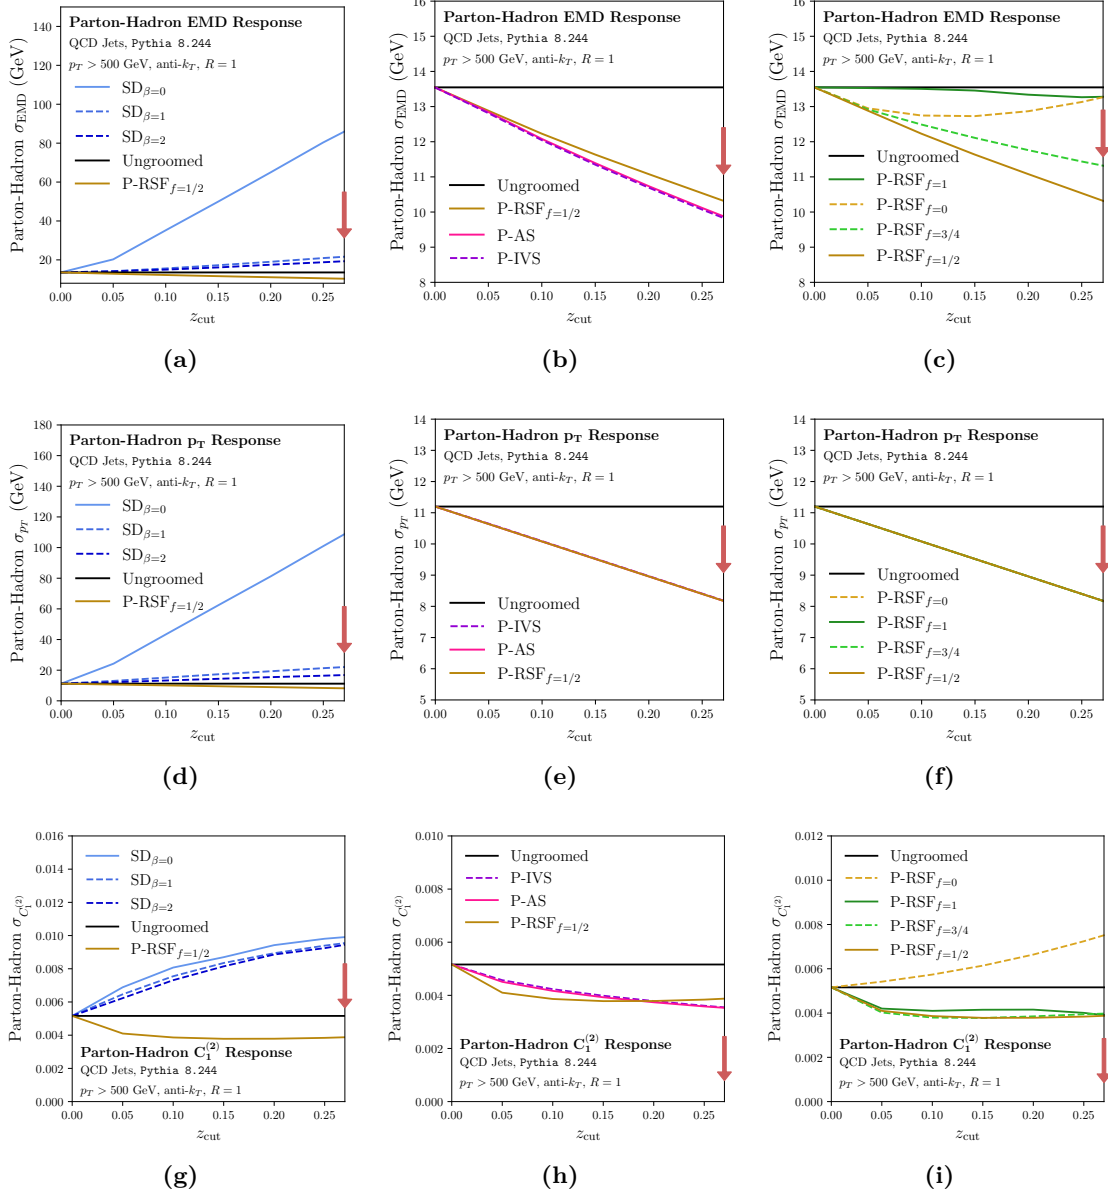
Figure 13. Per-jet hadronization responses of Piranha and traditional groomers for (top row) EMD, (middle row) Δp_T, and (bottom row) $C_1^{(2)}$; we compare P-RSF$_{f=1/2}$ to (left column) hard-cutoff groomers, (middle column) fully continuous groomers, and (right column) recursive subtractors. For brevity in this appendix, we focus on the variance of the shifts in each observable due to hadronization in groomed jets. The red arrows indicate the direction corresponding to better performance.
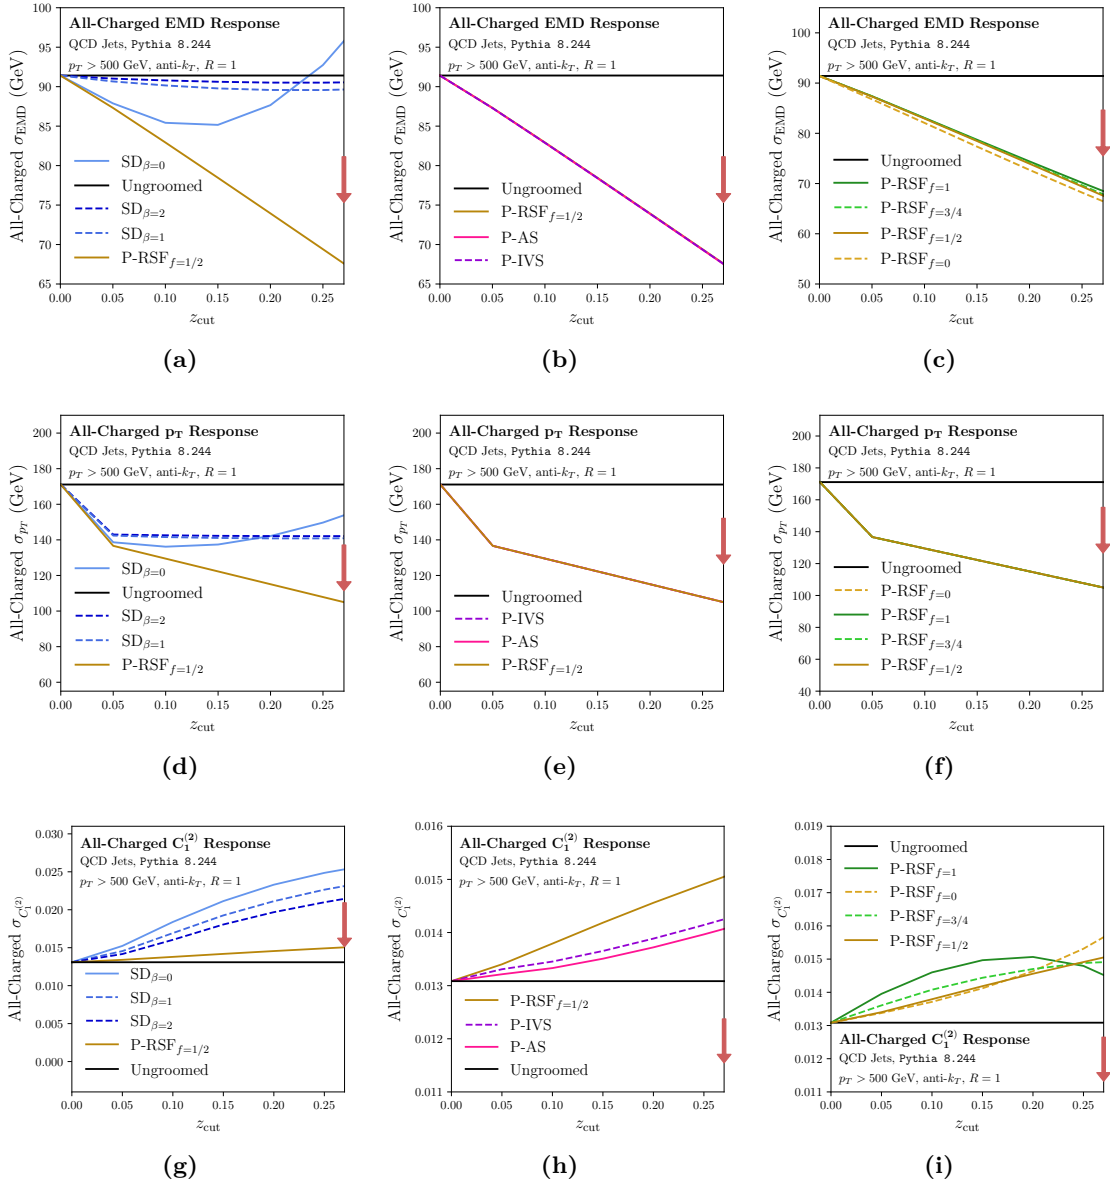
Figure 14. Same as figure 13, but for jet responses to the exclusion of neutral particles.

PIRANHA groomers we consider all have similar hadronization responses in the considered range of z_{cut}.

Figure 14 presents the all-charged groomed EMD and $\Delta C_1^{(2)}$ as discussed in section 4.3. The first column again shows that P-RSF$_{f=1/2}$ experiences significantly smaller fluctuations due to the exclusion of neutral particles than Soft Drop, and the second two columns show similar responses for all PIRANHA groomers.

B.2 Pileup

We now provide more detail regarding our PU mitigation procedure. As discussed in section 5.1, we use the dijet and minimum bias events of ref. [103] to study the effects of PU. We use the `GridMedianBackgroundEstimator` (GMBE) method of `FastJet` [117] to provide a rough estimate, ρ_{GMBE}, of the energy density of additive contamination due to pileup in each jet we would like to groom. However, since GMBE is designed to provide an estimate of the contaminating energy in an entire event,[22] we improve upon the GMBE procedure by optimizing a correction factor that scales the GMBE estimate for each groomer. We optimize the correction factor for each groomer by considering its action on 100,000 events, throwing away events where the presence of PU changes which jet within the dijet has more energy, and implementing the following procedure:

1. For each epoch i during our optimization, estimate the contaminating energy density due to PU within a jet,

$$\rho_{\text{PU}}[\mathcal{E}] = \frac{p_T^{(\text{PU in jet})}}{A_{\text{jet}}}, \tag{B.1}$$

   that needs to be applied by a groomer g to accurately remove PU contamination as

$$\rho_{\text{est.}}^{(\text{epoch } i)}[\mathcal{E}; g] = \rho_{\text{GMBE}}[\mathcal{E}] f_i(g), \tag{B.2}$$

   where we use parentheses to indicate that $\rho_{\text{est.}}$ and f are functions of the groomer g, and square brackets to indicate that the energy densities ρ are functionals of the energy flow $\mathcal{E}$. Since we are using jets clustered using the anti-k_t algorithm, we may take the jet area to be πR_{jet}^2 [124].

2. Introduce a new epoch every 1,000 events we consider, by adjusting our proposed correction factor f_i to

$$f_{i+1} = f_i + r_{\text{learn}} \left\langle \frac{\Delta \rho}{\rho} \right\rangle_i, \tag{B.3}$$

   where the angular brackets indicate an average over the jets considered during epoch i, $\Delta \rho$ indicates the difference in p_T density in a pileup subtracted event and a hard event,

$$\Delta \rho = \frac{\Delta p_T}{A_{\text{jet}}} = \frac{p_T^{(\text{groomed PU})} - p_T^{(\text{hard})}}{A_{\text{jet}}}, \tag{B.4}$$

   and r_{learn} is a learning rate which we set to $1/3$.

3. At the end of the final epoch, save the final correction factor $f_{\text{learned}}(g)$ for each groomer.

4. Perform a new analysis of the ability of each groomer to remove PU by estimating the optimal ρ required for PU-subtraction as

$$\rho_{\text{est.}}[\mathcal{E}; g] = \rho_{\text{GMBE}}[\mathcal{E}] f_{\text{learned}}(g). \tag{B.5}$$

   We recall for completeness that $z_{\text{cut}} = \rho_{\text{est.}} A_{\text{jet}} / p_T^{(\text{jet})}$.

---

[22]Note that `FastJet` also supports *Jet* Median Background Estimation (JMBE), which is more suited to the task of estimating the contaminating energy density in each jet. We were unable to find `Python` support for JMBE and used the method detailed here as a replacement.
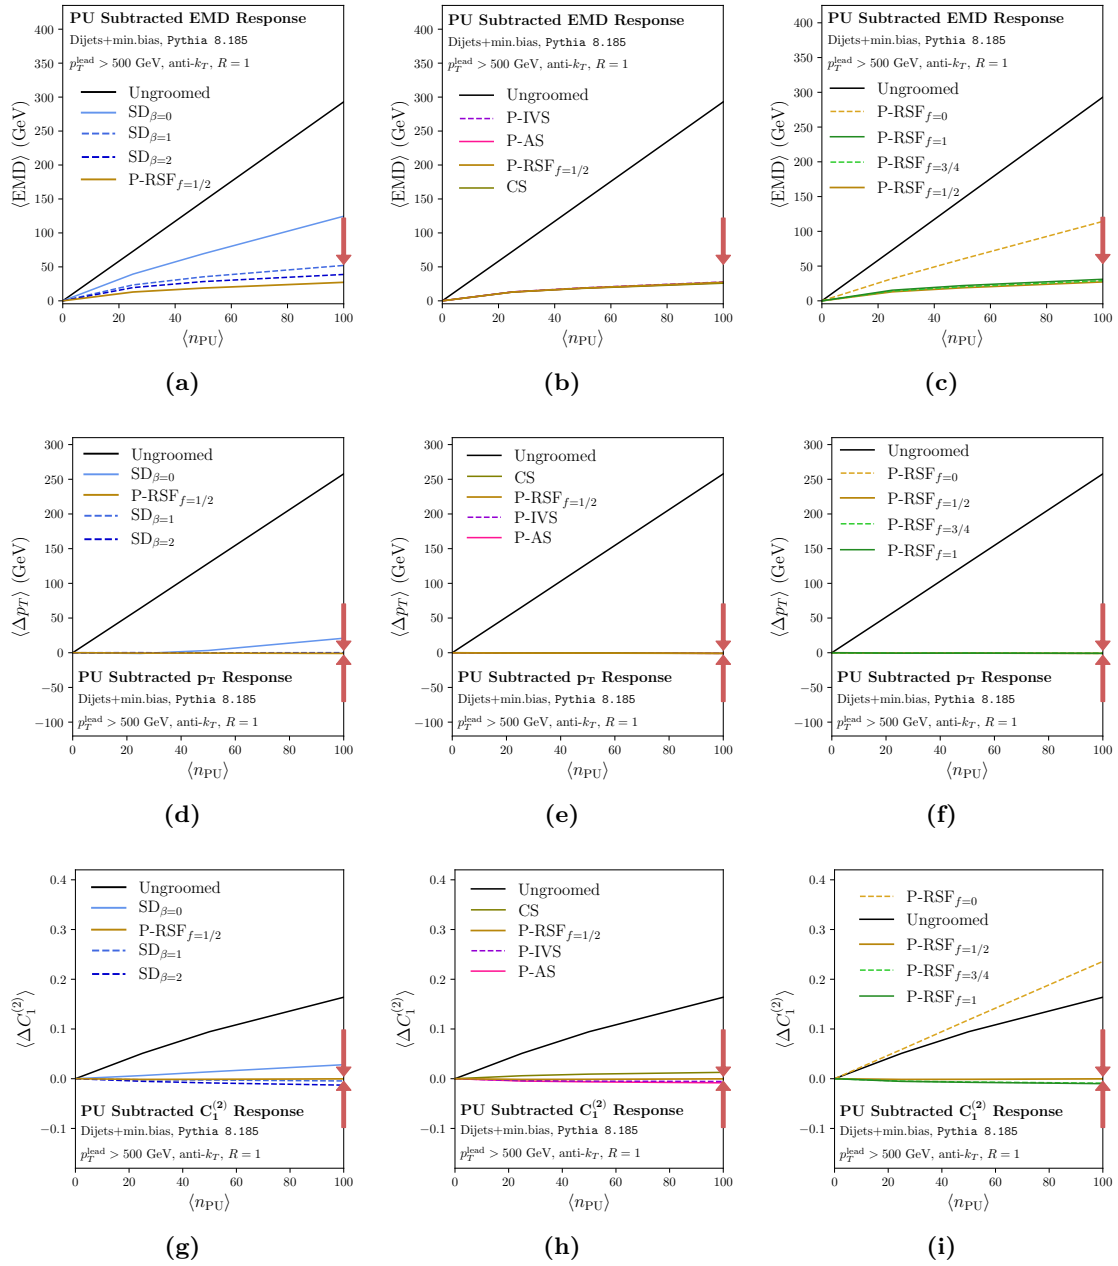
Figure 15. Average per-jet pileup-induced shifts in (top row) EMD, (middle row) Δp_T, and (bottom row) $C_1^{(2)}$; we compare P-RSF$_{f=1/2}$ to (left column) hard-cutoff groomers, (middle column) fully continuous groomers, and (right column) recursive subtractors, as discussed in section 5.1 and appendix B. The red arrows indicate the direction corresponding to better performance.

In figures 15 and 16, we present results on PU mitigation using the optimization procedure described above for a variety of groomers, as an extension of the discussion in section 5.1. We focus on the PU-subtracted EMD, the PU-subtracted Δp_T, and the PU-subtracted $\Delta C_1^{(2)}$ as a function of the average number of PU events $\langle n_{\rm PU} \rangle$.
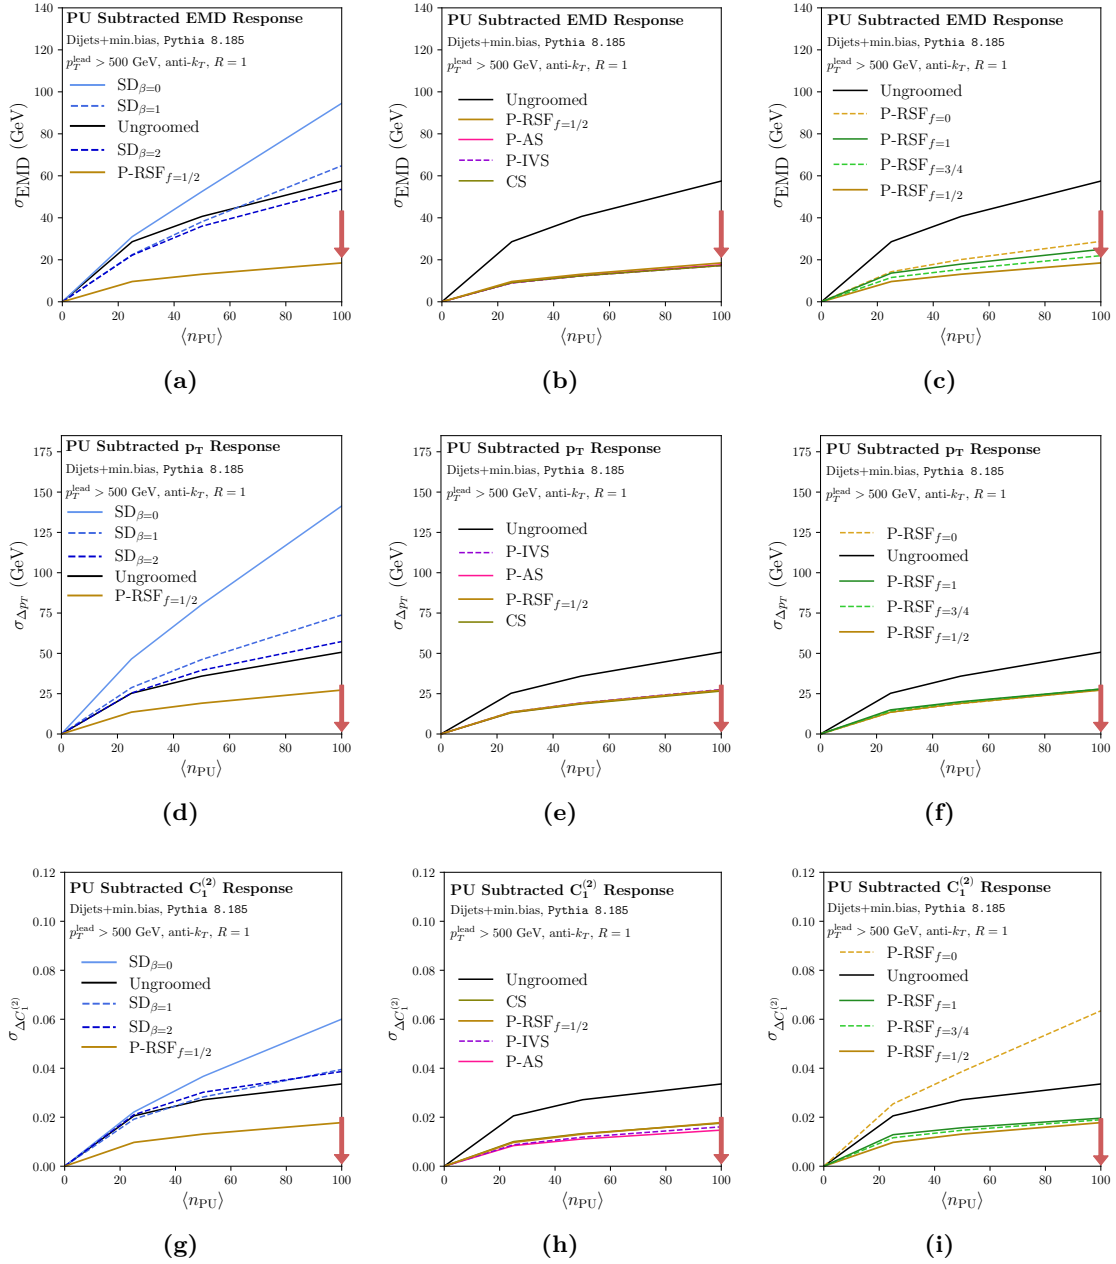
Figure 16. Same as figure 15, but for the standard deviation of the pileup-induced shifts.

The Piranha groomers displayed in figure 15, together with CS, tend to perform only slightly better than the traditional groomers, on average, in removing pileup; both Piranha and traditional groomers have been tuned so that $\langle \Delta p_T \rangle = 0$, and this tuning is reflected by the relatively small changes to the extensive properties and even the substructure of the PU-subtracted jets. The notable exception is P-RSF$_{f=0}$. Since P-RSF$_{f=0}$ preferentially grooms away hard radiation, it is poorly suited to grooming away the low-energy additive contamination that is contributed by pileup, and it performs more poorly
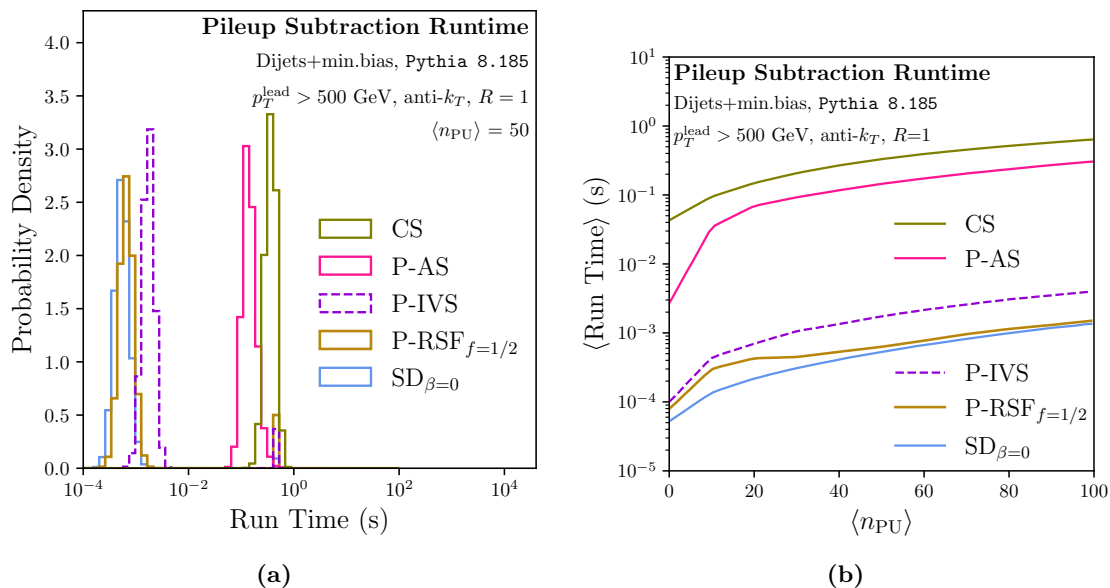
Figure 17. Runtimes for PIRANHA algorithms, and comparison to some traditional grooming and pileup mitigation algorithms. (a) shows the probability density for the logarithm of the runtime for each algorithm when $\langle n_{\mathrm{PU}} = 50 \rangle$, where n_{PU} is the average number of pileup events. (b) shows the average runtime for each algorithm as a function of $\langle n_{\mathrm{PU}} \rangle$. Runtimes are shown for the `Pythia 8.185` datasets of ref. [103], and we simulate pileup by layering a Poisson-distributed number of minimum-bias events over a hard dijet event.

than other PIRANHA groomers in reproducing both extensive and substructure properties of PU-subtracted jets.

However, figure 16 shows that PIRANHA groomers and CS are more *reliable* PU mitigators than traditional groomers: the standard deviations in the shifts of PU-subtracted extensive and substructure properties tend to be significantly smaller for PIRANHA groomers than for traditional groomers. Again, P-RSF$_{f=0}$ is a noticeable exception and is less reliable as a PU mitigator than either PIRANHA or traditional grooming methods.

We also briefly present results for the runtimes of the grooming algorithms discussed in this paper, including the Constituent Subtraction (CS) pileup mitigation algorithm, as a function of the amount of additive contamination in events with PU. In particular, figure 17 displays the properties of runtimes for grooming algorithms functioning as pileup mitigators as we vary the average number of layered minimum bias events, $\langle n_{\mathrm{PU}} \rangle$. The tree-based grooming strategies, such as Soft Drop and P-RS, are the fastest. P-IVS is an order of magnitude slower. Finally, P-AS and CS are the slowest, with runtimes of $\mathcal{O}(1\mathrm{s})$ even for a reasonable number of pileup events. The combination of high speed and high performance of balanced P-RS provides simple preliminary evidence for their utility in the removal of pileup.
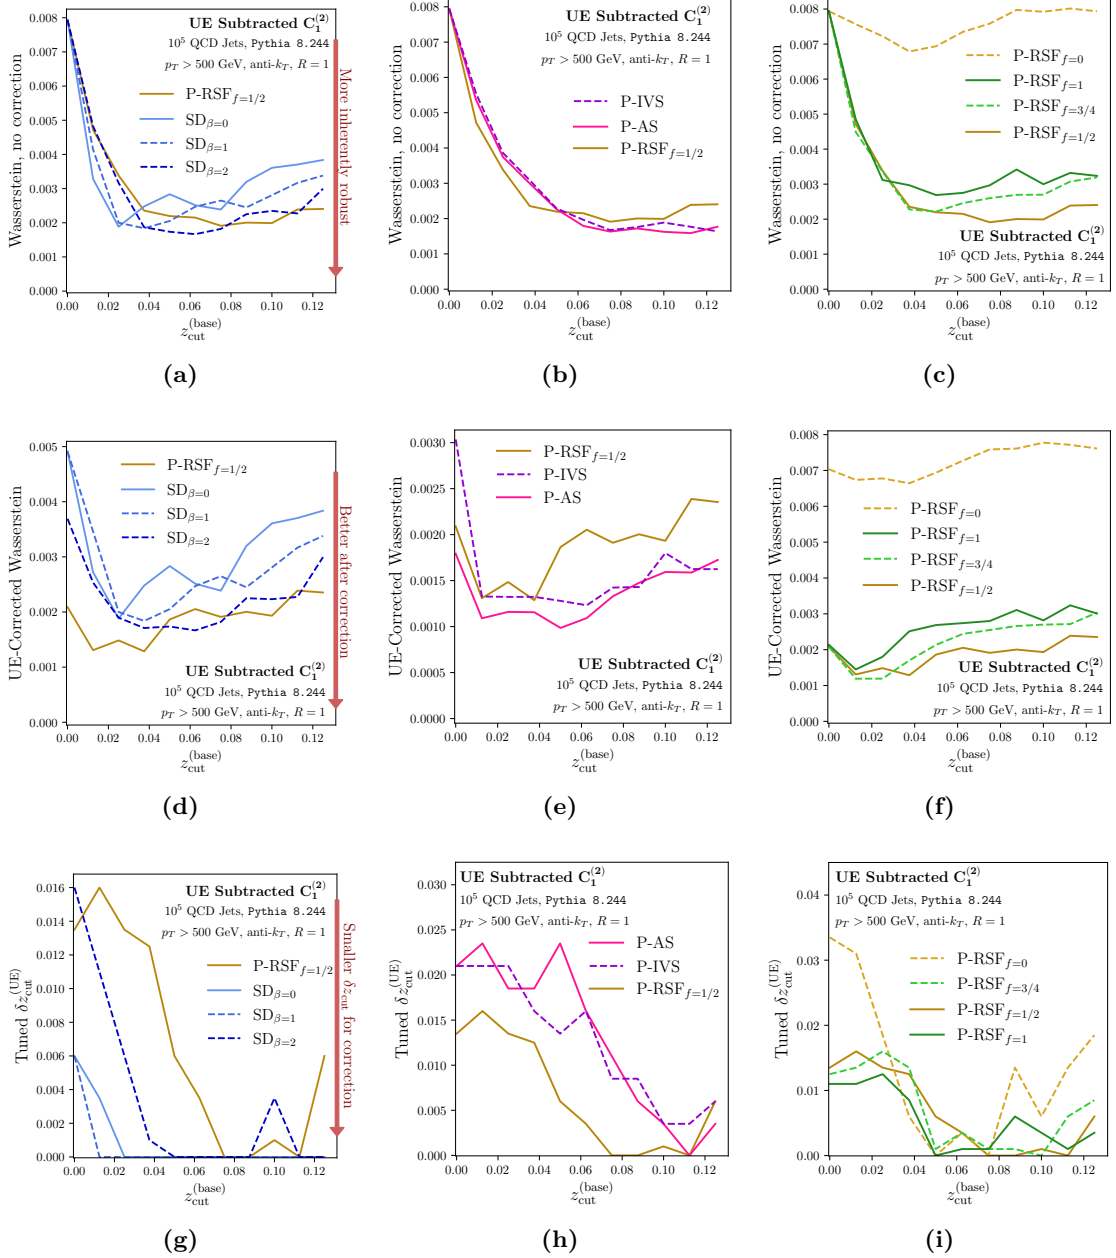
Figure 18. Wasserstein distance between jet $C_1^{(2)}$ distributions (top row) without UE correction and (middle row) with UE correction, and (bottom row) the optimal value of δz_{cut} for UE correction.

B.3 Underlying event

For the underlying event, we present results for the Wasserstein distances between substructure distributions with and without UE, discussed in section 5.2, as well as the additional grooming δz_{cut} required for optimal UE correction. We collect these results in figure 18.

In each of the plots, a lower value along the y-axis — of either the Wasserstein distance or the optimal additional δz_{cut} — indicates a desirable property for UE correction. A lower Wasserstein distance without UE correction indicates inherent robustness of the groomer to the presence of UE. A lower UE-corrected Wasserstein distance indicates that the groomer is able to better correct for the presence of UE. Finally, a lower δz_{cut} indicates greater robustness against UE: less additional grooming must be applied for UE correction.

Notably, figure 18d shows that the UE-corrected Wasserstein distance of P-RSF$_{f=1}$ is smaller than that of Soft Drop for several values of β_{SD} and a range of $z_{\text{cut}}^{(\text{base})}$, indicating that P-RSF$_{f=1}$ is a more delicate tool for the removal of additive contamination due to UE. However, P-IVS and P-AS perform even better at UE removal, as shown in figure 18e. It may be that simple improvements to P-RS that emulate the features of P-IVS and P-AS, such as the "Hard-Balanced Recursive Subtractors" mentioned briefly in section 3.4, may lead to even better performance in UE removal. Finally, while the P-RSF groomers generally exhibit similar performance in UE correction, P-RSF$_{f=0}$ performs poorly at UE correction and is extremely sensitive to the effects of UE; this is unsurprising, as P-RSF$_{f=0}$ deliberately grooms away hard radiation and leaves in soft radiation, such as the radiation contributed by UE.

Since the procedure we use to study the ability of each groomer to remove UE is based on a computationally expensive optimization of δz_{cut}, the plots we show in figure 18g have some unphysical, jagged behavior, especially in the second two rows. We expect that with a dedicated study with better statistics that the jagged behavior of these plots will be smoothed, but that the qualitative nature of our conclusions will remain unchanged.

C Towards resummed recursive subtraction

We would like to gain theoretical insight, where possible, into the properties of Piranha-groomed jets. To move towards a perturbative understanding of Piranha, we begin by studying the behavior of P-RSF at leading order (LO), deriving the results for P-RSF substructure presented in section 3.3, and describe why this is sufficient for generic Piranha groomers at LO. Beyond LO, we find difficulties in using existing resummation techniques because of the intricate correlations between subtractions of different emissions. Nonetheless, we can study the specific example of P-RSF$_{f=1}$ to obtain analytic resummed results up to leading logarithmic accuracy (LL) and numerical results beyond LL.

The angular-ordered, tree-based structure of P-RSF is an important feature that helps facilitate perturbative calculation. In particular, we recall that the branching structure of an angular-ordered tree of emissions is comparable to the history of coherent parton branchings in the parton model [123], and that the philosophy of local parton-hadron duality [145–148] asserts that energy flows are not dramatically changed in the presence

of hadronization. We may therefore approximate the probability of finding a physical particle within an angular-ordered jet tree by the probability of finding an unphysical, unconfined QCD parton with the same energy and angle. Each partonic emission in a jet is accompanied by a factor of the strong coupling α_s, enabling a perturbative analysis of the behavior of groomed observables.

We focus on quantifying the substructure of PIRANHA-groomed jets using the two-prong generalized energy correlation functions (ECFs) of ref. [39]. These were described by eqs. (3.2) and (4.2), and we repeat their approximate form in the central ($y = 0$) and narrow ($R_0 \ll 1$) limit here for convenience:

$$C_1^{(\varsigma)} \simeq \frac{1}{2} \sum_{i=1}^{M} \sum_{j=1}^{M} z_i z_j \left(\frac{\theta_{ij}}{R_0} \right)^\varsigma. \tag{C.1}$$

As before, z_i indicates the energy fraction of particle i, θ_{ij} indicates the opening angle between particles i and j, and R_0 is the jet radius.

The $C_1^{(\varsigma)}$ furnish a set of interesting but relatively simple observables to study in the context of PIRANHA. The behavior of the PIRANHA groomed jet p_T, for example, is relatively uninteresting: since the PIRANHA groomers discussed in this paper operate by subtracting a fixed amount of transverse momentum

$$\Delta p_T = z_{\text{cut}} \, p_T = \rho \, A_{\text{jet}}, \tag{C.2}$$

the changes to the transverse momentum of a jet due to the P-RSF grooming procedure are set trivially by the parameter z_{cut}. However, as discussed in section 3.3, refining calculations of observables involving the energy fraction z_g or angle r_g of the first emission to pass the grooming is a nuanced task that we will defer to future work.

We discuss the leading behavior of the distributions for P-RSF groomed ECFs in appendix C.1, with a brief discussion of the LO behavior of PIRANHA groomers in general. We discuss the behavior of P-RSF$_1$ groomed ECFs at resummed accuracy, comparing results obtained with perturbative QCD, leading logarithmic parton showers, and `Pythia 8.244`, in appendix C.3. We introduce the mathematical technology we use for calculations in perturbative QCD in appendix C.4, and our leading logarithmic parton shower in appendix C.5.

C.1 Substructure at fixed order

For convenience, we first briefly repeat some comments made in section 3.3 while discussing the behavior of P-RSF at lowest order in α_s (LO). At LO, the emission of a single parton with energy fraction z and angle θ is described approximately by the pseudo-probability distribution

$$\frac{\alpha_s}{\pi} p_i(z) \frac{1}{\theta} \approx \frac{2 C_{R_i} \alpha_s}{\pi} \frac{1}{z} \frac{1}{\theta}, \tag{C.3}$$

where $p_i(z)$ is a Dokshitzer-Gribov-Lipatov-Altarelli-Parisi (DGLAP) splitting function [131–133] and C_{R_i} is the quadratic Casimir for the SU(3) color representation R_i

of the mother parton, with $C_F = \frac{N_C^2-1}{2N_C} = \frac{4}{3}$ for quarks and $C_A = N_C = 3$ for gluons.[23] We use reduced splitting functions,

$$\overline{p}_i(z) = p_i(z) + p_i(1-z), \tag{C.4}$$

with $z \in (0, 1/2)$, which describe the probability distribution of the softest parton produced in the splitting.

The singular nature of the DGLAP splitting functions in the limits $z, \theta \to 0$ indicates that the phase space of partonic emissions is dominated by emissions with small energy fraction z and angle θ. The problems posed by these singularities are addressed systematically by resummation, and we review the resummation method we apply here (often called the CAESAR method [57, 135]) in appendix C.4. The non-singular pieces of these splitting functions are written explicitly in eqs. (C.20) and (C.21).

Because each emission is suppressed by a factor of α_s, the behavior of the P-RSF groomed ECFs of eq. (C.1), away from $C_1^{(\varsigma)} = 0$, is determined by the phase space distribution of two-parton jets whose particles survive the P-RSF grooming procedure, up to corrections proportional to α_s^2. If the two partons survive the grooming procedure, they can be described with the modified energy fractions

$$z'_{\text{soft}} = z - f\, z_{\text{cut}} > 0, \tag{C.5a}$$
$$z'_{\text{hard}} = 1 - z - (1-f)z_{\text{cut}} > 0, \tag{C.5b}$$

where z is the original energy fraction of the softer parton. The groomed jet ECF then takes the form

$$C_1^{(\varsigma)}(z,\theta) = (z - f\, z_{\text{cut}}) \left(\frac{\theta}{R_0}\right)^\varsigma \frac{(1 - z - (1-f)z_{\text{cut}})}{(1 - z_{\text{cut}})^2} \approx (z - f\, z_{\text{cut}}) \left(\frac{\theta}{R_0}\right)^\varsigma \frac{(1 - (1-f)z_{\text{cut}})}{(1 - z_{\text{cut}})^2}. \tag{C.6}$$

In the last step, we use that the phase space of the emission is dominated by regions with $z, \theta \ll 1$. The factors of $(1 - z_{\text{cut}})^2$ that appear in the denominators of eq. (C.6) emerge from the normalization of the groomed $C_1^{(\varsigma)}$ by the groomed jet p_T, as in the definition of eq. (C.1).

We also note that if either of the partons in our two-parton configuration at LO is entirely groomed away, the groomed value of $C_1^{(\varsigma)}$ is zero. The region of phase space where one parton is groomed away at LO conspires with virtual contributions and the one-parton configuration with no splittings to ensure that the $C_1^{(\varsigma)}$ distribution integrates to one. These contributions may quickly be calculated once we find the distribution away from the origin, $C_1^{(\varsigma)} > 0$, and we include them at the end of the following discussion.

In the DGLAP splitting functions, the energy fraction z takes values from 0 to 1. Since the $C_1^{(\varsigma)}$ are sensitive to the softest emission in the jet, we are instead interested in

---

[23]By focusing on the singular pieces of the splitting function in eq. (C.3), we implicitly focus on the probability that a *gluon* is emitted from the mother parton. For example, a quark may emit a gluon, so that the resulting final state consists of both a gluon and a quark, while a gluon may split into two gluons. In the language of splitting functions, a gluon may in principle split into two quarks; however, the associated splitting function does not have any singularities as $z \to 0$ and is therefore sub-leading to the singular behavior reflected in eq. (C.3).

restricting z to take values from 0 to $1/2$. We are therefore interested in the *reduced* splitting functions, $\bar{p}(z) = p(z) + p(1 - z)$; these have the same singular behavior as the unmodified DGLAP splitting functions near $z = 0$, but are only used in the range $z \in (0, 1/2)$. From these considerations, the fixed-order distribution of $C_1^{(\varsigma)}$ away from zero becomes

$$\rho_{i,\,\varsigma}^{(\text{LO})}(C > 0) = \frac{\alpha_s}{\pi} \int_0^{R_0} \frac{\mathrm{d}\theta}{\theta} \int_0^{1/2} \mathrm{d}z \, \bar{p}_i(z) \delta\left(C - C_1^{(\varsigma)}(z, \theta)\right), \qquad (\text{C.7})$$

where $C_1^{(\varsigma)}(z, \theta)$ is the functional form of the groomed ECF given in eq. (C.6).

Up to terms that are power-suppressed in C, z_{cut}, or both, the fixed-order distribution for the groomed observable away from $C = 0$ is

$$\rho_{i,\,\varsigma}^{(\text{LO})}(C > 0) = \frac{2\alpha_s C_{R_i}}{\varsigma \pi} \frac{1}{C} \left(2 \tanh^{-1}\left(1 - 2\tilde{C} - 2f\,z_{\text{cut}}\right) + B_i + \mathcal{O}\left(C, z_{\text{cut}}\right)\right) \qquad (\text{C.8a})$$

$$= \frac{2\alpha_s C_{R_i}}{\varsigma \pi} \frac{1}{C} \left(-\log\left(\tilde{C} + f\,z_{\text{cut}}\right) + B_i + \mathcal{O}\left(C, z_{\text{cut}}\right)\right) \qquad (\text{C.8b})$$

where $\tilde{C} = C \times (1 - z_{\text{cut}})^2 / (1 - (1 - f)z_{\text{cut}})$, C_{R_i} is the quadratic Casimir of the color representation of the initiating parton, and we have defined the factor B_i to capture behavior from non-singular components of the splitting functions: $B_q = -3/4$ for quark initiated jets and $B_g = -11/12 + n_f/(6C_A)$ for gluon initiated jets.

Including the $\mathcal{O}(\alpha_s^0)$ jet configuration with a single parton for which $C_1^{(\varsigma)} = 0$, the two-parton contributions for which $z < f z_{\text{cut}}$ and the softer emission is completely eliminated, and the virtual contributions at $\mathcal{O}(\alpha_s^1)$, we can write the full LO probability distribution for $C_1^{(\varsigma)}$ as

$$\rho_{i,\,\varsigma}^{(\text{LO})}(C) \approx \delta(C) + \frac{2\alpha_s C_{R_i}}{\varsigma \pi} \left[\frac{1}{C} \left(-\log\left(\tilde{C} + f\,z_{\text{cut}}\right) + B_i + \mathcal{O}\left(C, z_{\text{cut}}\right)\right)\right]_+^{(C_{\text{max}})}, \qquad (\text{C.9})$$

where $C_{\text{max}} = (1/2 - f z_{\text{cut}})(1 - (1 - f)z_{\text{cut}})/(1 - z_{\text{cut}})^2$ is the maximum value of $C_1^{(\varsigma)}$, and we have fixed the normalization of the distribution with *plus-function regularization*.[24] The plus-function regularized expression of eq. (C.9) may be derived by replacing $\bar{p}(z)$ in eq. (C.7) with its plus-function analog. Alternatively, one can recognize that the full distribution for $C_1^{(\varsigma)}$ must integrate to one even at LO, and that additional contributions to the LO distribution in eq. (C.8) can only come from the regions of phase space where $C_1^{(\varsigma)} = 0$; this is enough to uniquely specify the result of eq. (C.9) at LO.

Unlike the LO distributions for traditionally groomed jet ECFs, such as those of Soft Drop [3], the LO distributions for P-RSF$_{f=1/2}$ and P-RSF groomed $C_1^{(\varsigma)}$ do not exhibit piece-wise behavior. While the piece-wise behavior of Soft Drop observable distributions is smoothed out by all-orders effects [134], the smoothness of PIRANHA-groomed distributions at LO is a manifestation of their continuity.

---

[24]The plus-function regulation of the function $q(x)$ is a distribution $[q(x)]_+^{(x_0)}$ defined by its integral against any function $f(x)$ in a suitable space of test functions,

$$\int_0^{x_0} \mathrm{d}x [q(x)]_+^{(x_0)} f(x) = \int_0^{x_0} \mathrm{d}x \, q(x) \left(f(x) - f(0)\right). \qquad (\text{C.10})$$

See appendix B of ref. [149] and the presentation of ref. [150] for satisfying discussions.

C.2 The groomed energy fraction for a single emission

Next, we discuss the *groomed energy fraction*, z_g, of the first emission to survive the grooming procedure. In Soft Drop, z_g is the energy fraction of the first emission to pass the Soft Drop criterion $z > z_{\text{cut}}\theta^{\beta_{\text{SD}}}$. In P-RSF, z_g is the groomed energy fraction of the first emission that is not entirely removed by the grooming, which we call the *critical emission.*

The behavior of z_g for Soft Drop was derived to double-logarithmic accuracy by ref. [130] with the technology of Sudakov safety. Expanding their double-logarithmic result to $\mathcal{O}(\alpha_s)$, the z_g distribution for Soft Drop takes the approximate form

$$
\rho_{i,\,\text{SD}}(z_g;\beta_{\text{SD}}) = \begin{cases} \dfrac{2\alpha_s C_{R_i}}{\pi|\beta|}\,\overline{p}_i(z_g)\log(z_g/z_c)\,\Theta(z_{\text{cut}} < z_g) + \mathcal{O}(\alpha_s^2), & \beta_{\text{SD}} < 0; \\[2ex] \dfrac{\overline{p}_i(z_g)}{\int_{z_{\text{cut}}}^{1/2}\mathrm{d}z\,\overline{p}_i(z)}\,\Theta(z_{\text{cut}} < z_g), & \beta_{\text{SD}} = 0; \\[2ex] \sqrt{\dfrac{\alpha_s C_{R_i}}{\beta}}\,\overline{p}_i(z_g) + \dfrac{2\alpha_s C_{R_i}}{\beta}\,\overline{p}_i(z_g)\log\min(2z_{\text{cut}}, 2z_g) + \mathcal{O}(\alpha_s^{3/2}), & \beta_{\text{SD}} > 0, \end{cases}
$$

(C.11)

where i again indicates the flavor (quark or gluon) of the initiating parton, and we note that the expression for $\beta_{\text{SD}} = 0$ receives no $\mathcal{O}(\alpha_s)$ corrections at double-logarithmic accuracy. If $\beta_{\text{SD}} \leq 0$, the probability distribution has no support when $z_g < z_{\text{cut}}$ because z_g is determined by the Soft Drop criterion; the vanishing of the z_g distribution when $z_g < z_{\text{cut}}$ for $\beta_{\text{SD}} \leq 0$ therefore reflects the soft discontinuous behavior of Soft Drop. If $\beta_{\text{SD}} > 0$ there is instead a kink at $z_g = z_{\text{cut}}$ that similarly reflects Soft Drop's soft discontinuity.

The behavior of the P-RSF z_g distribution is more subtle, for reasons we discuss below. At the level of a single emission, however, the z_g distribution for Soft Drop may be translated directly to P-RSF because of the similarities between Soft Drop with $\beta = 0$ and P-RSF:[25]

$$
\rho_{i,\,\text{P-RSF}}(z_g; f) = (1 - z_{\text{cut}})\frac{\overline{p}_i\left[(1 - z_{\text{cut}})\,z_g + f z_{\text{cut}}\right]}{\int_{fz_{\text{cut}}}^{1/2}\mathrm{d}z'\,\overline{p}(z')} + \mathcal{O}(\alpha_s).
$$

(C.12)

$\rho_{i,\,\text{P-RSF}}(z_g; f)$ has support in the domain $z_g \in (0, (1/2 - f z_{\text{cut}})/(1 - z_{\text{cut}}))$ at this order of accuracy. The additional factors of $(1 - z_{\text{cut}})$ relative to eq. (C.11) capture the fact that z_g is normalized further after grooming: $z_g = (z_0 - f z_{\text{cut}})/(1 - z_{\text{cut}})$, where z_0 denotes the ungroomed softer energy fraction. The lower limit fz_{cut} on the integration variable z' reflects that we consider configurations for which the softer emission survives the grooming procedure.

Much like Soft Drop with $\beta_{\text{SD}} = 0$, the LO behavior of the P-RSF z_g distribution is determined by the fact that the first surviving emission must have an energy fraction

---

[25]In eq. (C.12), we choose z_g to be the groomed energy fraction of what had been the softer emission before grooming, regardless of which emission is softer after grooming. This avoids additional corrections if $f_{\text{soft}} < 1/2$, where more grooming is applied to the harder emission than the softer emission at each branch and it becomes possible that the harder of two ungroomed emissions becomes the softer of the two emissions after grooming. If we instead define z_g to be the softer *groomed* emission, there are corrections to this formula for $z_g > 1/2 - \mathcal{O}(z_{\text{cut}})$ when $f_{\text{soft}} < 1/2$. We also assume that $z_{\text{cut}} < 1/2$; if $f_{\text{soft}} < 1/2$ and $z_{\text{cut}} > 1/2$, there are additional $\mathcal{O}(z_{\text{cut}})$ corrections to eq. (C.12) due to the constraint that the harder ungroomed emission is not entirely groomed away.

greater than the assigned grooming. At LO, there is a single emission, and therefore the assigned grooming is simply $f z_{\text{cut}}$. Unlike Soft Drop, however, the z_g distribution for P-RSF has support down to $z_g = 0$: a two-parton event with $z = f z_{\text{cut}} + \delta$, $\delta \ll 1$, is mapped to a groomed two-parton event with $z_g = \delta$, while the two-parton event with $z = f z_{\text{cut}} - \delta$ is mapped to a one-parton event; though z_g is undefined in the latter case, where the groomed jet consists of only a single parton, the fact that z_g can become infinitesimally small is a mark of the continuity of P-RSF.

Improving the accuracy of the calculation of z_g for P-RSF will require a more detailed analysis which we leave to future work. An important complication emerges because there may be emissions before the critical emission that can soak up some of the grooming; the amount by which the critical emission is groomed is therefore not z_{cut} but a smaller effective value, $z_{\text{cut}}^{(\text{eff})} = z_{\text{cut}} - \Delta z_{\text{cut}}$, where Δz_{cut} depends on the energy fractions of earlier emissions. This leads to problems when following the approach of ref. [130], suited to the calculation of z_g for Soft Drop. Ref. [130] calculates the distribution of z_g in the framework of Sudakov safety by first marginalizing over the distribution of the angle of the critical emission, r_g. However, the distribution of r_g depends on the amount by which the critical emission is groomed, $z_{\text{cut}}^{(\text{eff})}$, while a leading logarithmic calculation of $z_{\text{cut}}^{(\text{eff})}$ depends on the value of r_g. The circular dependence of the distributions of r_g and $z_{\text{cut}}^{(\text{eff})}$ prevents us from applying the formalism of ref. [130] without more careful modification.

C.3 Substructure at leading logarithmic accuracy and beyond

The global nature of PIRANHA grooming algorithms leads to intricate correlations between the grooming of different final-state particles, and thus to complications in perturbative calculations beyond LO. For P-RSF$_{f=1/2}$, P-IVS, and P-AS, every particle in a jet is groomed by some amount, and the grooming of any particular final-state particle is correlated with the grooming of every other. This subtlety renders existing techniques for multiple emissions calculations less effective. PIRANHA-groomed observables depend on global information, and a more detailed understanding of PIRANHA-groomed correlators will require techniques that can elucidate global information associated with a jet.

While we do not get around this subtlety for the PIRANHA grooming procedures presented in this paper, we may use P-RSF to at least gain some intuition for the resummed behavior of recursive subtractors. In particular, we navigate around the subtleties associated with PIRANHA by studying the substructure of P-RSF with $f_{\text{soft}} = 1$, or P-RSF$_{f=1}$. Crucial to our analysis is that P-RSF$_{f=1}$ eventually terminates, such that it does not necessarily affect every particle within an event.[26] Due to this additional simplicity, we can study P-RSF$_{f=1}$ at leading logarithmic (LL) accuracy in slightly more detail, and find analytic expressions for observables related to the first emission to survive the grooming procedure, or *critical* emission.

---

[26]P-RSF$_{f=0}$ is another subtraction algorithm that eventually terminates, but it subtracts away hard radiation instead of soft radiation. Therefore, we focus on P-RSF$_{f=1}$ in the discussion below. We leave a more detailed survey of the analytic properties of P-RSF$_{f=1/2}$, P-RSF$_{f=0}$, and Recursive Subtraction more generally to future work.

The calculation of these resummed distributions is discussed in more detail in appendix C.4, whose methods may be used to show that the LL distribution for the energy fraction z and the angle θ of the first emission to survive the P-RSF$_1$ procedure is

$$\rho_{\text{crit}}^{(\text{LL})}(z_{\text{crit}}, \theta_{\text{crit}}) = \frac{2C_R \alpha_s}{\pi} \frac{1}{z_{\text{crit}}} \frac{1}{\theta_{\text{crit}}} \exp\left[-\frac{2C_R \alpha_s}{\pi} \log \theta_{\text{crit}} \log 2 z_{\text{cut}}\right]. \tag{C.13}$$

In eq. (C.13) and in the remainder of our discussion, we suppress the index i denoting the flavor of the initiating parton because our LL results depend unambiguously on the representation of the initiating parton through C_R.

Using this resummed distribution, the singular pieces of the splitting function, and dropping subleading terms in z_{cut}, we use the cumulative analog of eq. (C.7) at LL to find the cumulative distribution $\Sigma_{\varsigma,\,\text{crit}}^{(\text{LL})}(C)$ for the contribution of the critical emission to the groomed $C_1^{(\varsigma)}$ for general $f_{\text{soft}} = f$:

$$
\begin{aligned}
\Sigma_{\varsigma,\,\text{crit}}^{(\text{LL})}(C) = {} & C^{\eta_f} + \frac{2C_R \alpha_s}{\varsigma \pi} \left(\frac{C}{f\, z_{\text{cut}}}\right)^{\eta_f} \frac{1}{\eta_f(\eta_f - 1)} \\
& \times \Bigg[\left(\frac{f\, z_{\text{cut}}}{C}\right)^{\eta_f - 1} \; {}_2F_1\left(1, 1 - \eta_f, 2 - \eta_f, -\frac{C}{f\, z_{\text{cut}}}\right) \\
& \quad + \left(\frac{f\, z_{\text{cut}}}{C}\right)^{\eta_f} (\eta_f - 1) \log\left(1 + \frac{C}{f\, z_{\text{cut}}}\right) \\
& \quad - (f\, z_{\text{cut}})^{\eta_f - 1} \; {}_2F_1\left(1, 1 - \eta_f, 2 - \eta_f, 1 - 1/(2\, f\, z_{\text{cut}})\right) \\
& \quad - (f\, z_{\text{cut}})^{\eta_f} (\eta_f - 1) \log\left(1/(2\, f\, z_{\text{cut}})\right) \Bigg],
\end{aligned}
\tag{C.14}
$$

where we have defined

$$\eta_f = -\frac{2C_R \alpha_s}{\varsigma \pi} \log(2\, f\, z_{\text{cut}}). \tag{C.15}$$

To gain some intuition for this intimidating expression, we may also expand this to first order in our fixed coupling α_s to retrieve

$$\Sigma_{\varsigma\ \text{crit}}^{(\text{LL})}(C) \approx 1 + \frac{2C_R \alpha_s}{\varsigma \pi} \left(-\log(2\, f\, z_{\text{cut}}) \log C + \text{Li}_2\left(-\frac{C}{f\, z_{\text{cut}}}\right) - \text{Li}_2(1 - 1/(2\, f\, z_{\text{cut}}))\right). \tag{C.16}$$

Taking a derivative gives

$$\rho_{\varsigma\ \text{crit}}^{(\text{LL})}(C > 0) = \frac{\mathrm{d}}{\mathrm{d}C} \Sigma_{\varsigma\ \text{crit}}^{(\text{LL})}(C > 0) \approx -\frac{2C_R \alpha_s}{\varsigma\ \pi} \frac{1}{C} \log\left(2C + 2 f z_{\text{cut}}\right), \tag{C.17}$$

which more clearly reveals the LL contribution of the critical emission to the groomed value of $C_1^{(\varsigma)}$. Eq. (C.17) differs from the LO result of eq. (C.8) by a factor of 2 inside the logarithm simply because we have integrated z from 0 to $1/2$ and used only the singular parts of the splitting function in computing the LL result, while the LO result includes non-singular pieces of the splitting function. By noting that the probability distribution must integrate to 1, one can also produce an expression with end-point contributions at $C = 0$ which can then be compared to eq. (C.9).

Even for P-RSF with $f_{\text{soft}} = 1$, however, the subtractive nature of the grooming procedure leads to additional subtleties. Usually, the distribution of partonic emissions in the $\log z$–$\log\theta$ plane is approximately uniform, so that a particular emission will tend to have an exponentially higher contribution to the two-pronged substructure of a jet than any other. Since the first emission to survive an angular ordered hard-cutoff grooming procedure such as Soft Drop will have a non-negligible energy fraction $z > z_{\text{cut}}$ and a large angle, it is this first surviving emission, the critical emission, that contributes dominantly to $C_1^{(\varsigma)}$. This trick made the original calculation of Soft Drop groomed substructure quite simple [3].

In the case of P-RSF$_{f=1}$, however, the critical emission is partially groomed and may have a groomed energy fraction $z_{\text{P-RSF}_1} \sim z - z_{\text{cut}} \ll 1$. It is therefore possible for other emissions — the ungroomed *subsequent emissions* which are narrower than the critical emission and therefore untouched by the P-RSF$_{f=1}$ algorithm — to contribute greater values to $C_1^{(\varsigma)}$. The ungroomed *subsequent* emissions after the P-RSF$_{f=1}$ algorithm terminates can lead to larger contributions to substructure observables than they would in the case of Soft Drop.

As a result, performing resummed calculations of Soft Drop groomed substructure distributions that use only the critical emissions are more accurate than calculations for P-RSF$_{f=1}$ that only use the critical emission. The relative accuracy of the LL, single-emission result for Soft Drop and P-RSF$_{f=1}$ is demonstrated in figures 19a and 19c, which compares the fixed-order behavior of $C_1^{(2)}$ distributions for jets groomed with Soft Drop or P-RSF$_{f=1}$, respectively, using either the critical emission or all emissions in the jet. Our parton shower procedure is inspired by ref. [151] and is described in more detail in appendix C.5. Including multiple emissions does not dramatically change the Soft Drop substructure distribution, but leads to non-negligible changes in P-RSF$_{f=1}$ substructure distributions.

We may include effects from the running of the strong coupling constant and multiple emissions within the ungroomed jet to determine the behavior of the cumulative distribution of $C_1^{(\varsigma)}$ at an even higher level of accuracy, modified leading logarithmic (MLL) accuracy, as described in appendix C.4. While including these features in our calculation does not fully achieve next-to-leading-logarithmic (NLL) accuracy, the MLL calculation is an important piece of the full NLL result. To consider the effects of several emissions on the distribution of the groomed ECFs, we convolve the probability distributions of several emissions that contribute to the total observable $C_1^{(\varsigma)}$:

$$
\begin{aligned}
\rho_\varsigma^{(\text{MLL})}(C_{\text{tot}}) = \int_{f\,z_{\text{cut}}}^{1/2} \mathrm{d}z_{\text{crit}} \int_0^1 \mathrm{d}\theta_{\text{crit}} \; &\rho_{\text{crit}}(\theta_{\text{crit}})\rho_{\text{crit}}(z_{\text{crit}}|z_{\text{pre}}) \\
\times \int \mathrm{d}z_{\text{pre}} \; &\rho_{\text{pre}}(z_{\text{pre}}|\theta_{\text{crit}}) \int \mathrm{d}c_{\text{sub}} \; \rho_{\text{sub}}(c_{\text{sub}}|z_{\text{crit}},\theta_{\text{crit}}) \\
\times \; &\delta\left(C_{\text{tot}} - C_{1,\,\text{crit}}^{(\varsigma)}(\theta_{\text{crit}}, z_{\text{crit}}, z_{\text{pre}}) - c_{\text{sub}}\right).
\end{aligned}
\tag{C.18}
$$

Each factor of ρ in eq. (C.18) indicates a resummed probability distribution corresponding to a particular emission, discussed in greater detail in appendix C.4. At the MLL level of accuracy, these resummed probability distributions are all calculated with a running coupling α_s. We model non-perturbative effects by freezing the value of α_s below
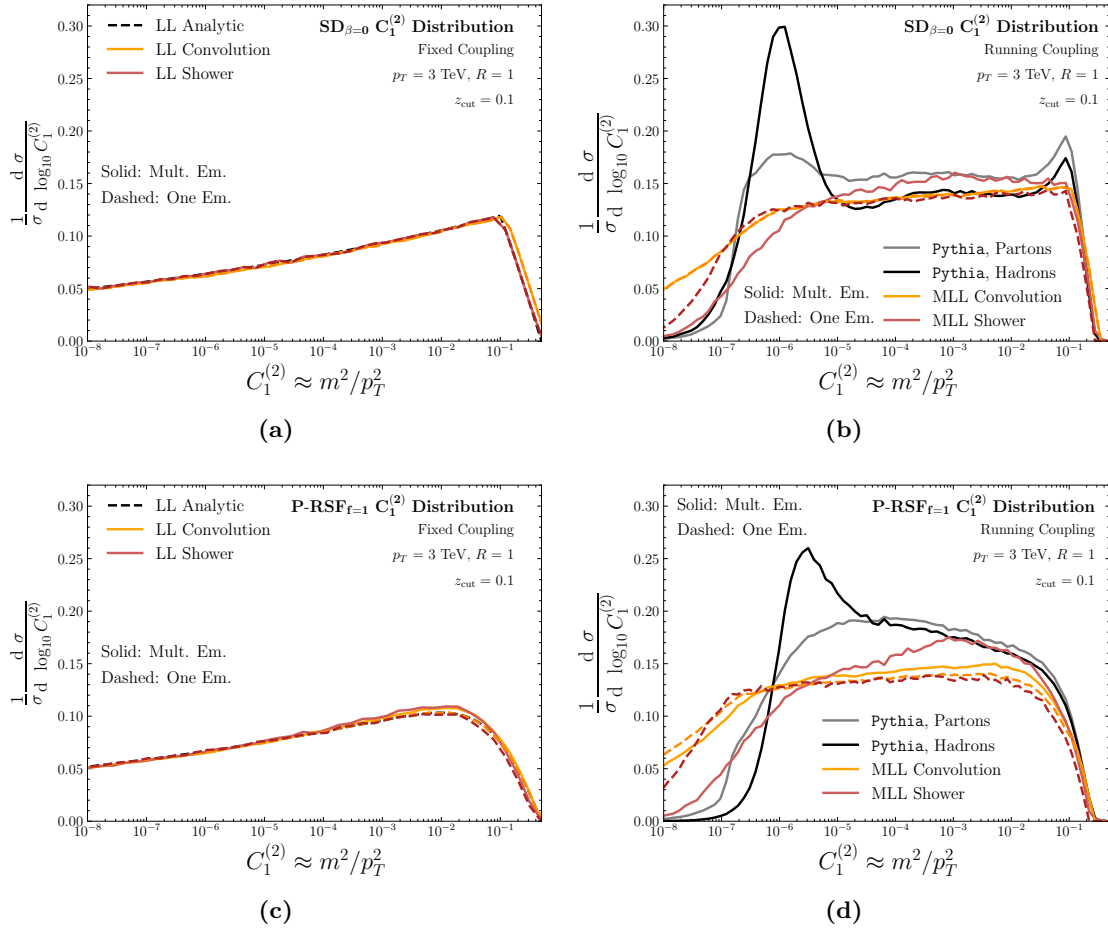
Figure 19. Distributions of the observable $C_1^{(2)} \approx m^2/p_T^2$ of eq. (C.1) for quark jets groomed with $\mathrm{SD}_{\beta=0}$ (top row) and $\mathrm{P\text{-}RSF}_{f=1}$ (bottom row), with and without multiple emissions, at (left column) LL accuracy and (right column) MLL accuracy. For each plot, we display the results of the convolution method of eq. (C.18) using numerical integration and perturbative QCD with 10^6 phase space samples as well as the parton shower algorithm of appendix C.5. At LL, we compare these results to the analytic prediction of eq. (C.14); at MLL, we compare them to parton-level and hadron-level results obtained with Pythia 8.244 using the anti-k_t algorithm. The figures demonstrate that $\mathrm{P\text{-}RSF}_{f=1}$ is less sensitive to hadronization corrections in the large observable regime, near $C_1^{(2)} \approx 10^{-2}$, and that $\mathrm{P\text{-}RSF}_{f=1}$ groomed substructure is more sensitive to the effects of multiple emissions than that of Soft Drop.

the non-perturbative scale of $\mu_{\mathrm{freeze}} = 1$ GeV, as in eq. (C.30). Variations in our choice of non-perturbative scale lead only to small changes to our results.

The expressions that describe multiple emissions within the groomed jet are:

- $\rho_{\mathrm{crit}}(\theta_{\mathrm{crit}})\rho(z_{\mathrm{crit}}|z_{\mathrm{pre}})$ represents the probability distribution for the emission angle θ_{crit} and energy fraction z_{crit} of the *critical* emission, or the first emission to survive the $\mathrm{P\text{-}RSF}_{f=1/2}$ procedure. This distribution is conditioned on the presence of earlier emissions that modify the grooming procedure.

- $\rho_{\text{pre}}(z_{\text{pre}}|\theta_{\text{crit}})$ encodes the probability that emissions that are wider than the critical emission, which we call *pre-critical* emissions, have an energy fraction z_{pre}. Any non-zero value of z_{pre} mitigates the grooming of the critical emission.

- $\rho_{\text{sub}}(c_{\text{sub}}|z_{\text{crit}}, \theta_{\text{crit}})$ is the probability distribution for emissions narrower than the critical emission to contribute a value c_{sub} to the overall observable $C_1^{(\varsigma)}$. We call these emissions *subsequent* emissions, and they are unmodified by the grooming procedure.

- Finally, $C_{1,\,\text{crit}}^{(\varsigma)}(\theta_{\text{crit}}, z_{\text{crit}}, z_{\text{pre}})$ is the contribution of the critical emission to $C_{1,\,\text{tot}}^{(\varsigma)}$, in the presence of pre-critical emissions with energy fraction z_{pre}.

We evaluate the convolution of eq. (C.18) through numerical integration and compare the resulting distributions to the results from running coupling parton showers and `Pythia 8.244` in figure 19. In particular, figures 19b and 19d compare our numerical MLL results to hadron-level results for anti-k_t jets in `Pythia 8.244` and demonstrate that this trend of comparatively large multiple-emissions effects in P-RSF$_{f=1}$ continues at MLL.

First, let us note the difference between our MLL results obtained with Monte Carlo integration, described in appendix C.4, and via parton showers, described in appendix C.5. Our Monte Carlo results use a frozen coupling below a non-perturbative scale $\mu_{\text{freeze}} = 1$ GeV (see, for example, eq. (C.30)). Our parton shower results both freeze the coupling at $\mu_{\text{freeze}} = 1$ GeV and stop the shower below a non-perturbative cutoff $\mu_{\text{cutoff}} = 1$ MeV, tuned by hand to match results from `Pythia 8.244`. The shower cutoff effectively sets α_s to zero below the scale set by μ_{cutoff}. Therefore, the plots in figures 19b and 19d effectively show the results obtained with three different models of non-perturbative effects, with `Pythia 8.244` the most accurate. Our MLL parton shower results in figures 19b and 19d are still normalized, but the normalization is not immediately evident due to the presence of zero bins that are not shown in the logarithmically-scaled figures. In passing from our rough parton shower model to the more accurate hadronization model of `Pythia`, we expect that it is these zero bins that contribute to the additional hadron-level peak shown in the `Pythia` data.

We also point out that in figure 19b, the parton-level `Pythia 8.244` distribution for $C_1^{(2)}$ also has an additional peak near the discontinuity in the parton-level LL result, in the regime of relatively large $C_1^{(2)} \approx 10^{-1}$. There is a comparable additional peak in the hadron-level `Pythia` distribution for $C_1^{(2)}$ at this relatively large value for the observable.

We also note that, while our numerical results for P-RSF$_{f=1}$ in the MLL figures agree roughly with `Pythia 8.244` in the regime of large $C_1^{(2)}$, the hadron-level `Pythia` distributions have an additional peak near $C_1^{(2)} \approx 3 \times 10^{-6}$ due to hadronization corrections to the observable at that scale. There is a similar additional peak in the hadron-level `Pythia` distribution for mMDT (SD$_{\beta=0}$) groomed $C_1^{(2)}$.

C.4 Summary of resummation: convolution

We now review the resummation procedure we apply in our LL and MLL results. In particular, we introduce the basics of the CAESAR formalism, named due to its use in the context of the Computer Automated Expert Semi-Analytical Resummer (CAESAR)

program [57].[27] To the best of our knowledge, this formalism was first introduced in the context of semi-analytic resummation of event shapes [135]. While we include an introduction here for completeness, a presentation of the concepts that appear in the following discussion can be found in more elegant or complete forms throughout existing literature [57, 111, 135, 155–161].

The distribution of energy fraction and angle of a partonic emission in a jet is described at fixed order by

$$\rho_{i,\text{ f.o.}}(z,\theta) = \frac{\alpha_s}{\pi} \bar{p}_i(z) \frac{1}{\theta}, \tag{C.19}$$

where $z \in (0, 1/2)$ is the energy fraction of the softer branch of the emission and $\bar{p}(z) = p(z) + p(1-z)$ indicates a reduced splitting function, modified from the splitting function $p(z)$ to describe the softer branch. We would like to regulate the singularities of this distribution by using techniques of resummation in perturbative QCD and thus gain greater predictive power over groomed jet observables. In this appendix, we discuss the resummed distributions used in the text and give some quick derivations.

To begin, we provide explicit expressions for the DGLAP splitting functions, which describe the probability for the emission of partonic radiation from quarks and gluons, at one loop accuracy:

$$p_{q \to qg}(z) = C_F \frac{1 + (1-z)^2}{z}, \tag{C.20}$$

$$p_{g \to gg}(z) = 2C_A \frac{1-z}{z} + C_A z(1-z) + T_F n_f (z^2 + (1-z)^2). \tag{C.21}$$

Here, z is the energy fraction of an emitted gluon relative to the total energy of the emitted gluon and an initial parton. p_q describes the distribution of z in the case that a gluon emitted by a quark, $q \to qg$, and p_g describes the distribution of z in the case $g \to gg$. $T_F = \frac{1}{2}$ is the Dynkin index of the fundamental representation of the standard model SU(3) color gauge group, and n_f indicates the number of quark flavors. In the following analysis, we take $n_f = 5$, including all quark flavors except for the top quark. As in the text, C_R is the quadratic Casimir for the representation R of SU(3). $C_R = C_F = \frac{4}{3}$ for quarks and $C_R = C_A = 3$ for gluons.

We also introduce some graphical notation to represent the probability of a splitting whose angle θ and softer energy fraction z lie within some region of phase space. In particular, we describe the probability of finding a splitting in some region in the $\log(z^{-1})$-$\log(\theta^{-1})$ plane, or Lund plane, with a diagram of the Lund plane with the corresponding region filled:

$$\iint_{\bigcirc} \frac{\alpha_s}{\pi} \frac{\mathrm{d}\theta}{\theta} \, \mathrm{d}z \, [\bar{p}_i(z)]_+^{(1/2)} \triangleq \quad \triangleq \quad , \tag{C.22}$$

where the shaded oval in the Lund plane above represents an arbitrary shape, and $[\bar{p}(z)]_+^{(1/2)}$ is a plus-function regulated *reduced* splitting function, $\bar{p}_i(z) = p_i(z) + p_i(1-z)$, using the

---

[27]The results of our presentation below are also a special case of the more general, observable-independent formalism of jet calculus [152] — closely connected to the resummation procedure used by parton showers — for which pedagogical introductions may be found in e.g. refs. [153, 154].

definition of plus-function regularization of eq. (C.10).[28] It follows, for example, that

 $= 0$ and $=$ $= -$. (C.23)

As depicted above, we may replace any vertical line — representing an integral over z at fixed θ — or any vertical strip — representing an integral over z and some range of θ — with zero, since the integral of a plus-function regulated distribution is zero.

In this paper, we are concerned with the jet ECFs $C_1^{(\varsigma)}$, which can be roughly calculated by adding contributions from all of the emissions within a jet, such that the contribution from each emission increases as z or θ increases and vanishes in the limit that z or θ vanishes. We call such observables *additive* observables, and we may treat the $C_1^{(\varsigma)}$ as additive observables in our leading order computations.

The calculation of distributions of additive observables is simplified by the approximately uniform distribution of partonic emissions within the Lund plane. In particular, the uniformity of emissions in the Lund plane implies that there is a single emission with splitting parameters z and θ that we expect to be exponentially higher than the values of z and θ of other emissions. Additive observables are therefore generically dominated by the contribution from a single such emission, and we may approximate the probability that an observable $\mathcal{O}$ is less than the sum of contributions from different emissions, $\mathbb{P}(\mathcal{O} < \sum_i x_i)$, by the probability that it is less than the dominant value, $\mathbb{P}(\mathcal{O} < X)$, with $X = \max_j x_j$. We refer to this as the *leading logarithmic* (LL) approximation.

Mathematically, the LL approximation for an additive observable $\mathcal{O}$, with contributions $x_i > 0$ from emissions i, asserts that there is a dominant contribution $X = \max_i x_i$ such that $\mathcal{O} = \sum_i x_i \approx X$. This approximation holds up to effects that are sub-leading in α_s and logarithms of the observable. In the calculation of cumulative distributions for $\mathcal{O}$, we may write without approximation that $\Theta(\mathcal{O} - \sum_i x_i) \leq \Theta(\mathcal{O} - X)$, where equality emerges when we have at most a single non-zero $x_j = X$. The LL approximation notes that x_i/X is exponentially small, and therefore that $\Theta(\mathcal{O} - \sum_i x_i) \approx \Theta(\mathcal{O} - X)$. We should also note that typically, calculations applying the LL approximation will also drop terms that are suppressed by powers of X since these are sub-dominant to the leading logarithmic terms in observable distributions.

We therefore define the LL cumulative probability distribution $\Sigma_{\mathcal{O}}^{(\mathrm{LL})}(X)$ of an additive observable $\mathcal{O}$ as the probability that no emission contributes a value greater than X to $\mathcal{O}$, up to sub-leading behavior in α_s and logarithms or powers of the observable. The region of phase space in which an observable contributes a value greater than X to the observable is called a *veto region*. We sum and integrate over all possible final-state emissions outside the veto region when calculating $\Sigma_{\mathcal{O}}^{(\mathrm{LL})}(x)$.

---

[28]We may use $[\bar{p}(z)]_+^{(1/2)}$ for the applications presented in this work because we are interested in observables whose dependence on energy is captured entirely by the energy fraction of the softer parton at any branching. If we would like to deal with observables that require more information than just the softer energy fraction, such as the flavor of the emitted partons, we would need to work with a plus-function regulated splitting function $[p(z)]_+^{(1)}$.

As an example, the approximate mass of a jet due to a single emission is given by $m^2 = Q^2 z\theta^2$, where Q is the total energy of the jet. Therefore, the jet mass distribution is given at leading logarithmic accuracy by

$$
\begin{aligned}
\Sigma_{m^2}^{(\mathrm{LL})}(m^2) &= \sum_{n=0}^{\infty} \frac{1}{n!} \left(\frac{\alpha_s}{\pi} \int_0^1 \frac{\mathrm{d}\theta}{\theta} \int_0^{\frac{1}{2}} \mathrm{d}z\, [\bar{p}_i(z)]_+^{[1/2]} \Theta\left(\frac{m^2}{Q^2} - z\theta^2 \right) \right)^n \\
&= \exp\left[-\frac{\alpha_s}{\pi} \int_0^1 \frac{\mathrm{d}\theta}{\theta} \int_0^{\frac{1}{2}} \mathrm{d}z\, \bar{p}_i(z) \Theta\left(z\theta^2 - \frac{m^2}{Q^2} \right) \right].
\end{aligned}
\tag{C.24}
$$

In graphical notation, the analogous derivation is

$$
\Sigma_{m^2}^{(\mathrm{LL})}(m^2) = \sum_{n=0}^{\infty} \frac{1}{n!} \left(\begin{array}{c}\includegraphics\end{array} \right)^n = \exp\left[-\begin{array}{c}\includegraphics\end{array} \right],
\tag{C.25}
$$

where in the final equality, we see that the LL cumulative distribution is one over the exponent of the veto region area. The area of the veto region for an observable is also called the *radiator* $R_{\mathcal{O}}$ for the observable.

A similar analysis yields the cumulative probability distribution function of the angle θ for the first emission in the P-RSF$_{f=1/2}$ procedure to satisfy $z > z_{\mathrm{cut}}$, which we call the *critical emission*:

$$
\Sigma_{\mathrm{crit},\theta}^{(\mathrm{LL})}(\theta_{\mathrm{crit}}) = \exp\left[-\begin{array}{c}\includegraphics\end{array} \right] = \exp\left[-\frac{\alpha_s}{\pi} \int_{\theta_{\mathrm{crit}}}^1 \frac{\mathrm{d}\theta}{\theta} \int_{z_{\mathrm{cut}}}^{1/2} \mathrm{d}z\, \bar{p}_i(z) \right], \tag{C.26}
$$

with a corresponding probability density

$$
\rho_{\mathrm{crit},\theta}^{(\mathrm{LL})}(\theta_{\mathrm{crit}}) = \frac{1}{\theta_{\mathrm{crit}}} \frac{\alpha_s}{\pi} \int_{z_{\mathrm{cut}}}^{1/2} \mathrm{d}z\, \bar{p}(z) \exp\left[-\frac{\alpha_s}{\pi} \int_{\theta_{\mathrm{crit}}}^1 \frac{\mathrm{d}\theta}{\theta} \int_{z_{\mathrm{cut}}}^{1/2} \mathrm{d}z\, \bar{p}_i(z) \right]. \tag{C.27}
$$

The resummed, joint probability distribution for the energy fraction and angle of the critical emission is therefore

$$
\rho_{\mathrm{crit}}(z,\theta) = \rho_{\mathrm{crit},z}(z|\theta)\rho_{\mathrm{crit},\theta}(\theta) = \alpha_s(\kappa_{\mathrm{crit}})\, \bar{p}(z)\frac{1}{\theta} \exp\left[-\frac{1}{\pi} \int_\theta^1 \frac{\mathrm{d}\theta'}{\theta'} \int_{z_{\mathrm{cut}}}^{1/2} \mathrm{d}z'\bar{p}_i(z')\, \alpha_s(\kappa') \right],
\tag{C.28}
$$

where the argument κ of the strong coupling accounts for the possible inclusion of running coupling effects. In particular, at LL we do not include effects due to the running of the coupling, and $\kappa = p_T R$ is a scale set only by the dynamics and parameters of the jet. When including running coupling effects, $\kappa \approx p_T R z\theta \approx p_\perp$ is approximately related to $p_\perp$, the momentum of the softer emission that is transverse to the momentum of the harder emission. This is because QCD correlations and cross sections that use α_s evaluated at a scale κ generically include logarithms of the form $\log(p_\perp/\kappa)$. To make these logarithms

small, one may choose $\kappa \sim p_\perp$, which then requires that we evaluate the strong coupling at the scale $p_\perp$ as well. At one loop, the running of α_s is captured by

$$\alpha_s^{1-\text{loop}}(\kappa) = \frac{\alpha_s(Q)}{1 + 2\alpha_s(Q)\beta_0 \log\left(\frac{\kappa}{Q}\right)}, \tag{C.29}$$

where $\beta_0 = (\frac{11}{3}C_A - \frac{4}{3}T_f n_f)/4\pi$ is a coefficient appearing in the QCD beta function: $\beta(\alpha_s) = -2\beta_0\alpha_s^2 + \mathcal{O}(\alpha_s^3)$. Calculations accounting for effects from the running of the QCD coupling and multiple emissions within a jet are called *modified leading logarithmic* (MLL) calculations.

In the numerical MLL results we present in figure 19, we also model non-perturbative QCD effects by *freezing* the coupling below a non-perturbative scale μ_{freeze},

$$\alpha_s^{1-\text{loop, frozen}}(\kappa; \mu_{\text{freeze}}) = \alpha_s^{1-\text{loop}}(\kappa)\Theta(\kappa > \mu_{\text{freeze}}) + \alpha_s^{1-\text{loop}}(\mu_{\text{freeze}})\Theta(\mu_{\text{freeze}} > \kappa). \tag{C.30}$$

Eq. (C.30) is a simplified analog of several motivated approaches to modeling the infrared structure of QCD for which α_s also approximately freezes below a non-perturbative scale.[29] We take $\mu_{\text{freeze}} = 1\,\text{GeV}$ in our numerical calculations, and we note that changing our choice of μ_{freeze} leads only to small effects in substructure distributions such as those presented in figure 19.

Let us now apply the technology reviewed in this appendix to distributions of P-RSF$_{f=1}$ groomed substructure observables. An efficient way to include the multiple emission effects that contribute to MLL accuracy is to consider P-RSF$_1$ so that the grooming affects only soft, wide-angle emissions and terminates before it affects narrow emissions forming the jet core. For P-RSF$_{f=1}$, we may divide up the emissions into three qualitative categories:

a) soft, wide angle *pre-critical emissions*, which are considered first in an angular-ordered grooming procedure and thus completely removed by P-RSF$_{f=1}$;

b) a single *critical emission*, which is the first soft emission to survive the grooming procedure but is nonetheless modified by the grooming; and

c) *subsequent emissions*, which are left unmodified by the grooming procedure.

If we have the radiator for an additive observable $\mathcal{O}$, it is simple to evaluate corrections due to the presence of multiple emissions. These corrections can be evaluated with Laplace transform methods [57] to find that at MLL accuracy,

$$\Sigma_{\mathcal{O}}^{(\text{MLL})}(x) = \frac{e^{-R_{\mathcal{O}}(x) - \gamma_E R'_{\mathcal{O}}(x)}}{\Gamma\left(1 + R'_{\mathcal{O}}(x)\right)}, \tag{C.31}$$
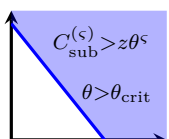
---

[29]The freezing of the effective value of α_s in the low-energy regime can be argued in several complementary contexts. The freezing of α_s is discussed by ref. [162] in the context of optimized perturbation theory [163], whose goal is to perform perturbative calculations that are minimally sensitive to changes in renormalization scheme, by refs. [164, 165] in the context of expanding perturbative results around the Caswell-Banks-Zaks fixed point $N_f = N_f^* \lesssim 11N_c/2$ of the QCD beta function [166, 167], and by ref. [168] in the context of holographic QCD. The freezing of α_s may be also seen qualitatively by comparing experimental data to field theoretic calculations [169–171], and is a special case of the more general program of using effective couplings to describe non-perturbative infrared effects [141, 172–175]. See ref. [176] for a recent review of models of the low-energy behavior of QCD with and without a frozen coupling.

where γ_E is the Euler-Mascheroni constant, and the prime denotes a derivative with respect to $\log(1/x)$.[30] These considerations lead us to the MLL expressions for the effects of multiple pre-critical and subsequent emissions on jet ECFs.

The cumulative distribution for the contribution of the subsequent emissions to the observable $C_1^{(\varsigma)}$ takes the form

$$\Sigma_{C_{\mathrm{sub}}^{(\varsigma)}}^{(\mathrm{MLL})}(C|\theta_{\mathrm{crit}}, z_{\mathrm{crit}}) = \frac{e^{-R_{\mathrm{sub}}(C)-\gamma_E R'_{\mathrm{sub}}(C)}}{\Gamma\left(1+R'_{\mathrm{sub}}(C)\right)} \approx e^{-R_{\mathrm{sub}}(C)} = \Sigma_{C_{\mathrm{sub}}^{(\varsigma)}}^{(\mathrm{LL})}(C|\theta_{\mathrm{crit}}), \qquad (\mathrm{C.32})$$

where $R_{\mathrm{sub}}(C)$ is shorthand for $R_{C_{\mathrm{sub}}^{(\varsigma)}}(C|\theta_{\mathrm{crit}})$, defined as

$$R_{C_{\mathrm{sub}}^{(\varsigma)}}(C|\theta_{\mathrm{crit}}) = \begin{matrix} C_{\mathrm{sub}}^{(\varsigma)} > z\theta^\varsigma \\ \theta > \theta_{\mathrm{crit}} \end{matrix} = \int_{C/\theta_{\mathrm{crit}}^\varsigma}^{1/2} \frac{\mathrm{d}c}{c} \int_c^{1/2} \mathrm{d}z\, \bar{p}_i(z) \frac{\alpha_s(\kappa)}{k\varsigma}. \qquad (\mathrm{C.33})$$

Of course, the value of z_{crit} will also have an impact on the subsequent emissions, but we may safely neglect these corrections, which contribute only by scaling $C_{\mathrm{sub}}^{(\varsigma)}$ by factors of the form $1 - z_{\mathrm{crit}}$.

The cumulative distribution for the energy fraction carried by pre-critical emissions is calculated similarly and takes the form

$$\Sigma_{z_{\mathrm{pre}}}^{(\mathrm{MLL})}(z|\theta_{\mathrm{crit}}) = \frac{e^{-R_{\mathrm{pre}}(z|\theta_{\mathrm{crit}})-\gamma_E R'_{\mathrm{pre}}(z|\theta_{\mathrm{crit}})}}{\Gamma\left(1+R'_{\mathrm{pre}}(z|\theta_{\mathrm{crit}})\right)} \approx e^{-R_{\mathrm{pre}}(z|\theta_{\mathrm{crit}})} = \Sigma_{z_{\mathrm{pre}}}^{(\mathrm{LL})}(z|\theta_{\mathrm{crit}}), \qquad (\mathrm{C.34})$$

where $R_{\mathrm{pre}}(z|\theta_{\mathrm{crit}}) = R_{z_{\mathrm{pre}}}(z|\theta_{\mathrm{crit}})$ is given by

$$R_{z_{\mathrm{pre}}}(z|\theta_{\mathrm{crit}}) = \int_{\theta_{\mathrm{crit}}}^1 \frac{\mathrm{d}\theta}{\theta} \int_z^{z_{\mathrm{cut}}} \mathrm{d}z'\, \bar{p}_i(z') \frac{\alpha_s(\kappa)}{\pi}. \qquad (\mathrm{C.35})$$

We need only use the singular pieces of the splitting functions to achieve LL accuracy, but, by definition, we must include the non-singular pieces of the splitting functions to achieve MLL accuracy. While the MLL analysis modifies the LL analysis by including additional, higher-order behavior, it is not quite sufficient to produce results at next-to-leading logarithmic (NLL) accuracy; a full NLL result requires a more intricate treatment including the effects of logarithms due to the constraint that radiation within the jet must lie within the jet radius. These logarithms emerge because jet observables depend on only subsets of the full phase (global) space of the outgoing radiation, and are known as non-global logarithms [56, 60, 65, 81, 144, 177–188]. Regardless, the MLL results presented in the paper provide a sense of the physics contained within NLL results.

---

[30] We note that additivity is a sufficient condition for eq. (C.31), but is not strictly necessary. For example, a weaker but still sufficient condition, named *recursive infrared-collinear (rIRC) safety*, was introduced by ref. [57]. The definition of rIRC safety is given in eqs. (3.4) and (3.5) of ref. [57], and an example of additivity is furnished by the thrust observable, discussed in section 3.2 of ref. [57].
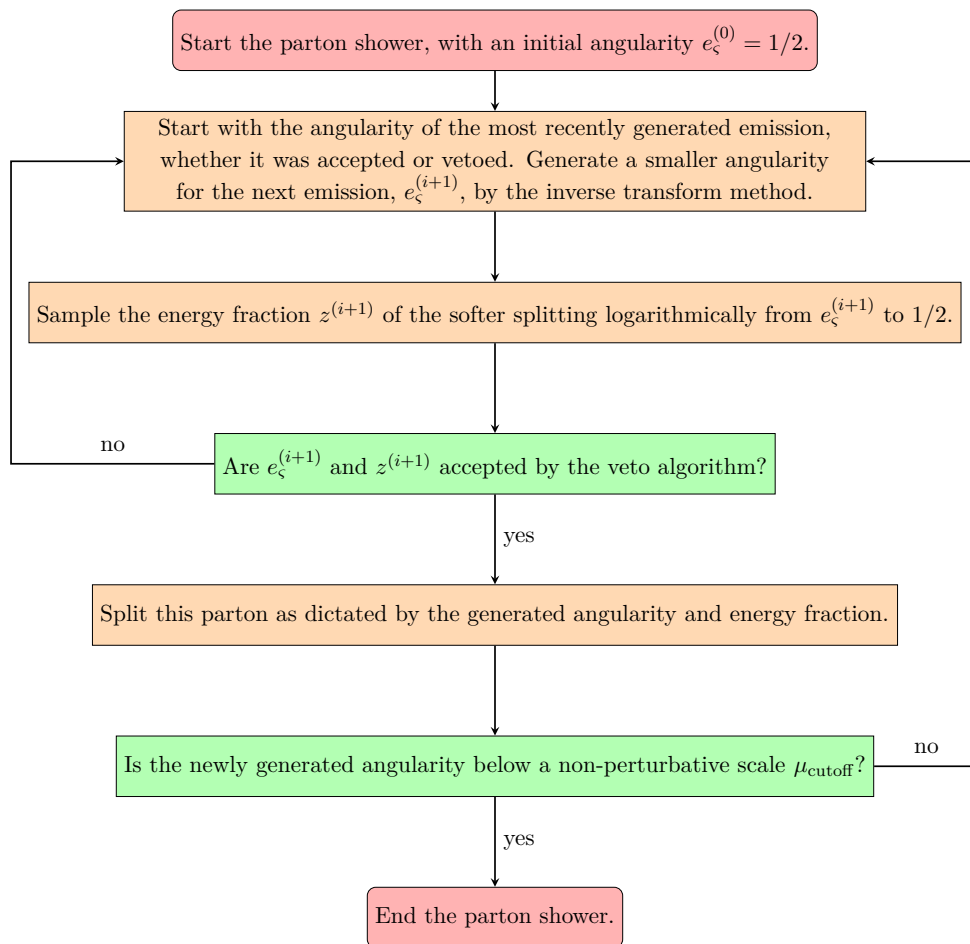
Figure 20. The basic code flow of the parton shower algorithm described in appendix C.5. The veto algorithm is only used when considering MLL effects and is skipped at LL. The computations of the algorithm are shown in orange, the decisions of the algorithm in green, and the start and stop of the algorithm in red.

C.5 Summary of resummation: parton shower algorithm

In this appendix, we describe the parton shower we use in the text to provide an additional check of our analytic results. Our parton shower, motivated by section 5 of ref. [151], is available on GitHub [189]. A concise flow of code for our parton shower at MLL accuracy is shown in figure 20.

In broad strokes, we order our parton shower by the angularity $e_\varsigma \approx z\theta^\varsigma$ of its emissions. The angularity e_ς associated with a single emission i with softer energy fraction z and opening angle θ takes the form $e_\varsigma^{(i)} = z\,(\theta/R)^\varsigma$, where R indicates the jet radius, which we set to 1.

To obtain a parton shower valid at LL accuracy, we implement the following procedure:

1. Initialize a jet consisting of a single parton. The splittings of this parton will give rise to all of the final-state particles in the jet at the parton level. We take the angularity

$e_\varsigma^{(0)} = R^\varsigma/2$ at this stage to be the maximum possible angularity for a single splitting within the jet radius R.

2. If the angularity $e_\varsigma^{(i)}$ of the parton under consideration is below a pre-determined cutoff scale, label it as a final-state parton. In our MLL parton shower, for example, the cutoff scale is taken to be a non-perturbative scale $\mu_{\text{cutoff}}/p_{T,\text{jet}} = 1\,\text{MeV}/p_{T,\text{jet}}$, where we have tuned the cutoff roughly by hand in order to get better agreement with substructure distributions obtained with `Pythia 8.244`. Otherwise, the parton splits into two daughter partons, and we continue to the next step.

3. The splitting of a parton is described by its angularity $e_\varsigma^{(f)}$. Using the techniques of the previous section, we may write the LL probability that there is no emission with angularity greater than $e_\varsigma^{(f)}$, given that there has been an emission with angularity $e_\varsigma^{(i)}$, as

$$\mathbb{P}\left(\text{no emission with } e_\varsigma > e_\varsigma^{(f)} \mid e_\varsigma < e_\varsigma^{(i)}\right) = \frac{\Sigma(e_\varsigma^{(f)})}{\Sigma(e_\varsigma^{(i)})} \tag{C.36}$$
$$= \exp\left[-\frac{C_R \alpha_s}{\varsigma\,\pi}\left(\log^2(2e_\varsigma^{(f)}) - \log^2(2e_\varsigma^{(i)})\right)\right].$$

We can therefore sample from the corresponding probability distribution by the inverse transform method: pick a random number r uniformly distributed between 0 and 1, and set it equal to the above cumulative distribution. The corresponding value of $e_\varsigma^{(f)}$ is

$$e_\varsigma^{(f)} = \frac{1}{2}\exp\left[-\sqrt{\log^2(2e_\varsigma^{(i)}) - \frac{\varsigma\,\pi}{C_R \alpha_s}\log(r)}\right], \tag{C.37}$$

which is taken to be the angularity of the splitting from this mother parton.

4. Given the angularity $e_\varsigma^{(f)}$ of our splitting, we may now sample from the associated distributions of the energy fraction z of the softer parton and the angle θ of the splitting. In particular, at leading logarithmic order, these are both logarithmically distributed. At this level of accuracy, we may therefore pick a uniformly distributed random variable r' and write

$$z = (2e_\varsigma^{(f)})^{r'}/2, \qquad \theta = (2e_\varsigma^{(f)})^{(1-r')/\varsigma}. \tag{C.38}$$

5. We may now continue this method recursively by returning to Step 2, allowing the partons we produced in Steps 3 and 4 to split. We can achieve results at LL and MLL accuracy if we allow only the harder of the two partons at any branch to split.

We can also extend this procedure to include running coupling effects and non-singular pieces of the splitting function, which contribute to MLL accuracy. We produce these higher-order effects by using veto method sampling [105, 151, 190]:

1. Generate an angularity, angle, and energy fraction with Steps 1–4 of the LL parton shower algorithm, with the coupling α_s of eqs. (C.36) and (C.37) replaced by a large fixed coupling $\hat{\alpha} = \frac{1}{2}$.

2. As dictated by the veto method, accept the generated variables with a probability given by the ratio $\rho^{(\text{MLL})}(z,\theta)/\hat{\rho}^{(\text{LL})}(z,\theta)$, where $\rho^{(\text{MLL})}$ indicates the MLL probability distribution function and $\hat{\rho}^{(\text{LL})}$ indicates the LL probability distribution function with the coupling $\hat{\alpha}$. In particular, generate a new random number r'' and accept the generated splitting if

$$r'' < \mathbb{P}_{\text{no veto}} = \frac{\bar{p}_i(z)}{p_{i,\,\text{gen}}(z)}\,\frac{\alpha_s(\kappa)}{\hat{\alpha}}, \tag{C.39}$$

where $p_{i,\,\text{gen}}(z) = 2C_{R_i}/z$ is the singular piece of the splitting function, used implicitly in drawing the angularity $e_\varsigma^{(f)}$ from an LL distribution in eq. (C.37). The choice $\hat{\alpha} = 1/2$ ensures that this probability is less than 1 if we freeze our coupling below the non-perturbative scale $\mu_{\text{freeze}} = 1\,\text{GeV}$.

3. If the veto method accepts the splitting, split the parton under consideration and continue the parton shower recursively by returning to Step 1. If the veto method rejects the splitting, we do *not* split the parton under consideration, but still return to Step 1 to generate a new angularity, this time starting the splitting algorithm at the scale of the rejected angularity. In particular, continue as dictated by eq. (C.37) with the coupling $\hat{\alpha}$, by setting $e_\varsigma^{(i)}$ to the angularity rejected by the veto algorithm in Step 2.

The final step is the key to the veto algorithm, and correctly takes into account the exponentiation of multiple emission contributions in the MLL case.

Additionally, figures 19a and 19c compare the LL ECF $C_1^{(2)} \approx m^2/p_T^2$ of eq. (C.1) for quark jets groomed using mMDT (Soft Drop with $\beta_{\text{SD}} = 0$) and P-RSF$_1$, obtained analytically, numerically, and using our parton shower procedure. The agreement between these calculations provides a simple validation for the results we obtain analytically, using numerical integration, and with parton shower methods for P-RSF$_1$ in the appendices.

References

[1] D. Krohn, J. Thaler and L.-T. Wang, *Jet trimming*, *JHEP* **02** (2010) 084 [arXiv:0912.1342] [INSPIRE].

[2] S.D. Ellis, C.K. Vermilion and J.R. Walsh, *Recombination algorithms and jet substructure: pruning as a tool for heavy particle searches*, *Phys. Rev. D* **81** (2010) 094023 [arXiv:0912.0033] [INSPIRE].

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

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

[5] 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].

[6] J. Tseng and H. Evans, *Sequential recombination algorithm for jet clustering and background subtraction*, *Phys. Rev. D* **88** (2013) 014044 [arXiv:1304.1025] [INSPIRE].

[7] M. Dasgupta, A. Powling, L. Schunk and G. Soyez, *Improved jet substructure methods: Y-splitter and variants with grooming*, *JHEP* **12** (2016) 079 [arXiv:1609.07149] [INSPIRE].

[8] M. Cacciari and G.P. Salam, *Pileup subtraction using jet areas*, *Phys. Lett. B* **659** (2008) 119 [arXiv:0707.1378] [INSPIRE].

[9] G. Soyez et al., *Pileup subtraction for jet shapes*, *Phys. Rev. Lett.* **110** (2013) 162001 [arXiv:1211.2811] [INSPIRE].

[10] P.T. Komiske, E.M. Metodiev, B. Nachman and M.D. Schwartz, *PileUp Mitigation with Machine Learning (PUMML)*, *JHEP* **12** (2017) 051 [arXiv:1707.08600] [INSPIRE].

[11] D. Bertolini, P. Harris, M. Low and N. Tran, *Pileup per particle identification*, *JHEP* **10** (2014) 059 [arXiv:1407.6013] [INSPIRE].

[12] M. Cacciari, G.P. Salam and G. Soyez, *The catchment area of jets*, *JHEP* **04** (2008) 005 [arXiv:0802.1188] [INSPIRE].

[13] J.M. Butterworth, A.R. Davison, M. Rubin and G.P. Salam, *Jet substructure as a new Higgs search channel at the LHC*, *Phys. Rev. Lett.* **100** (2008) 242001 [arXiv:0802.2470] [INSPIRE].

[14] 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].

[15] R. Alon et al., *A data-driven method of pile-up correction for the substructure of massive jets*, *Phys. Rev. D* **84** (2011) 114025 [arXiv:1101.3002] [INSPIRE].

[16] D. Krohn, M.D. Schwartz, M. Low and L.-T. Wang, *Jet cleansing: pileup removal at high luminosity*, *Phys. Rev. D* **90** (2014) 065020 [arXiv:1309.4777] [INSPIRE].

[17] ATLAS collaboration, *Pile-up subtraction and suppression for jets in ATLAS*, ATLAS-CONF-2013-083, CERN, Geneva, Switzerland (2013).

[18] ATLAS collaboration, *Performance of pile-up subtraction for jet shapes*, ATLAS-CONF-2013-085, CERN, Geneva, Switzerland (2013).

[19] CMS collaboration, *Measurement of the underlying event activity in pp collisions at $\sqrt{s} = 0.9$ and 7 TeV with the novel jet-area/median approach*, *JHEP* **08** (2012) 130 [arXiv:1207.2392] [INSPIRE].

[20] CMS collaboration, *Pileup jet identification*, CMS-PAS-JME-13-005, CERN, Geneva, Switzerland (2013).

[21] CMS collaboration, *Pileup mitigation at CMS in 13 TeV data*, 2020 *JINST* **15** P09018 [arXiv:2003.00503] [INSPIRE].

[22] ATLAS collaboration, *Studies of the impact and mitigation of pile-up on large-R and groomed jets in ATLAS at $\sqrt{s} = 7$ TeV*, ATLAS-CONF-2012-066, CERN, Geneva, Switzerland (2012).

[23] P. Hansen, J.W. Monk and C. Wiglesworth, *A wavelet based pile-up mitigation method for the LHC upgrade*, arXiv:1812.07412 [INSPIRE].

[24] J. Arjona Martínez et al., *Pileup mitigation at the Large Hadron Collider with graph neural networks*, *Eur. Phys. J. Plus* **134** (2019) 333 [arXiv:1810.07988] [InSPIRE].

[25] J. Thaler and L.-T. Wang, *Strategies to identify boosted tops*, *JHEP* **07** (2008) 092 [arXiv:0806.0023] [InSPIRE].

[26] J. Thaler and K. Van Tilburg, *Maximizing boosted top identification by minimizing N-subjettiness*, *JHEP* **02** (2012) 093 [arXiv:1108.2701] [InSPIRE].

[27] A. Hook, M. Jankowiak and J.G. Wacker, *Jet dipolarity: top tagging with color flow*, *JHEP* **04** (2012) 007 [arXiv:1102.1012] [InSPIRE].

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

[29] D.E. Soper and M. Spannowsky, *Finding top quarks with shower deconstruction*, *Phys. Rev. D* **87** (2013) 054012 [arXiv:1211.3140] [InSPIRE].

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

[31] CMS collaboration, *A Cambridge-Aachen (C-A) based jet algorithm for boosted top-jet tagging*, CMS-PAS-JME-09-001, CERN, Geneva, Switzerland (2009).

[32] CMS collaboration, *Search for high mass tt resonances in the all-hadronic mode*, CMS-PAS-EXO-09-002, CERN, Geneva, Switzerland (2009).

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

[34] ATLAS collaboration, *Reconstruction of high mass $t\bar{t}$ resonances in the lepton+jets channel*, ATL-PHYS-PUB-2009-081, CERN, Geneva, Switzerland (2009).

[35] ATLAS collaboration, *Prospects for top anti-top resonance searches using early ATLAS data*, ATL-PHYS-PUB-2010-008, CERN, Geneva, Switzerland (2010).

[36] Y. Cui, Z. Han and M.D. Schwartz, *W-jet tagging: optimizing the identification of boosted hadronically-decaying W bosons*, *Phys. Rev. D* **83** (2011) 074023 [arXiv:1012.2077] [InSPIRE].

[37] ATLAS collaboration, *Light-quark and gluon jets in ATLAS*, ATLAS-CONF-2011-053, CERN, Geneva, Switzerland (2011).

[38] CMS collaboration, *Studies of jet mass in dijet and W/Z + jet events*, *JHEP* **05** (2013) 090 [arXiv:1303.4811] [InSPIRE].

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

[40] M. Dasgupta, K. Khelifa-Kerfa, S. Marzani and M. Spannowsky, *On jet mass distributions in Z + jet and dijet processes at the LHC*, *JHEP* **10** (2012) 126 [arXiv:1207.1640] [InSPIRE].

[41] M. Backovic et al., *Measuring boosted tops in semi-leptonic $t\bar{t}$ events for the standard model and beyond*, *JHEP* **04** (2014) 176 [arXiv:1311.2962] [InSPIRE].

[42] ATLAS collaboration, *Performance of boosted top quark identification in 2012 ATLAS data*, ATLAS-CONF-2013-084, CERN, Geneva, Switzerland (2013).

[43] P.T. Komiske, E.M. Metodiev and J. Thaler, *An operational definition of quark and gluon jets*, *JHEP* **11** (2018) 059 [arXiv:1809.01140] [INSPIRE].

[44] 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].

[45] E.M. Metodiev and J. Thaler, *Jet topics: disentangling quarks and gluons at colliders*, *Phys. Rev. Lett.* **120** (2018) 241602 [arXiv:1802.00008] [INSPIRE].

[46] 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].

[47] ATLAS collaboration, *Measurement of jet shapes in top pair events at $\sqrt{s} = 7$ TeV using the ATLAS detector*, *PoS* **DIS2013** (2013) 142 [INSPIRE].

[48] B. Bhattacherjee et al., *Quark-gluon discrimination in the search for gluino pair production at the LHC*, *JHEP* **01** (2017) 044 [arXiv:1609.08781] [INSPIRE].

[49] S. Macaluso and D. Shih, *Pulling out all the tops with computer vision and deep learning*, *JHEP* **10** (2018) 121 [arXiv:1803.00107] [INSPIRE].

[50] S. Egan et al., *Long Short-Term Memory (LSTM) networks with jet constituents for boosted top tagging at the LHC*, arXiv:1711.09059 [INSPIRE].

[51] 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].

[52] J. Pearkes, W. Fedorko, A. Lister and C. Gay, *Jet constituents for deep neural network based top quark tagging*, arXiv:1704.02124 [INSPIRE].

[53] A. Butter, G. Kasieczka, T. Plehn and M. Russell, *Deep-learned top tagging with a Lorentz layer*, *SciPost Phys.* **5** (2018) 028 [arXiv:1707.08966] [INSPIRE].

[54] 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 [INSPIRE].

[55] Y.L. Dokshitzer, A. Lucenti, G. Marchesini and G.P. Salam, *On the QCD analysis of jet broadening*, *JHEP* **01** (1998) 011 [hep-ph/9801324] [INSPIRE].

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

[57] 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].

[58] A. Banfi and M. Dasgupta, *Problems in resumming interjet energy flows with k_t clustering*, *Phys. Lett. B* **628** (2005) 49 [hep-ph/0508159] [INSPIRE].

[59] S.D. Ellis, C.K. Vermilion and J.R. Walsh, *Techniques for improved heavy particle searches with jet substructure*, *Phys. Rev. D* **80** (2009) 051501 [arXiv:0903.5081] [INSPIRE].

[60] A. Banfi, M. Dasgupta, K. Khelifa-Kerfa and S. Marzani, *Non-global logarithms and jet algorithms in high-p_T jet shapes*, *JHEP* **08** (2010) 064 [arXiv:1004.3483] [INSPIRE].

[61] J.R. Walsh and S. Zuberi, *Factorization constraints on jet substructure*, arXiv:1110.5333 [INSPIRE].

[62] Y.-T. Chien, R. Kelley, M.D. Schwartz and H.X. Zhu, *Resummation of jet mass at hadron colliders*, *Phys. Rev. D* **87** (2013) 014010 [arXiv:1208.0010] [INSPIRE].

[63] H.-N. Li, Z. Li and C.-P. Yuan, *QCD resummation for light-particle jets*, *Phys. Rev. D* **87** (2013) 074025 [arXiv:1206.1344] [INSPIRE].

[64] T.T. Jouttenus, I.W. Stewart, F.J. Tackmann and W.J. Waalewijn, *Jet mass spectra in Higgs boson plus one jet at next-to-next-to-leading logarithmic order*, *Phys. Rev. D* **88** (2013) 054031 [arXiv:1302.0846] [INSPIRE].

[65] Y. Hatta and T. Ueda, *Resummation of non-global logarithms at finite N_c*, *Nucl. Phys. B* **874** (2013) 808 [arXiv:1304.6930] [INSPIRE].

[66] A.J. Larkoski, I. Moult and D. Neill, *Toward multi-differential cross sections: measuring two angularities on a single jet*, *JHEP* **09** (2014) 046 [arXiv:1401.4458] [INSPIRE].

[67] M. Procura, W.J. Waalewijn and L. Zeune, *Joint resummation of two angularities at next-to-next-to-leading logarithmic order*, *JHEP* **10** (2018) 098 [arXiv:1806.10622] [INSPIRE].

[68] ATLAS collaboration, *Jet reconstruction and performance using particle flow with the ATLAS detector*, *Eur. Phys. J. C* **77** (2017) 466 [arXiv:1703.10485] [INSPIRE].

[69] 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].

[70] L.G. Almeida et al., *Substructure of high-p_T jets at the LHC*, *Phys. Rev. D* **79** (2009) 074017 [arXiv:0807.0234] [INSPIRE].

[71] A.J. Larkoski, I. Moult and D. Neill, *Analytic boosted boson discrimination at the Large Hadron Collider*, arXiv:1708.06760 [INSPIRE].

[72] A.J. Larkoski, I. Moult and D. Neill, *Factorization and resummation for groomed multi-prong jet shapes*, *JHEP* **02** (2018) 144 [arXiv:1710.00014] [INSPIRE].

[73] J. Thaler and K. Van Tilburg, *Identifying boosted objects with N-subjettiness*, *JHEP* **03** (2011) 015 [arXiv:1011.2268] [INSPIRE].

[74] 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].

[75] A. Katz, M. Son and B. Tweedie, *Jet substructure and the search for neutral spin-one resonances in electroweak boson channels*, *JHEP* **03** (2011) 011 [arXiv:1010.5253] [INSPIRE].

[76] J. Gallicchio et al., *Multivariate discrimination and the Higgs + W/Z search*, *JHEP* **04** (2011) 069 [arXiv:1010.3698] [INSPIRE].

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

[78] CMS collaboration, *Identification of heavy-flavour jets with the CMS detector in pp collisions at 13 TeV*, 2018 *JINST* **13** P05011 [arXiv:1712.07158] [INSPIRE].

[79] L. Moore, K. Nordström, S. Varma and M. Fairbairn, *Reports of my demise are greatly exaggerated: N-subjettiness taggers take on jet images*, *SciPost Phys.* **7** (2019) 036 [arXiv:1807.04769] [INSPIRE].

[80] 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].

[81] M. Rubin, *Non-global logarithms in filtered jet algorithms*, *JHEP* **05** (2010) 005 [arXiv:1002.4557] [InSPIRE].

[82] ATLAS collaboration, *Measurement of jet-substructure observables in top quark, W boson and light jet production in proton-proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, *JHEP* **08** (2019) 033 [arXiv:1903.02942] [InSPIRE].

[83] CMS collaboration, *Performance of b tagging at $\sqrt{s} = 8$ TeV in multijet, ttbar and boosted topology events*, CMS-PAS-BTV-13-001, CERN, Geneva, Switzerland (2013).

[84] CMS collaboration, *Identifying hadronically decaying vector bosons merged into a single jet*, CMS-PAS-JME-13-006, CERN, Geneva, Switzerland (2013).

[85] G.D. Kribs, A. Martin, T.S. Roy and M. Spannowsky, *Discovering the Higgs boson in new physics events using jet substructure*, *Phys. Rev. D* **81** (2010) 111501 [arXiv:0912.4731] [InSPIRE].

[86] C.-R. Chen, M.M. Nojiri and W. Sreethawong, *Search for the elusive Higgs boson using jet structure at LHC*, *JHEP* **11** (2010) 012 [arXiv:1006.1151] [InSPIRE].

[87] C. Hackstein and M. Spannowsky, *Boosting Higgs discovery: the forgotten channel*, *Phys. Rev. D* **82** (2010) 113012 [arXiv:1008.2202] [InSPIRE].

[88] J.-H. Kim, *Rest frame subjet algorithm with SISCone jet for fully hadronic decaying Higgs search*, *Phys. Rev. D* **83** (2011) 011502 [arXiv:1011.1493] [InSPIRE].

[89] L.G. Almeida et al., *Three-particle templates for a boosted Higgs boson*, *Phys. Rev. D* **85** (2012) 114046 [arXiv:1112.1957] [InSPIRE].

[90] F. Pandolfi, *Search for the Standard Model Higgs boson in the $H \to ZZ \to l^+l^- q\bar{q}$ decay channel at CMS*, Ph.D. thesis, ETH, Zurich, Switzerland (2012) [DOI:10.1007/978-3-319-00903-2] [InSPIRE].

[91] CMS collaboration, *Search for the Higgs boson in the $b\bar{b}$ decay channel using the CMS detector*, *Nucl. Part. Phys. Proc.* **273-275** (2016) 733 [arXiv:1411.4362] [InSPIRE].

[92] CMS collaboration, *Search for the standard model Higgs boson in the dilepton plus photon channel in pp collisions at $\sqrt{s} = 13$ TeV*, CMS-PAS-HIG-17-007, CERN, Geneva, Switzerland (2018).

[93] 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].

[94] P.T. Komiske et al., *Exploring the space of jets with CMS open data*, *Phys. Rev. D* **101** (2020) 034009 [arXiv:1908.08542] [InSPIRE].

[95] P.T. Komiske, E.M. Metodiev and J. Thaler, *Metric space of collider events*, *Phys. Rev. Lett.* **123** (2019) 041801 [arXiv:1902.02346] [InSPIRE].

[96] A.H. Hoang, S. Mantry, A. Pathak and I.W. Stewart, *Nonperturbative corrections to soft drop jet mass*, *JHEP* **12** (2019) 002 [arXiv:1906.11843] [InSPIRE].

[97] V. Mateu, I.W. Stewart and J. Thaler, *Power corrections to event shapes with mass-dependent operators*, *Phys. Rev. D* **87** (2013) 014025 [arXiv:1209.3781] [InSPIRE].

[98] M. Wobisch and T. Wengler, *Hadronization corrections to jet cross-sections in deep inelastic scattering*, in the proceedings of the *Workshop on Monte Carlo generators for HERA physics (plenary starting meeting)*, (1998), p. 270 [hep-ph/9907280] [InSPIRE].

[99] ATLAS collaboration, *Impact of alternative inputs and jet grooming on large-R jet performance*, ATL-PHYS-PUB-2019-027, CERN, Geneva, Switzerland (2019).

[100] ATLAS collaboration, *Measurement of soft-drop jet observables in pp collisions with the ATLAS detector at* $\sqrt{s} = 13$ *TeV*, *Phys. Rev. D* **101** (2020) 052007 [arXiv:1912.09837] [INSPIRE].

[101] ATLAS collaboration, *Optimisation of large-radius jet reconstruction for the ATLAS detector in* 13 *TeV proton-proton collisions*, *Eur. Phys. J. C* **81** (2021) 334 [arXiv:2009.04986] [INSPIRE].

[102] P.T. Komiske, E.M. Metodiev and J. Thaler, *The hidden geometry of particle collisions*, *JHEP* **07** (2020) 006 [arXiv:2004.04159] [INSPIRE].

[103] G. Soyez, *Pileup mitigation at the LHC: a theorist's view*, *Phys. Rept.* **803** (2019) 1 [arXiv:1801.09721] [INSPIRE].

[104] T. Sjostrand and M. van Zijl, *A multiple interaction model for the event structure in hadron collisions*, *Phys. Rev. D* **36** (1987) 2019 [INSPIRE].

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

[106] M. Dasgupta, L. Magnea and G.P. Salam, *Non-perturbative QCD effects in jets at hadron colliders*, *JHEP* **02** (2008) 055 [arXiv:0712.3014] [INSPIRE].

[107] P. Kirchgaeßer, *Non-perturbative models for the simulation of hadronic collisions at the LHC*, Ph.D. thesis, KIT, Karlsruhe, Germany (2020) [DOI:10.5445/IR/1000120873] [INSPIRE].

[108] A. Moraes, C. Buttar and I. Dawson, *Prediction for minimum bias and the underlying event at LHC energies*, *Eur. Phys. J. C* **50** (2007) 435 [INSPIRE].

[109] CDF collaboration, *Study of the energy dependence of the underlying event in proton-antiproton collisions*, *Phys. Rev. D* **92** (2015) 092009 [arXiv:1508.05340] [INSPIRE].

[110] A.J. Larkoski and T. Melia, *A large-N expansion for minimum bias*, *JHEP* **10** (2021) 094 [arXiv:2107.04041] [INSPIRE].

[111] J. Baron et al., *Soft-drop grooming for hadronic event shapes*, *JHEP* **07** (2021) 142 [arXiv:2012.09574] [INSPIRE].

[112] S. Marzani, L. Schunk and G. Soyez, *The jet mass distribution after soft drop*, *Eur. Phys. J. C* **78** (2018) 96 [arXiv:1712.05105] [INSPIRE].

[113] P. Komiske, *PIRANHA GitHub repository*, https://github.com/pkomiske/Piranha (2022).

[114] V. Hartmann, *A geometry-based approach for solving the transportation problem with Euclidean cost*, arXiv:1706.07403.

[115] V. Hartmann and D. Schuhmacher, *Semi-discrete optimal transport — the case* $p = 1$, arXiv:1706.07650.

[116] D.P. Bourne, B. Schmitzer and B. Wirth, *Semi-discrete unbalanced optimal transport and quantization*, arXiv:1808.01962.

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

[118] P. Berta, M. Spousta, D.W. Miller and R. Leitner, *Particle-level pileup subtraction for jets and jet shapes*, *JHEP* **06** (2014) 092 [arXiv:1403.3108] [INSPIRE].

[119] T. Cai, J. Cheng, B. Schmitzer and M. Thorpe, *The linearized Hellinger-Kantorovich distance*, *SIAM J. Imaging Sci.* **15** (2022) 45 [arXiv:2102.08807].

[120] C. Sarrazin and B. Schmitzer, *Linearized optimal transport on manifolds*, arXiv:2303.13901.

[121] T. Cai, J. Cheng, N. Craig and K. Craig, *Linearized optimal transport for collider events*, *Phys. Rev. D* **102** (2020) 116019 [arXiv:2008.08604] [INSPIRE].

[122] T. Cai, J. Cheng, K. Craig and N. Craig, *Which metric on the space of collider events?*, *Phys. Rev. D* **105** (2022) 076003 [arXiv:2111.03670] [INSPIRE].

[123] J. Collins, *Foundations of perturbative QCD*, Cambridge University Press, Cambridge, U.K. (2013) [DOI:10.1017/9781009401845] [INSPIRE].

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

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

[126] A.H. Hoang, S. Mantry, A. Pathak and I.W. Stewart, *Extracting a short distance top mass with light grooming*, *Phys. Rev. D* **100** (2019) 074021 [arXiv:1708.02586] [INSPIRE].

[127] ATLAS collaboration, *A precise interpretation for the top quark mass parameter in ATLAS Monte Carlo simulation*, ATL-PHYS-PUB-2021-034, CERN, Geneva, Switzerland (2021) [INSPIRE].

[128] ATLAS and CMS collaborations, *Recent measurements of the top-quark mass and Yukawa coupling using the ATLAS and CMS detectors at the LHC*, *PoS* **EPS-HEP2021** (2022) 479 [INSPIRE].

[129] S. Marzani et al., *Fitting the strong coupling constant with soft-drop thrust*, *JHEP* **11** (2019) 179 [arXiv:1906.10504] [INSPIRE].

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

[131] V.N. Gribov and L.N. Lipatov, *Deep inelastic ep scattering in perturbation theory*, *Sov. J. Nucl. Phys.* **15** (1972) 438 [INSPIRE].

[132] 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 [INSPIRE].

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

[134] K. Benkendorfer and A.J. Larkoski, *Grooming at the cusp: all-orders predictions for the transition region of jet groomers*, *JHEP* **11** (2021) 188 [arXiv:2108.02779] [INSPIRE].

[135] A. Banfi, G.P. Salam and G. Zanderighi, *Semi-numerical resummation of event shapes*, *JHEP* **01** (2002) 018 [hep-ph/0112156] [INSPIRE].

[136] S. Brandt, C. Peyrou, R. Sosnowski and A. Wroblewski, *The principal axis of jets. An attempt to analyze high-energy collisions as two-body processes*, *Phys. Lett.* **12** (1964) 57 [INSPIRE].

[137] E. Farhi, *A QCD test for jets*, *Phys. Rev. Lett.* **39** (1977) 1587 [INSPIRE].

[138] S.D. Ellis et al., *Qjets: a non-deterministic approach to tree-based jet substructure*, *Phys. Rev. Lett.* **108** (2012) 182003 [arXiv:1201.1914] [INSPIRE].

[139] S.D. Ellis, A. Hornig, D. Krohn and T.S. Roy, *On statistical aspects of Qjets*, *JHEP* **01** (2015) 022 [arXiv:1409.6785] [INSPIRE].

[140] R. Corke and T. Sjostrand, *Interleaved parton showers and tuning prospects*, *JHEP* **03** (2011) 032 [arXiv:1011.1759] [INSPIRE].

[141] Y.L. Dokshitzer and B.R. Webber, *Calculation of power corrections to hadronic event shapes*, *Phys. Lett. B* **352** (1995) 451 [hep-ph/9504219] [INSPIRE].

[142] R. Field, *Min-bias and the underlying event at the LHC*, *Acta Phys. Polon. B* **42** (2011) 2631 [arXiv:1110.5530] [INSPIRE].

[143] J. Campbell, J. Huston and F. Krauss, *The black book of quantum chromodynamics: a primer for the LHC era*, Oxford University Press, Oxford, U.K. (2018) [DOI:10.1093/oso/9780199652747.001.0001] [INSPIRE].

[144] A.J. Larkoski, I. Moult and D. Neill, *Non-global logarithms, factorization, and the soft substructure of jets*, *JHEP* **09** (2015) 143 [arXiv:1501.04596] [INSPIRE].

[145] Y.I. Azimov, Y.L. Dokshitzer, V.A. Khoze and S.I. Troyan, *Similarity of parton and hadron spectra in QCD jets*, *Z. Phys. C* **27** (1985) 65 [INSPIRE].

[146] B. Andersson, P. Dahlqvist and G. Gustafson, *On local parton hadron duality*. 1. *Multiplicity*, *Z. Phys. C* **44** (1989) 455 [INSPIRE].

[147] M. Shifman, *The quark hadron duality*, *eConf* **C030614** (2003) 001 [INSPIRE].

[148] D. Neill, *The fragmentation spectrum from space-time reciprocity*, *JHEP* **03** (2021) 081 [arXiv:2010.02934] [INSPIRE].

[149] Z. Ligeti, I.W. Stewart and F.J. Tackmann, *Treating the b quark distribution function with reliable uncertainties*, *Phys. Rev. D* **78** (2008) 114014 [arXiv:0807.1926] [INSPIRE].

[150] M.A. Ebert and F.J. Tackmann, *Resummation of transverse momentum distributions in distribution space*, *JHEP* **02** (2017) 110 [arXiv:1611.08610] [INSPIRE].

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

[152] K. Konishi, A. Ukawa and G. Veneziano, *Jet calculus: a simple algorithm for resolving QCD jets*, *Nucl. Phys. B* **157** (1979) 45 [INSPIRE].

[153] Y.L. Dokshitzer, V.A. Khoze, A.H. Mueller and S.I. Troian, *Basics of perturbative QCD*, (1991) [INSPIRE].

[154] M. van Beekveld et al., *Collinear fragmentation at NNLL: generating functionals, groomed correlators and angularities*, arXiv:2307.15734 [INSPIRE].

[155] A.J. Larkoski, *An unorthodox introduction to QCD*, arXiv:1709.06195 [INSPIRE].

[156] A.J. Larkoski, *Jet physics from the ground up*, arXiv:2112.15122 [INSPIRE].

[157] T. Cohen, J. Doss and M. Freytsis, *Jet substructure from dark sector showers*, *JHEP* **09** (2020) 118 [arXiv:2004.00631] [INSPIRE].

[158] A. Banfi, G.P. Salam and G. Zanderighi, *Resummed event shapes at hadron-hadron colliders*, *JHEP* **08** (2004) 062 [hep-ph/0407287] [INSPIRE].

[159] A. Banfi, G.P. Salam and G. Zanderighi, *Phenomenology of event shapes at hadron colliders*, *JHEP* **06** (2010) 038 [arXiv:1001.4082] [INSPIRE].

[160] G. Luisoni and S. Marzani, *QCD resummation for hadronic final states*, *J. Phys. G* **42** (2015) 103101 [arXiv:1505.04084] [INSPIRE].

[161] C.W. Bauer and P.F. Monni, *A numerical formulation of resummation in effective field theory*, *JHEP* **02** (2019) 185 [arXiv:1803.07079] [INSPIRE].

[162] A.C. Mattingly and P.M. Stevenson, *QCD perturbation theory at low-energies*, *Phys. Rev. Lett.* **69** (1992) 1320 [hep-ph/9207228] [INSPIRE].

[163] P.M. Stevenson, *Optimized perturbation theory*, *Phys. Rev. D* **23** (1981) 2916 [INSPIRE].

[164] P.M. Stevenson, *The* $1612 - N_f$ *expansion and the infrared fixed point in perturbative QCD*, *Phys. Lett. B* **331** (1994) 187 [hep-ph/9402276] [INSPIRE].

[165] S.A. Caveny and P.M. Stevenson, *The Banks-Zaks expansion and 'freezing' in perturbative QCD*, hep-ph/9705319 [INSPIRE].

[166] W.E. Caswell, *Asymptotic behavior of non-Abelian gauge theories to two loop order*, *Phys. Rev. Lett.* **33** (1974) 244 [INSPIRE].

[167] T. Banks and A. Zaks, *On the phase structure of vector-like gauge theories with massless fermions*, *Nucl. Phys. B* **196** (1982) 189 [INSPIRE].

[168] S.J. Brodsky, G.F. de Teramond and A. Deur, *Nonperturbative QCD coupling and its* β-*function from light-front holography*, *Phys. Rev. D* **81** (2010) 096010 [arXiv:1002.3948] [INSPIRE].

[169] A. Deur, V. Burkert, J.P. Chen and W. Korsch, *Determination of the effective strong coupling constant* $\alpha_{s,g_1}(Q^2)$ *from CLAS spin structure function data*, *Phys. Lett. B* **665** (2008) 349 [arXiv:0803.4119] [INSPIRE].

[170] A. Deur, *The strong coupling constant at large distances*, *AIP Conf. Proc.* **1149** (2009) 281 [arXiv:0901.2190] [INSPIRE].

[171] D. Binosi et al., *Process-independent strong running coupling*, *Phys. Rev. D* **96** (2017) 054026 [arXiv:1612.04835] [INSPIRE].

[172] Y.L. Dokshitzer, G. Marchesini and B.R. Webber, *Dispersive approach to power behaved contributions in QCD hard processes*, *Nucl. Phys. B* **469** (1996) 93 [hep-ph/9512336] [INSPIRE].

[173] Y.L. Dokshitzer and B.R. Webber, *Power corrections to event shape distributions*, *Phys. Lett. B* **404** (1997) 321 [hep-ph/9704298] [INSPIRE].

[174] Y.L. Dokshitzer, A. Lucenti, G. Marchesini and G.P. Salam, *Universality of* $1/Q$ *corrections to jet-shape observables rescued*, *Nucl. Phys. B* **511** (1998) 396 [hep-ph/9707532] [INSPIRE].

[175] G.P. Korchemsky and G.F. Sterman, *Power corrections to event shapes and factorization*, *Nucl. Phys. B* **555** (1999) 335 [hep-ph/9902341] [INSPIRE].

[176] A. Deur, S.J. Brodsky and G.F. de Teramond, *The QCD running coupling*, *Nucl. Phys.* **90** (2016) 1 [arXiv:1604.08082] [INSPIRE].

[177] M. Dasgupta and G.P. Salam, *Resummed event shape variables in DIS*, *JHEP* **08** (2002) 032 [`hep-ph/0208073`] [INSPIRE].

[178] M. Dasgupta and G.P. Salam, *Accounting for coherence in interjet E_t flow: a case study*, *JHEP* **03** (2002) 017 [`hep-ph/0203009`] [INSPIRE].

[179] A. Banfi, G. Marchesini and G. Smye, *Away from jet energy flow*, *JHEP* **08** (2002) 006 [`hep-ph/0206076`] [INSPIRE].

[180] R.B. Appleby and M.H. Seymour, *Nonglobal logarithms in interjet energy flow with k_t clustering requirement*, *JHEP* **12** (2002) 063 [`hep-ph/0211426`] [INSPIRE].

[181] H. Weigert, *Nonglobal jet evolution at finite N_c*, *Nucl. Phys. B* **685** (2004) 321 [`hep-ph/0312050`] [INSPIRE].

[182] R. Kelley, M.D. Schwartz and H.X. Zhu, *Resummation of jet mass with and without a jet veto*, `arXiv:1102.0561` [INSPIRE].

[183] A. Hornig et al., *Non-global structure of the $O(\alpha_s^2)$ dijet soft function*, *JHEP* **08** (2011) 054 [*Erratum ibid.* **10** (2017) 101] [`arXiv:1105.4628`] [INSPIRE].

[184] R. Kelley, M.D. Schwartz, R.M. Schabinger and H.X. Zhu, *Jet mass with a jet veto at two loops and the universality of non-global structure*, *Phys. Rev. D* **86** (2012) 054017 [`arXiv:1112.3343`] [INSPIRE].

[185] M.D. Schwartz and H.X. Zhu, *Nonglobal logarithms at three loops, four loops, five loops, and beyond*, *Phys. Rev. D* **90** (2014) 065004 [`arXiv:1403.4949`] [INSPIRE].

[186] K. Khelifa-Kerfa and Y. Delenda, *Non-global logarithms at finite N_c beyond leading order*, *JHEP* **03** (2015) 094 [`arXiv:1501.00475`] [INSPIRE].

[187] A.J. Larkoski, I. Moult and D. Neill, *The analytic structure of non-global logarithms: convergence of the dressed gluon expansion*, *JHEP* **11** (2016) 089 [`arXiv:1609.04011`] [INSPIRE].

[188] A. Banfi, F.A. Dreyer and P.F. Monni, *Next-to-leading non-global logarithms in QCD*, *JHEP* **10** (2021) 006 [`arXiv:2104.06416`] [INSPIRE].

[189] S. Alipour-fard, *JETMONTECARLO GitHub repository*, https://github.com/samcaf/JetMonteCarlo (2022).

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