

MIT Open Access Articles

Pileup per particle identification

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Bertolini, Daniele, Philip Harris, Matthew Low, and Nhan Tran. "Pileup Per Particle Identification." J. High Energ. Phys. 2014, no. 10 (October 2014).

As Published: [http://dx.doi.org/10.1007/jhep10\(2014\)059](http://dx.doi.org/10.1007/jhep10(2014)059)

Publisher: Springer-Verlag

Persistent URL: <http://hdl.handle.net/1721.1/92526>

Version: Final published version: final published article, as it appeared in a journal, conference proceedings, or other formally published context

Terms of use: Creative Commons Attribution



Pileup per particle identification

Daniele Bertolini,^a Philip Harris,^b Matthew Low^c and Nhan Tran^d

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

^b*CERN, European Organization for Nuclear Research,
Geneva, Switzerland*

^c*Enrico Fermi Institute and Kavli Institute for Cosmological Physics, University of Chicago,
Chicago, IL 60637, U.S.A.*

^d*Fermi National Accelerator Laboratory (FNAL),
Batavia, IL 60510, U.S.A.*

E-mail: danbert@mit.edu, philip.coleman.harris@cern.ch,
mattlow@uchicago.edu, ntran@fnal.gov

ABSTRACT: We propose a new method for pileup mitigation by implementing “pileup per particle identification” (PUPPI). For each particle we first define a local shape α which probes the collinear versus soft diffuse structure in the neighborhood of the particle. The former is indicative of particles originating from the hard scatter and the latter of particles originating from pileup interactions. The distribution of α for charged pileup, assumed as a proxy for all pileup, is used on an event-by-event basis to calculate a weight for each particle. The weights describe the degree to which particles are pileup-like and are used to rescale their four-momenta, superseding the need for jet-based corrections. Furthermore, the algorithm flexibly allows combination with other, possibly experimental, probabilistic information associated with particles such as vertexing and timing performance. We demonstrate the algorithm improves over existing methods by looking at jet p_T and jet mass. We also find an improvement on non-jet quantities like missing transverse energy.

KEYWORDS: Jets, Hadronic Colliders

ARXIV EPRINT: [1407.6013](https://arxiv.org/abs/1407.6013)

Contents

1	Introduction	1
2	The algorithm	3
2.1	The local shape	4
2.2	Particle weights	6
2.3	Incorporating additional information	9
2.4	Choice of metric	9
3	Simulation details	10
4	Results	11
4.1	Jet kinematics	13
4.2	Jet shapes	14
4.3	Missing transverse energy	15
5	Summary and outlook	17

1 Introduction

Pileup, i.e. overlapping secondary proton-proton collisions on top of the primary interaction, will be a major challenge for the high luminosity LHC runs. Several methods for dealing with pileup are being successfully applied by ATLAS [1–6] and CMS [7–9] on present data. Current methods, however, will be strained both in upcoming LHC runs with expected pileup levels of $n_{\text{PU}} = 140$ or more, and at possible future hadron colliders. Recently, newer ideas for pileup mitigation have been proposed. A brief summary of the state-of-the-art is given below:

- *Four-vector area subtraction* [10, 11]: corrects the four-vector of a jet based on the characteristic pileup density of an event and on the jet area. It has been applied by ATLAS and CMS, but requires additional experimental tuning on top of area \times pileup density subtraction [2, 7].
- *Shape subtraction* [12]: a generalization of area subtraction from the jet four-vector to jet shapes, e.g. jet mass. Each shape is separately corrected using the same pileup density measure as area subtraction and using the susceptibility of a given shape to soft uniform contamination.
- *Grooming (e.g. filtering [13], pruning [14, 15], trimming [16], soft drop [17])*: systematically discards a subset of a jet’s constituents that are believed to obscure the signal process. Grooming can be used with subtraction.

- *Pileup jet identification* [2, 4]: removes entire jets that are identified as being composed primarily of pileup using both charged particle information and jet shapes.
- *Topological clustering* [18–20]: calorimeter cell signals are required to be several standard deviations above the typical noise level in the cells. These high significance cells are used as seeds to form local cell clusters used as inputs to jet algorithms.
- *Charged hadron subtraction (CHS)* [21]: removes charged particles from pileup based on the vertex from which particles originate. Four-vector area subtraction is then applied using the remaining particles.
- *Cleansing*¹ [22]: uses vertex information to remove charged pileup and rescales neutral pileup based on charged pileup composition in a subjet.
- *Constituent subtraction* [23, 24]: extends four-vector area subtraction and shape subtraction to particle level by representing the event density ρ with an area assigned to each particle, correcting the particles’s four-vector.

The methods listed above progressively move from a global approach towards a more local one. We note that, broadly speaking, the methods utilize three basic pieces of information to identify pileup: the event-wide pileup density, vertex information from charged tracks, and the local distribution of pileup with respect to particles from the leading vertex. As each technique has advantages and disadvantages it is unlikely that a single method alone will optimally remove pileup. It is therefore crucial to have a flexible framework to integrate the various pieces of information. We propose an algorithm to combine this information. This method uses both global information from the event, as well as local information to identify pileup at the particle level. As a shorthand, we refer to the method as PUPPI (PileUp Per Particle Identification). The algorithm is intended to remove pileup, rather than just correct jet quantities, to ultimately produce a consistent event interpretation.

It has been shown [22] that individually rescaling the four-momenta of particles in a jet (i.e. the jet’s constituents) is useful not only for correcting kinematic variables, but also for correcting jet shapes in an observable-independent way. Following the “jets without jets” paradigm [25], we propose a local approach, in which no clustering is performed and a weight is assigned to each individual particle. We then choose to use the weight to rescale the particle four-momentum. Ideally, particles coming from pileup would get a weight of zero and particles coming from the hard scatter would get a weight of one. This leads to a pileup-corrected event, where one can then proceed with jet finding without the need for further pileup correction. In fact, given the pileup-corrected event, event shapes can be measured with a reduced sensitivity to pileup. We show results for jet p_T , jet mass, and missing transverse energy and demonstrate that our algorithm improves over existing methods. We find the improvements on missing transverse energy reconstruction particularly relevant, as disentangling pileup contamination from missing energy is typically harder than for jet-based observables.

¹Depending on the precise definition of grooming, this may also be considered a groomer.

We anticipate that the PUPPI algorithm could potentially improve pileup mitigation for other jet and event shapes, as well as the identification of isolated photons and leptons. More generally, such a *per particle* approach may contribute valuable input into the design of future detectors by highlighting the complementary information measured by several detector subsystems.

The rest of the paper is organized as follows: in section 2 we describe the algorithm, in section 3 we describe the setup we used for generating Monte Carlo data, and in section 4 we present our results. Finally, we conclude and discuss future work in section 5.

2 The algorithm

Before discussing details, we describe qualitatively how the algorithm works. First we select a shape α , which attempts to locally distinguish parton shower-like radiation from pileup-like radiation, then compute it for each particle in an event. A basic handle to distinguish pileup and leading vertex particles is given by the p_T spectrum, with the pileup spectrum falling much faster. While we make use of this feature in the algorithm, the shape α itself attempts to exploit additional and complementary information with respect to the p_T of a single particle, as discussed in sections 2.1 and 2.4. Where tracking is available, we know the answer to whether charged particles are from the leading vertex (LV) or from a pileup vertex (PU). We can use the median and the RMS of the α values for charged pileup as an event-level characterization of the pileup distribution.

Next we assign a weight to each particle by comparing its α value to the median of the charged pileup distribution. The weight takes values between zero and one and indicates how much a particle is allowed to contribute to an event. Ideally, particles from the hard scatter would get a weight of one and pileup particles would get a weight of zero. Almost all pileup particles have α values within a few standard deviations of the median and are assigned small weights. On the other hand, α values that deviate far from the charged pileup are very uncharacteristic of pileup, and these particles are assigned large weights. As discussed in section 2.3, our method for computing weights allows for experimental information, such as vertexing and timing performance, to be smoothly incorporated.

Finally, we choose to apply the weights to rescale the particle's four-momentum. Particles with a very small weight or with a very small rescaled p_T are discarded. The final set of pileup-corrected particles can now be used as input to a jet algorithm or directly in the calculation of missing energy.

The rest of this section goes through the algorithm in detail. We use a $pp \rightarrow$ dijet sample at $\sqrt{s} = 14$ TeV generated with Pythia 8.176 [26] to show various distributions. The spectrum is generated such that the p_T of the $2 \rightarrow 2$ process is roughly flat across the range 15 – 500 GeV, in order to maintain reasonable statistics across different kinematic ranges. Pileup events are generated as zero-bias soft QCD events and overlaid onto the hard scatter event. Further details of the simulation are discussed in section 3.

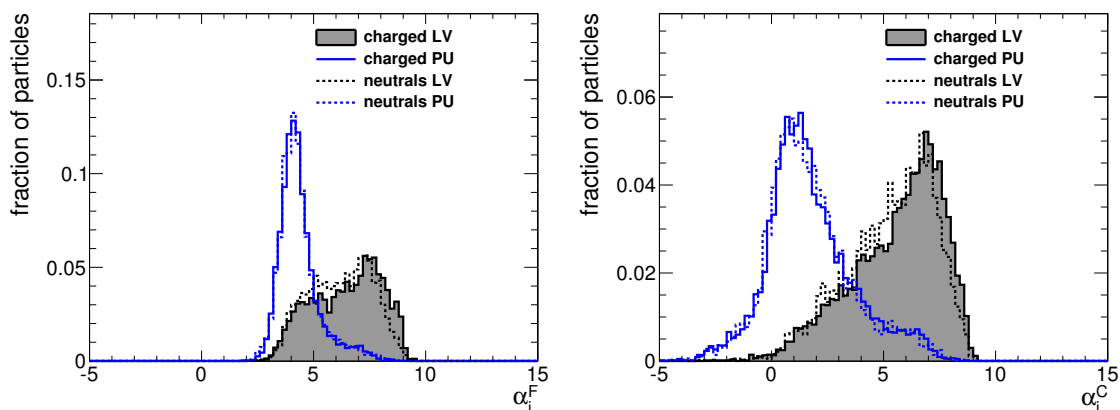


Figure 1. The distribution of α_i , over many events, for particles i from the leading vertex (gray filled) and particles from pileup (blue) in a dijet sample. For α_i^F (left) we sum over all particles as defined in eqs. (2.1) or (2.4), for α_i^C (right) we sum over charged particles from the leading vertex as defined in eq. (2.3). Both distributions consider only particles with a $p_T > 1$ GeV. Dotted and solid lines refer to neutral and charged particles respectively.

2.1 The local shape

For each particle i we define a shape

$$\alpha_i = \log \sum_{j \in \text{event}} \xi_{ij} \times \Theta(R_{\min} \leq \Delta R_{ij} \leq R_0), \tag{2.1}$$

where $\xi_{ij} = \frac{p_{Tj}}{\Delta R_{ij}}$.

Throughout the paper we use $\Theta(R_{\min} \leq \Delta R_{ij} \leq R_0)$ as a shorthand notation for $\Theta(\Delta R_{ij} - R_{\min}) \times \Theta(R_0 - \Delta R_{ij})$, where Θ is the Heaviside step function. ΔR_{ij} is the distance between particles i and j in $\eta\phi$ -space and p_{Tj} is the transverse momentum of particle j measured in units of GeV. R_0 defines a cone around each particle i , so that only particles within the cone enter the calculation of α_i . In addition, particles closer to i than R_{\min} are discarded from the sum, with R_{\min} effectively serving as a regulator for collinear splittings of particle i . Here we use $R_0 = 0.3$ and $R_{\min} = 0.02$.² Note that the logarithm is outside of the sum so it plays no role in the infrared-collinear behavior of the variable and just serves to rescale the range. The choice of ξ_{ij} is discussed in more detail in section 2.4.

Figure 1 (left) shows a sample distribution of α for particles from the leading vertex and pileup. Due to the collinear singularity of the parton shower, a particle i from a hard physics process is likely to be near other particles from the same parent process so that α_i tends to be larger. On the other hand, we expect pileup particles to have no shower-like structure and to be uncorrelated with particles from the leading vertex and so only to be spatially near by chance.³ So, α_i tends to be smaller if i is a pileup particle. In fact, this

²This choice of R_{\min} is related to typical detector resolutions, as discussed in more detail in section 3.

³It is worth noting that stochastic pileup jets can be found by jet algorithms. This is due to locally high pileup densities rather than a sequence of collinear parton branchings. As this results in a different radiation pattern on average, pileup jet removal uses differences in jet shapes to remove these pileup jets [4].

implies that the ideal version of eq. (2.1) would sum over particles from the leading vertex and ignore those from pileup. While we obviously do not know a priori which particles are from the leading vertex, we do have a handle on charged particles in the central ($|\eta| \lesssim 2.5$ for ATLAS and CMS) region of the detector. In that region, tracking information provides the ability to distinguish charged tracks originating from the leading vertex and charged tracks originating from pileup. Associating these tracks to particles can be done with the particle flow algorithm [21] which combines measurements from various detector subsystems to define individual candidates.⁴ Using particle flow, identified candidates can be sorted into three classes: neutral particles, charged hadrons from the leading vertex, and charged hadrons from pileup. Thus we can use charged particles from the leading vertex as a proxy for all particles from the leading vertex. To be explicit, in the central region the sum in eq. (2.1) can be decomposed as

$$\sum_j = \sum_{j \in \text{Ch,PU}} + \sum_{j \in \text{Ch,LV}} + \sum_{j \in \text{Neutral}} , \quad (2.2)$$

where Ch,PU refers to charged pileup, Ch,LV refers to charged particles from the leading vertex, and Neutral refers to all neutral particles both from pileup and the leading vertex. This leads to defining two versions of α for when tracking information is and is not available.

$$\alpha_i^C = \log \sum_{j \in \text{Ch,LV}} \xi_{ij} \Theta(R_{\min} \leq \Delta R_{ij} \leq R_0), \quad (2.3)$$

$$\alpha_i^F = \log \sum_{j \in \text{event}} \xi_{ij} \Theta(R_{\min} \leq \Delta R_{ij} \leq R_0). \quad (2.4)$$

Notice that $\alpha_i^F \equiv \alpha_i$ in eq. (2.1). Here it is renamed to stress the fact that we use this version of α_i in the forward region of the detector, as opposed to α_i^C which is used in the central region. Effectively, when tracking information is not available, we assume all particles in the sum are from the leading vertex. While there are noise contributions from pileup, these are suppressed relative to contributions from leading vertex particles by the p_{Tj} in the numerator. Thus the algorithm can still assign weights in regions where there is no tracking.

Figure 1 (right) shows the distributions of α^C . When there are no particles from the leading vertex around particle i to sum over, formally $\alpha_i \rightarrow -\infty$. In these cases the particle is assumed to be pileup and discarded from the event.⁵ The variable α^C has more discrimination power than α^F and is used in the central region of the detector.

There is a second advantage to be gained from tracking information. For central charged particles, we know the answer as to whether a particle is from pileup or not. Using only these particles, for a given event we can compute the distribution of both α^C and α^F and then assume that the neutral particles, for which we do not know the answer, belong

⁴The use of particles is not strictly necessary. In principle the algorithm can be performed with calorimeter cells and charged tracks as inputs. We discuss this later in section 5.

⁵The fact that, in practice, the appearance of a single isolated particle occurs much more frequently in pileup (with a moderately low number of pileup interactions) than in hard interactions, supports this assignment scheme.

to a distribution with the same properties. This assumes the distribution of α^F and α^C is the same for charged and neutral particles, and for central and forward particles. Neither of these assumptions is exact, but they both can be corrected if necessary. As an example, in figure 1, we show the distribution of α for neutral and charged particles separately and we find good agreement overall.

The quantities we use to characterize the distributions on an event-by-event basis are the median and the left-side RMS:⁶

$$\bar{\alpha}_{\text{PU}}^F = \text{median}\{\alpha_{i \in \text{Ch,PU}}^F\}, \quad \sigma_{\text{PU}}^F = \text{RMS}\{\alpha_{i \in \text{Ch,PU}}^F\}, \quad (2.5)$$

$$\bar{\alpha}_{\text{PU}}^C = \text{median}\{\alpha_{i \in \text{Ch,PU}}^C\}, \quad \sigma_{\text{PU}}^C = \text{RMS}\{\alpha_{i \in \text{Ch,PU}}^C\}. \quad (2.6)$$

The characterization of pileup contamination on an event basis is reminiscent of the area subtraction method where average information over an entire event is used to correct individual jets within the event [10, 11]. In the absence of vertex based discrimination, the median of α can be computed by taking the median over all particles as is done for the area subtraction method.

Because the computation of $\bar{\alpha}_{\text{PU}}^F$ and σ_{PU}^F is only done for charged pileup, it must be computed in the central region, even though these quantities are used to calculate the weights of particles in the forward region. Pileup density varies as a function of rapidity and the values of $\bar{\alpha}_{\text{PU}}^F$ and σ_{PU}^F do not account for this variation. A proper extrapolation can be performed by estimating the rapidity-dependence in a sample of minimum-bias events. The weights would then be calculated using the correction $\{\bar{\alpha}_{\text{PU}}^F, \sigma_{\text{PU}}^F\} \rightarrow f(y_i)\{\bar{\alpha}_{\text{PU}}^F, \sigma_{\text{PU}}^F\}$ where $f(y)$ is extracted from minimum-bias data.

2.2 Particle weights

Having introduced a variable with some separation power between shower-like radiation and pileup-like particles, we will use it to compute a weight for each particle. The ideal weight is one for leading vertex particles and zero for pileup particles. Since we are trying to estimate whether a particle is pileup or not given the available information, one can imagine that the weight may not be restricted to one and zero and can be a fractional value. Furthermore, even if one insists on assigning integer weights, in a detector environment neutral particles that are closer than the detector granularity will be treated as a single particle, leading to possible fractional weights.

In order to define weights, we first introduce the following quantity

$$\chi_i^2 = \Theta(\alpha_i - \bar{\alpha}_{\text{PU}}) \times \frac{(\alpha_i - \bar{\alpha}_{\text{PU}})^2}{\sigma_{\text{PU}}^2}, \quad (2.7)$$

where Θ is the Heaviside step function. Eq. (2.7) measures how far α_i fluctuates from the pileup median. Fluctuations below the median are considered to be pileup and are assigned a weight equal to zero, as defined below. On the contrary, large fluctuations

⁶The left-side RMS is computed with $\sum_{\alpha_i < \bar{\alpha}} (\alpha_i - \bar{\alpha})^2$ where $\bar{\alpha}$ is the median of the full distribution. We use the median and the left-side RMS because these are insensitive to tails in the charged pileup distributions that occur on the right side. These tails originate from pileup particles inside of jets.

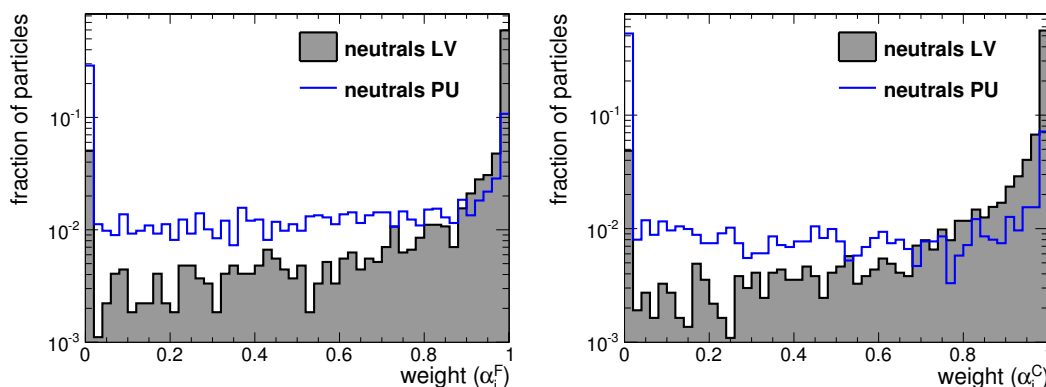


Figure 2. The distribution of weights from eq. (2.8), over many events, for neutral particles i with $p_T > 1$ GeV from the leading vertex (gray) and particles from pileup (blue) in a dijet sample. The weights are calculated using α_i^F (left) and α_i^C (right). In this sample, for weights from α_i^F , 30% (5%) of neutral PU (LV) particles have $w_i < 0.02$ while 10% (60%) have $w_i > 0.98$. For weights from α_i^C , 50% (5%) of neutral PU (LV) particles have $w_i < 0.02$ while 5% (55%) have $w_i > 0.98$.

above the median are very uncharacteristic of pileup and appropriately receive a weight close to 1. Any intermediate fluctuation above the median is assigned a fractional weight between zero and one. Whenever possible, the C variant of the quantities are used, and everywhere else the F variant is used. As seen in figure 1 the distribution of α for pileup looks roughly Gaussian-like. For this reason eq. (2.7) resembles a $\chi^2_{\text{NDF}=1}$ distribution, as the notation suggests. In fact, interpreting this distribution as a χ^2 distribution lends itself nicely into incorporating additional information, as is discussed in section 2.3. Particles are then assigned a weight given by

$$w_i = F_{\chi^2, \text{NDF}=1}(\chi_i^2), \quad (2.8)$$

where F_{χ^2} is the cumulative distribution function of the χ^2 distribution. As anticipated, particles with $\chi_i^2 = 0$ receive a weight $w_i = 0$. Figure 2 shows the weight distributions for particles both using α^F (left) and α^C (right). As expected, the weights are closer to their true value when computed from α^C .

At this point we could cut on the weight to decide whether or not to identify a particle as pileup and discard it from the event. In [22] it was found that rescaling the particles in subjets was able to correct kinematics and shapes. In light of this, we choose to use the weight in eq. (2.8) to rescale the particle's four-momentum. The complete algorithm proceeds as follows:

1. The values α_i^C and α_i^F are computed for all charged pileup in the event and the medians and left-side RMS's are computed.
2. All charged pileup particles are assigned a weight $w_i = 0$ and all charged leading vertex particles are assigned a weight $w_i = 1$.
3. The weights of all other particles are calculated using eq. (2.8).

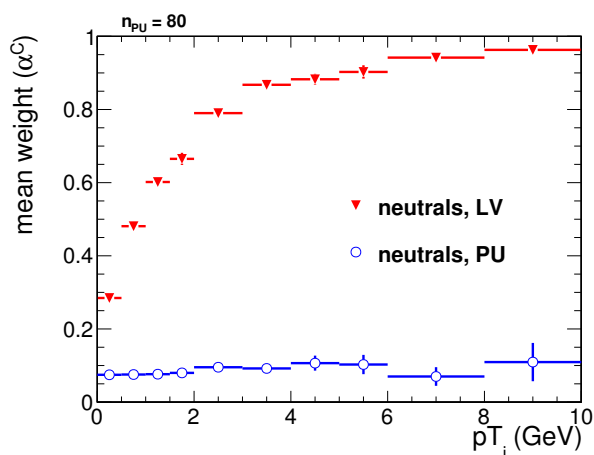


Figure 3. The mean weight, over many events, of neutral particles from the leading vertex (red) and pileup (blue) as a function of the particle’s p_T in a dijet sample.

4. The four-momentum of each particle is rescaled by its weight $p_i^\mu \rightarrow w_i \times p_i^\mu$.
5. Particles with small weights $w_i < w_{\text{cut}}$ or with low (rescaled) transverse momentum $p_{Ti} < p_{T,\text{cut}}$ are discarded.
6. The remaining set of rescaled particles is considered the pileup-corrected event.

Let us summarize the parameters of the algorithm. First, we have the cone size R_0 which specifies which particles are considered local. Neighboring particles inside a cone are the ones used to calculate α . We also have an R_{min} cutoff, such that neighboring particles with $\Delta R < R_{\text{min}}$ are not included in the computation of α . In our studies we use $R_0 = 0.3$ and $R_{\text{min}} = 0.02$. The choice of R_{min} is related to typical detector resolutions, as is discussed in section 3. Then we have a weight cut, w_{cut} , below which particles are deemed pileup and a p_T cut, $p_{T,\text{cut}}$. The precise choice of w_{cut} and $p_{T,\text{cut}}$ depends mildly both on the expected amount of pileup that will be encountered and detector parameters, such as calorimeter granularity. They can also, in general, be different for the central and forward regions. In our studies we use $w_{\text{cut}} = 0.1$, $p_{T,\text{cut}} \simeq 0.1 - 1.0$ GeV (the exact value will be described in section 4). We have checked that the performance of PUPPI algorithm depends weakly on the exact choice of these parameters, with a more significant degradation for much larger values of R_0 .

One may note that information from the distribution of particles from the leading vertex is primarily ignored. This is in contrast to matrix-element-like methods like shower deconstruction [27–29] which aim to optimize discrimination power by using as much signal and background information as possible. The specifics of the distributions for leading vertex particles depends on the sample, so we choose not to use the information from the distribution. In this way, the algorithm is not optimized for any specific signal, but rather looks for general features like a parton shower-like structure, and we expect it to behave consistently across a range of signal topologies.

2.3 Incorporating additional information

Many pileup removal algorithms are designed assuming a perfect detector and in many cases it is not straightforward to fold in information related to detector efficiencies or limitations. Using the PUPPI algorithm, experimental information can be used to directly modify the weight that is assigned to a particle. If one interprets the weight as a probability the particle is from the leading vertex (this will be discussed further in section 5), then vertex reconstruction efficiencies, for example, may affect this probability.

One advantage to the χ^2 approximation presented above is that it provides a scheme for calculating the weight based on experimental input. We further make the assumptions that the experimental information is Gaussian-distributed and independent both from the computation of α and other experimental information. Under these assumptions we can extend the $\chi^2_{\text{NDF}=1}$ approximation to a $\chi^2_{\text{NDF}=N}$ approximation

$$\chi^2_{i,\text{tot}} = \chi^2_i + \sum_{j=2}^N \chi^2_{i,j} = \chi^2_i + \chi^2_{i,\text{vertexing}} + \chi^2_{i,\text{calo depth}} + \dots \tag{2.9}$$

The weight is then appropriately adjusted to

$$w_i = F_{\chi^2, \text{NDF}=N}(\chi^2_{i,\text{tot}}). \tag{2.10}$$

Experimental information that may be useful includes tracking information, calorimeter depth information, and timing information.

2.4 Choice of metric

In separating pileup from leading vertex particles, it is necessary to identify features that distinguish between them. Here we consider leading vertex particles to originate from a parton shower. While the detailed jet structure will depend on the hard process, in the soft and collinear limit the parton shower is universal. In particular, it includes a soft and collinear singularity leading to the observed collinear structure of jets. Pileup, on the other hand, contains no hard scale and has no perturbative collinear structure. This motivates the use of a metric

$$\alpha_i = \log \sum_{j \in \text{event}} \frac{p_{Tj}}{\Delta R_{ij}^\beta} \times \Theta(R_{\min} \leq \Delta R_{ij} \leq R_0), \tag{2.11}$$

where this work uses $\beta = 1$. Particles from a parton shower are expected to have a small ΔR in relation to other particles from the shower, while pileup has no perturbative preference for small ΔR . The inclusion of p_{Tj} in the numerator is useful for the case where one sums over all particles. Here, the leading vertex contribution will dominate because the p_T spectrum of pileup falls much faster than leading vertex particles, resulting in the pileup contribution α being suppressed by p_T .

In eq. (2.11) β allows one to tune the relative importance of p_{Tj} vs. ΔR_{ij} . We have tried many metrics, including those not in the form of eq. (2.11), and find the one used here to be optimal.

One obvious question that may still arise in the choice of the metric is why p_{Ti} is not used. After all, we have already stated that the fact the p_T spectrum of pileup falls much faster than the p_T spectrum of leading vertex particles. Its exclusion from the metric is twofold. Firstly, it is already used in the algorithm. After weights are assigned and particles are rescaled, a cut on p_{Ti} is made. Secondly, we find that one of the reasons we find the weights useful as opposed to just a p_T cut is that the weights tend to not be strongly correlated with p_T in pileup, as shown in figure 3. In this way, α uses complementary information to just a p_{Ti} cut. In particular, in trying different metrics, we did try $\alpha_i = p_{Ti}$. We found its performance to be decent, however it degraded quicker than the ΔR -based metric when calorimeter cell discretization was introduced.

3 Simulation details

In order to study the performance of our algorithm and compare it to existing methods, we use a sample of $pp \rightarrow$ dijet at $\sqrt{s} = 14$ TeV, unless specified otherwise. Events are generated with Pythia 8.176 [26], tune 4C [30, 31]. The spectrum is generated such that the p_T of the $2 \rightarrow 2$ process is roughly flat across the range 15 – 500 GeV. This is done in order to maintain reasonable statistics across a range of jet p_T values and to demonstrate the method’s utility across different kinematic regimes. Pileup events are generated as zero-bias soft QCD events using Pythia and overlaid onto the hard scatter event. The number n_{PU} specifies the exact number of pileup interactions. We take as a baseline scenario $n_{PU} = 80$ pileup interactions overlaid and several results in section 4 consider this scenario. We also consider performance versus n_{PU} . Sections 4.1 and 4.2 show results on jet kinematics and shapes. In section 4.3 we show the algorithm’s performance on missing transverse energy (E_T^{miss}). In this section the sample used is $pp \rightarrow Z + \text{jets}$ at $\sqrt{s} = 14$ TeV, where the Z decays invisibly to neutrinos. In order to focus the performance study on pileup mitigation, underlying event is not included in the simulation.

We reconstruct particles in a naive detector simulation. The detector extends to $|\eta| < 5.0$ and includes a perfect tracker for $|\eta| < 2.5$. The perfect tracker exactly identifies if a charged hadron is from the leading vertex or from a pileup vertex (in contrast to a real tracker where misidentifications are possible) and has perfect spatial resolution. Neutral particles are discretized into calorimeter cells of size 0.1×0.1 in the $\eta\phi$ -plane.

Selecting an appropriate value for R_{min} is closely tied to the properties of the detector. The detector itself restricts cells from being closer than approximately $r_{\text{cell}} = 0.1$ from each other. Similarly, in a real detector the tracking efficiency degrades for distances closer than $r_{\text{track}} \lesssim 0.02$ from each other because for pairs of tracks closer than this distance it becomes possible that one of the tracks is lost. The distance between cells and tracks, on the other hand, is not necessarily regulated by the detector and could be as small as zero. Thus R_{min} directly regulates the cell-track distances and should be chosen as $R_{\text{min}} = \min(r_{\text{track}}, r_{\text{cell}}) = 0.02$ to be consistent with resolutions. We use $R_{\text{min}} = 0.02$ in our simulation for consistency with all inter-object distances and to mock-up the effect of track resolution.

Where particles are clustered into jets, we use Fastjet 3.0.5 [32] and the anti- k_T algorithm [33] with a radius of $R = 0.7$ (AK7). While smaller jet radii are more common

in phenomenological studies, larger jets receive more pileup contamination and are commonplace in substructure studies where correcting more than only a jet’s p_T becomes important [34–36]. We choose $R = 0.7$ as a compromise between these applications.

We define four particle collections from which we can derive algorithmic performance. They are:

- **LV**: only particles from the leading vertex.
- **PFlow**: all particles in the event including those from the leading vertex and pileup. These are the inputs that would be used in particle flow.
- **PFlowCHS**: all particles in the event except for charged particles from pileup (within the tracker volume). This corresponds to particle flow with charged hadron subtraction.
- **PUPPI**: the resulting rescaled particles from the algorithm described in section 2.

The **PFlowCHS** particle collection can be considered the current experimental state-of-the-art. We also apply four-vector subtraction to **PFlow** and **PFlowCHS** inputs as will be described in the following section, wherever jet quantities are shown.

In figure 4 we show a sample of an event display with $n_{PU} = 80$ for the four particle collections we consider. Particles from the leading vertex are drawn with filled squares and colored according to their p_T . Particles from pileup are drawn with unfilled, uncolored squares with their size logarithmically proportional to their p_T . The unfilled colored circles show anti- k_T $R = 0.7$ jets where the colors denote the p_T bin. The bins 25 – 50 GeV, 50 – 200 GeV, and > 200 GeV correspond to colors of magenta, cyan, and blue respectively.

The LV plot (top left) shows the original uncontaminated event. The **PFlow** plot (top right) shows the effect of all pileup particles being added to the event. The **PFlowCHS** plot (bottom left) shows a reduced pileup density in central region where charged pileup has been removed. The **PUPPI** plot (bottom right) is an event display that reproduces not only the hard jets from the LV collection, but also manages to capture features outside of the jets and remove a large portion of the pileup completely. The p_T of the jets from **PFlow** and **PFlowCHS** are area subtracted.

4 Results

In this section we study the performance of the **PUPPI** algorithm on several jet and event observables. Where jets are clustered using the **PFlow** collection, they are corrected using the “safe” modification of four-vector subtraction [37].^{7,8} Subtraction is also applied to jets clustered from the **PFlowCHS**. In the tracking region for **PFlowCHS**, charged pileup is already removed, so ρ is calculated only from neutral particles. In the forward region, ρ

⁷The results can differ based on the variant of four-vector subtraction used, however, the qualitative conclusions remain unchanged. In this work we use a modified version of four-vector subtraction presented in [37] which forbids negative masses by setting the mass of (sub)jets to zero in certain cases.

⁸We include corrections due to hadron masses following the method proposed in [12].

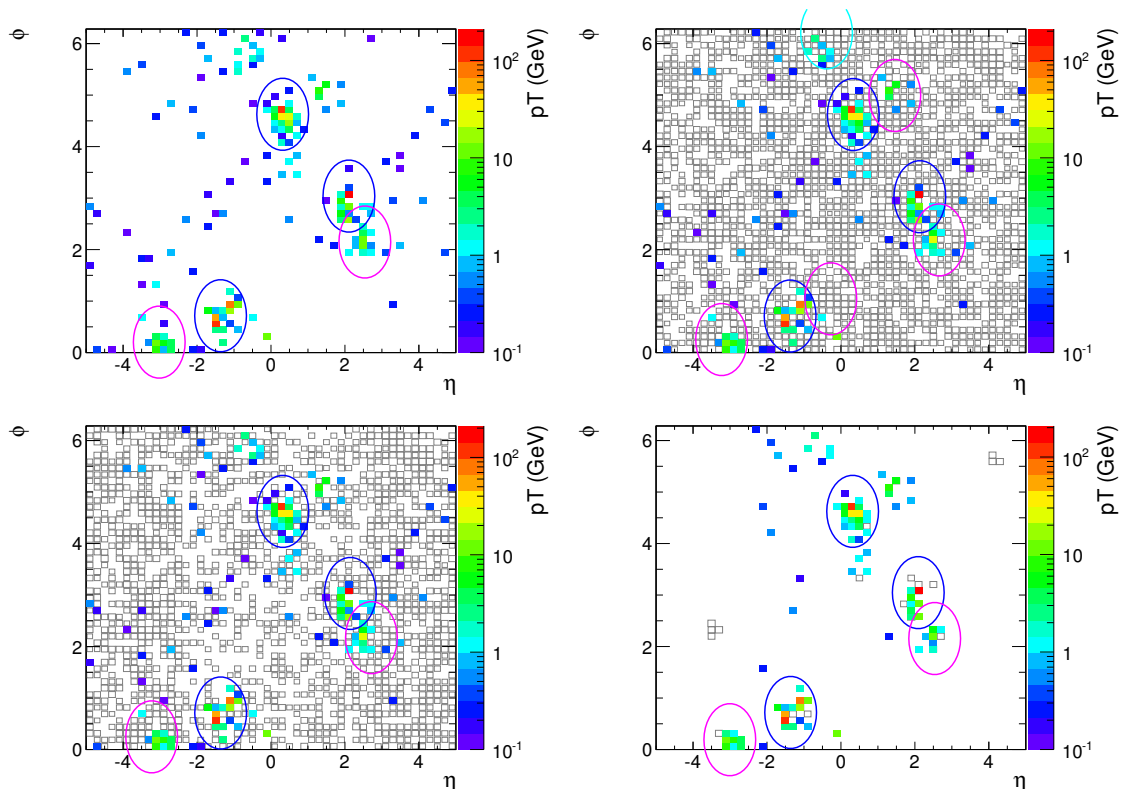


Figure 4. Event display for sample dijet event with 80 pileup interactions added. The particle collections shown are LV (top left), PFlow (top right), PFlowCHS (bottom left), and PUPPI (bottom right). Particles from the leading vertex are colored according to their p_T , while particles from pileup are uncolored and their size is logarithmically proportional to their p_T . The unfilled colored circles show anti- k_T $R = 0.7$ jets where the colors denote the p_T bin. The bins 25 – 50 GeV, 50 – 200 GeV, and > 200 GeV correspond to colors of magenta, cyan, and blue respectively. In the PFlow and PFlowCHS events, the average value of p_T among the pileup cells is ~ 0.7 GeV and ~ 0.4 GeV, respectively.

is computed using all particles. The jet clustering procedure is run separately on each particle collection.

For PUPPI we make the following parameter choices: $R_0 = 0.3$, $R_{\min} = 0.02$, $w_{\text{cut}} = 0.1$, and choose $p_{T,\text{cut}}$ according to

$$\text{central: } p_{T,\text{cut}} = 0.1 \text{ GeV} + n_{\text{PU}} \times 0.007 \text{ GeV}, \quad (4.1)$$

$$\text{forward: } p_{T,\text{cut}} = 0.2 \text{ GeV} + n_{\text{PU}} \times 0.011 \text{ GeV}. \quad (4.2)$$

In particular the $p_{T,\text{cut}}$ value has a weak dependence on the amount of pileup in the event and will depend on the granularity and energy resolution of a particular detector. We tune the values of this cut for our mock detector to minimize the offset between reconstructed observables and LV observables (see e.g. missing transverse energy in figure 10).

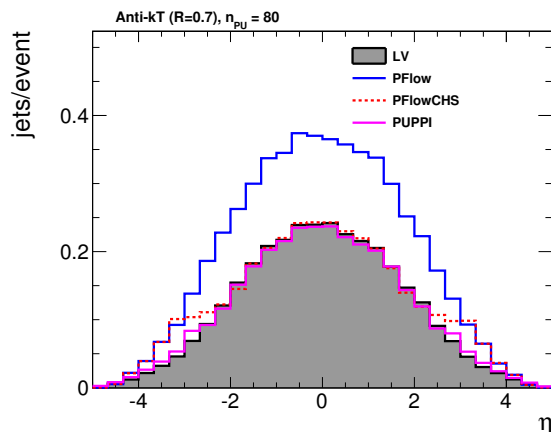


Figure 5. Jet multiplicity as a function of pseudorapidity for $n_{\text{PU}} = 80$.

4.1 Jet kinematics

We start by looking at the jet multiplicity as a function of pseudorapidity shown in figure 5 for $n_{\text{PU}} = 80$. Here all jets with $p_T > 25$ GeV after the pileup correction techniques are applied are considered. We see that in pseudorapidity regions where pileup correction is solely from subtraction the jet multiplicity tends to be too high. This is primarily from high density regions of pileup resulting in pileup jets. For PFlow this occurs across the full rapidity range, while for PFlowCHS this only occurs in the forward region where charged hadrons cannot be removed. The PUPPI jet distribution matches the LV distribution well across pseudorapidity.

Next we compare the jet p_T resolution across the methods. We define the resolution of an observable \mathcal{O} from the particle collection \mathcal{P} to be

$$\text{resolution}(\mathcal{O}_{\mathcal{P}}) = \text{RMS} \left\{ \frac{\mathcal{O}_{\mathcal{P}} - \mathcal{O}_{\text{LV}}}{\mathcal{O}_{\text{LV}}} \right\}. \quad (4.3)$$

Additionally, in plots where the resolution is cited as fitted σ , we adopt the common practice of fitting the distribution to a Gaussian and using the standard deviation as the resolution.

To compare jets from different collections, one needs a scheme to match jets. We consider jets from two collections matched if they are within $\Delta R = 0.3$ of each other. Figure 6 (left) shows the p_T resolution for central jets with p_T between 100 and 200 GeV. The p_T resolution of PUPPI is roughly 1.5 times better than PFlowCHS and 2.5 times better than PFlow. Figure 6 (right) shows the p_T resolution for forward jets with p_T between 25 and 50 GeV. In both cases the p_T response also tends to be more symmetric than PFlow and PFlowCHS. Despite the fact that there is no tracking information in the forward region, the PUPPI algorithm is able to maintain an improvement over subtraction even in the forward region. We also note that the improvement in PFlowCHS over PFlow in the forward region is due to the partial tracking information that PFlowCHS jets have near the tracker boundary.

Next we show the p_T resolution as a function of p_T for central jets in figure 7 (left). We show that the improvements found above hold across a wide kinematic range. In figure 7

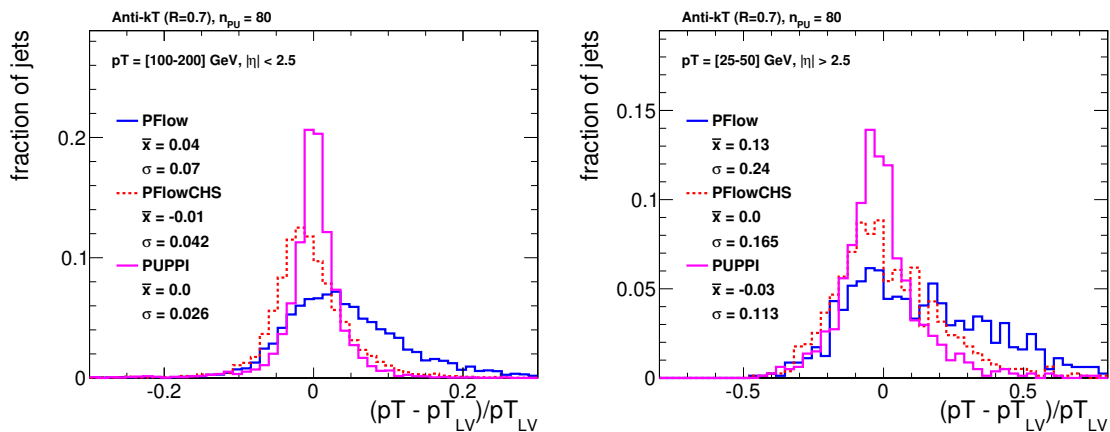


Figure 6. Jet p_T resolution with $n_{PU} = 80$ for jets with $100 \text{ GeV} < p_T < 200 \text{ GeV}$ and $|\eta| < 2.5$ (left) and jets with $25 \text{ GeV} < p_T < 50 \text{ GeV}$ and $|\eta| > 2.5$ (right).

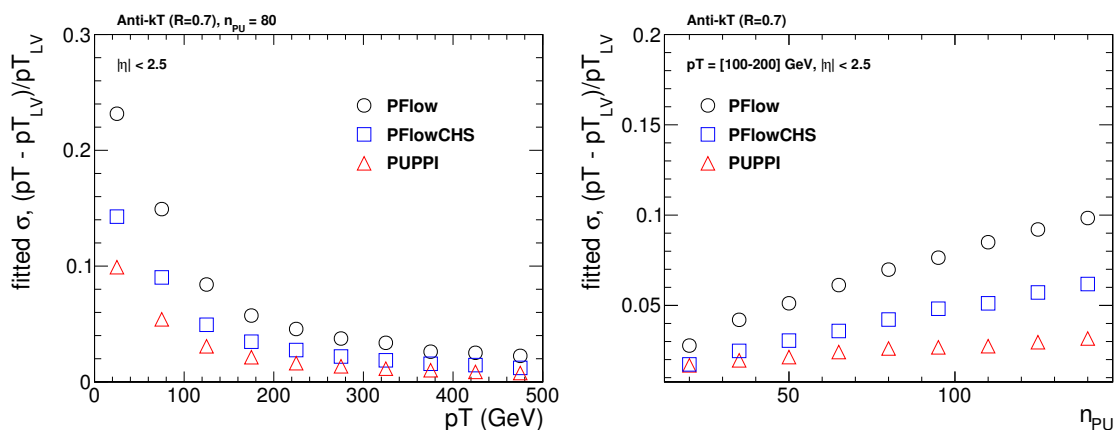


Figure 7. Jet p_T resolution vs. p_T (left) for $n_{PU} = 80$ for $|\eta| < 2.5$ and jet p_T resolution vs. number of pileup interactions (right) for jets with p_T between 100 and 200 GeV for $|\eta| < 2.5$.

(right) we show the p_T resolution as a function of number of pileup interactions. For low levels of pileup we see that the PUPPI algorithm does not offer much of an improvement over existing methods. This is for two reasons. Firstly, at low levels of pileup there is not much improvement to make. Secondly, in low pileup environments, there is less information available locally just due to the lack of pileup. This means the α distribution is not as well populated and the uncertainty on σ_{PU} is larger.

4.2 Jet shapes

Similar to our study of p_T distributions, we can study resolution and its pileup dependence for jet shapes. Here we show results for jet mass which is considered a reasonable proxy for generic jet shapes and is used in many applications such as boosted object tagging (see [34–36] and references therein).

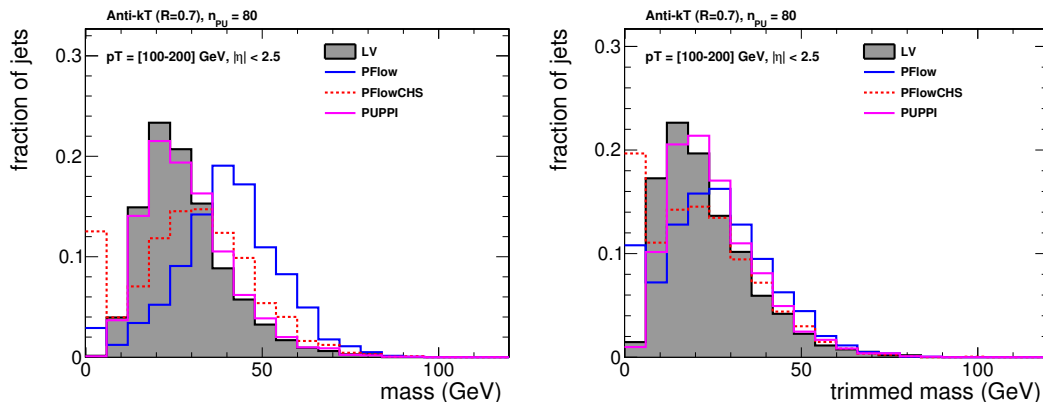


Figure 8. The single jet mass resolution for $n_{PU} = 80$ for jets with $100 \text{ GeV} < p_T < 200 \text{ GeV}$ and $|\eta| < 2.5$ for plain jet mass (left) and trimmed jet mass (right).

First we look at jet mass for central jets with $100 \text{ GeV} < p_T < 200 \text{ GeV}$. The distribution is shown in figure 8 (left). Here we see that PUPPI is not only able to correct the mean of the distribution, but also the distribution itself. Figure 8 (right) shows the results of PUPPI on trimmed mass. Trimming is performed on jets from all collections, including LV, using $r_{sub} = 0.2$ and $f_{cut} = 0.05$. For jets from PFlow and PFlowCHS subtraction is applied to the trimmed jet. Even with the application of grooming, PUPPI distributions do a consistent job of restoring distributions near to their LV distributions. We regard this as a positive indication that PUPPI is returning a consistent event interpretation.

In figure 9 (left) we show the mass resolution⁹ for jets with p_T between 100 GeV and 200 GeV at $n_{PU} = 80$. We find that the PUPPI jet mass resolution is improved with respect to the other inputs. Figure 9 (right) plots the mass resolution as a function of number of pileup interactions where the mass resolution from PUPPI is relatively stable as a function of n_{PU} .

4.3 Missing transverse energy

We now look at an event quantity, the missing transverse energy (E_T^{miss}), which is interesting from both a theoretical and an experimental point of view. From the theoretical perspective, missing transverse energy is one of the main signatures of new physics. For example, in R -parity conserving supersymmetry, every event in which superpartners are pair-produced the two lightest supersymmetric particles in the final state appear as missing transverse energy. Additionally for standard model measurements, E_T^{miss} plays a role in many analyses, such as the W mass measurement [38], the Higgs to WW discovery [39, 40] and the Higgs to $\tau\tau$ evidence [41, 42]. On the experimental side, E_T^{miss} is challenging because it compounds errors from the measurement of all objects in the event, both pileup and non-pileup alike. In the presence of pileup, the E_T^{miss} resolution rapidly degrades because the full energy of the additional pileup events is incorporated into the event [6, 9].

⁹For mass and missing transverse energy resolutions we use $\text{resolution}(\mathcal{O}_P) = \text{RMS}\{\mathcal{O}_P - \mathcal{O}_{LV}\}$ as opposed to eq. (4.3) to avoid divergent behavior as $\mathcal{O}_{LV} \rightarrow 0$.

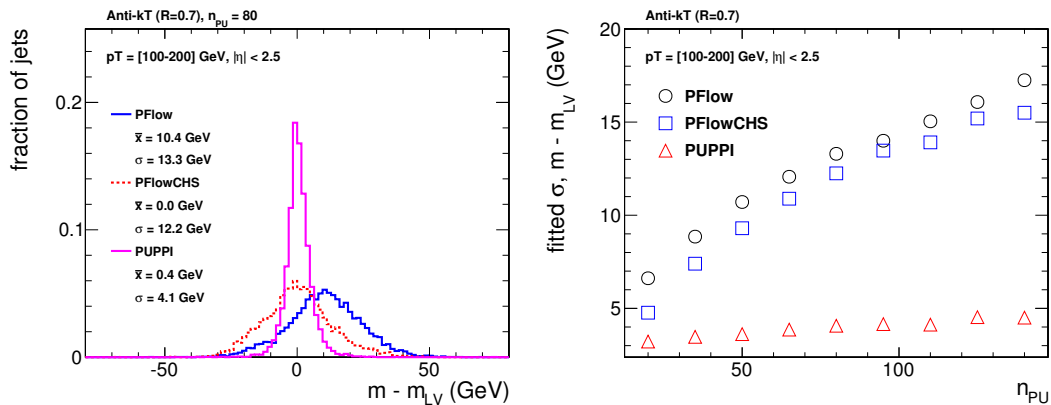


Figure 9. The single jet mass resolution for $n_{PU} = 80$, for jets with $100 \text{ GeV} < p_T < 200 \text{ GeV}$ and $|\eta| < 2.5$ (left) and jet mass vs. number of pileup interactions for jets with mass between 100 and 200 GeV.

Attempts at reducing the impact of pileup on the E_T^{miss} resolution are typically more difficult than on jets, because traditional approaches that work on jets breakdown. The pileup component of events has a natural tendency to have near zero E_T^{miss} . Applying a method that reconstructs E_T^{miss} from a fraction of the particles in the event, e.g. charged hadron subtraction, breaks the cancellation between neutral and charged pileup particles resulting in large distortions in E_T^{miss} measurements. In order to mitigate the effects of pileup in E_T^{miss} , both ATLAS and CMS have resorted to approaches that rely on combinations of various methods of calculating E_T^{miss} [6, 9, 43]. Such methods, either through a linear combination of different E_T^{miss} variables or through a boosted decision tree regression, can lead to a reduction of the pileup dependence on the E_T^{miss} resolution by a factor of roughly three. These calculations are typically quite elaborate and rely on the commissioning of 10 – 20 additional E_T^{miss} related variables.

To compare the performance of E_T^{miss} observables, we use a sample of events with large hadronic recoil and well-defined E_T^{miss} , in this case $pp \rightarrow Zj$ where the Z decays to neutrinos and has transverse momentum in the center of mass frame $p_T(Z) > 350 \text{ GeV}$. The missing transverse energy is constructed from negative vector sum of the particle transverse momenta

$$\vec{E}_T^{\text{miss}} = - \sum_{i \in \text{event}} \vec{p}_{Ti}, \quad (4.4)$$

where the length of this vector is denoted $E_T^{\text{miss}} = |\vec{E}_T^{\text{miss}}|$. Another related variable is the scalar sum of transverse energies

$$\sum E_T = \sum_{i \in \text{event}} |\vec{p}_{Ti}|. \quad (4.5)$$

We show the resolution in figure 10 (left), where we see that the PUPPI algorithm noticeably improves the $\sum E_T$ resolution over PFlow and PFlowCHS. While that fact that neither PFlow nor PFlowCHS are centered around zero is not an issue, the fact that PUPPI is

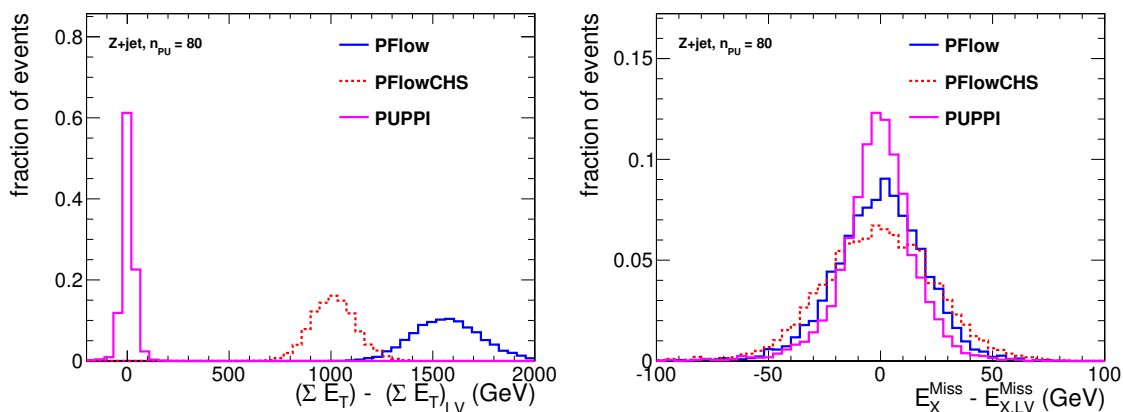


Figure 10. The resolution of $\sum E_T$ (left) and the resolution of E_x^{miss} (right) in $Z + \text{jets}$ events with $n_{PU} = 80$.

centered around zero supports the claim that applying PUPPI produces a consistent event interpretation without the need for further pileup correction.

To compare the E_T^{miss} resolution, we look at the resolution of the x -component of E_T^{miss} , shown in figure 10 (right). The relevance of this quantity to phenomenology is more directly seen, as this is one component of the E_T^{miss} vector. Both the length and direction of the E_T^{miss} vector are important discriminating variables in many new physics searches so it is plausible that a small signal could be washed out by poor E_T^{miss} resolution or non-unity E_T^{miss} response. We find that in our simplified set-up PUPPI displays improvements over PFlow and PFlowCHS. In fact, the resolution for PFlowCHS degrades the resolution with respect to PFlow. This effect is due to the observation above that the partial removal of pileup interactions can lead to larger E_T^{miss} resolution.

For the pileup-reduced E_T^{miss} computations in CMS [9], it was found that the key component to reducing the pileup dependence on the E_T^{miss} resolution resulted from the identification and (indirect) removal of pileup jets from the E_T^{miss} calculation. With PUPPI, pileup jet removal is naturally built into the algorithm, thereby allowing for a simplified approach to pileup mitigation in E_T^{miss} related quantities. We expect that given the algorithm’s flexibility in using experimental information, the improvement will persist in the full detector environment.

5 Summary and outlook

In this paper, we have introduced a new algorithm, PUPPI, for removing pileup contamination. This method employs a *per particle* approach and improves the reconstruction of not only jet quantities, but also of event-wide observables like missing transverse energy. PUPPI operates by using charged pileup to characterize the pileup in an event and then uses that knowledge to assign a weight to particles of unknown origin, like neutral hadrons or any particle in the forward region. The weight is used to rescale the particle’s four-momentum. The parameters of the algorithm are the size of the cone used to define neighboring par-

ticles R_0 , the minimum distance cutoff R_{\min} , the cut on the weights w_{cut} , and the cut on the rescaled transverse momentum $p_{T,\text{cut}}$.

By applying corrections at the particle level, before jet clustering, we can simultaneously perform pileup jet mitigation, and jet four-vector and shape corrections. We have shown the improvement of PUPPI over existing methods by studying jet p_T , mass, and missing transverse energy over a wide range of jet p_T and number of pileup interactions. Also, our method can be applied both in the central region of the detector (where tracking information is available) and in the forward region.

In this work we have introduced the simplest form of the algorithm. However, many modifications and extensions are worth further exploration. In particular, we have shown results for a single choice of metric, a particular weighting function, and a choice of how to use the weights. Further modifications considered for the metric can include a combination of discrimination power from a selection of metrics into a common multivariate discriminant. Preliminary studies with a boosted decision tree show modest improvements, although we leave it to future work to fully explore this avenue.

With regards to the particle weights, we have elected to allow fractional weights and chosen to use them to rescale four-vectors. It is not obvious that a four-vector rescaling is the optimal procedure to implement. As a simplification one could restrict weights to zero or one, in which case no rescaling is performed and particles are either kept or discarded. Taking a step in the opposite direction, one could interpret the weights as probabilities that a particle should be kept in the event. Given a probabilistic interpretation of weights, a natural approach would follow along the lines of Qjets [44, 45], where a given event would yield many event interpretations with particles either kept or discarded according to their probabilistic weight. All observables for a single event would then become distributions. We leave this study for future work.

Though we frame our studies within a “particle flow”-like set of inputs, it is not restricted only to inputs of this type. If we consider as inputs calorimeter clusters rather than particles, we can still similarly compute the distance of tracks to a given calorimeter cluster, i , within the tracking volume. Then for the forward region, we can consider nearby calorimeter clusters. The challenge is to identify pure pileup clusters; we expect this can be achieved using tracks from the non-leading vertices.

This method may also be applied to heavy ion events, where energy densities of the underlying event are similar to the levels of pileup studied. In this case, however, all particles originate from the leading vertex. One can use a modified version of PUPPI in which the leading vertex constraint is not applied in the algorithm.

Finally, given the performance of PUPPI on jet mass and E_T^{miss} we are optimistic that the PUPPI algorithm will be useful in improving pileup mitigation of other jet and event shapes, and more generally in the identification of other physics objects. For instance, given a pileup-corrected event, it is reasonable to expect that identifying isolated photons or leptons will be improved using a *per particle* weighting scheme.

While PUPPI was developed, another particle level pileup removal method called SoftKiller [46] has been proposed. Preliminary comparisons indicate comparable performances [47].

Acknowledgments

The authors are grateful to Jeff Berryhill, Matteo Cacciari, Dinko Ferencek, David Krohn, Andrew Larkoski, Filip Moortgat, Salvatore Rappoccio, Gavin Salam, Matthew Schwartz, Gregory Soyez, Jesse Thaler, and Lian-Tao Wang for useful discussions. We would also like to thank the organizers Filip Moortgat, Gavin Salam, and Ariel Schwartzman, as well as the participants of Mitigation of Pileup Effects at the LHC workshop for many fruitful discussions.

DB is partly supported by the U.S. Department of Energy under cooperative research agreement Contract Number DE-FG02-05ER41360, by the LHC-Theory Initiative Graduate Fellowship, and by Istituto Nazionale di Fisica Nucleare (INFN) through the Bruno Rossi Fellowship, PH is supported by CERN, ML is supported by NSERC of Canada, and NT is supported by the Fermi Research Alliance, LLC under Contract No. De-AC02-07CH11359 with the United States Department of Energy and the LPC-CMS Distinguished Researcher program operated through FNAL.

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

References

- [1] ATLAS collaboration, *Performance of jet substructure techniques for large- R jets in proton-proton collisions at $\sqrt{s} = 7$ TeV using the ATLAS detector*, *JHEP* **09** (2013) 076 [[arXiv:1306.4945](https://arxiv.org/abs/1306.4945)] [[INSPIRE](#)].
- [2] ATLAS collaboration, *Pile-up subtraction and suppression for jets in ATLAS*, *ATLAS-CONF-2013-083*, CERN, Geneva Switzerland (2013).
- [3] ATLAS collaboration, *Performance of pile-up subtraction for jet shapes*, *ATLAS-CONF-2013-085*, CERN, Geneva Switzerland (2013).
- [4] CMS collaboration, *Pileup jet identification*, *CMS-PAS-JME-13-005*, CERN, Geneva Switzerland (2013).
- [5] ATLAS collaboration, *Tagging and suppression of pileup jets with the ATLAS detector*, *ATLAS-CONF-2014-018*, CERN, Geneva Switzerland (2014).
- [6] ATLAS collaboration, *Pile-up suppression in missing transverse momentum reconstruction in the ATLAS experiment in proton-proton collisions at $\sqrt{s} = 8$ TeV*, *ATLAS-CONF-2014-019*, CERN, Geneva Switzerland (2014).
- [7] CMS collaboration, *Determination of jet energy calibration and transverse momentum resolution in CMS*, 2011 *JINST* **6** P11002 [[arXiv:1107.4277](https://arxiv.org/abs/1107.4277)] [[INSPIRE](#)].
- [8] CMS collaboration, *Studies of jet mass in dijet and $W/Z + jet$ events*, *JHEP* **05** (2013) 090 [[arXiv:1303.4811](https://arxiv.org/abs/1303.4811)] [[INSPIRE](#)].
- [9] CMS collaboration, *MET performance in 8 TeV data*, *CMS-PAS-JME-12-002*, CERN, Geneva Switzerland (2012).
- [10] M. Cacciari and G.P. Salam, *Pileup subtraction using jet areas*, *Phys. Lett. B* **659** (2008) 119 [[arXiv:0707.1378](https://arxiv.org/abs/0707.1378)] [[INSPIRE](#)].

- [11] M. Cacciari, G.P. Salam and G. Soyez, *The catchment area of jets*, *JHEP* **04** (2008) 005 [[arXiv:0802.1188](#)] [[INSPIRE](#)].
- [12] G. Soyez, G.P. Salam, J. Kim, S. Dutta and M. Cacciari, *Pileup subtraction for jet shapes*, *Phys. Rev. Lett.* **110** (2013) 162001 [[arXiv:1211.2811](#)] [[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] 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](#)].
- [15] 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](#)].
- [16] D. Krohn, J. Thaler and L.-T. Wang, *Jet trimming*, *JHEP* **02** (2010) 084 [[arXiv:0912.1342](#)] [[INSPIRE](#)].
- [17] A.J. Larkoski, S. Marzani, G. Soyez and J. Thaler, *Soft drop*, *JHEP* **05** (2014) 146 [[arXiv:1402.2657](#)] [[INSPIRE](#)].
- [18] ATLAS collaboration, *Jet energy measurement with the ATLAS detector in proton-proton collisions at $\sqrt{s} = 7$ TeV*, *Eur. Phys. J. C* **73** (2013) 2304 [[arXiv:1112.6426](#)] [[INSPIRE](#)].
- [19] W. Lampl et al., *Calorimeter clustering algorithms : description and performance*, [ATL-LARG-PUB-2008-002](#), CERN, Geneva Switzerland (2008).
- [20] ATLAS collaboration, T. Barillari et al., *Local hadronic calibration*, [ATL-COM-LARG-2008-006](#), CERN, Geneva Switzerland (2008).
- [21] CMS collaboration, *Particle-flow event reconstruction in CMS and performance for jets, taus and MET*, [CMS-PAS-PFT-09-001](#), CERN, Geneva Switzerland (2009).
- [22] 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](#)].
- [23] CMS collaboration, *Underlying-event subtraction for particle flow*, [CMS-DP-2013-018](#), CERN, Geneva Switzerland (2013).
- [24] 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](#)].
- [25] D. Bertolini, T. Chan and J. Thaler, *Jet observables without jet algorithms*, *JHEP* **04** (2014) 013 [[arXiv:1310.7584](#)] [[INSPIRE](#)].
- [26] T. Sjöstrand, S. Mrenna and P.Z. Skands, *A brief introduction to PYTHIA 8.1*, *Comput. Phys. Commun.* **178** (2008) 852 [[arXiv:0710.3820](#)] [[INSPIRE](#)].
- [27] D.E. Soper and M. Spannowsky, *Finding physics signals with shower deconstruction*, *Phys. Rev. D* **84** (2011) 074002 [[arXiv:1102.3480](#)] [[INSPIRE](#)].
- [28] D.E. Soper and M. Spannowsky, *Finding top quarks with shower deconstruction*, *Phys. Rev. D* **87** (2013) 054012 [[arXiv:1211.3140](#)] [[INSPIRE](#)].
- [29] D.E. Soper and M. Spannowsky, *Finding physics signals with event deconstruction*, *Phys. Rev. D* **89** (2014) 094005 [[arXiv:1402.1189](#)] [[INSPIRE](#)].

- [30] R. Corke and T. Sjöstrand, *Interleaved parton showers and tuning prospects*, *JHEP* **03** (2011) 032 [[arXiv:1011.1759](#)] [[INSPIRE](#)].
- [31] A. Buckley et al., *General-purpose event generators for LHC physics*, *Phys. Rept.* **504** (2011) 145 [[arXiv:1101.2599](#)] [[INSPIRE](#)].
- [32] M. Cacciari, G.P. Salam and G. Soyez, *FastJet user manual*, *Eur. Phys. J. C* **72** (2012) 1896 [[arXiv:1111.6097](#)] [[INSPIRE](#)].
- [33] M. Cacciari, G.P. Salam and G. Soyez, *The anti- k_t jet clustering algorithm*, *JHEP* **04** (2008) 063 [[arXiv:0802.1189](#)] [[INSPIRE](#)].
- [34] 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](#)].
- [35] A. Altheimer et al., *Jet substructure at the Tevatron and LHC: new results, new tools, new benchmarks*, *J. Phys. G* **39** (2012) 063001 [[arXiv:1201.0008](#)] [[INSPIRE](#)].
- [36] A. Altheimer et al., *Boosted objects and jet substructure at the LHC. Report of BOOST2012, held at IFIC Valencia, July 23–27 2012*, *Eur. Phys. J. C* **74** (2014) 2792 [[arXiv:1311.2708](#)] [[INSPIRE](#)].
- [37] M. Cacciari, G.P. Salam and G. Soyez, *On the use of charged-track information to subtract neutral pileup*, [arXiv:1404.7353](#) [[INSPIRE](#)].
- [38] CDF and D0 collaborations, T.A. Aaltonen et al., *Combination of CDF and D0 W -boson mass measurements*, *Phys. Rev. D* **88** (2013) 052018 [[arXiv:1307.7627](#)] [[INSPIRE](#)].
- [39] CMS collaboration, *Measurement of Higgs boson production and properties in the WW decay channel with leptonic final states*, *JHEP* **01** (2014) 096 [[arXiv:1312.1129](#)] [[INSPIRE](#)].
- [40] ATLAS collaboration, *Measurements of the properties of the Higgs-like boson in the $WW^{(*)} \rightarrow \ell\nu\ell\nu$ decay channel with the ATLAS detector using 25 fb^{-1} of proton-proton collision data*, [ATLAS-CONF-2013-030](#), CERN, Geneva Switzerland (2013).
- [41] ATLAS collaboration, *Evidence for Higgs boson decays to the $\tau^+\tau^-$ final state with the ATLAS detector*, [ATLAS-CONF-2013-108](#), CERN, Geneva Switzerland (2013).
- [42] CMS collaboration, *Evidence for the 125 GeV Higgs boson decaying to a pair of τ leptons*, *JHEP* **05** (2014) 104 [[arXiv:1401.5041](#)] [[INSPIRE](#)].
- [43] CMS collaboration, *Search for the standard model Higgs boson decaying to a W pair in the fully leptonic final state in pp collisions at $\sqrt{s} = 7\text{ TeV}$* , *Phys. Lett. B* **710** (2012) 91 [[arXiv:1202.1489](#)] [[INSPIRE](#)].
- [44] S.D. Ellis, A. Hornig, T.S. Roy, D. Krohn and M.D. Schwartz, *Q jets: a non-deterministic approach to tree-based jet substructure*, *Phys. Rev. Lett.* **108** (2012) 182003 [[arXiv:1201.1914](#)] [[INSPIRE](#)].
- [45] D. Kahawala, D. Krohn and M.D. Schwartz, *Jet sampling: improving event reconstruction through multiple interpretations*, *JHEP* **06** (2013) 006 [[arXiv:1304.2394](#)] [[INSPIRE](#)].
- [46] M. Cacciari, G.P. Salam and G. Soyez, *SoftKiller, a particle-level pileup removal method*, [arXiv:1407.0408](#) [[INSPIRE](#)].
- [47] *Mitigation of pileup effects at the LHC workshop*, <https://indico.cern.ch/event/306155/>, CERN, Geneva Switzerland May 16–18 2014.