

Fuzzy jets

Lester Mackey,^a Benjamin Nachman,^{b,c} Ariel Schwartzman^c and Conrad Stansbury^b

^a*Department of Statistics, Stanford University,
Stanford, CA 94305, U.S.A.*

^b*Department of Physics, Stanford University,
Stanford, CA 94305, U.S.A.*

^c*SLAC National Accelerator Laboratory, Stanford University,
2575 Sand Hill Rd, Menlo Park, CA 94025, U.S.A.*

E-mail: lmackey@stanford.edu, bnachman@cern.ch, sch@slac.stanford.edu,
chstan@stanford.edu

ABSTRACT: Collimated streams of particles produced in high energy physics experiments are organized using clustering algorithms to form *jets*. To construct jets, the experimental collaborations based at the Large Hadron Collider (LHC) primarily use agglomerative hierarchical clustering schemes known as sequential recombination. We propose a new class of algorithms for clustering jets that use infrared and collinear safe mixture models. These new algorithms, known as *fuzzy jets*, are clustered using maximum likelihood techniques and can dynamically determine various properties of jets like their size. We show that the fuzzy jet size adds additional information to conventional jet tagging variables in boosted topologies. Furthermore, we study the impact of pileup and show that with some slight modifications to the algorithm, fuzzy jets can be stable up to high pileup interaction multiplicities.

KEYWORDS: Jets

ARXIV EPRINT: [1509.02216](https://arxiv.org/abs/1509.02216)

Contents

1	Introduction	1
2	Mixture model jets	2
3	Clustering fuzzy jets: the EM algorithm	5
4	Comparisons with sequential recombination and jet tagging	6
4.1	Details of the simulation	6
4.2	Kinematic properties of fuzzy jets	8
4.3	New information from fuzzy jets	10
4.4	Fuzzy jets for tagging	11
5	Underlying event and pileup	14
5.1	Changing α for pileup suppression	14
5.2	Tower subtraction and the event jet: effective pileup correction	16
6	Conclusions	18
A	Wrapped gaussian	19
B	The EM algorithm	19
C	Controlling jet multiplicity with p_T	21
D	A leading order description of fuzzy jet σ	22

1 Introduction

As the result of a proton-proton collision at a hadron collider, hundreds of particles are created and detected [1, 2]. While some particles can be identified by their type, such as electrons [3, 4] and muons [5, 6], most of the detected particles are light hadrons produced in collimated sprays called *jets*. Jets are the consequence of high energy quarks or gluons fragmenting into colorless hadrons. Experimentally, jets are defined by clustering schemes which group together measured calorimeter energy deposits or reconstructed charged particle tracks. A jet algorithm is a clustering scheme that connects the measured objects with theoretical quantities that can be calculated and simulated. At a hadron collider, the natural coordinates for describing particles are p_T , y , and ϕ , where p_T is the magnitude of the momentum transverse to the proton beam, y is the rapidity, and ϕ is the azimuthal angle. Particles or calorimeter energy deposits are clustered using jet algorithms based on

distance metrics on their coordinates in $(p_T, \vec{\rho}) = (p_T, y, \phi)$. In order for a jet algorithm to be useful to experimentalists and theorists, the collection of jets should be IRC safe in the following sense:

1. Infrared safe (IR): if a particle i is added with $|p_T| \rightarrow 0$, the jets are unaffected.
2. Collinear safe (C): if a particle i with momentum p_i is replaced with two particles j and k with momenta $p_j + p_k = p_i$ such that $|\vec{\rho}_i - \vec{\rho}_j| = 0$, then the jets are unaffected.

The jet algorithms most widely used at hadron colliders fall into a class of schemes known as *sequential recombination* [7]. These IRC safe schemes require metrics d on momenta $d_{ij} = d(p_i, p_j) : (p_i, p_j) \rightarrow \mathbb{R}^+$, $d_{iB} = d(p_i) : p_i \rightarrow \mathbb{R}^+$ and proceed as follows:

1. Assign each particle as a proto-jet.
2. Repeat until there are no proto-jets left: let $(k, \ell) = \operatorname{argmin}_{i,j} d(p_i, p_j)$. If $d_{mB} < d_{k\ell}$ for $m = \operatorname{argmin}_i d(p_i)$, then declare proto-jet m a jet and remove it from the list. Otherwise, combine proto-jets k and ℓ into a new proto-jet with momentum $p_{\text{new}} = p_\ell + p_k$.

One common prescription is called the Cambridge-Aachen (C/A) algorithm [8, 9], which uses $d_{ij} = |\vec{\rho}_i - \vec{\rho}_j|^2 / R^2$ and $d_{iB} = 1$. The fixed quantity R is roughly the size of the jet in (y, ϕ) . By far, the most ubiquitous jet algorithm used at the Large Hadron Collider (LHC) is the anti- k_t algorithm [10] with $d_{ij} = \min(p_{T,i}^{-2}, p_{T,j}^{-2}) |\vec{\rho}_i - \vec{\rho}_j|^2 / R^2$ and $d_{iB} = p_{T,i}^{-2}$.

The purpose of this paper is to introduce a new paradigm for jet clustering, called *fuzzy jets*, based on probabilistic mixture modeling and to demonstrate its use in boosted topologies. section 2 introduces the statistical concept of a mixture model and describes the necessary modification to make the procedure IRC safe. Section 3 gives one efficient method for clustering fuzzy jets based on the Expectation-Maximization (EM) algorithm. Section 4 contains several examples comparing fuzzy jets with sequential recombination and section 5 describes how one might mitigate the impact of overlapping proton-proton collisions (pileup). We conclude in section 6 with some summary remarks and outlook for the future.

2 Mixture model jets

Mixture models [11] are a statistical tool for clustering which postulate a particular class of probability densities for the data to be clustered. Generically, for grouping n m -dimensional data points into k clusters, the mixture model density is

$$p(x_1, \dots, x_n | \theta) = \prod_{i=1}^n \left(\sum_{j=1}^k \pi_j f(x_i | \theta_j) \right), \tag{2.1}$$

where π_j is the unknown weight of cluster j such that $\sum_j \pi_j = 1$ and $f(x_i | \theta_j)$ is a probability density on m -dimensions with unknown parameters θ_j to be learned from the data. A




Figure 1. An example of the learned per-particle probability density specified in eq. (2.1) with $k = 3$ and Gaussian $f = \Phi$ in $m = 2$ dimensions. One cluster is associated with each component density $\Phi_i = \Phi(\cdot | \mu_i, \Sigma_i)$, where the dot \cdot is a placeholder for the function argument.

common choice for f is the normal density Φ with $\theta_j = (\mu_j, \Sigma_j)$ for μ_j the m -dimensional mean and Σ_j the $m \times m$ covariance matrix. In the mixture model paradigm, the θ_j are the cluster properties; in the Gaussian case, μ_j is the location of cluster j and Σ_j describes its shape in the m -dimensional space. When clustering with a finite mixture, the number of clusters k must be specified ahead of time,¹ which is dual to the usual use of sequential recombination² in which k is learned and the size of jets is specified ahead of time. The standard objective in (frequentist) mixture modeling is to select the parameters θ_j which maximize the likelihood (eq. (2.1)) of the observed dataset. Figure 1 illustrates what the learned event density might look like for $k = 3$ and Gaussian $f = \Phi$ in $m = 2$ dimensions.

An equivalent way of approaching mixture modeling is to view eq. (2.1) as the density used to generate the data. We view the data as having been drawn randomly from the density specified in eq. (2.1), with the following setup:

1. Throw n independent and identical k -sided dice with probability π_j to land on side $j = 1, \dots, k$ and label the outcomes $\lambda_1, \dots, \lambda_n$.
2. Independent of the others, data point $i \in \{1, \dots, n\}$ is drawn randomly from $f(\cdot | \theta_{\lambda_i})$.

Once θ and π are learned by minimizing eq. (2.1), we can compute $q_{ij} = \Pr(\lambda_i = j | x_i)$, the posterior probability that x_i was generated by $f(\cdot | \theta_j)$ or, intuitively, the posterior probability that x_i belongs to cluster j . The q_{ij} are the *soft assignments* of particles i to jet j and will play an important role in section 3 when we show how to maximize the likelihood in eq. (2.1). In particular, we can write, $q_{ij} = \pi_j f(x_i | \theta_j) / \sum_{j'} \pi_{j'} f(x_i | \theta_{j'})$. Jets produced with mixture modeling are called *fuzzy jets* because of the soft memberships —

¹There is a wealth of literature on the subject of choosing k , for a survey of methods, see [12]. The likelihood monotonically increases with k ; as alternatives to maximum likelihood, one can for instance look for kinks in the likelihood as a function of k [13].

²It is similar to the exclusive form of the k_T sequential recombination scheme [14]. The exclusive nature of the algorithm (and the minimization procedure used to find the jets) is similar to the XCono algorithm [15, 16] that became public as this manuscript was in its final preparation.

every particle can belong to every jet with some probability.³ This can be seen explicitly in figure 1 where the densities of all three clusters are everywhere nonzero, so $q_{ij} > 0$ for all j . The idea of probabilistic membership was recently studied in the context of the Q-jets algorithm [18] in which the same event is interpreted many times by injecting randomness into the clustering procedure. Unlike Q-jets, fuzzy jets allocates the soft membership functions deterministically throughout the clustering procedure. However, like Q-jets, there is an ambiguity in how to assign kinematic properties to the clustered jets. Fuzzy jets are defined by their shape (and location), not their constituents. This is in contrast to anti- k_t jets, which are defined by their constituents without an explicit shape determined from the clustering procedure. One simple assignment scheme is to define the momentum of a fuzzy jet j as

$$p_{\text{jet } j} = \sum_{i=1}^n p_i \left\{ \begin{array}{ll} 1 & j = \text{argmax}_k q_{ik} \\ 0 & \text{else} \end{array} \right\}. \tag{2.2}$$

In other words, this procedure assigns every particle to its most probable associated jet. This scheme will be known as the hard maximum likelihood (HML) scheme, but is not the only possible assignment algorithm. The dual problem in sequential recombination is the jet area, which must be defined [19], whereas the jet kinematics are the ‘natural’ coordinates.

We now specialize the likelihood in eq. (2.1) to the case of clustering particles into jets at a collider like the LHC. Consider a mixture model in two dimensions⁴ with $x_i = \rho_i$. The resulting mixture model (MM) jets are inherently not IR safe: particle p_T does not appear in the likelihood and therefore arbitrarily low energy particles can influence the clustering procedure. Therefore, we add a modification to the log likelihood:

$$\log \mathcal{L}(\{p_{T,i}, \rho_i\}|\theta) = \sum_{i=1}^n p_{T,i}^\alpha \log \left(\sum_{j=1}^k \pi_j f(\rho_i|\theta_j) \right), \tag{2.3}$$

where α is a weighting factor. Equation (2.3) is the log of eq. (2.1) with the term $p_{T,i}^\alpha$ inserted in the outer sum. For $\alpha > 0$, the resulting *modified* mixture model (mMM) jets are IR safe, and when $\alpha = 1$, the jets are C safe. Therefore, for $\alpha = 1$, the jets are IRC safe. Different choices of component densities f in eq. (2.3) give rise to different IRC safe MM jet algorithms. We have studied several possibilities for f , but for the remainder of this paper will specialize to (wrapped⁴) Gaussian⁵ $f = \Phi$. The resulting fuzzy jets are called modified Gaussian Mixture Model jets (mGMM) and are parameterized by the locations μ_j , the covariance matrices Σ_j , and the cluster weights π_j . We initialize $\pi_j = 1/k$ and $\Sigma_j = I$.

Since practical procedures for maximizing the modified likelihood in eq. (2.3) may converge to stationary points that are not globally optimal, the output of a fuzzy jet

³Soft assignments for jets during clustering was studied in the context of the ‘optimal jet finder’ [17] which maximizes a function of the soft assignments.

⁴One must take care in selecting a class of densities appropriate for the angular quantity ϕ . For more details on the wrapped Gaussian distribution and motivation for its use in this context, see appendix A.

⁵When f is a circular step function, the algorithm is related to the Snowmass iterative cone algorithm [20] via the ‘Snowmass Potential’ [21].

algorithm will depend on an initial setting of the cluster parameters θ and π . One simple procedure, used exclusively for the rest of the paper, is to seed fuzzy jets based on the output of a sequential recombination jet algorithm. This guarantees an IRC safe initial condition and therefore the entire procedure is IRC safe. We now discuss practically how one can find the maximum of the fuzzy jets likelihood.

3 Clustering fuzzy jets: the EM algorithm

One iterative procedure for maximizing the mixture model likelihood in eq. (2.1) is the *Expectation-Maximization* (EM) algorithm [22–24]. After initializing the cluster locations and prior density π , the following two steps are repeated:

Expectation Given the current values of θ_j , compute the fuzzy membership probabilities

$$q_{ij} = \pi_j \Phi(\vec{\rho}_i | \mu_j, \Sigma_j) / \sum_{j'} \pi_{j'} \Phi(\vec{\rho}_i | \mu_{j'}, \Sigma_{j'}).$$

Maximization Given q_{ij} , maximize the *expected modified complete log likelihood* over the parameters π, μ, Σ .

The expected modified complete log likelihood has the form

$$\sum_{i=1}^n \sum_{j=1}^k p_{Ti}^\alpha (q_{ij} \log \Phi(\vec{\rho}_i; \vec{\mu}_j, \Sigma_j) + q_{ij} \log \pi_j). \quad (3.1)$$

Note that the expected modified complete log likelihood is not the same as the expected modified log likelihood, shown in eq. (2.3). They differ in that the complete log likelihood has the second sum outside the logarithm while eq. (2.3) has the sum inside the logarithm. The power of the EM algorithm is that maximizing the complete log likelihood results in fixed point iteration to monotonically improve the original log likelihood. This desirable property of the EM algorithm is still true when $\alpha > 0$; for a proof, see appendix B. Many choices for f have closed form maxima for the M step; in the Gaussian $f = \Phi$ case outlined above, the updates are given by

$$\mu_j^* = \sum_{i=1}^n \tilde{q}_{ij} x_i \quad \Sigma_j^* = \sum_{i=1}^n \tilde{q}_{ij} (x_i - \mu_j)(x_i - \mu_j)^\top \quad \pi_j^* = \frac{1}{\sum_{i=1}^n p_{Ti}^\alpha} \sum_{i=1}^n p_{Ti}^\alpha \tilde{q}_{ij}, \quad (3.2)$$

where $\tilde{q}_{ij} = q_{ij} p_{Ti}^\alpha / \sum_{l=1}^n q_{il} p_{Ti}^\alpha$. The well-known k -means clustering algorithm [25] can be recovered as the limit of expectation-maximization in a Gaussian mixture model with $\Sigma = \sigma^2 I, \sigma^2 \rightarrow 0$. Figure 2 illustrates GMM clustering using the EM algorithm with $k = 2$ clusters. The EM algorithm readily accommodates constraints on the model parameters. One constraint we will consider throughout the rest of the paper is $\Sigma_j = \sigma_j^2 I$ for all j , which requires the curves of constant likelihood in (y, ϕ) to be circular. We will see in the next section that the learned value of σ_j is useful for distinguishing jets originating from different physics processes. Note that since the modified complete log likelihood is IRC safe, the EM algorithm does not break the IRC safety of the original log likelihood.

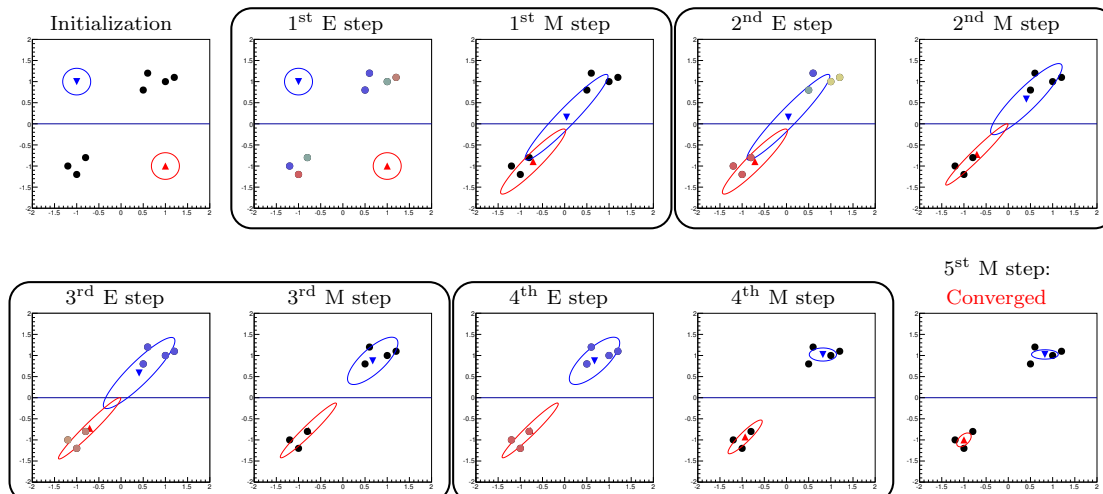


Figure 2. An illustration of of the EM algorithm for $k = 2$. The circles represent data points, the triangles represent the estimated cluster locations μ_j , and the ellipsoids are equidensity contours describing the shapes Σ_j of the learned cluster distributions. In the E-step, bluer colors correspond to higher value of $p_{i, \text{blue jet}}$.

4 Comparisons with sequential recombination and jet tagging

This section describes some numerical comparisons between sequential recombination and fuzzy jets. Section 4.1 summarizes the simulation details with some first event displays showing both fuzzy and sequential recombination jets. These two approaches to jet clustering are studied over an ensemble of events in section 4.2. A third subsection, section 4.3, illustrates that fuzzy jets captures new information about the hadronic final state, and in the fourth section, section 4.4, it is demonstrated that this new information can be used to classify the jet type.

4.1 Details of the simulation

In order to study fuzzy jets in a realistic scenario, we run Monte Carlo (MC) simulations. Three physics processes are generated using PYTHIA 8.170 [26, 27] at $\sqrt{s} = 8$ TeV. Hadronic W boson and top quarks are used for studying hard 2- and 3-prong type jets, respectively. To simulate high p_T hadronic W decays, W' bosons are generated to decay exclusively into a W and Z boson which subsequently decay into quarks and leptons, respectively. The p_T scale of the hadronically decaying W is set by the mass of the W' which is tuned to 800 GeV for this study so that the $p_T^W \lesssim 400$ GeV. In this p_T^W range, the W decay products are expected to merge within a cone of $R = 1.0$ where $\Delta R^2 = \Delta\phi^2 + \Delta\eta^2 \sim 4m_W^2/p_{T,W}^2$. A sample enriched in 3-prong type jets is generated with $Z' \rightarrow t\bar{t}$, where the Z' mass sets the energy scale of the hadronically decaying top quarks. In this analysis, we use $m_{Z'} = 1.0$ TeV, which sets $p_T^t \gtrsim 500$ GeV. To study the impact on signal versus background, QCD dijets are generated with a range of \hat{p}_T that is approximately in the same range as the relevant signal process. In all distributions, the QCD p_T spectrum is weighted to exactly match that of the signal to control for differences between signal and background due only to

the p_T spectrum differences. Pileup is simulated by overlaying additional independently generated minimum-bias interactions with each signal event. For the rest of this section, the number of pileup interactions $n_{\text{PU}} = 0$. See section 5 for studies of $n_{\text{PU}} > 0$.

For a comparison to fuzzy jets, anti- k_t jets are clustered using FASTJET [28] 3.0.3. The signal processes are chosen such that jets with radius parameter $R = 1$ are most appropriate in capturing the decay products of the heavy particles. The anti- k_t jets are trimmed [29] by re-clustering the constituents into $R = 0.3 k_t$ subjets and dropping those which have $p_T^{\text{subjet}} < 0.05 \times p_T^{\text{jet}}$. Anti- k_t jets are also used to seed the fuzzy jet clustering; the p_T threshold for this initialization is 5 GeV,⁶ and the impact of this choice is studied in appendix C. The choice of the parameters for the anti- k_t jet seeds is akin to the radius parameter R in the usual sequential recombination paradigm in that they can have a significant impact on the clustered jet properties. In complete analogy to the choice of R , the choice of seed jet parameters will depend on the targeted final state and the initial event conditions (e.g. pileup). Trimming is applied to all anti- k_t jets shown in subsequent sections and never applied to fuzzy jets.

At each iteration in the EM algorithm for fuzzy jet clustering, we record the value of the log likelihood. We make the choice to end the algorithm and declare convergence when the per iteration increase in this value is less than 10^{-6} for five consecutive iterations, or when a maximum of 100 iterations is reached. In practice most events converge after a much smaller number of iterations than this bound, with only a small fraction of events stopping for lack of convergence, and then only in high pileup scenarios ($n_{\text{PU}} > 80$).

To model the discretization and finite acceptance of a real detector, a calorimeter of towers with size 0.1×0.1 in (y, ϕ) extends out to $y = 5.0$. The total energy of the simulated particles incident upon a particular cell are added as scalars and the four-vector p_j of any particular tower j is given by

$$p_j = \sum_{i \text{ incident on } j} E_i(\cos \phi_j / \cosh y_j, \sin \phi_j / \cosh y_j, \sinh y_j / \cosh y_j, 1). \quad (4.1)$$

To simulate a particle flow reconstruction, the sum in eq. (4.1) contains only neutral particles for $|y| < 2.5$ and both charged and neutral particles for $2.5 < |y| < 5$. Charged particles within $|y| < 2.5$ are individually added to the list of inputs for clustering, unless they originate from a pileup collision. Anti- k_t jet momenta are corrected for pileup on average using area subtraction [19]. The median pileup density, ρ , is estimated by clustering hard scatter particles, neutral pileup particles, and charged pileup particles in the range $|y| < 2.5$ using k_t $R = 0.4$ jets in FASTJET with ghosted areas.

A representative event display for a $Z' \rightarrow t\bar{t}$ event is shown in figure 3. The top right plot in figure 3 shows the anti- k_t jets with $p_T > 5$ GeV as filled in (partial) circles. The filled area is determined by the jet area and there are deviations from circles only one a low p_T jet is close to a higher p_T jet. The two top quarks are depicted as red stars, each of which sits at the center of two high p_T jets. The top left plot in figure 3 shows mGMM

⁶This low threshold guarantees that there are enough seed jets around to capture the radiation from the underlying event. Another strategy could be to use the Event Jet (see section 5) even when there is no pileup.

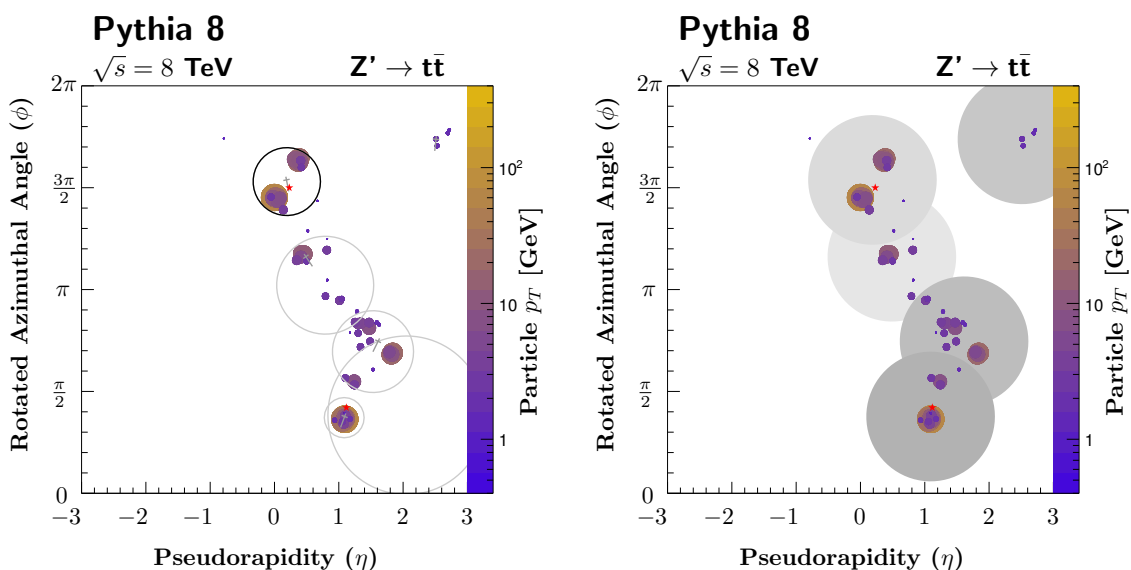


Figure 3. A representative event display for a $Z' \rightarrow t\bar{t}$ event. In the top left plot, gray circles show the location and size of mGMM fuzzy jets after clustering, with the size of the circle indicating $1\text{-}\sigma$ contours in the detector; the black circle indicates the highest p_T jet with HML particle assignment. The small filled colored circles are the particles, with the color and size indicating their energy. In each case, the events have been rotated in ϕ to place the truth top quark at $\phi = 3\pi/2$, which is indicated by a red star. Anti- k_t jet locations are shown with gray crosses in the left hand plot, the long tail of which points towards the mGMM jet for which it was a seed. In the top right plot, anti- k_t $R = 1.0$ jets passing a $5\text{ GeV } p_T$ cut are shown as discs under the particles indicating their active area, with centers the same as the crosses in the left hand side. Shades of gray in the anti- k_t discs have no scale and are meant to aid the eye, but go from low p_T (lighter) to high p_T (darker).

fuzzy jets. The fuzzy jets are depicted by their $1\text{-}\sigma$ contours. In contrast to the anti- k_t jets, fuzzy jets vary widely in radial size. Gray crosses in the top left plot indicate the locations of the anti- k_t jets shown in the top right plot. The long tail of the crosses point toward the fuzzy jet for which they were the seed. The two jets closest to the top quarks did not move a long distance from the seed location, though the size did change significantly from $R = 1$. The lowest p_T fuzzy jet moved a long distance from the seed to the final location.

Another new feature of fuzzy jets compared to anti- k_t jets is that they can overlap with each other. This is seen by the four jets with overlapping $1\text{-}\sigma$ contours in the top left plot of figure 3. Overlapping mGMM jets are an expression of structure inadequately captured with a single Gaussian shape. The ability to learn features at different scales in the same event without relying on a size parameter like the anti- k_t radius parameter can give mGMM fuzzy jets additional descriptive power over anti- k_t and other traditional jet algorithms. This particular event will be used again for reference in section 5 during a discussion on the performance of the technique in the presence of pileup interactions.

4.2 Kinematic properties of fuzzy jets

Jets clustered according to the mGMM algorithm capture similar hard jet locations and jet energy (under HML) as those clustered by anti- k_t $R = 1$. In figure 4, the p_T distribution




Figure 4. The jet p_T for the leading anti- k_t jet (left) and leading fuzzy jet under the HML particle assignment scheme (right). All the processes are re-weighted so that the anti- k_t p_T spectra are the same.

for the highest p_T jets for three different physics processes are plotted as given by anti- k_t $R = 1.0$ and mGMM jets. The anti- k_t p_T distributions are re-weighted so that all three processes have identical distributions in the left plot. On the right, the distributions are in good correspondence with those in the left plot, though there is a slight shift of the peak. Additionally, the (y, ϕ) locations of the highest p_T mGMM jets are in excellent correspondence with the locations of the anti- k_t jets as was already discussed in reference to figure 3.

The mGMM algorithm differs from the anti- k_t algorithm in how the size and structure of clustered jets. This was already shown qualitatively in figure 3: fuzzy jets come in a variety of sizes, and can overlap in complex ways. The matter is further complicated by the choice of particle assignment scheme for defining kinematic properties in the mGMM family of algorithms. The catchment area’s volume and shape of a fuzzy jet depends in general on the full set of learned jet locations and model parameters, Σ . In contrast, for anti- k_t jets, the catchment area is bounded from above by R and is only smaller when another high p_T jet is nearby. The nonlocality of the mGMM clustering model can be observed quantitatively by examining jet mass, given in eq. (4.2), which is sensitive to the distribution of energy within a jet. The jet mass distributions for both mGMM (HML assignment) and anti- k_t jets are shown in figure 5, with the same p_T weighting as in figure 4. Even though fuzzy jets learn the same core (i.e. p_T) for jets as anti- k_t , they do not learn the same mass. The white dashed lines in figure 5 mark the locations of the W boson and top quark masses at about 80 GeV and 175 GeV, respectively [30]. For both anti- k_t and fuzzy jets, there are clear peaks at the W mass for the boosted $W \rightarrow qq'$ from W' simulated events and at the top quark mass for $Z' \rightarrow t\bar{t}$ simulated events. However, there are clear differences in the shape of these distributions — in particular the fuzzy jets generally have a lower mass than their anti- k_t analogues. The W mass peak for W' events is more peaked for fuzzy jets, though there is also a low-mass contribution to the distribution. For Z' events, the top quark mass peak is less populated for fuzzy jets, which instead has shifted

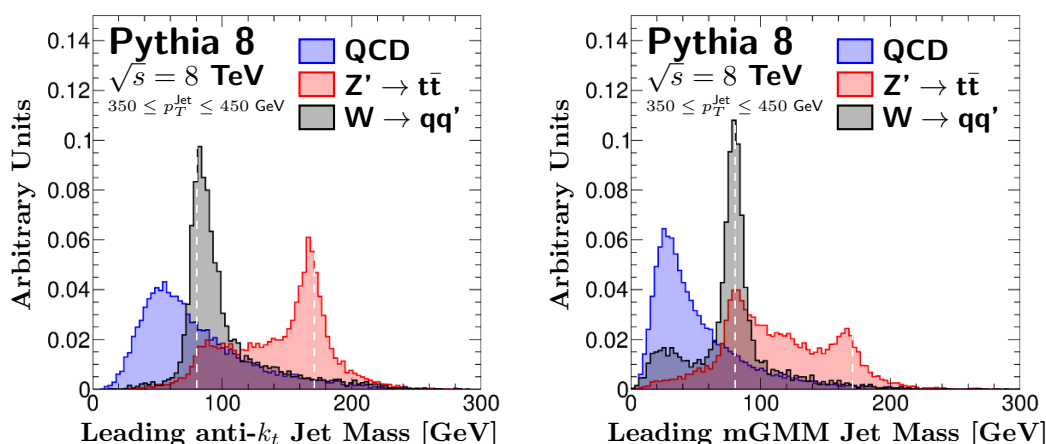


Figure 5. The jet mass for the leading anti- k_t (left) and leading fuzzy jet under the HML particle assignment scheme (right), in an anti- k_t leading jet p_T window of 350 to 450 GeV. All the processes are re-weighted so that the anti- k_t p_T distributions are the same. The dashed white lines mark $m_W = 80.4$ GeV and $m_t = 173.3$ GeV.

events to the W mass peak. This often happens when the tree-prong structure is learned by two (overlapping) fuzzy jets. The QCD multi-jet jet mass distribution is also qualitatively different between fuzzy jets and anti- k_t jets, with the former shifted to lower values of the mass

$$m_{\text{jet}}^2 = \left(\sum_{i \in \text{jet}} E_i \right)^2 - \left(\sum_{i \in \text{jet}} \vec{p}_i \right)^2. \quad (4.2)$$

4.3 New information from fuzzy jets

The properties Σ of a fuzzy jet can be useful in distinguishing jets resulting from different physics processes. In the simplest realization of mGMM jets already described above, $\Sigma = \sigma^2 I$, where σ is a measure of the size of the core of a jet. Although σ is a simple variable to construct from the wealth of data available after clustering with the mGMM algorithm, it captures at least some of the schematic differences in the likelihood for $Z' \rightarrow t\bar{t}$ and $W' \rightarrow WZ$ relative to a QCD multijet background (shown below).

The left plot of figure 6 also shows the average σ over all fuzzy jets in an event. The generic fuzzy jet is rather independent of the physics process and tends to be quite large. This is because fuzzy jets capturing hard radiation tend to be small, but most of the fuzzy jets needed to capture the sparse radiation pattern from the underlying event need to be large. In contrast, the σ for the leading⁷ mGMM jets are shown the right plot of figure 6 for each of the three physics processes. As expected, the decay relative size of the highest

⁷The distribution for σ for the sub-leading jet in the W' and Z' processes is qualitatively similar to that of the leading jet even though event-by-event they are largely uncorrelated. The σ for the sub-leading jet in QCD multijet events is systemically wider than for the leading jet, as expected from the discussions around figure 6.




Figure 6. The learned value of σ for the highest p_T jet under the HML scheme (left) and for all jets (right) for various physics processes.

p_T jets depends on the physics process. For the decay of a boosted heavy particle with mass m and p_T , the radial size of the decay products scales as $2m/p_T$ and thus since the p_T distribution in figure 6 is fixed, one would expect that the top quark jets have a larger σ than the W boson jets, which are in turn larger than the quark and gluon jets. This is reflected⁸ in the three peaks in the left plot of figure 6. The separation between the three physics processes is not 100% correlated with the naive scaling m/p_T of the corresponding leading anti- k_t jets. Figure 7 shows that there is a strong positive correlation between σ and the corresponding anti- k_t mass over p_T as expected. There are two peaks in the correlation for the $Z' \rightarrow t\bar{t}$ events because the anti- k_t mass spectrum has peaks at both the top mass, and the W boson mass. While the correlations between the fuzzy jet σ and the anti- k_t m/p_T are non-negligible, they are far from unity and thus there may be additional information contained in the fuzzy jet σ that is useful for tagging the flavour of a jet. Note that part of the new information in σ is resulting from the clustering procedure and not simply the definition of the observable. For example, computing σ from the constituents of an anti- k_t jet (i.e. running fuzzy jets on these constituents with $k = 1$) would result in

$$\sigma^2 = \frac{\sum_{i=1}^n p_{T,i} \Delta R^2}{\sum_{i=1}^n p_{T,i}}, \tag{4.3}$$

which is nearly the same as m/p_T .

4.4 Fuzzy jets for tagging

In this section, σ is compared with another class of jet substructure variables known to be useful for tagging: the N -subjettiness ratios [31]. N -subjettiness moments are defined

⁸At leading order, there is an exact relationship between σ and the jet mass — See appendix D. While beyond the scope of this study, it would be interesting to understand in more detail how σ compares to mass analytically (beyond leading order, with $k > 1$, etc.).




Figure 7. The left and right plots show the correlation between σ and the leading jet anti- k_t mass divided by p_T in an anti- k_t p_T window of 350 to 450 GeV for $Z' \rightarrow t\bar{t}$ and QCD events, respectively. Indicated in the lower right of each figure is the linear correlation between the variables.

over a set of N axes⁹, and calculated as:

$$\tau_N = \frac{1}{d_0} \sum_k p_{T,k} \min\{\Delta R_{1,k}, \Delta R_{2,k}, \dots, \Delta R_{N,k}\}, \quad (4.4)$$

where d_0 is the normalization

$$d_0 = \sum_k p_{T,k} R_0, \quad (4.5)$$

and R_0 is the radius of the jet. In practice, the useful variables for determining how much more i -pronged a jet is compared to j -pronged are the N -subjettiness ratios:

$$\tau_{ij} = \frac{\tau_i}{\tau_j}. \quad (4.6)$$

The variable τ_{21} is often used for the separation of W from QCD jets [32, 33] and is a measure of the compatibility of a jet with a 2-prong hypothesis compared to a 1-prong hypothesis. Low value of τ_{21} indicates that the jet likely has a 2-prong structure. Similarly, τ_{32} is useful for top tagging in that it measures whether a 3-prong structure is a better description of a jet relative to a 2-prong structure.

The rest of this section contains comparisons of the performance of σ relative to τ_{21} (τ_{32}) as well as m/p_T for separating W ($Z' \rightarrow t\bar{t}$) from QCD jets. In figures 8 and 9, a k -nearest neighbors classifier was trained with 2-fold cross validation¹⁰ in TMVA [35]. The left plots in figures 8 and 9 demonstrate an increase in performance for discriminating $Z' \rightarrow t\bar{t}$ from QCD relative to using τ_{32} or m/p_T (an anti- k_t analogue to the jet size)

⁹We use the “one-pass” k_t axes optimization technique, which uses an exclusive k_t algorithm to find N axes and then refines them by minimizing the N -subjettiness value.

¹⁰This is a standard procedure for selecting the free parameters of the algorithm (in this case, k) by dividing the training sample into a sub-training sample and a sub-test sample. For more details, see for instance ref. [34].

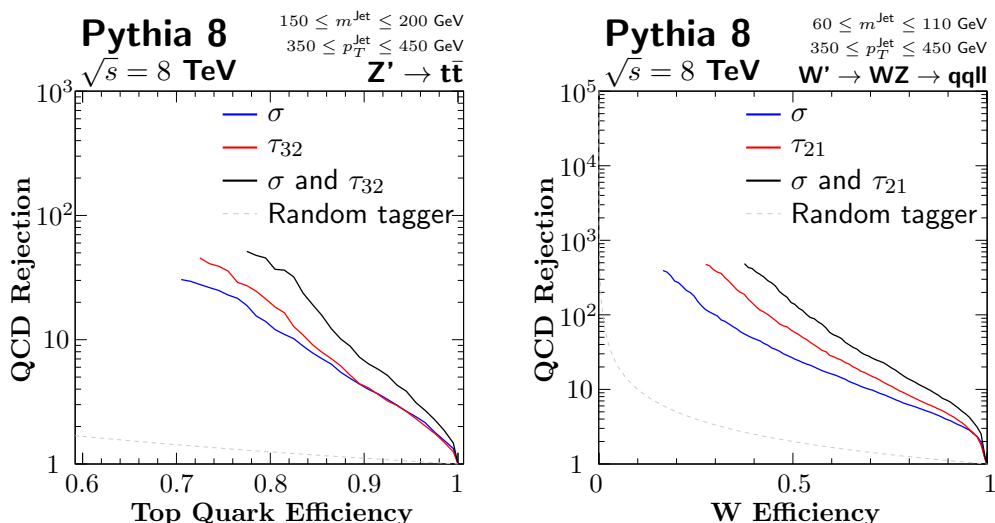


Figure 8. The tagging performance of σ relative to τ_{32} (τ_{21}) for distinguishing top quarks (W bosons) from a QCD background is shown on the left (right). The *random tagger* keeps a fixed fraction of all events, regardless of their origin and is a lower bound on the performance of any tagger.

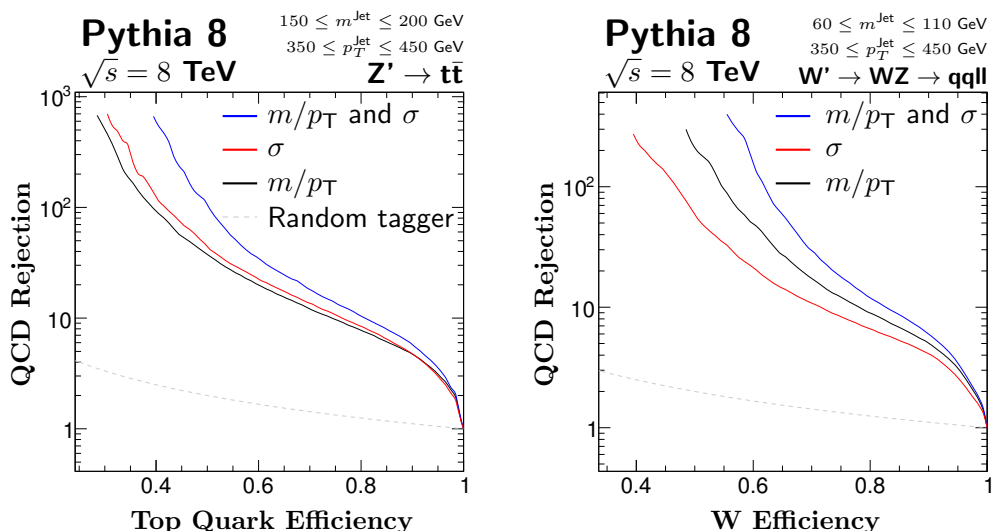


Figure 9. The tagging performance of σ relative to m/p_T (using anti- k_t $R = 1.0$ trimmed jets) for distinguishing top quarks (W bosons) from a QCD background is shown on the left (right). The *random tagger* keeps a fixed fraction of all events, regardless of their origin and is a lower bound on the performance of any tagger.

alone. The fuzzy jet σ is roughly equally useful to the N -subjettiness ratio (and m/p_T) at a signal efficiency of 0.85, and using both variables greatly improves background rejection. Similar results can be seen in the right plots of figures 8 and 9, where σ boosts background rejection relative to τ_{21} or m/p_T alone. In each case, the training and classification was performed in a mass window around the particles of interest, the top quark mass in the $Z' \rightarrow t\bar{t}$ sample and the W boson mass to discriminate $W \rightarrow qq'$ from QCD.

The comparisons of the fuzzy jet σ and N -subjettiness are intended to be an illustrative example. As discussed in the opening of this section, σ is just *one* variable that can be constructed by using mGMM clustered jets. Expanded studies of the various learned parameters could come up with additional variables, or the full learned parameter set could be thrown into an off the shelf classifier or machine learning model.

5 Underlying event and pileup

As with any new jet algorithm or jet variable, understanding the effect of pileup vertices from additional proton-proton collisions is essential to make meaningful statements about how the method will be applicable to real data analyses at the LHC. Studying pileup in the context of mGMM jets is complicated by the effective catchment area of the jets. For hierarchical-agglomerative algorithms like anti- k_t , the catchment area scales with the radius parameter. However fuzzy jets can have infinite catchment area because the likelihood for particle membership is nonzero for any finite distance and arrangement of Gaussian jets and particles. Furthermore, the catchment area can change depending on the other jets in an event. Although this effect also occurs in the hierarchical-agglomerative case, the effect is much more pronounced in the mGMM clustering algorithm, with some jets having finite catchment areas while others cluster infinite area.

The challenge of pileup for fuzzy jets is illustrated in figure 10, where the same event is shown with $n_{\text{PU}} = 0$, and with $n_{\text{PU}} = 40$. The event displays show the central region of the detector, where most of the decay products of the hard scatter lie. Qualitatively, it can be seen that the introduction of additional interaction vertices broadens all of the mGMM jets. This broadening clearly impacts the power of σ for differentiating QCD background from signal processes.

The next sections explore two methods for mitigating the impact of pileup in relation to fuzzy jets, illustrated with the variable σ .

5.1 Changing α for pileup suppression

In section 2, it was discussed that choosing $\alpha = 1$ in the likelihood (eq. (2.3)) guarantees IRC safety. With $\alpha = 1$, the mGMM algorithm treats hard structure and soft structure linearly in the particle or tower p_T . However, one can exploit the fact that σ is disproportionately a measure of the shape and extent of the leading jet hard structure to make the variable more resilient to the effects of pileup. In particular, choosing $\alpha > 1$ stabilizes σ at high n_{PU} because so long as the average input particle p_T due to pileup is significantly smaller than the p_T of the particles constituting the leading jet hard structure, the change in likelihood will be suppressed roughly according to $(p_{T,\text{hs}}/p_{T,\text{PU}})^\alpha$. An example of this effect is illustrated in figure 11, which shows the same event as in figure 10. The price for adjusting α is the loss of collinear safety. Varying α is not explored further, as section 5.2 demonstrates a method for dealing with pileup effectively that does not rely on moving α away from the IRC safe value of one.

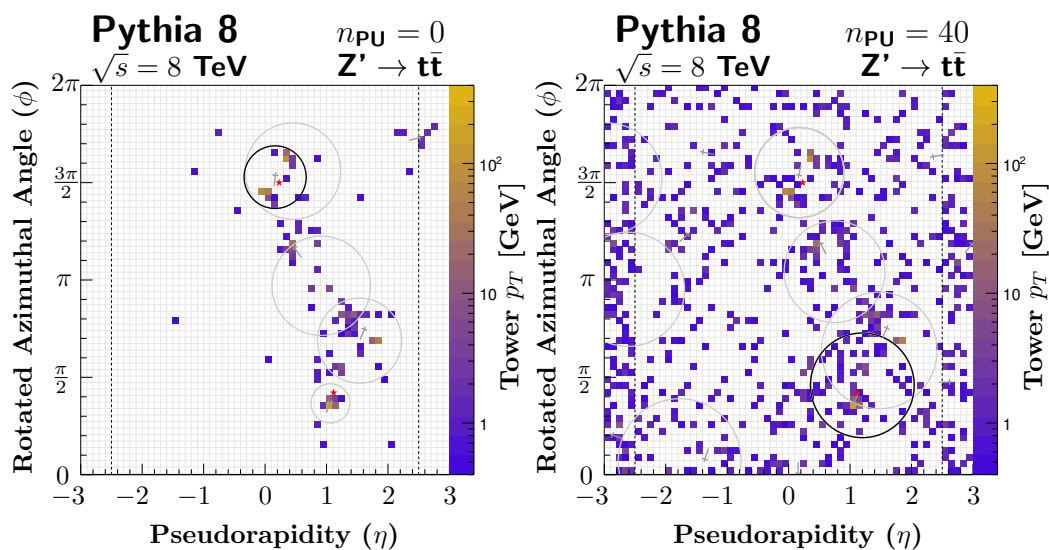


Figure 10. mGMM jets defined according to section 2 with an isotropic kernel are broadened as a result of the introduction of additional pp pileup vertices. The same hard scatter is clustered twice, on the left with $n_{\text{PU}} = 0$ and on the right with $n_{\text{PU}} = 40$. Vertical dashed lines at $\eta = \pm 2.5$ show the extent of a simulated tracker with the same η extent as that used at ATLAS and CMS. Charged pileup falling within the extent of the simulated tracker is discarded before clustering and the aggregation of particles into towers.

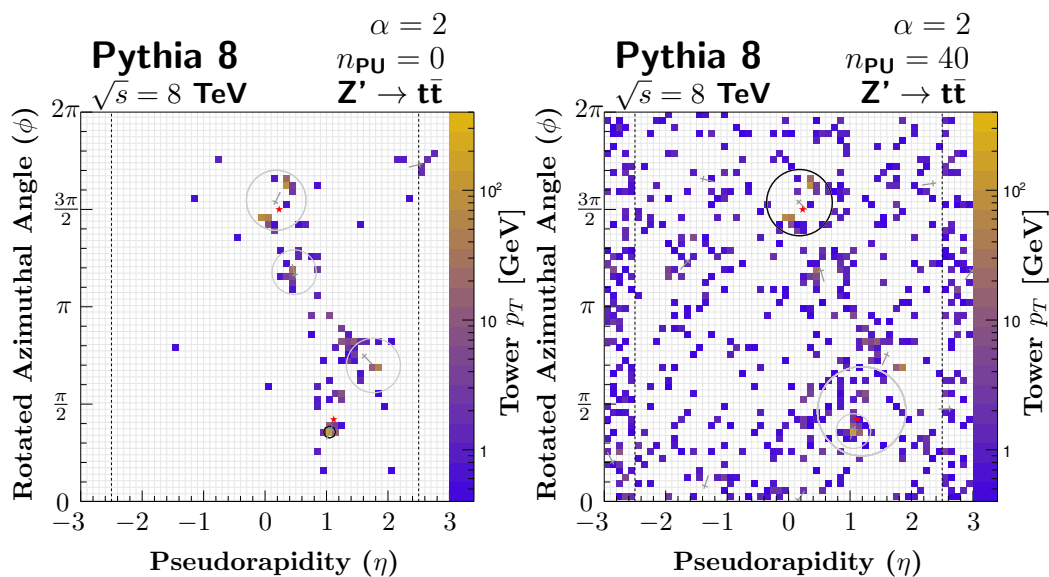


Figure 11. Clustering in the mGMM model with $\alpha = 2$. There is little broadening between the $n_{\text{PU}} = 0$ (left) and $n_{\text{PU}} = 40$ (right) cases, but jets at the locations of the tops in the event are substantially narrower than in the case where $\alpha = 1$, even with $n_{\text{PU}} = 0$ (compare to figure 10). Under the ML particle assignment, the $\alpha = 2$ algorithm identifies the other top as the highest p_T jet in the event, demonstrating the difficulty in dealing with fuzzy jet kinematics.

5.2 Tower subtraction and the event jet: effective pileup correction

Recent developments in pileup mitigation have led to several algorithms for correcting jet inputs before jet clustering begins. Such techniques include Pileup Per Particle Identification (PUPPI), Constituent Subtraction, and SoftKiller [36–38]. One simple input-correction scheme is to subtract from each calorimeter tower the estimated pileup p_T density per unit area multiplied by the size of the tower in the detector. As a first step, ρ is calculated in the same way as described in section 4.1. Tower momenta are then corrected according to eq. (5.1), where $p_{T,s}$ is the corrected momentum, $p_{T,o}$ is the original momentum, and A is the area of the tower.¹¹ Even though this procedure *undersubtracts*, it already mitigates a large amount of the impact of pileup (one reason a second step is introduced, below). In this study, all towers have area 0.1×0.1 in y - ϕ space.

$$p_{T,s} = \max(p_{T,o} - \rho A, 0). \tag{5.1}$$

While subtracting the average p_T background from towers before clustering is a relatively safe way of reducing the effect of pileup, at least when the p_T scales of the tower to tower fluctuations are small compared to the hard scatter p_T scale, it would still be helpful to systematically address the question of catchment areas. The mGMM clustering algorithm provides a natural framework in which to think about pileup, however, because the algorithm deals fundamentally with likelihoods, and the pileup likelihood is to leading order uniform over the detector (this is the motivation for the area-subtraction technique). This is the motivation for modifying the mGMM likelihood using a technique we call the *event jet*.

In addition to learning k mGMM jets throughout clustering, the event jet includes another background contribution to the likelihood which attempts to capture the intuition of a uniform contribution of particle likelihood due to pileup. Constraints are further imposed on the likelihood on the event level jet so that it has constant likelihood during the clustering process, making the necessary modifications to the algorithm procedures simpler.

Practically, the effect of the event jet can be parameterized through the introduction of an algorithmic parameter γ . Particle membership probabilities change according to eq. (5.2) with corresponding changes to the analytical M step for the Gaussian kernel type. The choice of γ is important, and it should reflect the fact that not all events are created equal in the sense that not all events have the same contributions due to pileup. Although there is no strict way of dealing with this issue, it is reasonable to replace γ by a meaningful combination of parameters which is sensitive to our estimates of the amount of pileup in a particular event. We have chosen to take $\gamma = \rho A \gamma_w$ where ρ is our estimate of the p_T density due to pileup, A is the calorimeter area, and γ_w is a parameter of the algorithm controlling the strength of the event jet. Initial studies with the event jet indicate that introducing a ρ dependent γ is much more effective than a ρ independent one.

$$q_{ij} \rightarrow \frac{q_{ij}}{\gamma + \sum_k p_{ik}} \tag{5.2}$$

¹¹A simple tower-based subtraction scheme is used here for illustration, but in principle the fuzzy-jet + event-jet algorithm can run after the application of any tower-based or particle-based constituent pileup subtraction technique.

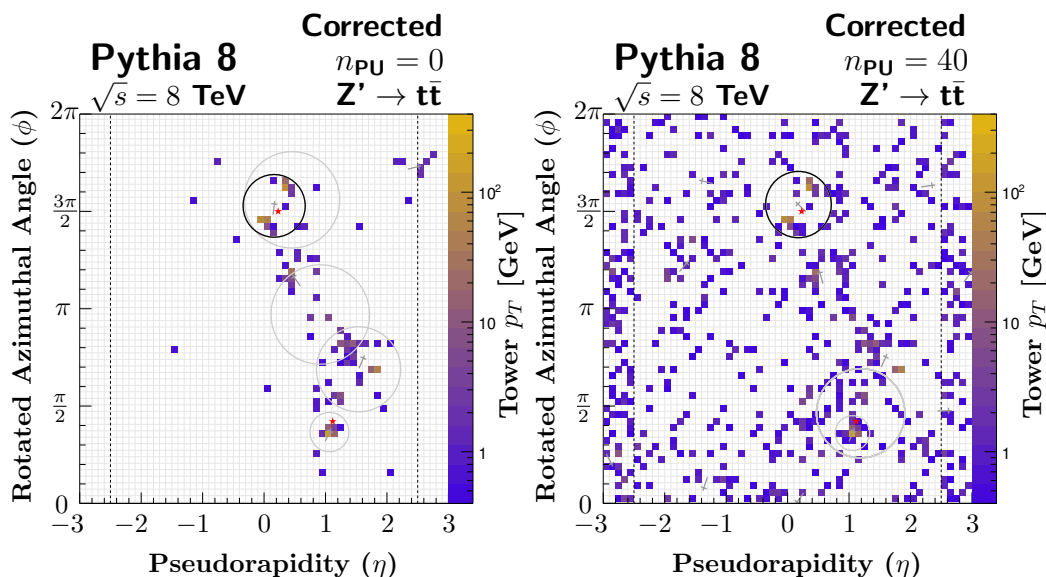


Figure 12. Jet correction using tower subtraction and the event jet with parameter $\gamma_w = 0.01$. The two leading p_T jets are almost identical in size in the left and right insets, which show the $n_{PU} = 0$ and $n_{PU} = 40$ cases respectively. Although many of the other jets change (including the migration of jets to higher $|\eta|$ as a result of the simulated tracker), those that give the σ and sub-leading σ variables are insensitive to the effect of pileup.

Studies of the pileup conditions similar to LHC Run I, with ~ 20 pileup interactions, indicate that with a 5 GeV p_T cut, $\gamma_w = 0.01$ provides reasonable stability¹² of the learned σ . This is demonstrated qualitatively in figure 12, in which the tower and event jet corrections are applied to the same event shown in figure 3 at both $n_{PU} = 0$ and $n_{PU} = 40$. Unlike any of the methods discussed previously, this method for correction maintains IRC safety, demonstrates very little jet broadening at $n_{PU} = 40$, and is not drastically different in its qualitative features by comparison to the standard mGMM algorithm. Note that the assignment of towers to jets under the HML scheme is impacted with the event jet because many towers belong to the event jet with higher probability than any of the other fuzzy jets. In particular, under the HML scheme, a tower is assigned to the event jet and not another jet if $\max_k p_{ik} < \gamma$. To preserve tower-to-jet assignments under pileup, a smaller value of γ_w should be chosen. The event jet is useful instead because it changes the dynamics of clustering, making jets less sensitive to soft radiation far away from the jet axis during the EM update steps, and therefore increasing the stability of the hard core that is eventually clustered.

A quantitative study of the pileup mitigation suggested qualitatively by figure 12 requires an ensemble of events. Figure 13 shows how the mean and standard deviation of learned σ evolve with n_{PU} . The uncorrected σ is shown in red downward pointing triangles while the tower subtraction and event jet corrections are shown in blue upward pointing

¹²This particular value was \sim optimal for both the W' and Z' signal processes and for n_{PU} ranging from 0 to 40. The optimal value of γ_w does depend on the p_T cut: changing the number of jets in the event modifies the scale for the likelihood and γ_w operates by rescaling the likelihood.

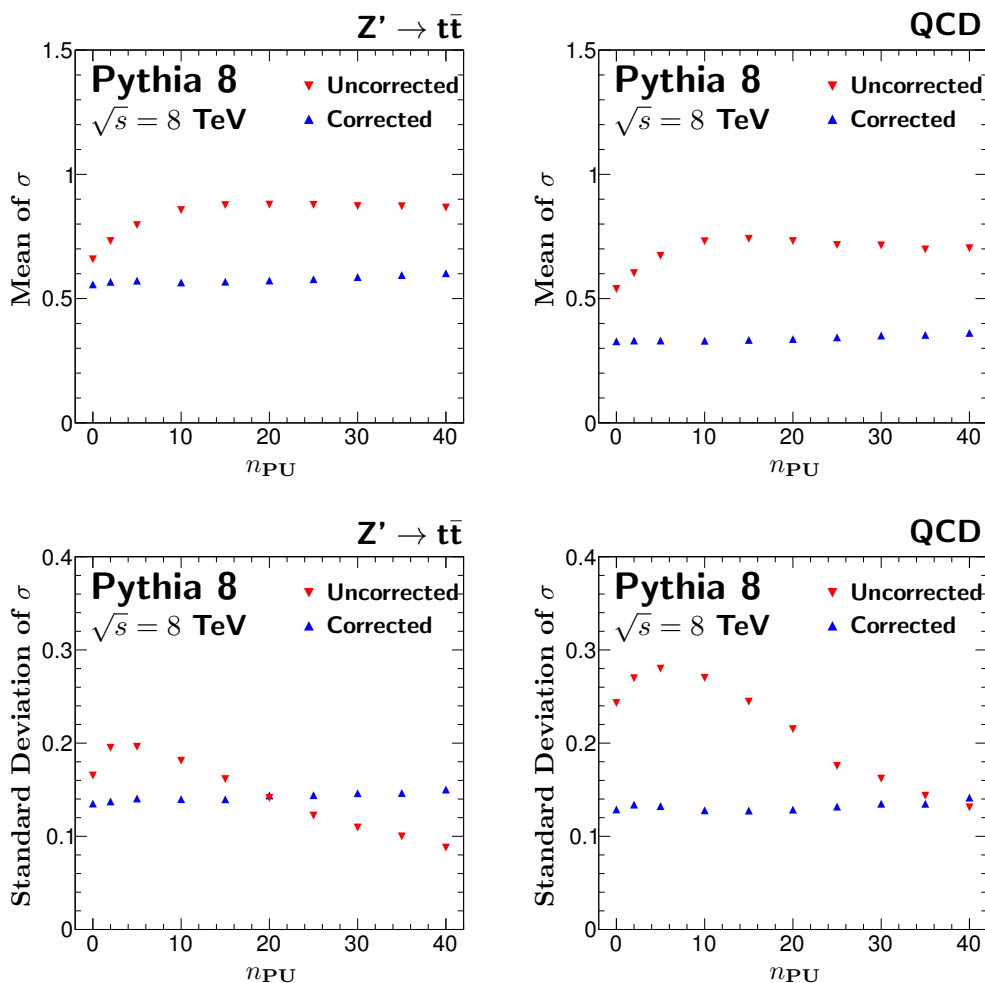


Figure 13. For both QCD and $Z' \rightarrow t\bar{t}$ samples, using the pileup correction (blue triangles) via the event jet and tower subtraction stabilizes the mean relative to the uncorrected samples (red inverted triangles), and prevents widening of the σ distribution in pileup conditions somewhat worse than during Run 1 at the LHC.

triangles. For both $Z' \rightarrow t\bar{t}$ and QCD, the pileup dependence is dramatically reduced with the tower subtraction and the event jet. The uncorrected mean σ increases as a function of n_{PU} as all of the fuzzy jets become the same size. The standard deviation of the uncorrected σ actually decreases beyond $n_{PU} \sim 5$ as all of the fuzzy jets become the same size. For modest levels of pileup, tower subtraction and event and the event jet maintain the mean and standard deviation of the σ distribution.

6 Conclusions

The modified mixture model algorithms provide a new way of looking at whole event structure. In contrast to the usual uses of hierarchical-agglomerative algorithms like anti- k_t , the number of seeds is fixed ahead of time and their properties are learned during the clustering process. The learned parameters provide a new set of handles for distinguishing

jets of different types. Even simple variables constructed out of the learned parameters of a mixture of isotropic Gaussian jets, like σ , offer complementary information to the n -subjettiness variables τ_{21} and τ_{32} for tagging W boson and top quark jets. Even though the variable σ is sensitive to changes in pileup conditions, small modifications to the fuzzy jets algorithm — correcting jet inputs and adding a pileup likelihood — can mitigate the impact of pileup.

Fuzzy jets are new paradigm for jet clustering in high energy physics. These IRC safe likelihood-based clustering schemes set the stage for many possibilities for future studies related to jet tagging, probabilistic clustering, and pileup suppression.

Acknowledgments

We would like to thank Jesse Thaler for useful discussions and helpful feedback on the manuscript. In addition, we thank Gavin Salam for useful comments on the algorithm description. This work is supported by the US Department of Energy (DOE) Early Career Research Program and grant DE-AC02-76SF00515. BN is supported by the NSF Graduate Research Fellowship under Grant No. DGE-4747 and by the Stanford Graduate Fellowship.

A Wrapped gaussian

In the EM algorithm described in section 3, there are explicit (and implicit) dependencies on the topology. For instance, if a Gaussian density is used to model ϕ , then, in the E step, a particle with ϕ_i near 2π will be deemed far from a cluster with location ϕ_j near 0. To avoid this undesirable behavior and enforce the equivalence of the angles 0 and 2π , we associate ϕ with a *wrapped Gaussian density* and y with a standard Gaussian density:

$$\Phi(y, \phi | \mu_\phi, \mu_y, \sigma^2) = \Phi_y(y | \mu_y, \sigma^2) \frac{1}{\sqrt{2\pi\sigma^2}} \sum_{I=-\infty}^{\infty} \exp \left[\frac{-(\phi - \mu_\phi(I))^2}{2\sigma^2} \right], \quad (\text{A.1})$$

where Φ_y is a normal distribution and $\mu_\phi(I) = \mu_\phi + 2\pi I$. In order to approximate the sum in eq. (A.1), we take only the leading contribution by choosing $\mu_\phi(I^*)$ for $I^* = \text{argmin}_I |\phi - \mu_\phi + 2\pi I|$. As other contributions are exponentially suppressed, this is a good approximation and recovers continuity near 0 and 2π . Figure 14 illustrates the improved clustering behavior that results when ϕ is modeled using the wrapped Gaussian approximation in place of the standard Gaussian density.

B The EM algorithm

This appendix contains two derivations: the modified EM algorithm updates in eq. (3.2) and the proof that the modified EM algorithm generically improves the original modified log likelihood eq. (2.3) with every iteration. Recall the expected modified complete log likelihood (mmCLL) from eq. (3.1):

$$\sum_{i=1}^n \sum_{j=1}^k p_{T_i}^\alpha (q_{ij} \log \Phi(\vec{\rho}_i; \vec{\mu}_j, \Sigma_j) + q_{ij} \log \pi_j).$$

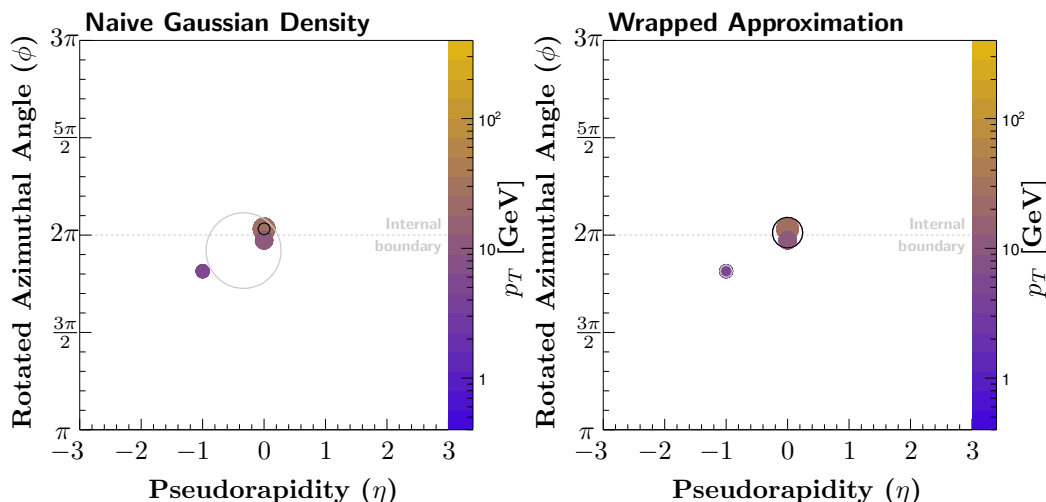


Figure 14. A three-particle event display illustrating the results of fuzzy jet clustering using a Gaussian density for ϕ (left) and a wrapped Gaussian density approximation for ϕ (right).

Viewing the mCLL as a function of $\vec{\mu}, \Sigma$ and π for fixed λ and $\vec{\rho}$ we can maximize. For π , we optimize

$$\sum_{i=1}^n \sum_{j=1}^k p_{Ti}^{\alpha} (q_{ij} \log \pi_j) + \lambda \left(\sum_{j=1}^k \pi_j - 1 \right),$$

where the last term is needed so that the optimal π^* is a probability. The derivative of this expression with respect to π_j is

$$\pi_j = -\frac{1}{\lambda} \sum_{i=1}^n p_{Ti}^{\alpha} q_{ij},$$

and then summing the equation over j and using $\sum_{j=1}^k q_{ij} = 1$ and the constraint equation $\sum_{j=1}^k \pi_j = 1$, we find that

$$\pi_j^* = \frac{1}{\sum_{i=1}^n p_{Ti}^{\alpha}} \sum_{i=1}^n p_{Ti}^{\alpha} q_{ij}$$

The updates for $\vec{\mu}$ and Σ follow from the standard derivation (by similarly taking derivatives of the mCLL with respect to components of these multi-dimensional objects) by noting that the only difference is that $q_{ij} \mapsto q_{ij} p_{Ti}^{\alpha}$ and there are no Lagrange multipliers needed unlike for π_j^* .

Finally, we prove the claim that the modified EM algorithm described in the body of the text monotonically improves the modified log likelihood in eq. (2.3). First, we note

that we can rewrite the (log) likelihood as

$$\begin{aligned}
 p_T^\alpha \log p(\rho|\theta) &= p_T^\alpha \log \left(\sum_{\lambda \in \{1,2,\dots,k\}} p(\rho, \lambda; \theta) \right) \\
 &= p_T^\alpha \log \left(\sum_{\lambda \in \{1,2,\dots,k\}} \frac{q(\lambda)p(\rho, \lambda; \theta)}{q(\lambda)} \right) \\
 &= p_T^\alpha \log \mathbb{E}_q \left[\frac{p(\rho, \lambda; \theta)}{q(\lambda)} \right] \\
 &\geq \mathbb{E}_q \left[p_T^\alpha \log \left(\frac{p(\rho, \lambda; \theta)}{q(\lambda)} \right) \right] \equiv \mathcal{L}(q, \theta),
 \end{aligned}$$

where the inequality in the last line follows from Jensen's inequality. Now, we are ready to prove the claim that $p_T^\alpha p(\rho|\theta^{(t)})$ improves monotonically with t , the index for the iteration of the EM algorithm. First, note that

$$\begin{aligned}
 \mathcal{L}(q, \theta) &= \mathbb{E}_q \left[p_T^\alpha \log \left(\frac{p(\rho, \lambda; \theta)}{q(\lambda)} \right) \right] \\
 &= \mathbb{E}_q [p_T^\alpha \log (p(\rho, \lambda; \theta))] - \mathbb{E}_q [p_T^\alpha \log (q(\lambda))],
 \end{aligned}$$

where the first term is the mCLL and the second term has no θ dependence and so maximize $\mathcal{L}(q, \theta)$ over θ is equivalent to maximize the mCLL over θ . Therefore, $\mathcal{L}(q^{(t+1)}, \theta^{(t)}) \leq \mathcal{L}(q^{(t+1)}, \theta^{(t+1)})$. By the inequality above, $\mathcal{L}(q^{(t+1)}, \theta^{(t+1)}) \leq p_T^\alpha p(\rho|\theta^{(t+1)})$. The E step can be recast as choosing

$$q^{(t+1)}(\lambda_i = j) = q_{ij}(\theta^{(t)}) = \mathbb{E}_{\theta^{(t)}} [q_{ij}] = p(\lambda|\rho, \theta^{(t)}).$$

This enforces:

$$\begin{aligned}
 \mathcal{L}(p(\lambda|\rho, \theta^{(t)}), \theta^{(t)}) &= \mathbb{E}_{p(\lambda|\rho, \theta^{(t)})} \left[p_T^\alpha \log \left(\frac{p(\rho, \lambda; \theta^{(t)})}{p(\lambda|\rho, \theta^{(t)})} \right) \right] \\
 &= \mathbb{E}_{p(\lambda|\rho, \theta^{(t)})} \left[p_T^\alpha \log (p(\rho; \theta^{(t)})) \right] \\
 &= p_T^\alpha \log (p(\rho; \theta^{(t)}))
 \end{aligned}$$

Putting this together with the bounds from the M step, we arrive at the desired result: $p_T^\alpha p(\rho|\theta^{(t)}) \leq p_T^\alpha p(\rho|\theta^{(t+1)})$, i.e., every step of the modified EM algorithm improves or leaves the same the original likelihood.

C Controlling jet multiplicity with p_T

In contrast to most uses of hierarchical-agglomerative clustering algorithms, the number of fuzzy jets is fixed before clustering begins. Whereas a single traditional jet can reasonably be considered to correspond to a parton in appropriate cases, mGMM jets should not be, as several mGMM jets can together express structure of what would be one or several

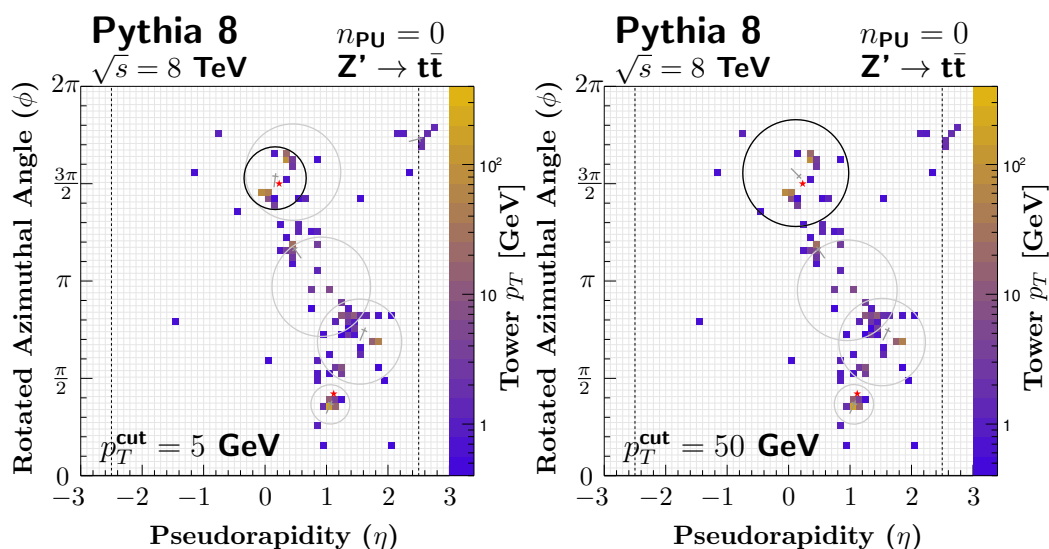


Figure 15. Changing the choice of the p_T cut used to select seeds can make a vast difference in the values of the constructed variables, like σ . In this event, clustered on the left with a cut of 5 GeV resulting in five jets, and on the right with a cut of 50 GeV resulting four jets. Fewer degrees of freedom in the four jet case means a much larger learned value for the σ variable.

jets according to another algorithm. The choice of the number of jets used in mGMM jet clustering therefore controls the expressive power of the algorithm to look at the event structure. In practice, choosing too many jets does not greatly affect the value of the leading learned σ variable, because the additional jets learn finer features of the event structure. On the other hand, choosing too few jets is often problematic as can be seen in figure 15 — the fuzzy jets need to grow in order to cover the full energy distribution in the event. Using anti- k_t jets as seeds for fuzzy jets has the feature that the number of fuzzy jets changes dynamically with the complexity of the event. The algorithm is not very sensitive to the exact locations of the anti- k_t jets — studies which randomly perturbed the initial jet locations inside a disc of radius 1.0 found that σ was robust to such fluctuations, even on an event by event basis. However, the p_T threshold for the seed anti- k_t jets can have a significant impact on the fuzzy jets as this alters the number of seeds. The p_T threshold for the anti- k_t seeds is typically lower than the p_T threshold one would use to consider anti- k_t jets alone because the fuzzy jets algorithm needs enough seeds to populate the low energy regions of the detector. One way of mitigating the impact of the p_T cut on the fuzzy jet clustering is to introduce an *event jet*, described in section 5.2.

D A leading order description of fuzzy jet σ

We have seen in section 4 that the fuzzy jet σ is correlated with $\rho = m/p_T$. We can build some intuition for this relationship by considering a leading order QCD calculation of σ . Consider an isolated quark jet with energy E which radiates a gluon with angle $\theta \ll 1$ from the jet axis and with energy fraction $z \ll 1$. Without loss of generality, suppose the

quark is moving in the $\phi = 0$ direction and the splitting happens in the $\phi = \pi/2$ direction so that the four vector of the quark is $q^\mu = E(1 - z)(1, 0, 0, 1)$, and the gluon four-vector is $g^\mu = Ez(1, \theta, 0, 1)$, to leading order. To this order, the jet mass is simply $m = Ez\theta^2$. What is σ ? Consider $k = 1$ and something like the event-jet applied so that we can treat this jet in isolation from other hadronic activity in the event. Since $k = 1$, the soft memberships are all one, i.e., $q_{i1} = 1$ and there is only one step of the EM algorithm. The anti- k_t jet has (y, ϕ) coordinates $(0, \theta)$, which could be used for the seed, but since $k = 1$, the seed is not used. The quark has coordinates $(0, 0)$, and the gluon has coordinates $(0, \theta)$. We can compute the fuzzy jet coordinates in the (single) M step:

$$\mu_y = 0 \tag{D.1}$$

$$\mu_\phi = \frac{0 \times E(1 - z) + \theta \times Ez}{E(1 - z) + Ez} = z\theta \tag{D.2}$$

$$\sigma^2 = \frac{(0 - z\theta)^2 \times E(1 - z) + (\theta - z\theta)^2 \times Ez}{2(E(1 - z) + Ez)} \tag{D.3}$$

$$= z\theta^2 + \mathcal{O}(\theta^2 z^2). \tag{D.4}$$

Therefore, to leading order and $k = 1$, the learned σ is the jet mass. For $k = 2$, there are enough degrees of freedom to resolve the substructure of the hard splitting and so the relationship between the jet mass and σ breaks down.

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, *Charged-particle multiplicities in pp interactions measured with the ATLAS detector at the LHC*, *New J. Phys.* **13** (2011) 053033 [[arXiv:1012.5104](https://arxiv.org/abs/1012.5104)] [[INSPIRE](#)].
- [2] CMS collaboration, *Charged particle multiplicities in pp interactions at $\sqrt{s} = 0.9, 2.36$ and 7 TeV*, *JHEP* **01** (2011) 079 [[arXiv:1011.5531](https://arxiv.org/abs/1011.5531)] [[INSPIRE](#)].
- [3] CMS collaboration, *Performance of electron reconstruction and selection with the CMS detector in proton-proton collisions at $\sqrt{s} = 8$ TeV*, *2015 JINST* **10** P06005 [[arXiv:1502.02701](https://arxiv.org/abs/1502.02701)] [[INSPIRE](#)].
- [4] ATLAS collaboration, *Electron and photon energy calibration with the ATLAS detector using LHC Run 1 data*, *Eur. Phys. J.* **C 74** (2014) 3071 [[arXiv:1407.5063](https://arxiv.org/abs/1407.5063)] [[INSPIRE](#)].
- [5] CMS collaboration, *Performance of CMS muon reconstruction in pp collision events at $\sqrt{s} = 7$ TeV*, *2012 JINST* **7** P10002 [[arXiv:1206.4071](https://arxiv.org/abs/1206.4071)] [[INSPIRE](#)].
- [6] ATLAS collaboration, *Measurement of the muon reconstruction performance of the ATLAS detector using 2011 and 2012 LHC proton-proton collision data*, *Eur. Phys. J.* **C 74** (2014) 3130 [[arXiv:1407.3935](https://arxiv.org/abs/1407.3935)] [[INSPIRE](#)].
- [7] S.D. Ellis and D.E. Soper, *Successive combination jet algorithm for hadron collisions*, *Phys. Rev.* **D 48** (1993) 3160 [[hep-ph/9305266](https://arxiv.org/abs/hep-ph/9305266)] [[INSPIRE](#)].

- [8] Y.L. Dokshitzer, G.D. Leder, S. Moretti and B.R. Webber, *Better jet clustering algorithms*, *JHEP* **08** (1997) 001 [[hep-ph/9707323](#)] [[INSPIRE](#)].
- [9] M. Wobisch and T. Wengler, *Hadronization corrections to jet cross-sections in deep inelastic scattering*, [hep-ph/9907280](#) [[INSPIRE](#)].
- [10] M. Cacciari, G.P. Salam and G. Soyez, *The anti- k_t jet clustering algorithm*, *JHEP* **04** (2008) 063 [[arXiv:0802.1189](#)] [[INSPIRE](#)].
- [11] G.J. McLachlan and D. Peel, *Finite mixture models*, Wiley series in probability and statistics. J. Wiley & Sons, New York U.S.A. (2000).
- [12] G. Milligan and M. Cooper, *An examination of procedures for determining the number of clusters in a data set*, *Psychometrika* **50** (1985) 159.
- [13] R. Tibshirani, G. Walther and T. Hastie, *Estimating the number of clusters in a data set via the gap statistic*, *J. R. Stat. Soc. B* **63** (2001) 411.
- [14] S. Catani, Y.L. Dokshitzer, M.H. Seymour and B.R. Webber, *Longitudinally invariant K_t clustering algorithms for hadron hadron collisions*, *Nucl. Phys. B* **406** (1993) 187 [[INSPIRE](#)].
- [15] I.W. Stewart, F.J. Tackmann, J. Thaler, C.K. Vermilion and T.F. Wilkason, *$XCone$: N -jettiness as an exclusive cone jet algorithm*, *JHEP* **11** (2015) 072 [[arXiv:1508.01516](#)] [[INSPIRE](#)].
- [16] J. Thaler and T.F. Wilkason, *Resolving boosted jets with $XCone$* , *JHEP* **12** (2015) 051 [[arXiv:1508.01518](#)] [[INSPIRE](#)].
- [17] D.Yu. Grigoriev, E. Jankowski and F.V. Tkachov, *Optimal jet finder*, *Comput. Phys. Commun.* **155** (2003) 42 [[hep-ph/0301226](#)] [[INSPIRE](#)].
- [18] S.D. Ellis, A. Hornig, T.S. Roy, D. Krohn and M.D. Schwartz, *$Qjets$: a non-deterministic approach to tree-based jet substructure*, *Phys. Rev. Lett.* **108** (2012) 182003 [[arXiv:1201.1914](#)] [[INSPIRE](#)].
- [19] M. Cacciari and G.P. Salam, *Pileup subtraction using jet areas*, *Phys. Lett. B* **659** (2008) 119 [[arXiv:0707.1378](#)] [[INSPIRE](#)].
- [20] J.E. Huth et al., *Snowmass 2001: Jet energy flow project*, [hep-ph/0202207](#).
- [21] S.D. Ellis, J. Huston, K. Hatakeyama, P. Loch and M. Tonnesmann, *Jets in hadron-hadron collisions*, *Prog. Part. Nucl. Phys.* **60** (2008) 484 [[arXiv:0712.2447](#)] [[INSPIRE](#)].
- [22] H. Hartley, *Maximum likelihood estimation from incomplete data*, *Biometrics* **14** (1958) 174.
- [23] A. Dempster, N. Laird and D. Rubin, *Maximum likelihood from incomplete data via the EM algorithm*, *J. Roy. Soc. Ser. B* **39** (1977) 1.
- [24] G. McLachlan and T. Krishnan, *The EM algorithm and extensions*, Wiley, New York U.S.A. (1997).
- [25] J. MacQueen, *Some methods for classification and analysis of multivariate observations*, in the proceedings of the *Berkeley symposium on mathematical statistics and probability*, University of California Press, Berkeley, U.S.A. (1967).
- [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] T. Sjöstrand, S. Mrenna and P.Z. Skands, *PYTHIA 6.4 physics and manual*, *JHEP* **05** (2006) 026 [[hep-ph/0603175](#)] [[INSPIRE](#)].

- [28] M. Cacciari, G.P. Salam and G. Soyez, *FastJet user manual*, *Eur. Phys. J. C* **72** (2012) 1896 [[arXiv:1111.6097](#)] [[INSPIRE](#)].
- [29] D. Krohn, J. Thaler and L.-T. Wang, *Jet trimming*, *JHEP* **02** (2010) 084 [[arXiv:0912.1342](#)] [[INSPIRE](#)].
- [30] PARTICLE DATA GROUP collaboration, K.A. Olive et al., *Review of particle physics*, *Chin. Phys. C* **38** (2014) 090001 [[INSPIRE](#)].
- [31] J. Thaler and K. Van Tilburg, *Identifying boosted objects with N -subjettiness*, *JHEP* **03** (2011) 015 [[arXiv:1011.2268](#)] [[INSPIRE](#)].
- [32] CMS collaboration, *Identifying hadronically decaying vector bosons merged into a single jet*, *CMS-PAS-JME-13-006* (2013).
- [33] ATLAS collaboration, *Performance of boosted W boson Identification with the ATLAS detector*, *ATL-PHYS-PUB-2014-004* (2014).
- [34] S. Jörg, A. Höcker, P. Speckmayer and H. Voss, *Current developments in TMVA — An outlook to TMVA4*, *PoS(ACAT08)063*.
- [35] A. Hocker et al., *TMVA — Toolkit for Multivariate Data Analysis*, *PoS(ACAT)040* [[physics/0703039](#)] [[INSPIRE](#)].
- [36] D. Bertolini, P. Harris, M. Low and N. Tran, *Pileup per particle identification*, *JHEP* **10** (2014) 059 [[arXiv:1407.6013](#)] [[INSPIRE](#)].
- [37] 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](#)].
- [38] M. Cacciari, G.P. Salam and G. Soyez, *SoftKiller, a particle-level pileup removal method*, *Eur. Phys. J. C* **75** (2015) 59 [[arXiv:1407.0408](#)] [[INSPIRE](#)].