

Jet substructure templates: data-driven QCD backgrounds for fat jet searches

Timothy Cohen,^a Martin Jankowiak,^b Mariangela Lisanti,^c Hou Keong Lou^c
and Jay G. Wacker^a

^aTheory Group, SLAC National Accelerator Laboratory,
Menlo Park, CA 94025, U.S.A.

^bInstitut für Theoretische Physik, Universität Heidelberg,
69120 Heidelberg, Germany

^cPhysics Department, Princeton University,
Princeton, NJ 08544, U.S.A.

E-mail: timcohen@slac.stanford.edu,
m.jankowiak@thphys.uni-heidelberg.de, mlisanti@princeton.edu,
hlou@princeton.edu, jgwacker@stanford.edu

ABSTRACT: QCD is often the dominant background to new physics searches for which jet substructure provides a useful handle. Due to the challenges associated with modeling this background, data-driven approaches are necessary. This paper presents a novel method for determining QCD predictions using templates — probability distribution functions for jet substructure properties as a function of kinematic inputs. Templates can be extracted from a control region and then used to compute background distributions in the signal region. Using Monte Carlo, we illustrate the procedure with two case studies and show that the template approach effectively models the relevant QCD background. This work strongly motivates the application of these techniques to LHC data.

KEYWORDS: Jets, Hadronic Colliders

ARXIV EPRINT: [1402.0516](https://arxiv.org/abs/1402.0516)

Contents

1	Introduction	1
2	Substructure templates	3
2.1	Separating kinematics and substructure	3
2.2	Using jet factorization	4
3	Implementing templates	7
3.1	Constructing templates	7
3.2	Dressing an event	9
3.3	Calculating errors	10
3.4	Measurement uncertainties	12
4	Applications	13
4.1	High multiplicity signals	13
4.2	Boosted three-jet resonances	16
5	Conclusions	17
A	A kernel smoothing primer	20
A.1	Deriving the AMISE bandwidth	21
A.2	Practical kernel smoothing	24
A.3	Validation	27
B	QCD Monte Carlo	30

1 Introduction

Analyzing the substructure of jets has proven useful for a multitude of new physics scenarios. A variety of substructure observables and techniques have been proposed [1–3], with applications ranging from boosted-object tagging [4–7] to high-multiplicity searches that take advantage of accidental substructure [8–10]. Simultaneously, tremendous progress has been made in modeling QCD at the analytical level and with Monte Carlo (MC) event generators [11–35]. These developments have been crucial in supporting efforts to perform Standard Model measurements, as well as searches for new physics, at the LHC. However, given the complexity of the LHC environment, data-driven background estimates for QCD remain a central component of any jet substructure study; new methods for determining these backgrounds are thus of great interest.

There are three complementary reasons why jet substructure is useful for new physics searches at the LHC. The first is that large-radius “fat” jets [36] naturally accommodate

grooming methods like trimming [37], pruning [38], and mass-drop filtering [4]. This is a significant advantage in the hadronic environment of the LHC, where the presence of the underlying event and pile-up makes mass reconstruction difficult. Second, teasing out the signatures of highly boosted states that undergo hadronic decays is difficult to impossible without jet substructure techniques. Maximizing sensitivity to very boosted final states will become increasingly important as rising trigger thresholds make low- p_T regions of phase space inaccessible. Finally, the fat jet paradigm provides an efficient way to separate signal from background in multijet events. For example, the total jet mass of an event (the sum of the individual fat jet masses) is a particularly useful discriminator [8] due to the simple fact that QCD does not tend to yield large-mass jets.

Given the many successes of the fat jet approach and the practical importance of improving methods for obtaining data-driven estimates for Standard Model rates, it is natural to focus on the specific case of the multijet background to fat jet searches. In [10, 39], a novel approach to computing multijet backgrounds was proposed that relies on the assumption that the properties of individual fat jets are independent of one another and universal — this is motivated by the approximate factorization of QCD jets [40]. Under this assumption, jet properties can be measured in signal-poor samples (e.g., inclusive dijets) and then extrapolated to the signal (e.g., multijet) region. This paper provides a concrete procedure for making such an extrapolation.

Typically, data-driven techniques require identifying a set of approximately uncorrelated variables to define control regions. Any non-negligible correlations require correction factors that are estimated with MC. Consequently, if one wants to perform cuts on non-trivially correlated variables, such as the mass and other substructure properties of a jet, the background estimate is no longer data-driven. In contrast, the template methods introduced in this paper naturally incorporate these correlations. As a result, searches can be performed with more handles than previously possible and with less dependence on MC.

Any realistic data-driven proposal must assess the statistical and systematic uncertainties associated with the background estimate. By taking advantage of kernel smoothing techniques, we develop a framework in which the statistical uncertainties associated with the template method can be reliably addressed. We demonstrate the feasibility of our approach by applying it to two example mock analyses.

The rest of the paper is organized as follows. Section 2 provides a general overview of the method. Section 3 continues the discussion with a more detailed and technical description. Section 4 presents two MC case studies to demonstrate how the method might be deployed in a realistic analysis. Finally, section 5 presents our conclusions. An appendix is included that gathers together many of the technical aspects of this study. A review of kernel smoothing is included in appendix A. The technical details of our MC generation and processing are discussed in appendix B. Code for implementing the approach presented in this paper is available at <http://sourceforge.net/projects/kernel smoother/>.

To keep track of the terminology associated with our proposal, new terms are emphasized with **bold red** typeface where they are first introduced.

2 Substructure templates

The approach advocated here is different from typical data-driven methods. In broad strokes, the usual strategy is to

1. define a **control region** that is expected to be signal-poor;
2. use this region to validate/normalize MC calculations;
3. use the validated MC to predict the backgrounds in the **signal region** for which a given search has been optimized.

The foundation of our approach, in contrast, involves constructing a data-driven **template** that models the substructure of fat jets. This template is then used to **dress** a sample of multi-fat jet events.¹ The selection criteria are then applied to the dressed sample to yield a background estimate in the signal region.

2.1 Separating kinematics and substructure

The templates $\hat{\rho}$ are transfer functions derived from data that encode the statistical distribution of output variables as a function of the input variable(s). In the limit of infinite statistics, $\hat{\rho}$ converges to the true distribution ρ . The outputs are jet substructure observables, such as the jet mass, while the inputs are kinematic variables. Assuming that the input variable is the jet’s transverse momentum p_T , then

$$\hat{P}(m | p_T) = \hat{\rho}(m | p_T) dm \tag{2.1}$$

is the data-driven estimate for the probability \hat{P} that a jet with the given p_T has a mass between m and $m + dm$. Here, “|” distinguishes input from output variables on the right and left sides, respectively, and hatted variables refer to estimated quantities.

Templates can be used to compute data-driven backgrounds through the following procedure:

1. determine a control region that is expected to be signal-poor;
2. map out templates using data in the control region;
3. generate a MC sample in the signal region, disregarding everything except for the distribution of the input variable(s);
4. convolve the templates derived from the control sample with the MC sample;
5. apply cuts to this data/MC hybrid dataset to obtain a data-driven background estimate.

¹Note that other examples of “dressing” have been implemented in the past. For example, in the context of charged Higgs boson searches, ATLAS makes a data-driven background estimate by taking a control sample with muons in the final state and replacing the muons with simulated taus, a technique that ATLAS refers to as “embedding” [41]. This is essentially the inversion of our recipe. CMS has also used a version of dressing in searches for boosted $t\bar{t}$ resonances. In particular, they rely on data in a control region to determine the properties of their top tagger, which is then propagated to the signal region [42, 43].

This article develops a concrete formalism for determining QCD backgrounds in this way, with particular care being paid to the statistical properties of the results. In more detail, the procedure is formulated as follows. First, a suitable control region is defined, which is referred to as the **training sample**. The fat jets in this sample are used to construct a histogram for the substructure observable(s) of interest as a function of kinematic observable(s). In order to obtain a continuous template from a limited number of fat jets, the histogram must be smoothed. This is done using kernel smoothing techniques, which are described in detail in section 3 and appendix A. The statistical error associated with the kernel smoothing procedure is folded into the final error estimate for the number of background events surviving a given set of cuts.

The next step requires the use of MC integration and is the analog of the extrapolation performed in other data-driven approaches. To begin, a **kinematic sample** is generated with event generator MC. The kinematic sample consists of events with N_j fat jets. This sample is then used to extract the kinematic distribution of the fat jet background — e.g.,

$$\frac{d^{N_j} \sigma(p_{T_i})}{dp_{T_1} \dots dp_{T_{N_j}}}. \quad (2.2)$$

It is worth emphasizing that any jet substructure information in the kinematic sample goes unused. Instead, each fat jet in the kinematic sample is dressed with substructure information using the template determined from the previous step; we refer to this as Monte Carlo integration, which results in the **dressed sample**. Lastly, cuts are performed on the dressed data set to obtain a background estimate and an associated smoothing variance. Note that this approach incorporates the correlations between the kinematics of the various fat jets into the final result. As shown in section 4 using two explicit mock analyses, this allows non-trivial correlations to propagate to the dressed sample and reproduces the predictions from MC, within statistical uncertainties.

To summarize, the proposed data-driven strategy is (this parallels the enumeration above):

1. determine a control region to obtain a training sample;
2. train a template $\hat{\rho}$;
3. generate a kinematic sample using MC;
4. perform the integration thereby convolving $\hat{\rho}$ with the kinematic sample;
5. apply cuts to the resulting set of dressed events to obtain a data-driven background estimate.

Figure 1 provides a pictorial representation of the procedure. The computation of smoothing variance is more involved and we postpone a detailed discussion to section 3.

2.2 Using jet factorization

Let us now take a moment to discuss the assumption of jet template independence in more detail. This assumption leads to a dramatic reduction in the dimension of the multijet




Figure 1. A pictorial representation of the procedure.

phase space and allows for data-driven background estimates to be made with minimal statistical error.

The jet properties we study here can be separated into two categories. Those in the first category, denoted by the d -dimensional vector \vec{x} , are the outputs of templates. Examples include jet mass, N -subjettiness, angularity, and other jet substructure observables. The physics of these observables is largely driven by the parton shower and, as such, is set by the kinematics of the hard parton. Physically, the corresponding templates can be mapped out by taking a parton-level event and showering it many times. The d observables in \vec{x} need not be independent of each other, and in most cases they are not. For instance, if a jet’s mass is anomalously large, the same jet will have non-trivial N -subjettiness moments.

The second category consists of kinematic variables, denoted by \vec{k} . In this paper, we only consider the jet p_T , although in principle the set \vec{k} could be augmented with other variables, such as the pseudo-rapidity of the jet or pileup conditions at the time of the event. These kinematic properties cannot be modeled as template outputs, since the kinematics of jets in an event are strongly correlated with one another.

An N_j -jet differential cross section is given as a function of the fat jet kinematics \vec{k}_i and substructure observables \vec{x}_i for $i = 1$ to N_j . It can be written as²

$$\frac{d^{2N_j}\sigma(\vec{x}_i, \vec{k}_i)}{d\vec{x}_1 \dots d\vec{x}_{N_j} d\vec{k}_1 \dots d\vec{k}_{N_j}} = \frac{d^{N_j}\sigma(\vec{k}_i)}{d\vec{k}_1 \dots d\vec{k}_{N_j}} \rho(\vec{x}_1, \dots, \vec{x}_{N_j} | \vec{k}_1, \dots, \vec{k}_{N_j}). \quad (2.3)$$

For a given \vec{k}_i , we assume that the substructure observable \vec{x}_i of the i^{th} fat jet is independent of all the other \vec{x}_j and \vec{k}_j , with $j \neq i$. Then the right-most term in eq. (2.3) splits into a

²Note that $\rho(\vec{x}_i | \vec{k}_i)$ is a legitimate (conditional) probability distribution, i.e. for each fixed value of the \vec{k}_i the distribution $\rho(\vec{x}_i | \vec{k}_i)$ yields unity when integrated over the \vec{x}_i .

product of N_j terms:

$$\frac{d^{N_j} \sigma(\vec{k}_i)}{d\vec{k}_1 \dots d\vec{k}_{N_j}} \rho(\vec{x}_1, \dots, \vec{x}_{N_j} | \vec{k}_1, \dots, \vec{k}_{N_j}) = \frac{d^{N_j} \sigma(\vec{k}_i)}{d\vec{k}_1 \dots d\vec{k}_{N_j}} \prod_{i=1}^{N_j} \rho_i(\vec{x}_i | \vec{k}_i). \quad (2.4)$$

This factorized form of the differential cross section allows for background estimates to be computed by a convolution of separate kinematic and substructure parts and forms the basis of our data-driven proposal. Note that the independence assumption underlying eq. (2.4) dramatically reduces the dimensionality of the configuration space that needs to be explored. Consider, for example, the case where $d = 1$ and $\mathbf{k} = p_T$. Then, instead of using MC to estimate the QCD prediction for a $2N_j$ -dimensional configuration space, only a N_j -dimensional configuration space needs to be explored. If multiple jet properties are being studied, $d > 1$, then the reduction in dimensionality is even more dramatic, with a $(d+1) \times N_j$ dimensional configuration space reduced to N_j dimensions. The additional cost involved is creating the $(d+1)$ -dimensional templates $\rho(\vec{x}|p_T)$. Of course, as d increases (keeping the size of the training sample fixed), the density estimation that underlies the construction of substructure templates becomes increasingly difficult and at some point the associated uncertainties will outweigh any gains from the reduced dimensionality. This is the price to pay for not including any *a priori* model for correlations between the various substructure observables; however, these minimal assumptions allow us to remain data-driven.

Soft QCD effects can violate the factorization in eq. (2.4). These include underlying event, pile-up, and $\mathcal{O}(N_c^{-2})$ corrections coming from perturbative QCD due to large-angle soft radiation. Jet grooming can minimize these effects. Throughout this paper, the fat jets are trimmed [37], as described in detail in appendix B. In general, it is important to consider the conditions under which eq. (2.4) is most likely to hold — for example, as a function of the fat jet radius.

There are many additional subleading effects that spoil the factorization in eq. (2.4). The most important of these is the quark-gluon composition of jets in an event. If the composition of the jets changes dramatically between the training sample and the signal region, then anomalous features may be introduced. In practice, quark- and gluon-initiated jets often have broadly similar features in the signal region and the composition does not change dramatically in many samples of interest. However, the templates taken from the two leading jets do not describe the third and fourth jets well. The introduction of quark and gluon templates may resolve this issue, as will be explored in future work [44].

Note that the three-momenta of the fat jets in the dressed data are automatically conserved because they are taken from MC. However, energy is not conserved since the fat jet masses are determined stochastically — the total energy of each dressed event is different from that of the kinematic event. In practice, we do not find that this is an issue for the cases studied here. The violation of energy conservation changes the center-of-mass energy of the event, but not appreciably because the ratio $m_j/p_T < 1$. These effects lead to correlations between the masses of the jets; how to account for this in the template approach remains an interesting open question.




Figure 2. The $p_T = 400$ GeV jet mass template is shown for three distinct bandwidths. The error band indicates the variance, while the bias can be seen in the mismatch between the solid curve and the dashed curves (variance and bias are defined later in this section). The lower panels give the fractional deviation between the solid and dashed curves.

3 Implementing templates

This section provides a detailed description of the procedure for generating substructure templates and calculating cut efficiencies. It is accompanied by appendix A, which reviews the basics of the statistical procedures used here. If the reader is new to kernel smoothing methods, we recommend that he or she take a moment to review appendix A before proceeding.

3.1 Constructing templates

The jets in the training sample are used to build a histogram. Technically, this is all that is needed to build the templates. However, due to the finite size of the training sample, the histograms might not be adequately populated and data-smoothing techniques are necessary.

The statistical technique of kernel smoothing is one way to approach the problem. In this procedure, every data point in the training sample is smoothed out to the adjoining region in the multi-dimensional space. The aggregate of these smooth contributions gives the final probability distribution function. Take, as an example, the jet mass distribution (solid purple) shown in figure 2. This distribution is derived by looking at the two leading jets with $p_T > 200$ GeV in the full MC sample. Now imagine that we want to reproduce this distribution using a statistically limited sample — say 1% of the total number of events in the full MC sample. Applying kernel smoothing to the statistically limited sample reproduces the full distribution, to within errors.

Take an event with N_j jets, each with particular kinematic and substructure properties described by the vectors \vec{k} and \vec{x} , respectively. For example, if we are interested in each jet's p_T , mass, and N -subjettiness ratio τ_{21} , then the i^{th} jet (ordered by p_T) is described by $\vec{k}_{j_i} = (p_T)_{j_i}$ and $\vec{x}_{j_i} = (m, \tau_{21})_{j_i}$. We unify \vec{k} and \vec{x} into the single D -dimensional vector \vec{z} , where D is the number of kinematic and substructure variables of interest. Then, the i^{th} jet in event j is described by

$$J_{ij}[\vec{z}] \equiv J_{ij}[\vec{x}, \vec{k}]. \quad (3.1)$$

To build a **training sample** for the i^{th} jet, we combine the i^{th} jet from *all* events in the sample. Said another way, the training sample for the i^{th} jet, denoted as T_i , is the set

$$T_i = \left\{ J_{i1}[\vec{z}_1], J_{i2}[\vec{z}_2], \dots, J_{iN_T}[\vec{z}_{N_T}] \right\}, \quad (3.2)$$

where N_T is the number of events with an i^{th} jet. The normalized histogram corresponding to T_i encodes a discrete probability distribution function with D dimensions. This probability distribution function may or may not be smooth, depending on the total number of events N_T . This is where kernel smoothing comes into play — it allows us to define a smooth probability distribution function in a statistically robust manner. In particular, we can define the smoothed template for the i^{th} jet, $\hat{\rho}_i(\vec{z})$, using

$$\hat{\rho}_i(\vec{z}) = \frac{1}{N_T} \sum_{J \in T_i} K_h(\vec{z} - \vec{z}(J)) \quad \text{and} \quad K_h(\vec{z}) = \frac{1}{\det h} K(h^{-1}\vec{z}), \quad (3.3)$$

where $\vec{z}(J)$ is the value of \vec{z} for jet J , K is the **kernel**, and h is a matrix describing the kernel width. The choice of kernel is not too critical and we simply adopt the canonical choice of a Gaussian:

$$K(\vec{z}) = \frac{1}{(4\pi)^{\frac{D}{2}}} \exp\left(-\frac{|\vec{z}|^2}{2}\right) \quad \text{and} \quad K_h(\vec{z}) = \frac{1}{(4\pi)^{\frac{D}{2}} \det h} \exp\left[-(2h^T h)^{-1} z^i z^j\right]. \quad (3.4)$$

The template generated by the kernel smoothing procedure, eq. (3.3), is simply a sum of Gaussians around each data point in the training sample. The degree to which this procedure appropriately smooths out the distribution depends on the matrix of widths, h_{ij} . Thus, the goal is to choose the smoothing parameters such that the estimated distribution function, $\hat{\rho}(\vec{z})$, is as close as possible to the true distribution function, $\rho(\vec{z})$. Selecting the width matrix h_{ij} of the kernel is one of the most important aspects of this procedure because an over/under-smoothed distribution may not accurately represent the true distribution function. A common practice is to use the ‘‘Asymptotic Mean Integrated Squared Error’’ (AMISE) metric described in appendix A. This gives rise to ‘‘Silverman’s Rule-of-Thumb’’ [45], where h_{ij} is

$$h_{ij}^{\text{AMISE}} \simeq c \hat{\sigma}_{ij} N_T^{-\frac{1}{D+4}}. \quad (3.5)$$

Here, c is an $\mathcal{O}(1)$ constant and $\hat{\sigma}_{ij}$ is an estimator for the square root of the covariance matrix of ρ . For distributions that are not vastly different from the normal distribution, this provides a useful starting point for determining the kernel width.

Figure 2 illustrates the effects of choosing different values of c . All three panels show the result of applying kernel smoothing to the full MC sample; this serves as an approximation to the “true” distribution.³ Overlaid are templates that have been generated using only 1% of the full sample in order to demonstrate how kernel smoothing approximates the true distribution for different values of the kernel width.⁴ The width is varied from the optimal value $c = c_{\text{AMISE}}$ [middle, green] to an over-smoothed value $c = 1.5 c_{\text{AMISE}}$ [left, red], and an under-smoothed value $c = 0.7 c_{\text{AMISE}}$ [right, orange].

The smoothing procedure introduces two competing sources of error in the template, **bias** and **variance**, whose magnitudes are set by the amount of smoothing. These are defined as follows:

$$\text{bias} \equiv b(\vec{z}) = \rho(\vec{z}) - \hat{\rho}(\vec{z}); \quad (3.6)$$

$$\text{variance} \equiv v^2(\vec{z}) = \langle \hat{\rho}(\vec{z})^2 \rangle - \langle \hat{\rho}(\vec{z}) \rangle^2, \quad (3.7)$$

where $\langle \dots \rangle$ is the expectation value with respect to the true distribution ρ . The variance is interpreted as the scatter in $\hat{\rho}(\vec{z})$ when estimated from different samples drawn from $\rho(\vec{z})$, while the bias is a measure of the systematic displacement of the estimated distribution from the true distribution as a consequence of smoothing.

As discussed in more detail in appendix A, the presence of bias is problematic and must be corrected for. Throughout, we use a **corrected template** $\hat{\rho}^*$, which has substantially reduced bias. To compute $\hat{\rho}^*(\vec{z})$, all that is needed is an estimate for $b(\vec{z})$. We adopt an estimator $\hat{b}(\vec{z})$ that is formed by comparing the twice-smoothed template

$$\hat{\hat{\rho}}(\vec{z}) = \int d^D z' \hat{\rho}(\vec{z}') K_h(\vec{z} - \vec{z}') \quad (3.8)$$

with the once-smoothed template $\hat{\rho}(\vec{z})$:

$$\hat{b}(\vec{z}) = \hat{\hat{\rho}}(\vec{z}) - \hat{\rho}(\vec{z}). \quad (3.9)$$

This estimator is valid to leading order in the kernel width [47], as shown in eq. (A.24). The corrected template is then defined as

$$\hat{\rho}^*(\vec{z}) = \hat{\rho}(\vec{z}) - \hat{b}(\vec{z}). \quad (3.10)$$

It is this corrected template that will be used to make background estimates in the signal region. This procedure is justified in appendix A, where the corrected template is shown to give consistent results for a test probability distribution.

3.2 Dressing an event

The previous section showed how to construct a template from a training sample. With this template in hand, one can calculate the efficiency for a kinematic event in the signal

³The non-trivial shape of these distributions results from jet trimming, see [46].

⁴It is important to mention one subtlety in this figure. The solid curve is the corrected jet mass template $\hat{\rho}^*$ (see eq. (3.10)). The dashed curves are the uncorrected templates $\hat{\rho}$ (see eq. (3.3)) and demonstrate the bias inherent in the smoothing procedure.

region to pass a given set of cuts. We will now describe this calculation in detail. Note that throughout the rest of this paper, unless stated otherwise, our substructure templates are universal. That is, we do not make any distinctions between, say, the 1st and 2nd jets (as ordered by p_T) in the training sample. Rather, a single template is constructed making use of all the jets in the training sample.

The corrected template $\hat{\rho}^*(\vec{x})$ constructed in the previous section does not distinguish between the kinematic variables \vec{k} and substructure variables \vec{x} . However, we wish to view the substructure variables as functionally dependent upon the kinematic input variables. That is, we want to convert the joint probability distribution $\hat{\rho}^*(\vec{x}, \vec{k})$ into the conditional probability distribution $\hat{\rho}^*(\vec{x}|\vec{k})$ using

$$\hat{\rho}^*(\vec{x}|\vec{k}) = \frac{\hat{\rho}^*(\vec{x}, \vec{k})}{\hat{\rho}^*(\vec{k})} = \frac{\hat{\rho}^*(\vec{x}, \vec{k})}{\int d^d x' \hat{\rho}^*(\vec{x}', \vec{k})} \quad \text{such that} \quad \int d^d x \hat{\rho}^*(\vec{x}|\vec{k}) = 1 \quad (3.11)$$

for each \vec{k} . The template $\hat{\rho}^*(\vec{x}|\vec{k})$ is normalized so that it acts as a probability distribution function for the substructure information of a jet with given kinematics \vec{k} .

Given a multi-fat jet event in the signal region, obtained from MC, we can apply cuts to the first and second jets, with specified kinematics \vec{k}_1 and \vec{k}_2 , respectively. The efficiency, $\hat{\epsilon}^*$, that this event passes the cuts is

$$\hat{\epsilon}^*[C] = \int_C d^d x_1 d^d x_2 \hat{\rho}^*(\vec{x}_1|\vec{k}_1) \hat{\rho}^*(\vec{x}_2|\vec{k}_2), \quad (3.12)$$

where $C = C[\vec{x}_1, \vec{x}_2]$ is the region of the (\vec{x}_1, \vec{x}_2) configuration space selected by the cuts. For instance, if \vec{x} is the mass of a jet, then

$$C[\vec{x}_1, \vec{x}_2] = \{m_1 + m_2 > M_J\} \quad (3.13)$$

is the cut requiring that the sum of the two individual jet masses is larger than M_J . While the integrand of eq. (3.12) is a product of templates, the region of integration may be a complicated function of multiple variables.

3.3 Calculating errors

This procedure is trivially generalized to a *set* of kinematic events, thus yielding a background prediction with error bars. The most straightforward use of the template is to start with a sample of jets, of size N_e , where only the kinematics are specified. A dressed MC sample may be created by dressing each event with the template multiple times. Said another way, the integral in eq. (3.12) is performed by Monte Carlo integration; for each kinematic event, the template is used to produce n_I dressed events. It is sufficient to take $n_I \sim \mathcal{O}(10^3)$. Schematically, to dress events with two kinematic objects, one converts each MC event into n_I dressed events

$$(J[\vec{k}_1], J[\vec{k}_2])_i \quad \longrightarrow \quad \left\{ \left((J[\vec{k}_1, \vec{x}_{1\alpha}], J[\vec{k}_2, \vec{x}_{2\alpha}])_i, w_\alpha^*, w_\alpha \right); \alpha = 1, \dots, n_I \right\}, \quad (3.14)$$

where each dressed event, labeled by α , has fully specified jet substructure, i.e., specific values of $\vec{x}_{1\alpha}$ and $\vec{x}_{2\alpha}$ that have been drawn uniformly from their domain. Since the $\vec{x}_{i\alpha}$ are

drawn uniformly (and not from the template itself), each dressed event is also associated with a corrected weight w_α^* and an uncorrected weight w_α :

$$w_\alpha^* = \hat{\rho}^*(\vec{x}_{1\alpha}|\vec{k}_1) \hat{\rho}^*(\vec{x}_{2\alpha}|\vec{k}_2) \quad \text{and} \quad w_\alpha = \hat{\rho}(\vec{x}_{1\alpha}|\vec{k}_1) \hat{\rho}(\vec{x}_{2\alpha}|\vec{k}_2). \quad (3.15)$$

The normalizations of w^* and w are irrelevant, because they cancel when calculating efficiencies (see eq. (3.19)). The corrected weight is used for the background prediction, while the uncorrected weight is used to estimate the total bias for the predicted number of events.

To determine the final number of events that survive the cuts for a given search, one first computes the sum of corrected weights for each kinematic event before and after cuts:

$$W_i^* = \sum_{\alpha=1}^{n_I} w_{i\alpha}^* \quad \text{and} \quad W_i^*[C] = \sum_{\alpha=1}^{n_I} w_{i\alpha}^* \Theta(c_{i\alpha}), \quad (3.16)$$

where

$$c_{i\alpha} = \begin{cases} 1 & \text{if } (\vec{x}_1, \vec{x}_2)_{i\alpha} \in C[\vec{x}_1, \vec{x}_2] \\ 0 & \text{otherwise.} \end{cases} \quad (3.17)$$

The estimate for the efficiency that a given event i passes the cuts is

$$\hat{\epsilon}_i^* \equiv \frac{W_i^*[C]}{W_i^*}. \quad (3.18)$$

Note that this is just the discrete version of eq. (3.12). These efficiencies are then summed to obtain the final background prediction. The predicted number of background events after cuts, $\hat{N}_e^*[C]$, is

$$\hat{N}_e^*[C] = \sum_{i=1}^{N_e} \hat{\epsilon}_i^*[C]. \quad (3.19)$$

The statistical uncertainty in the estimate $\hat{N}_e^*[C]$ is challenging to determine because of the implicit dependence of $\hat{N}_e^*[C]$ on the training sample. This leads to important correlations between the cut efficiencies, $\hat{\epsilon}^*[C]$, for different kinematic events. A convenient technique for computing errors makes use of the bootstrap (see e.g., [47–50]). An ensemble of datasets (resampled from the training sample) is used to calculate the corresponding set of templates. This yields an ensemble of estimates for $\hat{N}_e^*[C]$ whose variance is calculated directly. While this procedure is computationally expensive, it results in a robust estimate of the statistical uncertainty, as will be demonstrated in appendix A.

In more detail, for an ensemble of resampled datasets, a corresponding set of templates $\{\hat{\rho}^*_b\}$ is computed, with b ranging from 1 to N_{BS} . These, in turn, are used to compute an ensemble of background predictions $\mathcal{N}_e^*[C] = \{\hat{N}_e^*[C]_1, \dots, \hat{N}_e^*[C]_{N_{\text{BS}}}\}$ using eq. (3.19). The **smoothing variance** is given by:

$$\hat{\sigma}_V^2 = \frac{1}{N_{\text{BS}} - 1} \sum_{b=1}^{N_{\text{BS}}} (\hat{N}_e^*[C]_b - \langle \mathcal{N}_e^*[C] \rangle_{\text{BS}})^2 \quad \text{with} \quad \langle \mathcal{N}_e^*[C] \rangle_{\text{BS}} = \frac{1}{N_{\text{BS}}} \sum_{b=1}^{N_{\text{BS}}} \hat{N}_e^*[C]_b. \quad (3.20)$$

In practice, we find that $N_{\text{BS}} \sim \mathcal{O}(100)$ is sufficient. Eq. (3.20) provides a good estimate of the statistical uncertainty in $\hat{N}_e^*[C]$ provided that the bias is under control (see appendix A

for more details). To make sure this is the case, we estimate the size of the residual bias $\hat{\sigma}_B$ inherent in the estimate $\hat{N}_e^*[C]$ by comparing $\hat{N}_e^*[C]$ to $\hat{N}_e[C]$, where the latter estimate uses the uncorrected weights:

$$\hat{\epsilon}_i[C] = \frac{W_i[C]}{W_i} \quad \text{and} \quad \hat{N}_e[C] = \sum_{i=1}^{N_e} \hat{\epsilon}_i[C], \quad (3.21)$$

with

$$W_i = \sum_{\alpha=1}^{n_I} w_{i\alpha} \quad \text{and} \quad W_i[C] = \sum_{\alpha=1}^{n_I} w_{i\alpha} \Theta(c_{i\alpha}). \quad (3.22)$$

The contribution of each kinematic event, i , to the bias is the difference between the two efficiencies $\hat{\epsilon}_i^*[C]$ and $\hat{\epsilon}_i[C]$. Note that this difference is a signed quantity, and the biases from different kinematic events may add (in)coherently. A measure of the total bias $\hat{\sigma}_B$ is

$$\hat{\sigma}_B = |\hat{N}_e[C] - \hat{N}_e^*[C]| = \left| \sum_{i=1}^{N_e} (\hat{\epsilon}_i[C] - \hat{\epsilon}_i^*[C]) \right|. \quad (3.23)$$

The probability distribution function for the bias is unknown *a priori* and is in general not Gaussian. In order for $\hat{\sigma}_V$ to be a reliable estimate of the statistical uncertainty in $\hat{N}_e^*[C]$, we must have $\hat{\sigma}_B \ll \hat{\sigma}_V$. In practice, we choose a kernel width such that $\hat{\sigma}_B$ is a factor of 2 to 5 less than $\hat{\sigma}_V$. Because bias correction is done at the template-level, $\hat{\sigma}_B$ is an imprecise measure of the residual bias in $\hat{N}_e^*[C]$. Consequently, $\hat{\sigma}_B$ is taken to be unsigned and can be interpreted as a systematic error.

In addition to the statistical fluctuation of the templates, there is an additional systematic error σ_{sys} coming from the mismatch between the template in the control region and the associated template that would be derived from a background-only sample in the signal region. For the examples presented in the next section, we compute $\hat{\sigma}_V$ and $\hat{\sigma}_B$ explicitly. These results demonstrate that σ_{sys} is subdominant to the statistical errors associated with the template procedure, validating the approximation that the templates are independent.

3.4 Measurement uncertainties

The preceding discussion made no mention of measurement uncertainties. If one were interested in comparing substructure templates extracted from data to analytical calculations, then the size and nature of any detector effects would be an important consideration in designing an appropriate template procedure. Indeed, the presence of these uncertainties introduces a number of complications (see e.g., [51]). In our case, templates are used to make background estimates that are then directly compared to data in the signal region. Consequently, the substructure variables \vec{x} are implicitly taken to be the measured values; the templates are derived using the measured data and extrapolated to make a prediction for cuts on measured variables. Furthermore, as long as detector-smearing is smooth, the kernel width can be taken to be smaller than the inherent detector resolution (and it will be in the limit of large statistics, see eq. (3.5)). This is not a problem, since what is being estimated is the convolution of the underlying kinematic distributions with the detector response.

The measurement uncertainties for the input variables \vec{k} are relevant if the kinematics are being modeled by MC. In this situation, it is important that the kinematic sample be passed through a reliable detector simulator prior to the integration step so that all of the input and output variables correspond to quantities measured by the detector. On the other hand, if the signal contamination in the region of interest is small (in the absence of any substructure cuts), it might be possible to take the kinematic sample directly from data, in which case these issues can be avoided.

4 Applications

This section is devoted to illustrating the data-driven template approach for two different searches. The smoothing variance $\hat{\sigma}_V$ for 8 TeV LHC data is estimated and shown to be under control. As we will see, the MC and template predictions are in excellent agreement.

The first example is a search for high multiplicity signals (see section 4.1). Fat jets from signal processes should exhibit accidental substructure [8–10] that would be suppressed for QCD backgrounds. We demonstrate that the jet mass and N -subjettiness of the two leading jets of a four-jet sample can be predicted using templates derived from a three-jet sample. For a second case study, the template method is applied to an existing ATLAS search for pair-produced gluinos undergoing an R -parity-violating (RPV) decay to three light quarks each (see section 4.2). The observed correlation between the masses of the two leading jets is captured using the template approach.

Because we do not have access to actual LHC data, we rely on MC events. Weighted events were generated using a variation on the binning procedure outlined in [52]. Specifically, two-, three-, and four-parton QCD events were generated using `Madgraph5` [53]. These events were subsequently showered using `Pythia8` [54]; matching was done using the MLM procedure. The particle-level events were grouped into cells to simulate the finite resolution of the calorimeter. Trimming was applied to these pixels to reduce sensitivity to pileup effects (which we do not model), and N -subjettiness was computed using the `FastJet-contrib` plugin [55, 56]. For a full description of the simulation framework, see appendix B.

4.1 High multiplicity signals

This section demonstrates the application of template techniques to the types of high-multiplicity searches advocated in [8–10]. This case study does not reproduce these searches in exact detail. Instead, we study a simplified version of the event-subjettiness search in [9] in which the substructure of only the two leading jets is considered. Despite the modification, this example is still complex enough to constitute a real test of the template methodology.

As a preselection, each event is required to have at least $N_j = 4$, $R = 1.0$ anti- k_T [57] jets with $p_T > 200$ GeV. Cuts are placed on observables that depend on the fat jets' substructure (but not on their kinematics): the sum of jet masses

$$M_J = \sum_{i=1}^2 m_i \tag{4.1}$$

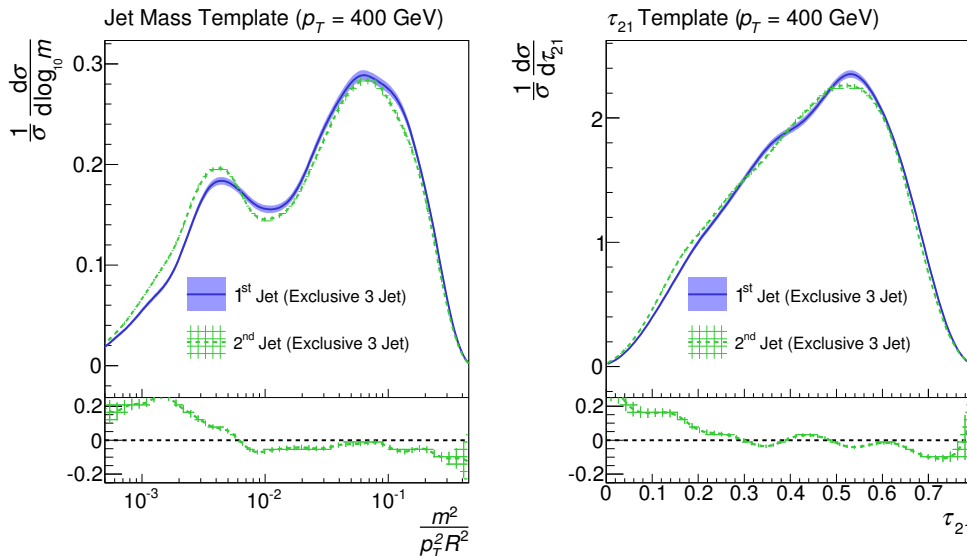


Figure 3. The plot on the left (right) compares the jet mass (τ_{21}) templates between the first and second jets in the exclusive three-jet bin. The lower panels show the deviation between the two templates, normalized to the template of the leading jet. The error bands on the templates are statistical, calculated using the bootstrapped ensemble of templates.

and the ‘event-subjettiness’

$$T_{21} = \left[\prod_{i=1}^2 (\tau_{21})_i \right]^{1/2}, \tag{4.2}$$

where τ_{ij} is the N -subjettiness ratio τ_j/τ_i [55, 56]. Note that there is no impediment to replacing event-subjettiness by one of the subjet counting observables proposed in [10], although the discrete nature of the latter would require modifications to the kernel smoothing procedure.

Requiring four fat jets is already extremely efficient at reducing the Standard Model background [8]. An exclusive three-jet sample is therefore expected to be signal-poor, making it an ideal training sample. The two leading jets in each event are used to fill a 3-dimensional template $\hat{\rho}^*(m, \tau_{21} | p_T)$.

A good set of coordinates for the template is:

$$\hat{\rho}^* = \hat{\rho}^* \left(-\log_{10} \left(\frac{m}{p_T} \right), \tau_{21}, \ln \left(\frac{p_T}{200 \text{ GeV}} \right) \right). \tag{4.3}$$

The log-transformed variables are appropriate given the logarithmic evolution of the strong coupling constant with respect to the fat jet p_T and the form of the collinear singularity that governs the generation of mass in QCD jets.

The bandwidth is chosen to produce a small total bias, $\hat{\sigma}_B \lesssim \hat{\sigma}_V$, while keeping the error $\hat{\sigma}_V$ as small as possible. Figure 3 shows the resulting templates, while figure 4 compares the 3-jet template with the 4-jet prediction for a fixed p_T slice. These figures demonstrate that the substructure variables for the two leading jets in the exclusive 3-jet

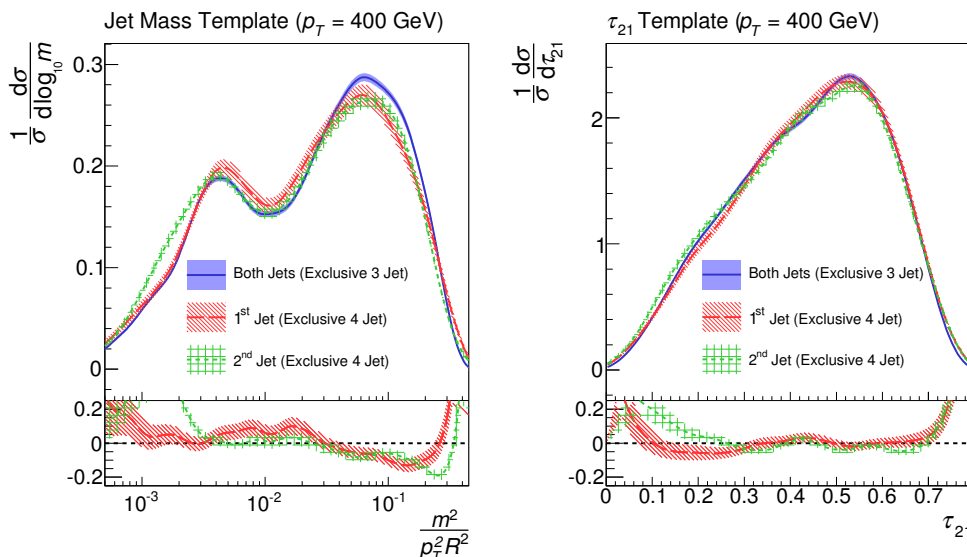


Figure 4. The plot on the left (right) compares the jet mass (τ_{21}) templates between different exclusive jet bins. The exclusive three-jet template (blue) is obtained from the two leading jets. The error bands only include the template errors, computed using the bootstrapped ensemble of templates. The lower panels show the deviation between the templates, normalized to the three-jet template.

and 4-jet samples are equivalent within 10% in the high-mass region and a broad range of τ_{21} . We have checked that this holds for a wide range of p_T choices.

With the template in hand, the next step is to dress a kinematic sample, i.e., the integration step. An inclusive 4-jet cut, with each $R = 1.0$ anti- k_T jet required to have $p_T > 200$ GeV, fully characterizes the kinematic sample. The leading and subleading jets (but not the 3rd or 4th jets) in the kinematic sample are then dressed with the substructure template. The result is a dressed 4-jet sample in which the leading and subleading jets are associated with values of τ_{21} and fat jet mass $m_{1,2}$.

Using the dressed sample, one can compute the fraction of dressed events passing the cuts. We apply a preselection that $m_j > 20$ GeV for each jet, ensuring that the results are IR-safe [35]. Another preselection cut of $p_T > 250$ GeV is also applied to ensure insensitivity to the template boundary. The resulting differential distributions for M_J (with $T_{21} < 0.3$) and T_{21} (with $M_J > 250$ GeV) are shown in figure 5. These distributions are in excellent agreement with the MC. Several additional cuts are summarized in table 1; the template predictions agree with the MC to within template errors.

Note that only the first and second jets in the kinematic sample were dressed in this case study. Empirically, we find that the template derived from the two leading jets in the exclusive three-jet sample accurately models the two leading jets in the four-jet sample. While the template for the third jet is qualitatively different, it does provide a good model of the third jet in the four-jet sample, although this statement suffers from large statistical uncertainties given the size of the MC sample. The difference in the third-jet template is driven in large part by the average quark/gluon content of the third jet. It is likely that by




Figure 5. A comparison between the template estimate and the MC for the differential distributions for M_J (with $T_{21} < 0.3$) and T_{21} (with $M_J > 250$ GeV) for the high multiplicity case study. The MC error bands show the statistical uncertainty for the weighted event sample. The template error bands are given by $\hat{\sigma}_V$ in each bin (note that the errors are correlated).

c	M_J cut [GeV]	T_{21} cut	MC	Template $\pm \hat{\sigma}_V \pm \hat{\sigma}_B$
0.37	500	0.3	20.3 ± 2.2	$19.2 \pm 2.3 \pm 0.6$
0.52	750	0.3	0.86 ± 0.10	$0.96 \pm 0.19 \pm 0.05$
0.37	500	0.6	45.8 ± 3.5	$45.2 \pm 3.7 \pm 1.3$
0.52	750	0.6	1.67 ± 0.14	$1.90 \pm 0.19 \pm 0.13$

Table 1. Expected number of background events for the high multiplicity case study, assuming an integrated luminosity of 1 fb^{-1} . The amount of smoothing used for each background estimate is specified in the first column (in units of c_{AMISE}). The errors for the template estimates are given by $\hat{\sigma}_V$, computed using eq. (3.20). The template and MC predictions agree to within the calculated errors. $\hat{\sigma}_B$ is not normally distributed and therefore cannot be simply combined with $\hat{\sigma}_V$.

incorporating quark and gluon information (extracted from MC) into the determination of the templates, the procedure could be generalized to incorporate the third and even fourth jets into the analysis. Exploring such a generalization is left for future work [44].

4.2 Boosted three-jet resonances

The second case study is based on the ATLAS search for boosted gluinos undergoing RPV decays to three light quarks [58]. The ATLAS analysis proceeds by clustering the event into fat $R = 1.0$ anti- k_T jets and narrow $R = 0.4$ jets. Several search regions are considered, each with a combination of cuts on narrow jet quantities (multiplicity, scalar sum p_T) and fat jet quantities (multiplicity, jet masses). In addition, a substructure requirement of $\tau_{32} < 0.7$ is imposed on both fat jets in order to select on the expected three-pronged structure of gluino decays. The case study presented here mirrors the fat jet analysis,

but ignores the narrow jet selections. Note that we compute N -subjettiness using the “min-axes” algorithm, whereas ATLAS used the “ k_T -axes” algorithm.

To proceed with a background estimate as outlined in the previous two sections, we must first define a training sample from which we can construct a substructure template. As a preselection on the MC sample, each event is required to have at least two fat jets with $p_T > 320$ GeV. Because the signal region consists of events with two high-mass fat jets, the training sample is defined by requiring at least one low mass jet with $m_j < 140$ GeV. This ensures that signal contamination in the training sample is small. That is, for every pair of leading and subleading fat jets (j_1, j_2) the three-dimensional template $\hat{\rho}^*(m, \tau_{32} | p_T)$ is filled with j_2 whenever $m_{j_1} < 140$ GeV. The procedure is then repeated for (j_2, j_1) . The template is formed using the coordinates

$$\hat{\rho}^* = \hat{\rho}^* \left(-\log_{10} \left(\frac{m}{p_T} \right), \tau_{32}, \ln \left(\frac{p_T}{320 \text{ GeV}} \right) \right), \quad (4.4)$$

as in the previous study.

The next step is to define a kinematic sample. In general, the natural way to define the kinematic sample is to use the same cuts as the signal region; apart from substructure cuts, the two samples are defined identically. In the present case, this corresponds to choosing a sample generated by `MadGraph5` at $\sqrt{s} = 8$ TeV, with each event required to have at least two fat jets each with $p_T > 350$ GeV and no requirements on the fat jets’ masses or values of τ_{32} . In the integration step, both of the fat jets in the kinematic sample are dressed with the template $\hat{\rho}^*(m, \tau_{32} | p_T)$. This results in a dressed data set in which each event consists of a pair of fat jets, with given values of p_T , m , and τ_{32} .

To assess the efficacy of the template approach, the dressed data set is compared to the full MC sample. One interesting test is to compare the correlations present in the m_{j_1} vs. m_{j_2} plane in the kinematic sample with the results from the dressed data set. Figure 6 shows that the resulting jet mass distributions are indeed reproduced by the template estimate. This figure also provides a good qualitative match with the relevant plot published by ATLAS [58]. The dominant correlations between the two jet masses are captured by virtue of the p_T dependence of the $\hat{\rho}^*$ template.

As in the previous study, the dressed data set can also be used to estimate yields in the signal region after imposing substructure cuts. The predicted cross sections are listed in table 2 for a selection of cuts. There is excellent agreement between the MC and template predictions. One expects the smoothing variance at the LHC to be reduced when a large dataset is available for training the templates. Note that as the training sample increases, statistical uncertainties decrease more slowly than the familiar $1/\sqrt{N_T}$, see eq. (3.5).

5 Conclusions

Recent years have seen a revolution in fat jet techniques that could allow for the discovery of new physics in signal regions that would otherwise have remained obscured by QCD backgrounds. Realizing the full potential of this statement assumes that we can reliably




Figure 6. The jet mass distributions of the two leading jets obtained directly from MC (left) and estimated using template methods (right). The data-driven technique captures the dominant correlations between the jet masses.

c	m_J cut [GeV]	τ_{32} cut	MC	Template $\pm \hat{\sigma}_V \pm \hat{\sigma}_B$
1.0	400	0.5	220 ± 26	$210 \pm 14 \pm 6$
1.0	600	0.5	3.70 ± 0.32	$4.70 \pm 0.98 \pm 0.05$
1.0	500	0.6	67.0 ± 3.4	$73.0 \pm 10.0 \pm 2.6$
1.0	600	0.6	9.1 ± 0.7	$11.0 \pm 2.0 \pm 0.3$

Table 2. Expected number of background events for the RPV case study, assuming an integrated luminosity of 20 fb^{-1} . The amount of smoothing used for each background estimate is specified in the first column (in units of c_{AMISE}). The errors for the template estimates are given by $\hat{\sigma}_V$, computed using eq. (3.20). The template and MC predictions agree to within the calculated errors. $\hat{\sigma}_B$ is not normally distributed and therefore cannot be simply combined with $\hat{\sigma}_V$.

model the relevant backgrounds. This strongly motivates developing new methods for data-driven background estimates, particularly for the challenging QCD background.

An attractive feature of the fat jet paradigm is that it offers a natural division of phase space into inputs (kinematics) and outputs (substructure variables) as a consequence of jet factorization. This work takes advantage of this division to develop a novel method of estimating QCD backgrounds from data. The method makes use of jet substructure templates that are obtained from a control region and used to dress a kinematic sample in the signal region. The implementation chosen here depends on kernel smoothing techniques, which are used to obtain a statistical error for the background prediction. For the two case studies in section 4, the QCD MC was treated as mock data and the predictions from kernel smoothing were shown to match the MC predictions. This illustrates that the systematic error introduced by treating the two leading jets as independent and identical is subdominant to the statistical error.

It is worth contrasting our data-driven proposal with existing approaches, such as the “ABCD” method. Ideally, the ABCD method uses a highly uncorrelated pair of variables




Figure 7. Quark and gluon templates for jet mass (left) and τ_{21} (right). A 20 GeV jet mass cut is indicated by the vertical grey line.

in order to define signal and control regions. If correlations are non-negligible, MC is used to calculate a correction factor that is folded into the background extrapolation. Thus we should consider two cases in comparing the two methods. In the first case, uncorrelated variables can be found, and ABCD yields a data-driven background estimate without relying on MC. Because the systematics associated with template methods and the ABCD method are complementary, each provides a valuable crosscheck on the background determination. In the second case, uncorrelated variables cannot be found, and any ABCD estimate is strongly dependent on MC. Template-based methods, however, may still be able to provide a data-driven estimate. This latter situation is expected to hold in many jet substructure searches of interest, e.g., for searches involving cuts on jet mass and additional substructure observables (because these observables are in general highly correlated).

This paper is just the start of exploring applications of the template approach. One important area of study is how templates depend on the quark/gluon composition of jets. The techniques in this paper can be generalized to include separate templates for quarks and gluons, which can be viewed as a discrete input label. To explore how the templates depend on quark/gluon content, we generated two MC samples of pure-quark and pure-gluon dijets and created the associated templates, shown in figure 7. Once a 20 GeV mass cut is imposed on each jet (dashed gray line), the quark and gluon mass distributions are very similar. However, there is clearly information that could potentially be incorporated into the procedure. Using separate quark/gluon templates will be a crucial step in applying these methods to third (or higher) jets in the event, which tend to be gluon jets.

There are other directions that deserve further study. It will be important to understand how the template estimates depend on the fat jet radius, and the impact of overlapping jets and pile-up. While the jet factorization assumption facilitates straightforward data-driven estimates, it would be interesting to explore what happens when this

condition is relaxed. This might include systematically improving the template method with input from perturbative QCD calculations. Finally, the template technique can be generalized to many different new physics searches, with different kinematic inputs and substructure variables.

Template functions provide a powerful way to estimate backgrounds in regions where direct MC calculations are prohibitively hard. This allows extrapolation of both MC and data into signal-rich regions. A deviation would be the first sign of new physics or, at worst, an indication of the breakdown of jet factorization or novel QCD effects not included in MC. No matter the result, we will learn something from the application of template methods to data.

Acknowledgments

We thank G. Salam, M. Swiatlowski, and B. Tweedie for helpful discussions, as well as J. Thaler for feedback on the draft. TC and JGW are supported by DoE contract number DE-AC02-76SF00515. HKL is supported by the DOE SCG Fellowship. TC and ML thank the Galileo Galilei Institute and the INFN, the Aspen Center for Physics, and the NSF Grant No. PHY-1066293 for partial support during the completion of this work.

A A kernel smoothing primer

This appendix provides a detailed introduction to kernel smoothing and presents the key derivations that justify the procedures used in the main body of the paper. It is divided into three parts: appendix A.1 reviews the basic formulae and the main properties of kernel smoothing; appendix A.2 details the various algorithmic choices that underlie our template procedure; and appendix A.3 validates the procedure on a test probability distribution. We chose our particular prescription for its relative simplicity, computational efficiency, and the fact that it works well in MC. More sophisticated variations are possible, including for example, adaptive kernel smoothing (which uses a variable bandwidth) and boundary kernels (which improve performance near the edge of the domain).

A basic statistical problem is to estimate a probability distribution $\rho(\vec{z})$ from a random sample of finite size. The simplest strategy is to bin the data into a histogram. Histograms, however, are intrinsically non-smooth, and the resulting estimates can be very sensitive to the choice of binning scheme, especially when the sample size is small. Furthermore, the statistical errors associated with histograms can be difficult to compute.

A better alternative is to use kernel smoothing techniques [45, 47, 59, 60]. Kernel smoothing replaces a histogram with a probability estimate $\hat{\rho}(\vec{z})$, where \vec{z} denotes a D -dimensional vector that can include kinematic and/or substructure variables. The weighting function, or kernel $K(\vec{z})$, must balance between the degree of locality and the smoothness of the resulting estimate.

Given a training sample with N_T independent data points $\{\vec{z}_1, \vec{z}_2, \dots, \vec{z}_{N_T}\}$ that have been obtained from the true distribution $\rho(\vec{z})$, one can obtain an estimate $\hat{\rho}$ of the true

density ρ with the following master formula:

$$\hat{\rho}(\vec{z}) = \frac{1}{N_T} \sum_i^{N_T} K_h(\vec{z} - \vec{z}_i), \quad \text{where} \quad K_h(\vec{z}) = \frac{1}{\det h} K(h^{-1}\vec{z}). \quad (\text{A.1})$$

Here, K is the smoothing kernel (often chosen to be a Gaussian) and h is a non-degenerate matrix that controls the smoothness of the resulting $\hat{\rho}$ by changing the width of the kernel. For a Gaussian kernel,⁵

$$K(\vec{z}) = \frac{1}{(4\pi)^{\frac{D}{2}}} \exp\left(-\frac{|\vec{z}|^2}{2}\right) \quad \text{and} \quad K_h(\vec{z}) = \frac{1}{(4\pi)^{\frac{D}{2}} \det h} \exp\left[-(2h^T h)_{ij}^{-1} z^i z^j\right]. \quad (\text{A.2})$$

The density estimate is normalized, $\int \hat{\rho} d^D z = 1$, as long as the kernel is normalized to unity, $\int K_h d^D z = 1$. After fixing K to be a Gaussian, one still has to choose the bandwidth matrix h , which should be large enough that the resulting estimate is reasonably smooth, but small enough that important features are not washed-out. As we will see below, the choice of h is typically more important than the choice of kernel, as it controls the size of the bias and variance. Choosing the best value of h is challenging because different measures result in different choices of h (see e.g., [61]). The following subsections present a concrete prescription for choosing h that leads to good statistical behavior.

A.1 Deriving the AMISE bandwidth

To select an optimal bandwidth h , one must specify a metric by which to measure optimality. A common choice is the Mean Integrated Squared Error (MISE), defined as the expectation of the square of the error term:

$$\text{MISE} = \int d^D z \langle (\hat{\rho}(\vec{z}) - \rho(\vec{z}))^2 \rangle \quad (\text{A.3})$$

$$= \int d^D z \langle (\langle \hat{\rho} \rangle - \rho)^2 \rangle + \int d^D z \langle (\hat{\rho} - \langle \hat{\rho} \rangle)^2 \rangle \quad (\text{A.4})$$

$$\equiv \int d^D z b(\vec{z})^2 + \int d^D z v^2(\vec{z}), \quad (\text{A.5})$$

where the notation $\langle \dots \rangle$ denotes the expectation value with respect to the true distribution; for any statistical quantity $f(\vec{z}_1, \dots, \vec{z}_{N_T})$ one defines

$$\langle f(\vec{z}_1, \vec{z}_2, \dots, \vec{z}_{N_T}) \rangle = \int \prod_i^{N_T} d^D z_i \rho(\vec{z}_i) f(\vec{z}_1, \vec{z}_2, \dots, \vec{z}_{N_T}). \quad (\text{A.6})$$

In eq. (A.5), the MISE is split into two terms. The first is identified as the bias $b(\vec{z})$ squared and arises from the fact that kernel smoothing systematically smooths out peaks and valleys. The second term yields the variance $v^2(\vec{z})$, which encodes statistical fluctuations about the mean estimate $\langle \hat{\rho}(\vec{z}) \rangle$. Eq. (A.5) encapsulates the classic bias-variance trade-off

⁵Because the Gaussian has infinite support, care should be taken when using this formula near (i.e., within a few bandwidths of) the boundary of the domain of \vec{z} .

in statistics: in order to reduce bias, one must reduce the amount of smoothing and hence increase the variance; to reduce the variance, one has to smooth out more features and hence increase the bias.

Computing the bias and variance is difficult because each depends explicitly on the true distribution ρ one is trying to estimate. Given a large sample size, it usually suffices to consider leading-order expressions in the asymptotic limit of large N_T . As the number of data points approaches infinity, the optimal bandwidth will approach zero, $h^{\text{AMISE}} \rightarrow 0$. As a result, it is useful to expand the MISE around the limit $N_T \rightarrow \infty$ and $h \rightarrow 0$ to derive an asymptotic expression for the optimal bandwidth, h^{AMISE} . Consider the average estimate $\langle \hat{\rho}(\vec{z}) \rangle$. In the large N_T limit, the summation in eq. (A.1) becomes a continuum integral:

$$\langle \hat{\rho}(\vec{z}) \rangle \rightarrow \frac{1}{\det h} \int d^D z' K(h^{-1}(\vec{z} - \vec{z}')) \rho(\vec{z}'). \quad (\text{A.7})$$

Introducing $\vec{s} = h^{-1}(\vec{z} - \vec{z}')$, one has $\vec{z}' = \vec{z} - h\vec{s}$ and $d^D z' = \det h d^D s$ so that

$$\langle \hat{\rho}(\vec{z}) \rangle = \int d^D s K(\vec{s}) \rho(\vec{z} - h\vec{s}). \quad (\text{A.8})$$

Expanding the integrand to second order in h and writing out all the indices explicitly yields

$$\begin{aligned} \langle \hat{\rho}(\vec{z}) \rangle &\simeq \int d^D s \left(K(\vec{s}) \rho(\vec{z}) - K(\vec{s}) (h s)^i \frac{d\rho(\vec{z})}{dz^i} + \frac{1}{2} K(\vec{s}) (h s)^i (h s)^j \frac{d^2\rho(\vec{z})}{dz_i dz_j} \right) \\ &= \rho(\vec{z}) + \frac{h_{ik} h_{jl}}{2} \frac{d^2\rho(\vec{z})}{dx_i dx_j} \int d^D s s_k s_l K(\vec{s}), \end{aligned} \quad (\text{A.9})$$

where a symmetric kernel with $K(s) = K(-s)$ has been assumed. For the Gaussian kernel, which satisfies $\int d^D s s_i s_j K(\vec{s}) = \delta_{ij}$, the leading-order bias is

$$b(\vec{z}) \equiv \langle \hat{\rho}(\vec{z}) \rangle - \rho(\vec{z}) = \frac{(h^T h)_{ij}}{2} \frac{d^2\rho(\vec{z})}{dz_i dz_j}. \quad (\text{A.10})$$

The interpretation is that the smoothing bias is most pronounced at peaks and valleys where the second derivative is large. Local maxima (minima) are estimated to be lower (higher) due to smoothing.

A similar calculation can be performed for the variance:

$$v^2(\vec{z}) = \langle \hat{\rho}^2(\vec{z}) \rangle - \langle \hat{\rho}(\vec{z}) \rangle^2 \quad (\text{A.11})$$

$$\begin{aligned} &= \frac{1}{N_T^2} \left[\int \prod_k^{N_T} d^D z_k \rho(\vec{z}_k) \sum_{i,j}^{N_T} K_h(\vec{z} - \vec{z}_i) K_h(\vec{z} - \vec{z}_j) \right. \\ &\quad \left. - \left(\int \prod_k^{N_T} d^D z_k \rho(\vec{z}_k) \sum_i^{N_T} K_h(\vec{z} - \vec{z}_i) \right)^2 \right]. \end{aligned} \quad (\text{A.12})$$

In the first summation, the $i = j$ terms yield N_T integrals over K_h^2 . For $i \neq j$, each of the $N_T^2 - N_T$ integrals factorizes into the square of $\int d^D z' K_h(\vec{z} - \vec{z}')$, as \vec{z}_i and \vec{z}_j are

independent random variables drawn from the same distribution. In the second summation, independence of the z_i yields a further N_T^2 terms $(\int d^D z' K_h(\vec{z} - \vec{z}'))^2$. Gathering everything together, one has

$$v^2(\vec{z}) = \frac{1}{N_T^2} \left[N_T \int d^D z' \rho(\vec{z}') K_h^2(\vec{z} - \vec{z}') - N_T \left(\int d^D z' \rho(\vec{z}') K_h(\vec{z} - \vec{z}') \right)^2 \right]. \quad (\text{A.13})$$

Until this point, the computation is exact. To proceed further, consider the limit $h \rightarrow 0$, where $K_h(\vec{z}) \rightarrow \delta(\vec{z})$. The first term dominates because it contains a factor of $\delta^2(\vec{z})$. Therefore, in the asymptotic limit:

$$v^2(\vec{z}) \simeq \frac{1}{N_T} \int d^D z' \rho(\vec{z}') K_h^2(\vec{z} - \vec{z}') \quad (\text{A.14})$$

$$= \frac{1}{N_T \det h^2} \int d^D z' \rho(\vec{z}') K^2(h^{-1}(\vec{z} - \vec{z}')) \quad (\text{A.15})$$

$$= \frac{1}{N_T \det h} \int d^D s \rho(\vec{z} - h \vec{s}) K^2(s) \simeq \frac{\rho(\vec{z})}{N_T \det h} \int d^D s K^2(s), \quad (\text{A.16})$$

where $\rho(\vec{z} - h \vec{s}) \simeq \rho(\vec{z})$ to leading order in the last line. The variance is directly proportional to the true density in the neighborhood of \vec{z} . For a Gaussian kernel, the integral evaluates to $\int d^D s K^2(s) = (4\pi)^{-\frac{D}{2}}$. The leading-order variance in the asymptotic limit is then

$$v^2(\vec{z}) = \frac{\rho(\vec{z})}{N_T (4\pi)^{\frac{D}{2}} \det h}. \quad (\text{A.17})$$

Armed with asymptotic formulae for the bias and variance, we can compute the Asymptotic MISE (AMISE) by plugging eq. (A.10) and eq. (A.17) into eq. (A.5):

$$\text{AMISE} \equiv \int d^D z \left[\left(\frac{(h^T h)_{ij}}{2} \frac{d^2 \rho(\vec{z})}{dz_i dz_j} \right)^2 + \frac{\rho(\vec{z})}{N_T (4\pi)^{\frac{D}{2}} \det h} \right]. \quad (\text{A.18})$$

The optimal bandwidth is determined by minimizing the AMISE; this derivation yields a bandwidth h^{AMISE} that is correct to leading order in N_T^{-1} . To explicitly minimize this expression, it is useful to rewrite the bias. Using integration by parts, eq. (A.10) is rewritten as

$$\frac{1}{4} (h^T h)_{ij} (h^T h)_{kl} \int d^D z \frac{\rho(\vec{z}) d^4 \rho(\vec{z})}{dz_i dz_j dz_k dz_l} \equiv \frac{1}{4} (h^T h)_{ij} (h^T h)_{kl} S_{ijkl}, \quad (\text{A.19})$$

where S_{ijkl} is a totally symmetric rank-4 tensor. Any such tensor can be diagonalized by an orthogonal transformation; it is possible to perform a coordinate transformation $\vec{z} \rightarrow \vec{y}$ such that $S_{ijkl} = \frac{1}{4} (\delta_{ij} \delta_{kl} + \delta_{ik} \delta_{jl} + \delta_{il} \delta_{jk}) (4\pi)^{-\frac{D}{2}}$, where the normalization factor assumes ρ has been transformed into standard normal form. Defining $H_{ij} \equiv (h^T h)_{kl} \frac{dy_i}{dz_k} \frac{dy_j}{dz_l}$, the AMISE can be rewritten as

$$\text{AMISE} = \frac{1}{(4\pi)^{\frac{D}{2}} \det \frac{d\vec{y}}{d\vec{z}}} \left(\frac{1}{16} [2 \text{tr} H^2 + (\text{tr} H)^2] + \frac{1}{N_T \sqrt{\det H}} \right). \quad (\text{A.20})$$

The minimum occurs when H_{ij} is proportional to the unit matrix δ_{ij} . The optimal bandwidth in the asymptotic limit is then identified as⁶

$$H_{ij}^{\text{AMISE}} = \delta_{ij} \left(\frac{4}{2+D} \right)^{\frac{2}{D+4}} N_T^{-\frac{2}{D+4}} \quad \text{and} \quad h_{ij}^{\text{AMISE}} = \frac{dy_i}{dz_j} \left(\frac{4}{2+D} \right)^{\frac{1}{D+4}} N_T^{-\frac{1}{D+4}}. \quad (\text{A.21})$$

Even in the asymptotic limit, the expression for the optimal bandwidth depends on the true distribution $\rho(\vec{z})$ via $\frac{dy_i}{dz_j}$. This dependence cannot be removed without making further assumptions. For a Gaussian distribution, the required transformation is simply $y_i = \sigma_{ij} z_j$, where σ_{ij} is the square root of the covariance matrix. For more general distributions that are not too far from the Gaussian distribution, one expects the correct transformation $z \rightarrow y$ to roughly correspond to rotating away the correlations and scaling the width of the distribution to unity. Hence, the following choice of the bandwidth should be near optimal for most reasonable probability distributions:

$$h_{ij}^{\text{AMISE}} \simeq c \sigma_{ij} N_T^{-\frac{1}{D+4}}. \quad (\text{A.22})$$

Here, σ_{ij} is the square root of the covariance matrix computed from the distribution $\rho(\vec{z})$, and c is an N_T -independent $\mathcal{O}(1)$ constant; eq. (A.22) is known as Silverman’s Rule-of-Thumb [45]. In practice, σ_{ij} is replaced by an estimate $\hat{\sigma}_{ij}$ from data. As we will see in appendix A.3, variance estimation is more reliable than bias estimation in practice. It is therefore advantageous to choose a value for c that deliberately undersmooths $\hat{\rho}$, so that the bias is smaller than the variance and the error estimate can be trusted.

A.2 Practical kernel smoothing

This subsection presents the specific algorithms and formulae used in this paper. In the applications we consider, kernel estimates $\hat{\rho}(\vec{z})$ are typically constructed from very large samples, and it is not always feasible to keep the \vec{z}_i for each event in memory.⁷ To limit memory usage, the dataset is first binned with a small bin size $h_b \ll h$, and then kernel smoothing is performed on the binned dataset. Additionally, to save computational time, the kernel estimate is precomputed at each bin location and linear interpolation is used to compute $\hat{\rho}(\vec{z})$ for intermediate values of \vec{z} .

More explicitly, let \vec{q} denote the center of a given bin and let $\vec{q}(\vec{z})$ be the function that sends \vec{z} to the center of the corresponding bin. Given a training sample $\{\vec{z}_i\}$ of size N_T , denote the number of \vec{z}_i in each bin \vec{q} as $n_{\vec{q}}$. The master formula for kernel smoothing, eq. (A.1), then becomes

$$\hat{\rho}(\vec{q}') = \frac{\kappa}{N_T} \sum_i^{N_T} K_h(\vec{q}' - \vec{q}(\vec{z}_i)) = \frac{\kappa}{N_T} \sum_{\vec{q}} n_{\vec{q}} K_h(\vec{q}' - \vec{q}), \quad (\text{A.23})$$

⁶ h_{ij} is ambiguous up to an orthogonal transformation $h_{ij} \rightarrow O_{ik} h_{kj}$. This ambiguity is unimportant, as only $h^T h$ is relevant for Gaussian kernels.

⁷All events need to be stored in memory to implement the Fast Fourier Transform [62–64] used in our C++ code.

where κ is a normalization constant. Eq. (A.23) only defines $\hat{\rho}$ at the center of each bin; to extend the domain of $\hat{\rho}$ to all values of \vec{z} , multi-dimensional linear interpolation is performed. In principle, the constant κ can be set to normalize $\hat{\rho}(\vec{z})$ to unity; when computing cut efficiencies this constant drops out, so it can be ignored in practice.

Combined with the rule-of-thumb for bandwidth selection given in eq. (A.22), the calculation of $\hat{\rho}(\vec{z})$ is now straightforward. Estimating the bias $\hat{b}(\vec{z})$ and variance $\hat{v}^2(\vec{z})$ is less trivial. One could in principle substitute $\hat{\rho}(\vec{z})$ for $\rho(\vec{z})$ into eq. (A.10) and eq. (A.17) and use the resulting estimates, which are correct to leading order in h . We find, however, that it is preferable to use a different approach, which we now describe.⁸

The bias can be estimated by performing kernel smoothing twice and comparing the result with that obtained from performing kernel smoothing once. Specifically,⁹

$$\begin{aligned} \hat{b}(\vec{z}) &= \int d^D z' \hat{\rho}(\vec{z} - \vec{z}') K_h(\vec{z}') - \hat{\rho}(\vec{z}) \\ &\simeq \frac{(h^T h)_{ij}}{2} \frac{d^2 \hat{\rho}(\vec{z})}{dz_i dz_j} \simeq \frac{(h^T h)_{ij}}{2} \frac{d^2 \rho(\vec{z})}{dz_i dz_j}. \end{aligned} \tag{A.24}$$

This procedure is equivalent to eq. (A.10) in the asymptotic limit [47]. We find that eq. (A.24) yields relatively stable results in the test simulation discussed in appendix A.3. Note that the kernel smoothing in eq. (A.24) is performed using the same algorithm outlined above.¹⁰

We use the bootstrap [47–50] to estimate the variance. The flexibility of the bootstrap is particularly well-suited to our case, where the variance needs to be propagated through complicated kinematic cuts. For a given statistic ϵ (e.g., ϵ can be the efficiency for the cuts of interest) and estimator $\hat{\epsilon}$, the bootstrap relies on creating a set of resamples of the observed dataset. Each resampled dataset then yields a value for $\hat{\epsilon}$, the distribution of which is used to measure properties of the estimator $\hat{\epsilon}$ (e.g., its variance).

To compute the variance, one first converts a given dataset $\{\vec{z}_i\}$ into a set of bin counts $\{n_{\vec{q}}\}$. One then fluctuates each element of the set $\{n_{\vec{q}}\}$ a total of N_{BS} times, thus creating N_{BS} fluctuated datasets, denoted by $\{n_{\vec{q}}\}_b$, with $b \in \{1, \dots, N_{\text{BS}}\}$. Each element of $\{n_{\vec{q}}\}_b$ is sampled from the Poisson distribution:¹¹

$$\{n_{\vec{q}}\}_b \sim \text{Poisson}(n_{\vec{q}}). \tag{A.25}$$

⁸For example, because eq. (A.10) requires an estimation of the second derivative of ρ , it involves the additional complication of selecting a bandwidth appropriate for derivative estimation.

⁹In principle, eq. (A.24) depends on two distinct bandwidths. In practice, we find that it is sufficient to take these to be equal as determined using eq. (A.22). It should be kept in mind, however, that different bias estimators are possible and may perform better on real data.

¹⁰One might worry that binning the data yields additional sources of bias not accounted for in eq. (A.24). As long as $h_b \ll h$ and linear interpolation is applied, however, this additional bias is negligible. In practice, we find that h_b as large as $h/2$ still yields a consistent estimate. Of course, it is preferable to take h_b as small as allowed by memory limitations.

¹¹For N_T data points $\{\vec{z}_i\}$ the standard bootstrap procedure generates each resampled dataset by drawing exactly N_T samples from the empirical distribution (i.e., from the observed dataset) with replacement. Our Poisson fluctuation procedure (which we choose for computational convenience) is not identical to the standard bootstrap, but in practice we find that it provides robust estimates of the variance. The accuracy of this prescription deserves further scrutiny in any future experimental study.

For each fluctuated dataset, one can compute a corresponding $\{\hat{\rho}(\vec{z})\}_b$ using the same binned kernel smoothing procedure as in eq. (A.23):

$$\{\hat{\rho}(\vec{z})\}_b = \frac{\kappa_b}{\mathcal{N}_b} \sum_{\vec{q}} \{n_{\vec{q}}\}_b K_h(\vec{q}' - \vec{q}), \quad (\text{A.26})$$

where, as before, κ_b is a normalization constant, and \mathcal{N}_b is the total number of events in the b^{th} fluctuated dataset. Finally, we obtain an estimate of the smoothing variance:

$$\hat{v}^2(\vec{z}) = \frac{1}{N_{\text{BS}} - 1} \sum_{b=1}^{N_{\text{BS}}} \left[\{\hat{\rho}(\vec{z})\}_b - \langle \{\hat{\rho}(\vec{z})\} \rangle_{\text{BS}} \right]^2; \quad \langle \{\hat{\rho}(\vec{z})\} \rangle_{\text{BS}} = \frac{1}{N_{\text{BS}}} \sum_{b=1}^{N_{\text{BS}}} \{\hat{\rho}(\vec{z})\}_b. \quad (\text{A.27})$$

In practice, $N_{\text{BS}} \simeq 100$ is sufficient for a reasonable estimate of the variance.

For our application, it is important that density estimates $\hat{\rho}(\vec{z})$ be associated with a confidence interval. However, due to the presence of asymptotic bias (see especially the discussions in [49, 50]), a naive application of the smoothing variance in eq. (A.27) might lead us to construct incorrect confidence intervals that are centered at $\langle \hat{\rho}(\vec{z}) \rangle$ instead of at $\rho(\vec{z})$. There are two approaches to this problem, both of which can be used to ensure the construction of properly centered confidence intervals. The first is to deliberately under-smooth the data so that bias is negligible in comparison to the root variance. The second is to explicitly correct for the bias. In practice, we find that it is necessary to use both approaches simultaneously. Relying solely on undersmoothing is inadequate because background estimates rely on estimates of $\rho(\vec{z})$ at a variety of different values of \vec{z} ; to ensure small bias over the entire range of relevant values of \vec{z} would lead to too much undersmoothing. Relying solely on bias correction is inadequate because of the difficulty of constructing an accurate bias estimator. By making use of both approaches simultaneously, it is possible to mitigate the problems of asymptotic bias. In practice, for any specified cut, we choose a bandwidth such that $\hat{\sigma}_V \gtrsim 2 \hat{\sigma}_B$, where $\hat{\sigma}_B$ is a measure of the total bias of $\hat{N}_e^*[C]$, see eq. (3.23). It is worth stressing the need to pay careful attention to these issues, because they play a central role in determining the robustness of the final statistical uncertainties associated with the template procedure. Note that more sophisticated approaches that, e.g. make use of adaptive kernel smoothing or implement bias correction at the level of $\hat{e}[C]$, may allow for an improved treatment of asymptotic bias.

To introduce explicit bias correction, we define the bias-corrected density estimate $\hat{\rho}^*(\vec{z})$:

$$\hat{\rho}^*(\vec{z}) \equiv \hat{\rho}(\vec{z}) - \hat{b}(\vec{z}). \quad (\text{A.28})$$

Note that $\int d^D z \hat{\rho}^*(\vec{z}) = 1$ as long as $\hat{b}(\vec{z})$ is the difference of two probability distributions as computed in eq. (A.24). The resulting $\hat{\rho}^*$ is not necessarily positive everywhere. We find that this does not cause any problems for the cut efficiencies computed in this paper, but one should be aware of this issue. Eq. (A.27) can also be used to define the bias-corrected variance via the substitution¹²

$$\hat{\rho} \rightarrow \hat{\rho}^* \implies v^2 \rightarrow v^{2*}. \quad (\text{A.29})$$

¹²The variance itself is not bias-corrected, but rather v^{2*} represents the dispersion about the bias-corrected density estimate $\hat{\rho}^*(\vec{z})$.




Figure 8. Results of the first validation test with 1000 trial experiments. Left: $\rho(x)$ is the true distribution (blue), $\hat{\rho}(x)$ is the kernel estimate from a single trial experiment (green), and $\langle \hat{\rho}(x) \rangle$ is the average over all experiments (red). Right: the bias-corrected results for the same set of trial experiments.

Using the bias-corrected $\hat{\rho}^*(\vec{z})$, one can construct a bias-corrected estimator $\hat{\epsilon}^*$, which has substantially reduced bias. Note that ϵ^* is still a biased estimator, but the bias is now of higher order in the bandwidth h . The cost of reducing bias is an increase in the variance, which can be reliably computed with the bootstrap.

Given the ensemble of bias-corrected estimators, the corresponding statistical error is computed with the bootstrap. As discussed in section 3.3, each of the $\{\hat{\rho}^*\}_b$ is used as an input to generate an ensemble of bootstrapped results $\{N_e^*[C]_1, \dots, N_e^*[C]_{N_{BS}}\}$ for a given set of cuts C . The variance of the final prediction is estimated using this ensemble:

$$\hat{\sigma}_V^2 \simeq \frac{1}{N_{BS} - 1} \sum_{b=1}^{N_{BS}} (\mathcal{N}_e^*[C]_b - \langle \mathcal{N}_e^*[C] \rangle_{BS})^2 \quad \text{where} \quad \langle \mathcal{N}_e^*[C] \rangle_{BS} = \frac{1}{N_{BS}} \sum_{b=1}^{N_{BS}} \mathcal{N}_e^*[C]_b. \tag{A.30}$$

This is to be distinguished from the variance of the templates themselves, $v^2(\vec{z})$.

The next subsection presents a validation study showing that the bias correction and variance estimates outlined above are statistically robust and lead to sound confidence intervals.

A.3 Validation

This subsection presents a validation study for our particular implementation of kernel smoothing. The first test evaluates the accuracy of the estimates $\hat{\rho}(\vec{z})$ by applying kernel smoothing to a known one-dimensional probability distribution $\rho(x)$ with $x \in [0, 1]$. The second test evaluates the data-driven procedure for estimating cut efficiencies by applying cuts to the joint probability distribution $\rho(x_1, x_2) = \rho(x_1)\rho(x_2)$. Both tests use the distribution $\rho(x) \sim x(1-x)(\cos(8x) + 1.5)$, illustrated in figure 8.




Figure 9. The bias and variance for the first validation test. Left: $b(x)$ is the true bias (blue), $\hat{b}(x)$ is the bias estimate for a single trial experiment (green), and $\langle \hat{b}(x) \rangle$ is the estimated bias averaged over all trial experiments (red). Right: the corresponding results for the square root of the bias-corrected variance, $\sqrt{v^{2*}(x)}$.

The first test is formulated as follows:

1. Do $N_{\text{trials}} = 1000$ trial experiments, with each experiment containing $N = 500$ events, $\{x_1, x_2, \dots, x_N\}$, drawn from $\rho(x)$.
2. For each experiment i , compute the estimates $\hat{\rho}_i(x)$, $\hat{\rho}_i^*(x)$, $\hat{b}_i(x)$ and $\sqrt{\hat{v}_i^{2*}(x)}$ using the methods outlined in appendix A.2.
3. Compute the averages $\langle \hat{\rho}(x) \rangle$, $\langle \hat{\rho}^*(x) \rangle$, $\langle \hat{b}(x) \rangle$ and $\langle \sqrt{\hat{v}^{2*}(x)} \rangle$ from the 1000 experiments.
4. Using the true distribution $\rho(x)$, compute $b(x)$ and $\sqrt{v^{2*}(x)}$.

The resulting kernel estimates are illustrated in figure 8. The bandwidth is chosen with Silverman’s Rule-of-Thumb with a moderate amount of undersmoothing ($c \simeq 0.9 c_{\text{AMISE}}$ and $h \simeq 0.04$ in eq. (A.22)). The bin width is $h_b = 0.0025 \ll h$. The bias-corrected version on the right shows substantial reduction in the average bias (i.e., red matches blue). The variance, however, increases as can be seen by the increased error of the sample estimate, shown in green.

The performance of the bias and variance estimators is shown in figure 9. The left panel shows systematic underestimation of the bias. The bias, however, is subdominant to the variance as a consequence of undersmoothing. Since the bias is generally of the correct sign and order-of-magnitude, bias correction leads to confidence intervals that have improved probability coverage (as we will see explicitly below). Compared to the bias, the variance can be more reliably estimated, as seen in the right panel. It is likely that the bias estimator can be improved, which would be an important goal for any experimental study.

The second test goes a step further and allows us to evaluate the statistical performance of our procedure for estimating cut efficiencies. In particular it constitutes a test of




Figure 10. The results of the second validation test, which evaluates the probability coverage of the confidence intervals corresponding to the cut efficiencies $\hat{\epsilon}$ and $\hat{\epsilon}^*$ with $x_{\text{cut}} = 1.4$ (left) and 1.8 (right). The statistics s and s^* are defined in eq. (A.31). The fact that s^* is approximately Gaussian with mean zero and unit variance indicates that the corresponding confidence intervals are sound.

the smoothing variance $\hat{\sigma}_V^2$ and the probability coverage of the corresponding confidence interval. It differs from the cut computations in section 4 only in that: i) the kinematic sample is effectively of size 1 and ii) the dimension is $D = 1$.

The test is formulated as follows:

1. Do $N_{\text{trials}} = 5000$ trial experiments, with each experiment containing $N = 500$ events, $\{x_1, x_2, \dots, x_N\}$, drawn from $\rho(x)$.
2. For each trial experiment i , compute $\hat{\rho}_i(x)$ and $\hat{\rho}_i^*(x)$.
Let $\epsilon[C]_i \equiv \int_C dx_1 dx_2 \hat{\rho}_i(x_1) \hat{\rho}_i(x_2)$, where C is given by the region $x_1 + x_2 > x_{\text{cut}}$. Compute $\hat{\epsilon}[C]_i$ and $\hat{\epsilon}^*[C]_i$.
3. For each trial experiment i , calculate the bias estimator $\hat{\sigma}_{Bi} = \hat{\epsilon}[C]_i - \hat{\epsilon}^*[C]_i$, as in eq. (3.23), and use the bootstrap to calculate $\hat{\sigma}_{Vi}$, as in eq. (3.20).
4. Using the true distribution $\rho(x)$, compute $\epsilon[C]$ and form the statistics

$$s_i \equiv \frac{\hat{\epsilon}[C]_i - \epsilon[C]}{\hat{\sigma}_{Vi}} \quad \text{and} \quad s_i^* \equiv \frac{\hat{\epsilon}^*[C]_i - \epsilon[C]}{\hat{\sigma}_{Vi}}. \quad (\text{A.31})$$

Note that $\hat{\sigma}_{Vi}$ does not include a bias correction when constructing s_i .

5. Finally, calculate the variance of the estimates $\{\hat{\epsilon}^*[C]_i\}$:

$$\sigma_V^2 = \frac{1}{N_{\text{trials}} - 1} \sum_{i=1}^{N_{\text{trials}}} (\hat{\epsilon}^*[C]_i - \langle \hat{\epsilon}^*[C] \rangle_{\text{trials}})^2 \quad \langle \hat{\epsilon}^*[C] \rangle_{\text{trials}} = \frac{1}{N_{\text{trials}}} \sum_{i=1}^{N_{\text{trials}}} \hat{\epsilon}^*[C]_i. \quad (\text{A.32})$$

The distributions for s and s^* are shown in figure 10. Note that a robust confidence interval should result in a statistic s (s^*) that is approximately Gaussian with zero mean




Figure 11. The performance of $\hat{\sigma}_B$ and $\hat{\sigma}_V$ in the 5000 trial experiments of the second validation test. Plotted is the distribution of deviations for the estimator $\hat{\sigma}_B = |\hat{\epsilon}[C] - \epsilon^*[C]|$ from the true bias $\sigma_B = |\hat{\epsilon}[C] - \epsilon[C]|$ (left). Also shown is the deviation of the estimator $\hat{\sigma}_V$ from the true root variance σ_V (right).

and unit variance. As can be seen in the figure, bias correction significantly improves the Gaussianity of s^* (which corresponds to the bias-corrected estimate $\hat{\epsilon}^*$) as compared to s (which corresponds to the estimate $\hat{\epsilon}$). For $x_{\text{cut}} = 1.4$ and with $h = 0.06$ the bias is roughly a factor of 20 smaller than the variance, but the difference between s and s^* is still non-negligible. For $x_{\text{cut}} = 1.8$, h needs to be significantly smaller with $h = 0.015$ in order to achieve a subdominant bias. The bias in this case is roughly a factor of two smaller than the variance, and the bias correction results in a sizable shift from s to s^* . For both cuts, bias correction improves the probability coverage of the confidence interval for $\hat{\epsilon}^*$ (especially its centeredness) and thus justifies the error bars given for the background estimates in section 4.

Finally, the performance of the estimators $\hat{\sigma}_B$ and $\hat{\sigma}_V$ is shown in figure 11. As is apparent, variance estimation is much more reliable than bias estimation. The bias estimator $\hat{\sigma}_B$ (as opposed to the bias estimator $\hat{b}(\vec{z})$ used to define $\hat{\rho}^*(\vec{z})$) need not be very accurate because it is only used to ensure that we are in a statistical regime where the bias is subdominant so that confidence intervals have good probability coverage. Note that the performance of $\hat{\sigma}_B$ illustrated in the figure is for the particular cut value of $x_{\text{cut}} = 1.4$; the agreement between $\hat{\sigma}_B$ and σ_B varies as a function of x_{cut} , although the overall level of agreement shown here is representative.

B QCD Monte Carlo

This section describes the simulated data utilized in our studies. Up to four final-state partons were simulated at the generator-level, which were then merged with a shower. The final hadron-level output was pixelated to model the finite granularity of the calorimeter.¹³

¹³This more closely resembles the setup for ATLAS than for CMS, where particle-flow jets are used instead. We expect that this will have little to no impact on our results.

bin	H_T [GeV]	$2j$		$3j$		$4j$	
		σ [pb]	N_{events}	σ [pb]	N_{events}	σ [pb]	N_{events}
1	0–50	9.0×10^7	2.8×10^6	0	0	0	0
2	50–100	1.0×10^8	4.1×10^6	6.8×10^6	3.0×10^6	3.2×10^4	8.3×10^5
3	100–200	4.3×10^6	3.0×10^6	6.5×10^6	3.1×10^6	2.8×10^6	3.3×10^6
4	200–400	1.1×10^5	1.6×10^6	2.9×10^5	1.8×10^6	1.2×10^6	4.2×10^6
5	400–800	1.9×10^3	8.2×10^5	6.9×10^3	8.4×10^5	9.5×10^4	5.2×10^6
6	800–1400	2.7×10^1	4.9×10^5	1.1×10^2	4.8×10^5	3.6×10^3	4.9×10^6
7	1400–2000	6.3×10^{-1}	4.0×10^5	2.7×10^0	3.8×10^5	1.3×10^2	5.1×10^6
8	2000–3000	4.5×10^{-2}	3.9×10^5	1.9×10^{-1}	3.8×10^5	1.0×10^1	5.4×10^6
9	3000–4000	1.0×10^{-3}	4.2×10^5	4.4×10^{-3}	4.2×10^5	2.2×10^{-1}	4.6×10^6
10	4000–5000	2.6×10^{-5}	4.3×10^5	1.1×10^{-4}	4.6×10^5	5.2×10^{-3}	5.4×10^6
11	5000–6000	4.4×10^{-7}	4.9×10^5	1.9×10^{-6}	5.6×10^5	7.8×10^{-5}	4.8×10^6
12	> 6000	2.7×10^{-9}	5.7×10^5	1.2×10^{-8}	6.9×10^5	3.8×10^{-7}	4.2×10^6

Table 3. The number of MC events after matching in the different H_T and jet multiplicity bins. The associated matched cross sections for each bin are also provided.

Finally, these pixels were clustered into fat anti- k_T jets with $R = 1.0$. This level of detail in the simulation is sufficient for demonstrating the viability of our data-driven technique. Note that if we were considering backgrounds for searches involving missing energy, it would be important to also simulate detector-smearing effects.

In more detail, parton-level QCD events were generated using **Madgraph5** version 5.1.5.10 at $\sqrt{s} = 8$ TeV and merged with the parton shower via the MLM “MadGraph style” prescription using **PYTHIA** version 8.1.76 [54]. To more accurately model the large phase space for multijet production at the LHC, we employed a variation of the weighting technique discussed in [52]. Specifically, events were binned in both jet multiplicity (2 exclusive, 3 exclusive, and 4 inclusive jets) and in H_T , see table 3.

The post-PYTHIA hadron-level information was then passed to a code based on **FastJet** version 3.0.2 [65, 66]. The granularity of the calorimeter was simulated by creating a grid of cells in the η - ϕ plane with size 0.1×0.1 for $|\eta| < 2.5$ and 0.2×0.2 for $2.5 < |\eta| < 4.5$; any particles with $|\eta| > 4.5$ were discarded. For each event, the energy of all particles that fell within that cell was summed. This gave the energy of the cell for the given event. The energy and position of these cells were combined into light-like four-vectors that were used as inputs to the **FastJet** routines.

Next, the cells were clustered into $R = 1.0$ anti- k_T jets. In order to reduce sensitivity to the impact of pileup, the jets were groomed using the trimming algorithm with $f_{\text{cut}} = 0.05$ and $R_{\text{trim}} = 0.3$. The fat jets used in our study are the output of this trimming procedure. N -subjettiness was computed using the **FastJet** package [55, 56] with the ‘min_axes’ algorithm and $\beta = 1$.

One of the main goals of section 4 is to demonstrate that the error bars associated with the smoothing procedure are reasonably small. As discussed above, the width of the kernel scales with the number of events in the training sample, N_T , and the total AMISE scales as $N_T^{-4/(D+4)}$ (see eq. (A.20)). As a result, we must understand the size of the training sample available at the LHC in order to make realistic projections for the statistical uncertainties of the template procedure. For very inclusive triggers, the event rate can be so large that it is not possible to write all events that pass selection cuts to tape. To deal with this, the collaborations perform prescaling, i.e. they record only a fraction of the events and reweight them to account for events that went unrecorded. Prescaling is relevant because the weighted events come with larger statistical uncertainties than the naive expectation of \sqrt{N} .

The training samples we consider are taken from low-multiplicity exclusive events. We assume these events are obtained through a dijet trigger, which requires an enormous prescaling factor for low- p_T cuts. To avoid these issues, we choose a p_T cut high enough that prescaling is not necessary. The p_T threshold is computed as follows:

1. for a given threshold p_T , compute the two fat jet cross section;
2. assume that the instantaneous luminosity at the LHC is $10^{34} \text{ cm}^{-2}/\text{s}$, and compute the prescaling factor required to keep the trigger rate below 200 Hz;
3. choose the lowest p_T threshold such that the prescaling factor is roughly 1.

This procedure yields a p_T threshold of 200 GeV. Since both our mock analyses stay above this p_T threshold, it is not necessary for us to take prescaling explicitly into account.

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

References

- [1] A. Abdesselam et al., *Boosted objects: a probe of beyond the standard model physics*, *Eur. Phys. J. C* **71** (2011) 1661 [[arXiv:1012.5412](https://arxiv.org/abs/1012.5412)] [[INSPIRE](#)].
- [2] 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](https://arxiv.org/abs/1201.0008)] [[INSPIRE](#)].
- [3] A. Altheimer et al., *Boosted objects and jet substructure at the LHC*, [arXiv:1311.2708](https://arxiv.org/abs/1311.2708) [[INSPIRE](#)].
- [4] J.M. Butterworth, A.R. Davison, M. Rubin and G.P. Salam, *Jet substructure as a new Higgs search channel at the LHC*, *Phys. Rev. Lett.* **100** (2008) 242001 [[arXiv:0802.2470](https://arxiv.org/abs/0802.2470)] [[INSPIRE](#)].
- [5] D.E. Kaplan, K. Rehermann, M.D. Schwartz and B. Tweedie, *Top tagging: a method for identifying boosted hadronically decaying top quarks*, *Phys. Rev. Lett.* **101** (2008) 142001 [[arXiv:0806.0848](https://arxiv.org/abs/0806.0848)] [[INSPIRE](#)].
- [6] T. Plehn, M. Spannowsky, M. Takeuchi and D. Zerwas, *Stop reconstruction with tagged tops*, *JHEP* **10** (2010) 078 [[arXiv:1006.2833](https://arxiv.org/abs/1006.2833)] [[INSPIRE](#)].

- [7] Y. Cui, Z. Han and M.D. Schwartz, *W-jet tagging: optimizing the identification of boosted hadronically-decaying W bosons*, *Phys. Rev. D* **83** (2011) 074023 [[arXiv:1012.2077](#)] [[INSPIRE](#)].
- [8] A. Hook, E. Izaguirre, M. Lisanti and J.G. Wacker, *High multiplicity searches at the LHC using jet masses*, *Phys. Rev. D* **85** (2012) 055029 [[arXiv:1202.0558](#)] [[INSPIRE](#)].
- [9] T. Cohen, E. Izaguirre, M. Lisanti and H.K. Lou, *Jet substructure by accident*, *JHEP* **03** (2013) 161 [[arXiv:1212.1456](#)] [[INSPIRE](#)].
- [10] S. El Hedri, A. Hook, M. Jankowiak and J.G. Wacker, *Learning how to count: a high multiplicity search for the LHC*, *JHEP* **08** (2013) 136 [[arXiv:1302.1870](#)] [[INSPIRE](#)].
- [11] R.K. Ellis, K. Melnikov and G. Zanderighi, *Generalized unitarity at work: first NLO QCD results for hadronic W + 3 jet production*, *JHEP* **04** (2009) 077 [[arXiv:0901.4101](#)] [[INSPIRE](#)].
- [12] C.F. Berger et al., *Precise predictions for W + 3 jet production at hadron colliders*, *Phys. Rev. Lett.* **102** (2009) 222001 [[arXiv:0902.2760](#)] [[INSPIRE](#)].
- [13] A. Bredenstein, A. Denner, S. Dittmaier and S. Pozzorini, *NLO QCD corrections to $pp \rightarrow t\bar{t}b\bar{b} + X$ at the LHC*, *Phys. Rev. Lett.* **103** (2009) 012002 [[arXiv:0905.0110](#)] [[INSPIRE](#)].
- [14] G. Bevilacqua, M. Czakon, C.G. Papadopoulos, R. Pittau and M. Worek, *Assault on the NLO wishlist: $pp \rightarrow t\bar{t}b\bar{b}$* , *JHEP* **09** (2009) 109 [[arXiv:0907.4723](#)] [[INSPIRE](#)].
- [15] T. Binoth et al., *Next-to-leading order QCD corrections to $pp \rightarrow b\bar{b}b\bar{b} + X$ at the LHC: the quark induced case*, *Phys. Lett. B* **685** (2010) 293 [[arXiv:0910.4379](#)] [[INSPIRE](#)].
- [16] H. Ita et al., *Precise predictions for Z + 4 jets at hadron colliders*, *Phys. Rev. D* **85** (2012) 031501 [[arXiv:1108.2229](#)] [[INSPIRE](#)].
- [17] J. Alwall et al., *MadGraph/MadEvent v4: the new web generation*, *JHEP* **09** (2007) 028 [[arXiv:0706.2334](#)] [[INSPIRE](#)].
- [18] M.L. Mangano, M. Moretti, F. Piccinini, R. Pittau and A.D. Polosa, *ALPGEN, a generator for hard multiparton processes in hadronic collisions*, *JHEP* **07** (2003) 001 [[hep-ph/0206293](#)] [[INSPIRE](#)].
- [19] T. Gleisberg et al., *Event generation with SHERPA 1.1*, *JHEP* **02** (2009) 007 [[arXiv:0811.4622](#)] [[INSPIRE](#)].
- [20] J. Alwall et al., *Comparative study of various algorithms for the merging of parton showers and matrix elements in hadronic collisions*, *Eur. Phys. J. C* **53** (2008) 473 [[arXiv:0706.2569](#)] [[INSPIRE](#)].
- [21] S. Catani, F. Krauss, R. Kuhn and B.R. Webber, *QCD matrix elements + parton showers*, *JHEP* **11** (2001) 063 [[hep-ph/0109231](#)] [[INSPIRE](#)].
- [22] V. Hirschi et al., *Automation of one-loop QCD corrections*, *JHEP* **05** (2011) 044 [[arXiv:1103.0621](#)] [[INSPIRE](#)].
- [23] S. Hoche, F. Krauss, M. Schonherr and F. Siegert, *Automating the POWHEG method in SHERPA*, *JHEP* **04** (2011) 024 [[arXiv:1008.5399](#)] [[INSPIRE](#)].
- [24] S. Alioli, P. Nason, C. Oleari and E. Re, *A general framework for implementing NLO calculations in shower Monte Carlo programs: the POWHEG BOX*, *JHEP* **06** (2010) 043 [[arXiv:1002.2581](#)] [[INSPIRE](#)].
- [25] P. Nason and B. Webber, *Next-to-leading-order event generators*, *Ann. Rev. Nucl. Part. Sci.* **62** (2012) 187 [[arXiv:1202.1251](#)] [[INSPIRE](#)].

- [26] P. Nason, *A new method for combining NLO QCD with shower Monte Carlo algorithms*, *JHEP* **11** (2004) 040 [[hep-ph/0409146](#)] [[INSPIRE](#)].
- [27] S. Frixione and B.R. Webber, *Matching NLO QCD computations and parton shower simulations*, *JHEP* **06** (2002) 029 [[hep-ph/0204244](#)] [[INSPIRE](#)].
- [28] M. Dasgupta, K. Khelifa-Kerfa, S. Marzani and M. Spannowsky, *On jet mass distributions in Z +jet and dijet processes at the LHC*, *JHEP* **10** (2012) 126 [[arXiv:1207.1640](#)] [[INSPIRE](#)].
- [29] Y.-T. Chien, R. Kelley, M.D. Schwartz and H.X. Zhu, *Resummation of jet mass at hadron colliders*, *Phys. Rev. D* **87** (2013) 014010 [[arXiv:1208.0010](#)] [[INSPIRE](#)].
- [30] T.T. Jouttenus, I.W. Stewart, F.J. Tackmann and W.J. Waalewijn, *Jet mass spectra in Higgs + one jet at NNLL*, *Phys. Rev. D* **88** (2013) 054031 [[arXiv:1302.0846](#)] [[INSPIRE](#)].
- [31] M. Rubin, *Non-global logarithms in filtered jet algorithms*, *JHEP* **05** (2010) 005 [[arXiv:1002.4557](#)] [[INSPIRE](#)].
- [32] I. Feige, M.D. Schwartz, I.W. Stewart and J. Thaler, *Precision jet substructure from boosted event shapes*, *Phys. Rev. Lett.* **109** (2012) 092001 [[arXiv:1204.3898](#)] [[INSPIRE](#)].
- [33] A.J. Larkoski, G.P. Salam and J. Thaler, *Energy correlation functions for jet substructure*, *JHEP* **06** (2013) 108 [[arXiv:1305.0007](#)] [[INSPIRE](#)].
- [34] M. Dasgupta, A. Fregoso, S. Marzani and A. Powling, *Jet substructure with analytical methods*, *Eur. Phys. J. C* **73** (2013) 2623 [[arXiv:1307.0013](#)] [[INSPIRE](#)].
- [35] A.J. Larkoski and J. Thaler, *Unsafe but calculable: ratios of angularities in perturbative QCD*, *JHEP* **09** (2013) 137 [[arXiv:1307.1699](#)] [[INSPIRE](#)].
- [36] M.H. Seymour, *Searches for new particles using cone and cluster jet algorithms: a comparative study*, *Z. Phys. C* **62** (1994) 127 [[INSPIRE](#)].
- [37] D. Krohn, J. Thaler and L.-T. Wang, *Jet trimming*, *JHEP* **02** (2010) 084 [[arXiv:0912.1342](#)] [[INSPIRE](#)].
- [38] 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](#)].
- [39] Y. Bai, A. Katz and B. Tweedie, *Pulling out all the stops: searching for RPV SUSY with stop-jets*, *JHEP* **01** (2014) 040 [[arXiv:1309.6631](#)] [[INSPIRE](#)].
- [40] J.C. Collins, D.E. Soper and G.F. Sterman, *Factorization of hard processes in QCD*, *Adv. Ser. Direct. High Energy Phys.* **5** (1988) 1 [[hep-ph/0409313](#)] [[INSPIRE](#)].
- [41] ATLAS collaboration, *Data-driven estimation of the background to charged Higgs boson searches using hadronically-decaying τ final states in ATLAS*, *ATLAS-CONF-2011-051* (2011).
- [42] CMS collaboration, *Search for anomalous $t\bar{t}$ production in the highly-boosted all-hadronic final state*, *JHEP* **09** (2012) 029 [[arXiv:1204.2488](#)] [[INSPIRE](#)].
- [43] CMS collaboration, *Searches for new physics using the $t\bar{t}$ invariant mass distribution in pp collisions at $\sqrt{s} = 8$ TeV*, *Phys. Rev. Lett.* **111** (2013) 211804 [[arXiv:1309.2030](#)] [[INSPIRE](#)].
- [44] T. Cohen, M. Jankowiak, M. Lisanti, H.K. Lou and J.G. Wacker, work in progress.
- [45] B.W. Silverman, *Density estimation for statistics and data analysis*, Chapman & Hall, London U.K. (1986).

- [46] M. Dasgupta, A. Fregoso, S. Marzani and G.P. Salam, *Towards an understanding of jet substructure*, *JHEP* **09** (2013) 029 [[arXiv:1307.0007](#)] [[INSPIRE](#)].
- [47] D. Scott, *Multivariate density estimation: theory, practice, and visualization*, Wiley, New York U.S.A. (1992).
- [48] B. Efron, *Bootstrap methods: another look at the jackknife*, *Ann. Statist.* **7** (1979) 1.
- [49] P. Hall, *The bootstrap and Edgeworth expansion*, Springer Series in Statistics, Springer (1995).
- [50] J.L. Horowitz, *The bootstrap*, *Handbook of Econometrics* **5** (2001) 3159.
- [51] L.A. Stefanski and R.J. Carroll, *Deconvoluting kernel density estimators*, *Statistics* (1990) 169.
- [52] A. Avetisyan et al., *Methods and results for standard model event generation at $\sqrt{s} = 14$ TeV, 33 TeV and 100 TeV proton colliders (a Snowmass whitepaper)*, [arXiv:1308.1636](#) [[INSPIRE](#)].
- [53] J. Alwall, M. Herquet, F. Maltoni, O. Mattelaer and T. Stelzer, *MadGraph 5: going beyond*, *JHEP* **06** (2011) 128 [[arXiv:1106.0522](#)] [[INSPIRE](#)].
- [54] 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](#)].
- [55] J. Thaler and K. Van Tilburg, *Identifying boosted objects with N -subjettiness*, *JHEP* **03** (2011) 015 [[arXiv:1011.2268](#)] [[INSPIRE](#)].
- [56] J. Thaler and K. Van Tilburg, *Maximizing boosted top identification by minimizing N -subjettiness*, *JHEP* **02** (2012) 093 [[arXiv:1108.2701](#)] [[INSPIRE](#)].
- [57] M. Cacciari, G.P. Salam and G. Soyez, *The anti- k_t jet clustering algorithm*, *JHEP* **04** (2008) 063 [[arXiv:0802.1189](#)] [[INSPIRE](#)].
- [58] ATLAS collaboration, *Search for pair production of massive particles decaying into three quarks with the ATLAS detector in $\sqrt{s} = 7$ TeV pp collisions at the LHC*, *JHEP* **12** (2012) 086 [[arXiv:1210.4813](#)] [[INSPIRE](#)].
- [59] T. Hastie, R. Tibshirani and J. Friedman, *The elements of statistical learning: data mining, inference, and prediction*, 2nd ed., Springer (2008).
- [60] K.S. Cranmer, *Kernel estimation in high-energy physics*, *Comput. Phys. Commun.* **136** (2001) 198 [[hep-ex/0011057](#)] [[INSPIRE](#)].
- [61] R. Cao, A. Cuevas and W.G. Manteiga, *A comparative study of several smoothing methods in density estimation*, *Comput. Stat. Data Anal.* **17** (1994) 153.
- [62] J.W. Cooley and J.W. Tukey, *An algorithm for the machine calculation of complex Fourier series*, *Math. Comput.* **19** (1965) 297.
- [63] M. Frigo and S.G. Johnson, *The design and implementation of FFTW3*, *Proc. IEEE* **93** (2005) 216.
- [64] <http://www.fftw.org>.
- [65] M. Cacciari, G.P. Salam and G. Soyez, *FastJet user manual*, *Eur. Phys. J. C* **72** (2012) 1896 [[arXiv:1111.6097](#)] [[INSPIRE](#)].
- [66] M. Cacciari and G.P. Salam, *Dispelling the N^3 myth for the k_t jet-finder*, *Phys. Lett. B* **641** (2006) 57 [[hep-ph/0512210](#)] [[INSPIRE](#)].