RECEIVED: August 8, 2014

REVISED: October 22, 2014

ACCEPTED: November 25, 2014

PUBLISHED: December 19, 2014

Quantifying the power of multiple event interpretations

Yang-Ting Chien,^a David Farhi,^b David Krohn,^b Andrew Marantan,^b
David Lopez Mateos^b and Matthew Schwartz^b

^aTheory Division T-2, Los Alamos National Laboratory, MS B283,
Los Alamos, NM, 87545 U.S.A.

^bDepartment of Physics, Harvard University,
Oxford Street, Cambridge, MA, 02138 U.S.A.

E-mail: ytchien@lanl.gov, farhi@physics.harvard.edu,
dkrohn@physics.harvard.edu, amarantan@physics.harvard.edu,
david.lopez@cern.ch, schwartz@physics.harvard.edu

ABSTRACT: A number of methods have been proposed recently which exploit multiple highly-correlated interpretations of events, or of jets within an event. For example, Qjets reclusters a jet multiple times and telescoping jets uses multiple cone sizes. Previous work has employed these methods in pseudo-experimental analyses and found that, with a simplified statistical treatment, they give sizable improvements over traditional methods. In this paper, the improvement gain from multiple event interpretations is explored with methods much closer to those used in real experiments. To this end, we derive and study a generalized extended maximum likelihood procedure, and find that using multiple jet radii can provide substantial benefit over a single radius in fitting procedures. Another major concern we address is that multiple event interpretations might be exploiting similar information to that already present in the standard kinematic variables. We perform multivariate analyses (boosted decision trees) on a set of standard kinematic variables, a single observable computed with several different cone sizes, and both sets combined. We find that using multiple radii is still helpful even on top of standard kinematic variables (providing a 12% improvement at low p_T and 20% at high p_T). These results suggest that including multiple event interpretations in a realistic search for Higgs to $b\bar{b}$ would give additional sensitivity over traditional approaches.

KEYWORDS: Jets

ARXIV EPRINT: [1407.2892](https://arxiv.org/abs/1407.2892)

Contents

1	Introduction	1
2	Monte Carlo simulation and event selection	3
3	Significance measures	4
4	Maximum likelihood fits	5
4.1	Merged histogram likelihood fit	6
4.2	Multidimensional likelihood fits	8
4.3	Results	9
5	Comparison to standard observables	11
6	Conclusions	12

1 Introduction

Both the ATLAS and CMS collaborations have recently released their full Run 1 analyses of the search for the $H \rightarrow b\bar{b}$ decay mode [1, 2]. With only Run 1 data, neither search was capable of finding this process at Standard Model rates. With the additional statistics from Run 2 data, $H \rightarrow b\bar{b}$ will surely be observed. However, a precision measurement providing a meaningful extraction of the bottom and top Yukawa couplings with implications for Beyond the Standard Model physics will require these searches to increase their sensitivity beyond what might be achievable with currently used experimental techniques. Most of the proposed improvements involve looking in special kinematic regions where backgrounds are smaller [3] or computing new, physically-motivated observables [4, 5].

In [6], a qualitatively new way to construct observables called “Qjets” was proposed. Instead of comparing a single observable between signal and background, Qjets proposed to look at the sensitivity of an observable to multiple interpretations of that observable, generated by small variations in some parameter. The intuition behind Qjets is as follows. In traditional jet clustering, a jet is constructed as the union of particles most likely to have been radiated from a single hard parton, according to some measure. For example, the Cambridge-Aachen algorithm uses an angular distance measure and the k_T algorithm a relative transverse momentum measure. Naturally, different algorithms can produce slightly different jets. Different parameters, like jet size R , also lead to slightly different jets. For very clean events, with well-separated clusters of energy, all the algorithms will agree. However, in realistic situations, there can be large variation. Qjets exploits this variation by accepting that the optimal set of particles, according to some measure, may not be that much better than slightly sub-optimal sets. By combining the many different but reasonable interpretations of an event, a more robust picture of an event emerges, and consequently signal and background can be more effectively distinguished.

In the original Qjets proposal, the multiple interpretations were generated by perturbing around the original jet clustering algorithm: instead of always merging the two closest

particles during jet clustering, the Qjets algorithm considers merging more distant pairs. This generates a distribution of highly-correlated observables for each jet in each event. The width of this distribution for pruned jet mass [7], which was called *volatility* in [6], provides a strong signal-to-background discriminant in boosted Higgs or boosted W boson searches. Volatility has been measured in experiment [8, 9] with similar results and discrimination power to simulation. An application of the Qjets method to event reconstruction was proposed in [10]. In [11], a simpler and faster way of using multiple event interpretations was proposed: simply compute the same observable using different values of the jet size R . Both [10] and [11] computed the reach of the $H \rightarrow b\bar{b}$ search by combining the multiple interpretations, finding as much as a 46% improvement in significance over using a single interpretation.

The method used to estimate significance improvement in [6, 10, 11] is a natural generalization of cut-and-count for a single observable. Normally, an event either passes a set of cuts ($z = 1$) or does not ($z = 0$). With multiple interpretations, a fraction z of the interpretations pass the cuts. This fraction z is an observable, measurable in data and computable with Monte Carlo. In [6, 10, 11], the 1-dimensional distributions of z for signal and background were used to estimate the probability that a given set of events could be explained by a fluctuation of the background only. The procedure is reviewed below in section 3. Using this method, multiple event interpretations were shown in [6, 10, 11] to give significant improvement to search reaches.

One drawback of the method used in [6, 10, 11] is that it presupposes a knowledge of the background cross section. Many LHC analyses try to avoid taking cross sections from theory. Instead, they often use control regions to establish background normalizations, which are not necessarily precisely known in the specific regions of phase space exploited by the analysis. These control regions are typically defined to have minimal overlap with the signal regions. When using multiple event interpretations, however, it is generally not possible to define non-overlapping regions, since a single event can potentially cover a large range of values for the observable of interest. One goal of this paper is to generalize the extended maximum likelihood procedure, which fits to signal and background cross sections separately, to observables based on multiple event interpretations.

An important question about multiple event interpretations is whether the improvement is due to features of the events which are already accounted for in the current search strategies. For example, most analyses (including the $H \rightarrow b\bar{b}$ searches) use many kinematic variables in the event to maximize significance, often combined with sophisticated multivariate methods like neural networks and boosted decision trees. It has not been proven so far that the improvement observed when using multiple event interpretations is independent of what can be obtained by exploiting the kinematic features of the event exploited by multivariate discriminants. We also address this concern, by showing that multiple event interpretations can indeed improve the significance of a search when combined with kinematic variables.

This paper is organized as follows. Section 2 describes the simulation and event selection we used. Section 3 gives a quick introduction to the statistical methods we discuss and their relative merits. In section 4 we first review the typical extended maximum likelihood (EML) fit, emphasizing those aspects which break down when using multiple

Improvement (over $R=0.5$ only)	$p_T^Z < 120$ GeV		$p_T^Z > 120$ GeV	
	xs-based	EML	xs-based	EML
1 R , fraction in window	0.83	0.74	0.80	0.71
12 R 's, fraction in window	0.98	0.97	0.92	0.88
1 R , $m_{b\bar{b}}$	1.00	1.00	1.00	1.00
5 R 's, $m_{b\bar{b}}$ merged	0.94	1.08	0.94	1.06
4 R 's, 3 bins	0.99	1.00	1.16	1.20
2 R 's, full	1.10	1.14	1.35	1.38
2 R 's, BDT	1.04	1.08	1.30	1.34
12 R 's, BDT	1.19	1.30	1.52	1.41
12 kinematic	1.33	1.50	1.35	1.29
12 kinematic + 12 R 's	1.39	1.68	1.67	1.55

Table 1. Relative significances are shown for various collections of jet sizes and the different methods discussed in the text. This table is explained in section 4.3.

event interpretations. Then we derive the modifications in the EML formalism necessary to account for the statistical correlations among the different interpretations of the same event. Section 4.2 describes a two-dimensional extension of the likelihood fit that avoids modifications of the EML fit at the expense of adding complexity, and compares the performance of this extension to the results of the previous section. Section 5 compares the performance of multivariate analyses including kinematic information and multiple event interpretations, to understand whether multiple event interpretations indirectly make use of kinematic information already exploited in current LHC analyses. We summarize the results from all these methods in table 1.

2 Monte Carlo simulation and event selection

We generate signal and background processes for proton-proton collisions at $\sqrt{s} = 8$ TeV using MadGraph 5.1 [12], interfaced to Pythia 6.4 [13] to simulate the parton shower and non-perturbative effects such as underlying event and hadronization. $ZH \rightarrow e^+e^-b\bar{b}$ events are generated and used for the signal, and $Z(\rightarrow e^+e^-) + b\bar{b}$ events are used as background. High statistics are produced for these samples, but in quoting expected significances, the signal and backgrounds are normalized to 25 fb^{-1} . This normalization is also used for generating toy models used for estimating the uncertainties in the likelihood fits. Jets are clustered from stable particles with lifetimes above 10 ps (excluding neutrinos) using the anti- k_t algorithm with different R parameters. Unlike in [10], these interpretations are built using a deterministic method, similar to the telescoping jets approach introduced in [11]. The event selection is a simplified version of that in [1] and requires $83 \text{ GeV} < m_{ee} < 99 \text{ GeV}$, $p_T(b_{\text{lead}}) > 45 \text{ GeV}$, $p_T(b_{\text{sublead}}) > 25 \text{ GeV}$ and $p_T^{\text{lep}} > 25 \text{ GeV}$. A jet is defined as a b -jet if it contains any decay products from the original b -quark. Studies are

performed in two kinematic regions defined by the transverse momentum of the vector boson ($p_T^Z < 120$ GeV and $p_T^Z > 120$ GeV). The invariant mass distribution of the two leading jets in p_T that are labeled as b -quarks is used for our significance estimates.

3 Significance measures

In this section, we quickly review the essential differences between a cross section based significance calculation, like the ones used in [6, 10, 11], and the extended likelihood method.

The approach used in [6, 10, 11] starts with an observable z , defined as the fraction of event interpretations satisfying a given set of cuts. In a normal analysis, z would be either 1 (event passes cuts) or 0 (event fails cuts). The insight of Qjets was that one can use multiple event interpretations to make z a rational number. One can then compute in Monte Carlo the probabilities $\rho_S(z)$ and $\rho_B(z)$ for finding certain z 's (see figure 3 of [10] for some $\rho(z)$ distributions). Then the probability that the data cannot be accounted for by a fluctuation of the background (measured in standard deviations of the signal away from the background) is given by

$$\text{significance} = \frac{N_{\text{data}} - N_{\text{expected}}^B}{\delta N_{\text{expected}}^B} = \frac{\int_0^1 dz \frac{\rho_S(z)}{\rho_B(z)} [\rho_{\text{data}}(z) - \rho_B(z)]}{\sqrt{\int_0^1 dz \frac{\rho_S(z)^2}{\rho_B(z)}}} \quad (3.1)$$

where ρ_{data} is the observed probability. See section 5 of [10] for a derivation of this formula. In a simulation, we replace data-minus-expected-background in the numerator by signal: $\rho_{\text{data}}(z) - \rho_B(z) \rightarrow \rho_S(z)$. For terminological clarity, we call the estimate of significance using this method the *cross section (xs)-based significance*. It is cross section based since one needs to know how much background there should be in order to see a fluctuation above this amount.

The xs-based significance can be applied to any observable, not just this z variable. It corresponds to an analysis which provides a weight function $w(z) = \rho_S(z)/\rho_B(z)$, and measures a final observable N which is the sum of the weight of all observed events. In particular, it can be applied to multidimensional data, using $w(z_1, z_2, \dots)$ if the multidimensional distributions $\rho_S(z_1, z_2, \dots)$ and $\rho_B(z_1, z_2, \dots)$ are known. We show the relative significance using various choices of z_i in table 1.

Many LHC analyses are instead based on likelihood fits. In a likelihood fit, the distribution of some observable z in the data is used to fit the signal and background cross sections:

$$\rho_{\text{data}}(z) = (\sigma_S \pm \delta_S) \rho_S(z) + (\sigma_B \pm \delta_B) \rho_B(z). \quad (3.2)$$

While the xs-based significance is computed using the normalized probability distributions for signal and background and the expected background cross section, likelihood fits extract both signal and background cross sections directly from data. A likelihood fit which takes into account the Poisson fluctuations of signal and background is called an *extended maximum likelihood (EML) fit*; in the rest of this paper all references to “likelihood” fits include this extension as appropriate. In section 4 we show how extended maximum likelihood can be computed from events with multiple interpretations.

In likelihood fits the ultimate observable of interest is σ_S . For a discovery, we need the probability that $\sigma_S > 0$. The analog of the significance in eq. (3.1) is the expected value of σ_S divided by the variance in the measured value of σ_S on data sets with no signal. In other words:

$$\text{significance} = \frac{\langle \sigma_S \rangle_{\text{signal+background}}}{\langle \delta \sigma_S \rangle_{\text{background only}}}. \quad (3.3)$$

Note that while eqs. (3.1) and (3.3) are both measures of the probability that a signal is there, the two methods have different priors so they cannot be directly compared. In the xs-based significance, the analysis knows the expected background rate, and counts any difference in the data from the background as observed signal. For an EML significance, both the signal rate and the background rate are fit. Since the xs-based analysis has access to more information it will often give higher significances. The likelihood fit is procedurally closer to experimental conditions, where backgrounds are never known with perfect accuracy and are estimated from sidebands.

The difference in priors between xs-based and EML can be illustrated by a simple example. Suppose we had only a single bin, and so the entirety of the data is the fact that N events were measured. The xs-based approach would report an observed value N , subtract the expected N_B (an input from theory), and declare an observed signal rate of $N_S = N - N_B$, with a corresponding significance $N_S/\sqrt{N_B}$. The likelihood fit, however, would simply fail, for it would be attempting to fit $N = \sigma_S + \sigma_B$ for both σ_S and σ_B , and a single bin *cannot* fit two parameters. Thus the error would be infinite, and the significance would not be meaningful.

4 Maximum likelihood fits

Many methods exist to find the cross section fits σ_S and σ_B in eq. (3.2) by maximizing the likelihood of the observed distributions. In this section we explore how to apply such methods to multiple-interpretation data. When using multiple interpretations, each event $k = 1 \cdots K$ gives us not a single number x_k , as is the usual case, but a series of numbers x_k^i , with $i = 1 \cdots I$ indexing the interpretations. We consider two approaches to constructing a likelihood fit from these observables x_k^i , which we call the merged histogram likelihood fit and the multidimensional likelihood fit.

For the *merged histogram fit*, we combine all of the interpretations into a single one-dimensional distribution, as though we had not K events but $K \times I$ events. Then we can apply the usual fitting technology to this distribution, as we might for, say, an invariant mass distribution in a conventional analysis. This approach will converge upon the correct fit values in the limit of infinite statistics. However, because the distribution contains both highly-correlated contributions (from a single event) and statistically-uncorrelated ones (from different events), the existing technology to estimate the errors in the fit parameters will give vastly incorrect errors. We show how to correctly estimate the errors with this method in section 4.1.

In the *multidimensional likelihood fit*, we treat the x_k^i as I observables in K events. This is possible when the multiple interpretations are indexed, as in telescoping jets where

they correspond to different R , but not when they are generated by adding randomness to a jet algorithm, as in the original Qjets proposal. If the I interpretations are distinguishable, we can then do a fit to the K data points in an I -dimensional space. This is no different from a regular fit to multi-dimensional data, and so the usual fitting technology can be applied. As in other multidimensional fits, the high-dimensionality of the space will quickly saturate the statistics either of data or simulation. To control for this, we consider three approaches. We first try severely limiting the number of interpretations, down to $I = 2$. This allows us to map out the full I -dimensional distribution over events and produce a nearly exact maximum likelihood solution. Second, we try limiting the number of bins, but taking I larger. Third, we try using multivariate methods, in particular boosted decision trees, in place of the exact likelihood. The multidimensional approaches are explored in section 4.2. Numerical results are summarized in table 1 and discussed in section 4.3.

4.1 Merged histogram likelihood fit

In this section we derive appropriate formulae for uncertainty estimation when highly correlated data from multiple event interpretations are merged into a single distribution. We first review general features of how likelihood fits are done, and then discuss how things are modified with multiple event interpretations.

In an extended maximum likelihood (EML) fit, a set of unknown parameters $\{a_\alpha\}$ are estimated by maximizing the likelihood density over a sample of N data values, x_n :

$$\mathcal{L} = \left[\prod_n^N P(x_n; \{a_\alpha\}) \right] e^{-\mathcal{N}(\{a_\alpha\})}, \quad (4.1)$$

where $P(x; \{a_\alpha\})$ is the probability density for the data value x when $\mathcal{N}$ data points are expected. It should be noted that N , the actual number of observed events, is generally not equal to $\mathcal{N}$, due to Poisson fluctuations ($\mathcal{N}$ can take non-integer values). Since the normalization has to be constrained to the number of expected events, the following is true:

$$\int P(x; \{a_\alpha\}) dx = \mathcal{N}(\{a_\alpha\}), \quad (4.2)$$

where P is taken as a continuous probability density function in this equation and not evaluated at each data point. This shows one way that the EML fit differs from the standard maximum likelihood fit, for which this normalization is 1. The EML thus allows to fit for the shape of a distribution as well as for the normalization. The EML is thus widely used in searches, where the total number of events collected is subject to Poisson fluctuations.

The difference between the standard maximum likelihood and the EML lies on the normalization condition, represented by the $e^{-\mathcal{N}}$ term [14], which guarantees that the normalization of $P(x_n; \{a_\alpha\})$ does not increase arbitrarily in the maximization procedure, and that that normalization obeys Poisson statistics for $\mathcal{N}$ expected events. The log-likelihood for this function takes the form

$$\ln \mathcal{L} = \sum_{n=1}^N \ln P(x_n; \{a_\alpha\}) - \mathcal{N}(\{a_\alpha\}). \quad (4.3)$$

The fitted parameters are then taken to be the set $\{\hat{a}_\alpha\}$ which maximize this likelihood:

$$0 = \frac{d \ln \mathcal{L}}{da_\alpha} (\{\hat{a}_\alpha\}). \quad (4.4)$$

The errors on these parameters come from incoherent Poisson fluctuations on the number of observed events in each bin, $\delta n(x)$. Thus the total error on each parameter, δa_α , is given by the quadrature sum of the errors due to each bin fluctuation:

$$\delta a_\alpha^2 = \sum_{\text{bin } x} [\delta a_\alpha (\delta n(x))]^2. \quad (4.5)$$

The error due to fluctuations of a single bin can be computed by expanding the minimization equation:

$$0 = \frac{d \ln \mathcal{L}}{da_\alpha} (n(x) + \delta n, \hat{a} + \delta \hat{a}) = \frac{d \ln \mathcal{L}}{dn(x) da_\alpha} \delta n(x) + \sum_\beta \frac{d \ln \mathcal{L}}{da_\alpha da_\beta} \delta \hat{a}_\beta. \quad (4.6)$$

$d \ln \mathcal{L} / dn(x)$ is just the contribution to $\ln \mathcal{L}$ from a single event in bin x , which is given by $\ln P$. From this we can compute the total error on parameters a_α :

$$\delta a_\alpha^2 = \sum_{\text{bin } x} \left(\delta n(x) \sum_\beta M_{\alpha\beta}^{-1} \frac{d \ln P}{da_\beta} \right)^2. \quad (4.7)$$

Where $M_{\alpha\beta}$ is the second derivative matrix of $\ln \mathcal{L}$. Assuming $n(x)$ has Poisson errors $\delta n(x) = \sqrt{n(x)}$, we get a final formula for the errors:

$$\delta a_\alpha^2 = \sum_{\text{event } n} \sum_{\beta, \gamma} M_{\alpha\beta}^{-1} \frac{d \ln P(x_n)}{da_\beta} M_{\alpha\gamma}^{-1} \frac{d \ln P(x_n)}{da_\gamma} = N (M^{-1} C M^{-1})_{\alpha\alpha}. \quad (4.8)$$

Where C is the covariance matrix: $C_{\alpha\beta} \equiv \langle \partial_{a_\alpha} \ln P \partial_{a_\beta} \ln P \rangle$.

In the simple case of only a single a , this reduces to:

$$\delta \hat{a} = \frac{\sqrt{\sum_n \left(\frac{\partial}{\partial a} \ln P(x_n; \hat{a}) \right)^2}}{\frac{\partial^2 \ln \mathcal{L}(\hat{a})}{\partial a^2}}. \quad (4.9)$$

This formalism can be extended to understand the use of the EML for multiple event interpretations. For multiple event interpretations, correlations between points can be resolved by expressing the sum over K events as a sum over the multidimensional parameter space of $\vec{x} = \{x^1, x^2, \dots, x^I\}$, where the new index refers to each event interpretation. For the analysis where all interpretations are combined in the same x distribution, $P(x) = \sum_i P_i(x)$ and $P_i(x)$ is the projection of the multidimensional distribution $P(\vec{x})$ along the interpretation i .

In this approach the likelihood function to be maximized is:

$$\ln \mathcal{L} = \sum_{k=1}^K \sum_{i=1}^I \ln P(x_k^i; \{a_\alpha\}) - I N(\{a_\alpha\}). \quad (4.10)$$

Computing the errors from this point of view is complicated because the x_m^i are correlated. But we can recast it like this:

$$\ln \mathcal{L} = I \left(\sum_{k=1}^K \ln \tilde{P}(\vec{x}_k; \{a_\alpha\}) - \mathcal{N}(\{a_\alpha\}) \right). \quad (4.11)$$

Where $\ln \tilde{P}(\vec{x}) \equiv \frac{1}{I} \sum_i \ln P(x^i)$. Now eq. (4.11) looks exactly like eq. (4.3), and the derivation of the errors is exactly the same, except using $\tilde{P}$. Thus we have:

$$\delta a_\alpha^2 = K (M^{-1} C M^{-1})_{\alpha\alpha} \quad (4.12)$$

$$M_{\alpha\beta} = \frac{d \ln \mathcal{L} / I}{da_\alpha da_\beta} \quad (4.13)$$

$$C_{\alpha\beta} = \frac{1}{I^2} \left\langle \frac{d \sum_i \ln P(x^i)}{da_\alpha} \frac{d \sum_i \ln P(x^i)}{da_\beta} \right\rangle. \quad (4.14)$$

And in the case of a single fit parameter a :

$$\delta \hat{a} = \frac{\sqrt{\sum_k \left(\frac{\partial}{\partial a} \sum_i \ln P(x_k^i; \hat{a}) \right)^2}}{\frac{\partial^2 \ln \mathcal{L}(\hat{a})}{\partial a^2}}. \quad (4.15)$$

In the case of interest, we have a signal distribution $\rho_S(x)$ and a background distribution $\rho_B(x)$ (each normalized so $\int \rho(x) dx$ is the number of interpretations expected to pass the cuts). We are fitting to a predicted distribution $P = N_S \rho_S + N_B \rho_B$, so we have two parameters N_S and N_B . The likelihood is:

$$\ln \mathcal{L} = \sum_{k=1}^K \sum_{i=1}^I \ln(N_S \rho_S(x_k^i) + N_B \rho_B(x_k^i)) - R(N_S + N_B). \quad (4.16)$$

And then the errors are given by:

$$\delta N_\alpha^2 = N (M^{-1} C M^{-1})_{\alpha\alpha} \quad (4.17)$$

$$M_{\alpha\beta} = \sum_{k=1}^K \frac{1}{I} \sum_{i=1}^I \frac{-\rho_\alpha(x_k^i) \rho_\beta(x_k^i)}{(N_S \rho_S(x_k^i) + N_B \rho_B(x_k^i))^2} \quad (4.18)$$

$$C_{\alpha\beta} = \left\langle \left(\frac{1}{I} \sum_{i=1}^I \frac{\rho_\alpha(x_k^i)}{N_S \rho_S(x_k^i) + N_B \rho_B(x_k^i)} \right) \left(\frac{1}{I} \sum_{i=1}^I \frac{\rho_\beta(x_k^i)}{N_S \rho_S(x_k^i) + N_B \rho_B(x_k^i)} \right) \right\rangle. \quad (4.19)$$

Using eqs. (4.17)–(4.19) and (3.3), the improvement in the significance of the search can be estimated for an analysis using a likelihood fit.

4.2 Multidimensional likelihood fits

Rather than merging all the interpretations into a single distribution, it would clearly be smarter to keep the interpretations separate and exploit their correlations. Unfortunately, computing the exact likelihood from an I -dimensional space would require an exponentially large data set. For example, with $I = 10$ interpretations, and only $B = 5$ bins in each direction, we would need 10 million events just to have each bin populated with 1 event. This is of course just the usual curse of dimensionality for multivariate fits, which is present in any analysis with correlated observables. A popular approach is to replace

Figure 1. Two-dimensional distributions of invariant masses reconstructed with the anti- k_t algorithm with $R = 0.5$ and $R = 1.2$ for events with $p_T^Z > 120$ GeV for the signal ((a)) and background ((b)). Colors indicate relative cross sections, with darker colors being more probable. Normalization is arbitrary. Since it is very unusual for the jet with $R = 0.5$ to have larger mass than the jet with $R = 1.2$, there are essentially no statistics in the lower-right halves of these plots.

the exact likelihood with boosted decision trees (BDTs), neural networks, or other sophisticated algorithms. Another approach is just to take I and B small enough so that the dimensionality is not intractable. In this section, we compare these alternatives.

If we take $I = 2$, then it is possible to populate the 2 dimensional space quite well. For example, the 2D distributions for signal and background $m_{b\bar{b}}$ distributions for two different choices of R are shown in figure 1. The figure shows the low p_T^Z sample clustered with anti- k_t using $R = 0.5$ and $R = 1.2$. The results at high p_T^Z are qualitatively comparable except for an overall shift at higher masses for the background. The two invariant masses are clearly correlated, but also clearly not 100% correlated. To quantify the correlation, the linear correlation coefficients among some representative R values are shown in figure 2. A picture emerges in which, particularly at high p_T^Z , the correlation between small R and large R interpretations is quite different for signal and background, with smaller correlations in the signal. The 2-dimensional data can be run through a regular likelihood fitting procedure, with no modification since each event is only contributing a single point. We show results in table 1.

As an alternative to taking $I = 2$, we can produce a statistically tractable fitting problem by instead reducing the number of bins B . For instance, we consider $B = 3$ and $I = 4$, where we bin every interpretation into $m_{b\bar{b}} < 110$ GeV, $110 \text{ GeV} < m_{b\bar{b}} < 140$ GeV, and $140 \text{ GeV} < m_{b\bar{b}}$. Results for this approach are also shown in table 1.

4.3 Results

Table 1 summarizes the results obtained with the different methods described in previous sections. The baseline is the third line of the table; for the xs-based results it uses $\rho(m_{b\bar{b}})$ (see section 3) and for the EML-based result uses a fit to the $m_{b\bar{b}}$ distribution. The first two lines use the fraction of interpretations in a window of $110 \text{ GeV} < m_{b\bar{b}} < 140$ GeV to

Figure 2. Linear correlation coefficients for $m_{b\bar{b}}$ among different R values for signal ((a), (c)) and background ((b), (d)), with $p_T^Z < 120$ GeV ((a), (b)) and $p_T^Z > 120$ GeV on bottom ((c), (d)).

define z as in [11]. These two rows list the xs-based significance and the EML significance, both evaluated on the observable z given by the fraction of interpretations that lie in the window. For 1 R , z is either 0 or 1, and this is equivalent to a simple cut-and-count experiment. The use of the full $m_{b\bar{b}}$ distribution (from 50 GeV to 200 GeV with an overflow bin on each side), even with just 1 R (the third row) is more powerful than the use of multiple event interpretations for the simple cut-and-count model (the second row). When using 5 R 's ($R = 0.4, 0.5, 0.6, 0.7$ and $R = 0.8$) and pooling all interpretations in the same $m_{b\bar{b}}$ distribution (the row labeled “merged”), there is a small gain when using the likelihood fit, despite some features of the distribution being washed out in the procedure.

The next row shows the results when using 3 $m_{b\bar{b}}$ bins ($m_{b\bar{b}} < 110$ GeV, 110 GeV $< m_{b\bar{b}} < 140$ GeV and $m_{b\bar{b}} > 140$ GeV), to reduce the dimensionality, and 4 R 's ($R = 0.4, 0.6, 0.8, 1.0$). Using 3 bins provides non-negligible gains at high p_T , but does not manage to perform better than a single interpretation at low p_T . This indicates that at low p_T the added bins do not help tell apart the different interpretations but that a p_T -dependent choice of binning might be worth further exploration.

The method that gives the second highest significance is the one using the full 2D distributions of $m_{b\bar{b}}$ computed with $R = 0.5$ and $R = 1.2$. At high p_T , we find 35% and 38% improvements with the xs-based and EML fit respectively. As long as the two radii are far enough apart that the correlations between the reconstructed invariant masses

is not very high, the choice of the two radii does not impact significantly the observed improvement. The use of boosted decision trees¹ to combine the two radii does not do quite as well as the full 2D $m_{b\bar{b}}$ distribution, suggesting that there is a loss of information in the construction of the BDT, at least in our implementation (the TMVA default).

The improvement in both the xs-based and the EML-based approaches is highest using the most radii which we combine using BDTs. We find up to a 52% improvement for xs-based or 41% for EML based in the high p_T sample by combining 12 R 's over using just a single R . Two additional entries in the table refer to the use of kinematic variables in the BDTs and are discussed in section 5.

5 Comparison to standard observables

So far, we have seen that combining measurements of $m_{b\bar{b}}$ computed by clustering with the anti- k_t algorithm with different R 's can have a sizable improvement over using $m_{b\bar{b}}$ with a single R . It is natural to wonder whether the improvement is due to the exploitation of properties of the events which could be exploited equally well using more traditional kinematic variables. For example, as R increases, the jet momentum increases. Since the signal and background have different momentum dependence, one might imagine that the same gain could be realized simply by including the p_T of the $b\bar{b}$ system into the analysis. To explore whether multiple R values leads to improvement in the $HZ \rightarrow b\bar{b}\ell^+\ell^-$ search, we compare the improvement using multiple event interpretations (multiple R 's) to the improvement from standard kinematic observables.

For the kinematic variables, we take the set from the HZ LHC study in [4], table 4. All observables are computed using jets reconstructed with the anti- k_t algorithm with $R = 0.7$. The observables we consider are constructed from either the hadrons (b -jets):

$$m_{b\bar{b}}, \quad p_T^{b_1}, \quad p_T^{b_2}, \quad |p_T^{b_1}| + |p_T^{b_2}|, \quad |\eta^{b_1} - \eta^{b_2}|, \quad |\eta^h - \eta^{b_2}|, \quad \text{and} \quad |p_T^h - p_T^{b_2}| \quad (5.1)$$

or the leptons (from the Z decay):

$$p_T^{\ell_1}, \quad p_T^{\ell_2}, \quad p_T^Z \quad \text{and} \quad m_{\ell\ell} \quad (5.2)$$

and one variable dependent on the leptons and b -jets:

$$|\eta^h - \eta^Z|. \quad (5.3)$$

In these expressions, η is rapidity, p_T is transverse momentum, b_1 refers to the harder of the two b -jets and b_2 to the softer, and similarly for the leptons.

These 12 variables are all input to a BDT and used to compute the significance just as we have for multiple event interpretations. We also try combining these kinematic variables with the 11 remaining masses in our study (computed using $R = 0.4$ to $R = 1.5$ in steps of 0.1). Results are shown on the last two lines of table 1. We can draw a few interesting qualitative conclusions from this analysis. First, we see that the kinematic variables work

¹We use the default BDT parameters in the TMVA package. The significance is computed using the probability-distribution-functions produced from TMVA from the BDT classifier [15].

relatively better at low p_T than high p_T . At high p_T angular differences are smaller and there are, thus, fewer handles to distinguish signal from background. This is unlike the multiple event interpretations, for which improvements are more significant at high p_T . Second, we see that using multiple event interpretations (multiple R 's) still gives serious benefit on top of all of the kinematic variables. The improvement is more significant at high p_T as expected from the previous observation.

6 Conclusions

In previous work [6, 10, 11] multiple event interpretations, in particular the reclustering of an event using different jet sizes, were shown to give sizable improvement in the potential significance for a $H \rightarrow b\bar{b}$ search in association with a W or Z boson. In this paper we have attempted to refine those analyses using methodologies as close as possible to those used in experimental analyses at the LHC. For this purpose, a new expression of the likelihood has been developed to account for correlations across events populating several bins in one dimension. The improvement in the significance of the search with this treatment has been shown to be sizable. We find as much as $\approx 41\%$ improvement when using 12 R 's over a single R when the $b\bar{b}$ system has $p_T > 120$ GeV and $\approx 30\%$ improvement for $p_T < 120$ GeV.

We also explored whether the improvement from multiple event interpretations carries overlapping information to traditional kinematic variables or complementary information. To answer this question, we took 12 kinematic variables that have been demonstrated to be nearly optimal in a multivariate analysis [4] (and some of which were used in a recent CMS search) for $H \rightarrow b\bar{b}$ and compared their efficacy to what we get from just using multiple R 's and what we get by combining them. We find that adding the m_{bb} at multiple interpretations gives a 12% improvement at low p_T and 20% improvement at high p_T over the kinematic variables at a single R . These improvements are particularly encouraging, since the phase space explored in this paper does not include boosted topologies and thus cannot benefit from otherwise highly successful jet substructure techniques [3, 16, 17]. This phase space is, however, quite relevant for finding the $H \rightarrow b\bar{b}$ decay. In summary, these results strongly suggest that multiple interpretations can help in searches with realistic statistical methods.

Acknowledgments

MDS and DF are supported in part by the DOE under grant DE-SC003916. YTC is supported by the US Department of Energy, Office of Science. AM is supported in part by a DoD NDSEG fellowship.

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

References

- [1] ATLAS collaboration, *Search for the bb decay of the Standard Model Higgs boson in associated W/ZH production with the ATLAS detector*, CERN Geneva, [ATLAS-CONF-2013-079](#) (2013).
- [2] CMS collaboration, *Search for the standard model Higgs boson produced in association with W or Z bosons, and decaying to bottom quarks for LHCp 2013*, CERN Geneva, [CMS-PAS-HIG-13-012](#) (2013).
- [3] J.M. Butterworth, A.R. Davison, M. Rubin and G.P. Salam, *Jet substructure as a new Higgs search channel at the LHC*, *Phys. Rev. Lett.* **100** (2008) 242001 [[arXiv:0802.2470](#)] [[INSPIRE](#)].
- [4] J. Gallicchio, J. Huth, M. Kagan, M.D. Schwartz, K. Black and B. Tweedie, *Multivariate discrimination and the Higgs + W/Z search*, *JHEP* **04** (2011) 069 [[arXiv:1010.3698](#)] [[INSPIRE](#)].
- [5] J. Gallicchio and M.D. Schwartz, *Seeing in Color: Jet Superstructure*, *Phys. Rev. Lett.* **105** (2010) 022001 [[arXiv:1001.5027](#)] [[INSPIRE](#)].
- [6] S.D. Ellis, A. Hornig, T.S. Roy, D. Krohn and M.D. Schwartz, *Qjets: A Non-Deterministic Approach to Tree-Based Jet Substructure*, *Phys. Rev. Lett.* **108** (2012) 182003 [[arXiv:1201.1914](#)] [[INSPIRE](#)].
- [7] 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](#)].
- [8] ATLAS collaboration, *Performance and Validation of Q -Jets at the ATLAS Detector in pp Collisions at $\sqrt{s} = 8$ TeV in 2012*, CERN Geneva, [ATLAS-CONF-2013-087](#) (2013).
- [9] CMS collaboration, *Identifying Hadronically Decaying Vector Bosons Merged into a Single Jet*, CERN Geneva, [CMS-PAS-JME-13-006](#) (2013).
- [10] D. Kahawala, D. Krohn and M.D. Schwartz, *Jet Sampling: Improving Event Reconstruction through Multiple Interpretations*, *JHEP* **06** (2013) 006 [[arXiv:1304.2394](#)] [[INSPIRE](#)].
- [11] Y.-T. Chien, *Telescoping Jets: Multiple Event Interpretations with Multiple R 's*, *Phys. Rev. D* **90** (2014) 054008 [[arXiv:1304.5240](#)] [[INSPIRE](#)].
- [12] J. Alwall, M. Herquet, F. Maltoni, O. Mattelaer and T. Stelzer, *MadGraph 5: Going Beyond*, *JHEP* **06** (2011) 128 [[arXiv:1106.0522](#)] [[INSPIRE](#)].
- [13] T. Sjöstrand, S. Mrenna and P.Z. Skands, *PYTHIA 6.4 Physics and Manual*, *JHEP* **05** (2006) 026 [[hep-ph/0603175](#)] [[INSPIRE](#)].
- [14] R.J. Barlow, *Extended maximum likelihood*, *Nucl. Instrum. Meth. A* **297** (1990) 496 [[INSPIRE](#)].
- [15] TMVA CORE DEVELOPER TEAM collaboration, J. Therhaag, *TMVA: Toolkit for multivariate data analysis*, *AIP Conf. Proc.* **1504** (2009) 1013 [[INSPIRE](#)].
- [16] 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](#)] [[INSPIRE](#)].
- [17] Y. Cui, Z. Han and M.D. Schwartz, *Top condensation as a motivated explanation of the top forward-backward asymmetry*, *JHEP* **07** (2011) 127 [[arXiv:1106.3086](#)] [[INSPIRE](#)].