

Lawrence Berkeley National Laboratory

LBL Publications

Title

Tagging partially reconstructed objects with jet substructure

Permalink

<https://escholarship.org/uc/item/9n5527pz>

Authors

Freytsis, Marat
Volansky, Tomer
Walsh, Jonathan R

Publication Date

2017-06-01

DOI

10.1016/j.physletb.2016.08.044

Peer reviewed



Tagging partially reconstructed objects with jet substructure



Marat Freytsis^{a,*}, Tomer Volansky^b, Jonathan R. Walsh^{c,d}

^a Department of Physics, Harvard University, Cambridge, MA, 02138, USA

^b Raymond and Beverly Sackler School of Physics and Astronomy, Tel-Aviv University, Tel-Aviv 69978, Israel

^c Ernest Orlando Lawrence Berkeley National Laboratory, University of California, Berkeley, CA 94720, USA

^d Berkeley Center for Theoretical Physics, University of California, Berkeley, CA 94720, USA

ARTICLE INFO

Article history:

Received 2 November 2015
 Received in revised form 14 July 2016
 Accepted 3 August 2016
 Available online 24 August 2016
 Editor: B. Grinstein

Keywords:

QCD jets
 Hadronic top decays
 Jet substructure
 Top tagging

ABSTRACT

We present a new tagger which aims at identifying partially reconstructed objects, in which only some of the constituents are collected in a single jet. As an example, we focus on top decays in which either part of the hadronically decaying W or the b jet is soft or falls outside of the top jet cone. We construct an observable to identify remnant substructure from the decay and employ aggressive jet grooming to reject QCD backgrounds. The tagger is complementary to existing ones and works well in the intermediate boost regime where jet substructure techniques usually fail. It is anticipated that a similar tagger can be used to identify non-QCD hadronic jets, such as those expected from hidden valleys.

© 2016 The Authors. Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>). Funded by SCOAP³.

1. Introduction

The successful discovery of the Higgs boson at the Large Hadron Collider (LHC) [1,2] together with null results in new physics searches stresses the need to improve and develop more sophisticated techniques to search for rare or exotic phenomena. Considerations related to the hierarchy problem typically predict new physics which couples to the top quark and which is often charged under QCD. As a consequence, such new physics may be buried under immense hadronic background.

Interesting classes of models that may escape detection are ones that admit partially reconstructed objects. These occur when, for example, constituents from the decay or shower of a heavy particle are soft, not confined to a single jet, or escape detection altogether. Two examples are hadronic tops in which either part of the hadronically decaying W or the b -quark falls outside of the top-jet cone, and hidden valleys [3,4], in which some of the hidden-sector particles are stable and escape detection while the others decay to hadrons. In both cases, some of the information on the hard parton which initiates the jet is missing and the jet no longer has an obvious hard scale to distinguish it from QCD. As a consequence it is hard to identify such events and with existing tools new physics of this kind will go unnoticed.

To improve on the situation, a better understanding of the substructure of such jets is needed. In recent years, significant progress has been made in developing tools for studying the substructure of hadronic events in order to disentangle new physics signals from background (see Ref. [5–7] and references therein). Most existing techniques are effective when used on boosted objects, and are therefore useful at discovering heavy new physics particles that decay hadronically. On the other hand, partially reconstructed jets are best studied in the intermediate regime of moderate boost. The production rates in this regime are often enhanced by many orders of magnitude and hence may play a crucial role in discovering new physics. The available tools, however, have significantly reduced sensitivity in this region of parameter space [8] despite its discovery potential. This letter focuses on this moderate-boost region, demonstrating its utility by studying partially reconstructed tops. The case of hidden valleys will be presented elsewhere.

Top quarks are a standard testing ground for boosted techniques and are considered a standard candle for comparison between methods [5,6]. The LHC experiments have used a variety of techniques to identify boosted top quarks, and used them to search for new physics signals [10–16]. In the case of fully reconstructed boosted tops, the jets typically admit 3 subjets and counting them provides a powerful discrimination against QCD background. However, when only two of the top decay products are reconstructed as clear subjets, the resulting substructure exhibits only 2 subjets and does not reconstruct the top. This sample is significantly more

* Corresponding author.

E-mail address: freytsis@uoregon.edu (M. Freytsis).

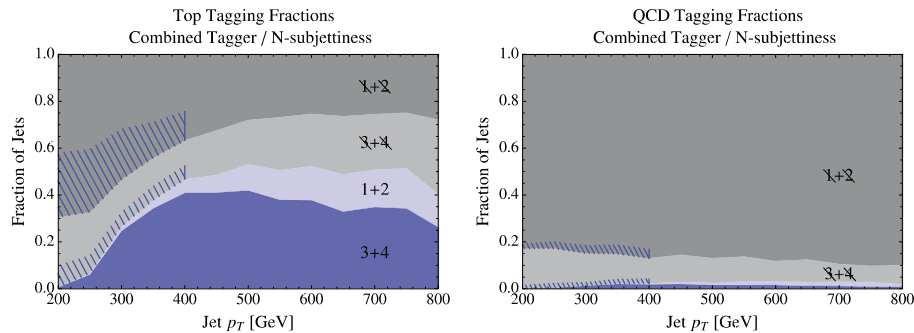


Fig. 1. Subject fractions for top (**left**) and QCD (**right**) jets as a function of jet p_T . The number of subjects is calculated using the Johns Hopkins top tagger [9], and we group the subjects into two categories: 1 or 2 (“1+2”) subjects and 3 or 4 (“3+4”) subjects, with plain and slashed numbers indicating tagged and untagged jets, respectively. The shaded regions correspond to the N -subjettiness top tagger, with the tagged fractions shown in blue and the untagged in gray. For $p_T < 400$ GeV, we show, hatched in blue, the additional fraction of untagged jets from N -subjettiness that are tagged by our top tagger. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

difficult to separate from background since there is no resonance mass to cut on, and due to the large QCD background in which a hard gluon splitting yields 2 subjects.

Despite the above, we argue below that partially reconstructed tops may be efficiently identified. Our key observation is that such jets have radiation patterns that differ from a vast majority of QCD jets, even though the partially reconstructed tops may not have kinematic invariants on which to base selection. We introduce an observable that is sensitive to these differences, and utilize aggressive jet grooming to further remove QCD jets with multiple subjects which may fake a partially reconstructed top. The combination results in a robust top tagging tool which works well in the intermediate-boost regime and that is complementary to existing techniques. While we focus on top quarks as a case study, the variables we consider are useful beyond this example, and may be used to identify other partially reconstructed objects in the intermediate boosted regime.

2. Kinematics properties of unreconstructed decays

The tagging of boosted heavy particles is a balance between two competing effects. On the one hand, the direct production cross-section of heavy objects (which decay into colored final states) falls with increasing momentum, while on the other, tagging with substructure techniques becomes more efficient at higher boost. Consequently, there is often a kinematic regime at intermediate boost where the cross section is large compared to the boosted regime but where most taggers do not effectively function [8]. These unreconstructed decays can be targeted with some simple considerations about their kinematics, as we now explain.

When a boosted particle such as a Z undergoes a $1 \rightarrow 2$ hadronic decay, some fraction of those decays can be reconstructed in a single jet. The decays can be characterized by the angle θ between the forward decay product and the boost direction in the rest frame of the Z . When $\theta \sim \pi/2$, the two quarks from the decay will have similar momenta in the lab frame and the subsequent hadronic final states can be reconstructed in a fat jet. For $\theta \sim 0$, the decay products become asymmetric, with one softer than the other, and further apart. As a consequence, the soft particle is emitted away from the jet direction, which lies along the hard decay product, and reconstruction is challenging. Thus, while symmetric decays can be reconstructed in a single jet, asymmetric ones may not be. Furthermore, these asymmetric decays mimic QCD-like topologies with a hierarchy of parton energies, and so the ability to differentiate these already challenging decays from QCD

backgrounds degrades. Thus, the efficiency to tag boosted Z bosons significantly decreases in this region of phase space.

Consider now the more complex case of a top quark which undergoes two successive $1 \rightarrow 2$ decays. Two main parameters are of interest to understand the top reconstruction efficiency: θ_W , the angle of the W to the top boost direction in the top rest frame, and θ_q , the angle of the forward quark from the W decay (defined in the W rest frame) to the W boost direction in the top rest frame. When $\theta_W \sim 0$, the W will carry most of the top quark’s momentum and the b quark will be soft and at a wide angle. Such a decay will lead to the boosted W being easily reconstructed, but the top will not be. When $\theta_W \sim \pi$, the W has a small boost in the lab frame and hence the two quarks from the W decay are not reconstructed in the same jet. Similarly, if $\theta_q \sim 0$, then the W decay products are asymmetric in energy and at wide angles, also leading to an unreconstructed top. It is only when θ_W and θ_q are both not extremal that the top has a large reconstruction efficiency. Furthermore, when the top is not reconstructed, it is often one of the quarks from the W decay that is lost, meaning the b and the other quark are clustered in the same jet, which is a configuration without any resonance mass to select on.

The above considerations hold in both the intermediate- and highly-boosted regimes. As an illustration, in Fig. 1 we show the fraction of tagged and untagged tops (*left*) and the fraction of mistagged and vetoed QCD jets (*right*), using the N -subjettiness top tagger [17] operating at the 50% efficiency point. (Unless explicitly stated otherwise, a jet tagged by N -jettiness will be taken to mean a jet tagged at this efficiency point, computed over the p_T range shown as defined in Ref. [18], i.e. with $\beta = 1$ and a mass window of $m_J \in (160, 240)$ GeV.) Each sample is further divided according to the number of subjects (as defined by the Johns Hopkins (JHU) tagger [9]) found in the event. We note that the fraction of 4 subjects is numerically small for both tops and QCD, and top jets additionally only rarely have a single subjet. We see that as the jet p_T decreases, the rate of untagged jets significantly increases, and a non-negligible fraction of untagged jets with ≤ 2 -subjects persists. The hashed regions for $p_T < 400$ GeV in each plot indicate the fraction of untagged jets (by N -subjettiness) that are tagged by the tagger we describe below.

From the above discussion we identify two interesting types of unreconstructed tops in which not all constituents are found in a single jet: the ones in which the b jet is soft and the ones in which the W is soft. In the former, the W is boosted into a single jet but the b is at a wide angle, while in the latter we typically find that the b is clustered with one of the quarks from the W decay and the other quark is at a wide angle. In both cases, one expects roughly equal energy sharing between the quarks in the same jet,

which is uncharacteristic of QCD jets that preferentially have soft splittings. Therefore, a useful kinematic handle is the intrinsic two-prong nature of unreconstructed top jets, in which neither of the two subjects are soft. We will focus our tagger around observables and techniques that exploit this feature.

A variety of different methodologies have been adopted in order to tag the hadronic decays of boosted tops [9,19–26,17,27,18,28–32]. All of these search strategies are focused on tops where the decay may be reconstructed in a single jet, typically by identifying 3 well-separated subjects. Although combinations of techniques show improved tagging power, they are often limited by the fact that the subsets of tops being tagged are significantly overlapping (to a greater degree than the QCD jets that fake these tops). However, the tagger discussed below targets a kinematically distinct sample of tops and therefore is largely orthogonal to other techniques, meaning that the combination with other taggers shows increased power in top tagging.

3. Tagging partially reconstructed objects

The common currency for jet substructure is subjects: isolated clusters of high- p_T radiation in the jet. Most substructure tagging methods count the number of subjects in a jet (sometimes subject to some additional cuts per subject) and select those matching the expected number from the boosted decay, such as 3 for the top quark (or 4 if allowing a hard radiated gluon). Techniques to count subjects typically use declustering algorithms (e.g., [33,9,23,34]) or shape measurements (e.g., [17,18,28,32]). More general jet grooming methods [22,24,35–37] can also be used to shape the jets, but must be supplemented with a subject counting method like those above. An additional focus of jet substructure has also been the removal of pileup contamination in jets [38–42]. With partially reconstructed objects in jets leaving their effect in the distribution of soft radiation, the application of such techniques to the problem at hand is more likely to limit discrimination than help. In top tagging, present approaches explicitly remove cases where the top is not fully reconstructed, i.e., when there are fewer than 3 explicitly identified subjects.

Of course, having 3 subjects does not guarantee tagging. In the intermediate p_T regime, 3- and 4-subjet events are often misidentified and do not correspond to the actual top decay products. Consequently, additional (typically mass) cuts applied to combinations of subjects in existing top-taggers are designed to veto such events. Conversely, in the boosted regime, as the p_T is raised, the subjects increasingly correspond to the true decay products, but these additional cuts are still necessary to control QCD backgrounds. The above is illustrated in Fig. 1 where it is apparent that a sizable sample of top jets, with any number of subjects, is not tagged, with the situation significantly worse at lower p_T .

The partially reconstructed tops we target in this paper are precisely those tops missed by conventional top taggers. This sample is dominated by events where either the b quark or one of the quarks from the W is soft or escapes the jet. In the latter case, due to the sequential nature of the top decay, the phase space ensures that such configurations are approximately 3 times more likely to contain a b quark and one of the quarks from the decay of the W rather than a W alone. To tag such events, it is best to assume as little about the precise substructure of the jet as possible.

Our first variable is therefore one that allows us to remove the QCD sample with no hard splittings while making no assumptions about the number of hard subjects in the signal sample. We consider 1-subjettiness with variable angular weighting exponent β [17],

$$\tau_1^{(\beta)} = \frac{1}{p_{T,J}} \min_{\hat{n}} \sum_{i \in \text{jet}} p_{T,i} \Delta R_{n,i}^\beta. \quad (1)$$

For small angles, this sum can be written as $\sum_i z_i \theta_i^\beta$, with angles defined with respect to the axis that minimizes the variable, and z_i the fraction of the jet p_T carried by the i th particle. As detailed in Ref. [43], for $\beta = 1$ this sum is always minimized when the axis is aligned with the most energetic particle and if particle splittings are strongly ordered in angle. Alternatively, for $\beta = 2$, this is minimized when the axis is aligned with the jet axis itself in the narrow jet limit.

In the decay of a boosted object, if at least two subjects are reconstructable, the $\tau_1^{(1)}$ axis precisely tracks the hardest subject, while the $\tau_1^{(2)}$ axis continues to fall on the jet axis. The two axes are then expected to be $O(1)$ apart. Conversely, in the case of a QCD jet with no hard splittings, the jet axis can only move away from the hardest particle when that particle undergoes a collinear splitting. This distance is expected to only be $O(\theta_c)$ from the hardest particle in the jet, with θ_c the angle of the widest collinear splitting in the shower. We then expect that this distance to still be $O(\theta_c) \approx m_J/p_T \ll 1$. All of this remains true for an arbitrary number of subjects.

This suggests that a cut on a minimal separation between the $\tau_1^{(1)}$ and $\tau_1^{(2)}$ axes will be able to eliminate QCD jets with no hard splittings in its parton shower, while remaining agnostic about the nature of the signal. We call this variable the *jet lean*, or ΔR_{τ_1} , the distance between the $\tau_1^{(1)}$ and $\tau_1^{(2)}$ axes. (As noted above, the $\tau_1^{(2)}$ axis and jet axis can be interchanged with no loss of performance. For a single emission off a single hard parton, jet lean is identical to $\tau_1^{(1)}$ itself [43]. But with a multi-prong structure, as is the case for tops, jet lean picks up an additional factor of the ratio of jet to average subject p_T , providing improved discrimination compared to using $\tau_1^{(1)}$.) Similar variables have also been considered in the discrimination of boosted colored octets in Ref. [44]. The discriminating power of jet lean is illustrated on the left of Fig. 2, which shows its distribution for the tagged and untagged top and QCD samples. A clear separation is visible. Demanding that jet lean be above a certain value will eliminate the QCD background coming from jets that get their mass from a wide, diffuse parton shower with no discernible substructure.

The remaining QCD sample with 2 or more subjects can be reduced further, but as we desire to remain agnostic as to the number of reconstructed subjects in the signal sample, we need to exploit the kinematic properties of the subject splittings. Coming from on-shell decays, top decay products tend to have fairly democratic energy sharing among subjects, while QCD jets, dominated by singularities of the splitting functions even in the case of hard splittings, tend to produce a hierarchy in subject p_T . While one can check for this condition by simply declustering a jet once, and looking at the p_T ratio of the two subjects, superior discrimination is achieved using jet grooming algorithms [24,35]. This allows us to more efficiently isolate the leading hard subject. Because of the presence of soft radiation in the jet from the unidentified components of the top decay, grooming the subjects before measuring their p_T can be more effective than it otherwise would be.

Here we use trimming, which reclusters the jet into subjects of small radius R_{sub} . Only subjects with a p_T fraction of the jet p_T greater than a parameter f_{cut} are kept. By running the trimming procedure with a much larger value of f_{cut} (≈ 0.20) than typically used, one may veto softer subjects if their p_T fractions are small enough, instead of simply removing soft radiation which is the usual goal of trimming. We call this aggressive use of trimming, *chopping*. Chopping then preferentially removes softer subjects from QCD and converts some of the 2-subjet QCD sample into 1-subjet

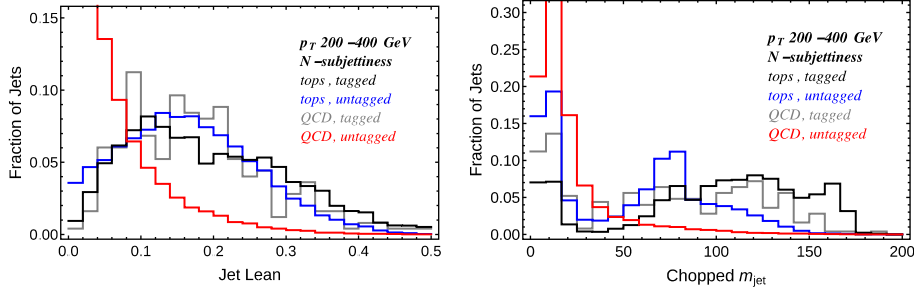


Fig. 2. Distributions of variables used in our tagger: jet lean (**left**) and the chopped jet mass (**right**). These variables are shown for samples of tagged and untagged top and QCD jets. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)




Fig. 3. Two-dimensional distributions of our top tagger variables, jet lean and chopped jet mass, for top jets tagged (**left**) and untagged (**middle**) by N -subjettiness, as well as untagged QCD jets (**right**). For each bin in the two-dimensional distribution, the area is proportional to the number of events in the bin. A dashed line on each plot indicates the linear cut for the 50% efficiency point for the combined tagger with N -subjettiness. The sample of QCD jets which pass the conventional N -subjettiness tagger are indistinguishable from the tagged top sample in our variables. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

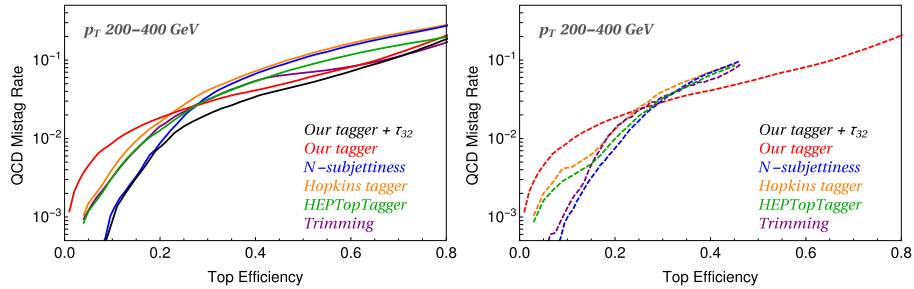


Fig. 4. Plots of top tagger performance in the intermediate p_T range of 200–400 GeV. The data is taken from the boosted top benchmark samples generated for the BOOST2010 [5] workshop and the performance is shown for the trimming tagger [35] (purple), Johns Hopkins tagger [9] (orange), HEPTopTagger [25] (green), N -subjettiness tagger [18] (blue), the tagger studied in this paper (red) and the tagger of this paper in conjunction with the cut on τ_3/τ_2 (black). **Left:** taggers with all cuts optimized for the subsample. **Right:** same as left but with an additional hard cut of $m_{\text{jet}} > 120$ GeV on all jets for all taggers other than our own. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

jets. A cut on the chopped jet mass will then remove these jets. The behavior of this variable is illustrated in the right plot of Fig. 2, where the distribution of the chopped jet mass is shown for the tagged and untagged top and QCD samples.

To summarize, our tagger consists of a two-step procedure with parameter values,

$$\Delta R_{\tau_1} \gtrsim O(0.1), \quad \text{chopped } m_J \gtrsim O(50 \text{ GeV}), \quad (2)$$

giving good separation between boosted top and QCD jets. This combination allows us to simultaneously veto massive QCD jets with no discernible substructure and purify the remaining sample of jets with hard splittings that are more QCD-like, without imposing a particular subjet counting or mass window on the signal.

4. Efficiency studies and results

We wish to characterize the behavior of our top tagger in the intermediate p_T range, $200 \text{ GeV} < p_T < 400 \text{ GeV}$, as this is where

conventional top taggers start to lose efficiency. The effect of the two cuts separately is displayed in Fig. 2, and together in Fig. 3, for QCD and tops that have passed and failed the conventional implementation of the N -subjettiness top tagger. (Distributions for tagged QCD jets are identical to those of tagged tops within statistics.) Fig. 3 also displays the cut corresponding to a linear optimization at the 50% top efficiency point. The optimal linear Fisher discriminant [45] is constructed from the two variables, while scanning over the input parameters controlling calculation of the chopped mass. As noted above, Fig. 2 shows that the distribution of both conventionally tagged and untagged tops looks markedly different from the majority of QCD events, demonstrating that the variables presented here are sensitive to discriminating structure that other taggers miss. We use the boosted top benchmark Monte Carlo samples generated for the BOOST2010 workshop to perform our analysis [5].

In Fig. 4, we show the efficiency of the top tagger described above in the intermediate p_T range, along with corresponding

performance from several different types of conventional taggers. In all cases, optimized cuts are found following the procedure described above, with a cut on the appropriate linear combination of output variables of the tagger used, while input parameters (where present) are scanned over a reasonable range. We carry out the above in two different cases. Initially, we do not employ the conventional hard lower cut on the total jet mass, designed to remove isolated boosted W -jets from the top sample. We then compare this with the imposition of an $m_J > 120$ GeV cut on all conventional taggers. The performance in both cases is comparable with the exception of the trimming tagger at high purity, but in the second case conventional taggers cannot achieve an efficiency better than $\approx .46$ due to the hard cut. While the HEPTOPTAGGER algorithm [25] has been advocated in the past as having better performance in the intermediate boost regime compared to other approaches, we find that at the low p_T 's we focus on, after optimization, its performance is only somewhat better than the Hopkins tagger, and not better than our advocated combination with N -subjettiness. (This should be contrasted with the HEPTOPTAGGER2 algorithm [46], with which we do not compare performance, as it is a much more involved tagger than any ones discussed here, with many additional input variables and possibilities to tune for specific subsamples.) We see that unless one desires a high purity sample, our tagger on its own outperforms the conventional approaches for such a moderately boosted sample.

While preferentially tagging poorly reconstructed tops on its own might be of formal interest, in practical applications it would be rare to wish to tag tops only in the intermediate-boost regime. Instead, the hope would be to recover some untagged tops in addition to those already tagged by more conventional methods. Since we use the N -subjettiness variable as part of our method, it is natural to look at the performance of running our tagger in parallel with the more conventional cut on τ_3/τ_2 , which for the 50% efficiency point corresponds to the cut,

$$\frac{\tau_3}{\tau_2} + 1.1 \left(\frac{\Delta m}{m} \right)_+ - 1.8 \left(\frac{\Delta m}{m} \right)_- \lesssim 0.7, \quad (3)$$

with $(\Delta m/m)_\pm$ being the fractional deviation above (below) the top mass for the given jet. Thus, we also show the efficiency for a double tagger for which events which pass either the conventional τ_3/τ_2 cut or our new cuts are considered tagged. It turns out that a substantial overlap in QCD mistags in both methods is present. This allows for notably better tagging performance than either approach alone, particularly in the 0.25–0.4 efficiency range. In fact, the combined tagger shows at least a 50% improvement in QCD mistag reduction over previously available methods in the entire 0.25–0.5 efficiency range.

5. Discussion

In this note we have presented a method for differentiating (partially) boosted hadronically decaying objects without having to specify a particular topology or mass window for grouping of sub-jets. Our approach allows us to extend the range of scenarios in which substructure methods can be applied, going beyond conventional approaches that excel when all decay products are cleanly identifiable. As an example, we considered hadronically decaying top quarks with intermediate boosts. Here, our method alone proves more powerful when larger top efficiencies are required, and shows even more significant improvement when combined with a conventional top tagging method (N -subjettiness). Additional applications, which we leave to future work, could involve reconstruction of objects decaying both to hadrons and invisible products, for which conventional substructure approaches are entirely inapplicable.

Acknowledgements

We thank David Curtin, Gilly Elor, Matt Schwartz, Jesse Thaler, Jay Wacker for useful discussions. MF thanks Tel-Aviv University for its hospitality while this work was in preparation. The work of TV is supported by the US-Israel Binational Science Foundation, the EU-FP7 Marie Curie, CIG fellowship and by the I-CORE Program of the Planning Budgeting Committee and the Israel Science Foundation (Grant No. 1937/12). The work of JRW is supported by the Office of Science, Office of High Energy Physics, of the U.S. Department of Energy (DOE) under contract DE-AC02-05CH11231. The work of MF is supported by the DOE under grant DE-SC003916 and the National Science Foundation (NSF) under grant No. PHY-1258729.

References

- [1] G. Aad, et al., Observation of a new particle in the search for the standard model higgs boson with the atlas detector at the Lhc, Phys. Lett. B 716 (2012) 1–29, <http://dx.doi.org/10.1016/j.physletb.2012.08.020>, arXiv:1207.7214.
- [2] S. Chatrchyan, et al., Observation of a new boson at a mass of 125 GeV with the Cms experiment at the Lhc, Phys. Lett. B 716 (2012) 30–61, <http://dx.doi.org/10.1016/j.physletb.2012.08.021>, arXiv:1207.7235.
- [3] M.J. Strassler, K.M. Zurek, Echoes of a hidden valley at hadron colliders, Phys. Lett. B 651 (2007) 374–379, <http://dx.doi.org/10.1016/j.physletb.2007.06.055>, arXiv:hep-ph/0604261.
- [4] M.J. Strassler, K.M. Zurek, Discovering the Higgs through highly-displaced vertices, Phys. Lett. B 661 (2008) 263–267, <http://dx.doi.org/10.1016/j.physletb.2008.02.008>, arXiv:hep-ph/0605193.
- [5] A. Abdesselam, E.B. Kuutmann, U. Bitenc, G. Brooijmans, J. Butterworth, et al., Boosted objects: a probe of beyond the standard model physics, Eur. Phys. J. C 71 (2011) 1661, <http://dx.doi.org/10.1140/epjc/s10052-011-1661-y>, arXiv:1012.5412.
- [6] A. Altheimer, S. Arora, L. Asquith, G. Brooijmans, J. Butterworth, et al., Jet substructure at the tevatron and LHC: new results, new tools, new benchmarks, J. Phys. G 39 (2012) 063001, <http://dx.doi.org/10.1088/0954-3899/39/6/063001>, arXiv:1201.0008.
- [7] A. Altheimer, A. Arce, L. Asquith, J. Backus Mayes, E. Bergeaas Kuutmann, et al., Boosted objects and jet substructure at the LHC. Report of BOOST2012, held at IFIC Valencia, 23rd–27th of July 2012, Eur. Phys. J. C 74 (3) (2014) 2792, <http://dx.doi.org/10.1140/epjc/s10052-014-2792-8>, arXiv:1311.2708.
- [8] T. Plehn, P. Schichtel, D. Wiegand, Where boosted significances come from, Phys. Rev. D 89 (5) (2014) 054002, <http://dx.doi.org/10.1103/PhysRevD.89.054002>, arXiv:1311.2591.
- [9] D.E. Kaplan, K. Rehermann, M.D. Schwartz, B. Tweedie, Top tagging: a method for identifying boosted hadronically decaying top quarks, Phys. Rev. Lett. 101 (2008) 142001, <http://dx.doi.org/10.1103/PhysRevLett.101.142001>, arXiv:0806.0848.
- [10] Search for BSM $t\bar{t}$ production in the boosted all-hadronic final state, 2011.
- [11] G. Aad, et al., A search for $t\bar{t}$ resonances in lepton+jets events with highly boosted top quarks collected in pp collisions at $\sqrt{s} = 7$ TeV with the ATLAS detector, J. High Energy Phys. 1209 (2012) 041, [http://dx.doi.org/10.1007/JHEP09\(2012\)041](http://dx.doi.org/10.1007/JHEP09(2012)041), arXiv:1207.2409.
- [12] G. Aad, et al., Search for resonances decaying into top-quark pairs using fully hadronic decays in pp collisions with ATLAS at $\sqrt{s} = 7$ TeV, J. High Energy Phys. 1301 (2013) 116, [http://dx.doi.org/10.1007/JHEP01\(2013\)116](http://dx.doi.org/10.1007/JHEP01(2013)116), arXiv:1211.2202.
- [13] S. Chatrchyan, et al., Search for anomalous $t\bar{t}$ production in the highly-boosted all-hadronic final state, J. High Energy Phys. 1209 (2012) 029, [http://dx.doi.org/10.1007/JHEP03\(2014\)132](http://dx.doi.org/10.1007/JHEP03(2014)132), arXiv:1204.2488.
- [14] Performance of boosted top quark identification in 2012 ATLAS data, 2013.
- [15] G. Aad, et al., Performance of jet substructure techniques for large- R jets in proton-proton collisions at $\sqrt{s} = 7$ TeV using the ATLAS detector, J. High Energy Phys. 1309 (2013) 076, [http://dx.doi.org/10.1007/JHEP09\(2013\)076](http://dx.doi.org/10.1007/JHEP09(2013)076), arXiv:1306.4945.
- [16] Search for top-Higgs resonances in all-hadronic final states using jet substructure methods.
- [17] J. Thaler, K. Van Tilburg, Identifying boosted objects with N -subjettiness, J. High Energy Phys. 1103 (2011) 015, [http://dx.doi.org/10.1007/JHEP03\(2011\)015](http://dx.doi.org/10.1007/JHEP03(2011)015), arXiv:1011.2268.
- [18] J. Thaler, K. Van Tilburg, Maximizing boosted top identification by minimizing N -subjettiness, J. High Energy Phys. 1202 (2012) 093, [http://dx.doi.org/10.1007/JHEP02\(2012\)093](http://dx.doi.org/10.1007/JHEP02(2012)093), arXiv:1108.2701.
- [19] J. Thaler, L.-T. Wang, Strategies to identify boosted tops, J. High Energy Phys. 0807 (2008) 092, <http://dx.doi.org/10.1088/1126-6708/2008/07/092>, arXiv:0806.0023.

- [20] L.G. Almeida, S.J. Lee, G. Perez, G.F. Sterman, I. Sung, et al., Substructure of high- p_T jets at the LHC, Phys. Rev. D 79 (2009) 074017, <http://dx.doi.org/10.1103/PhysRevD.79.074017>, arXiv:0807.0234.
- [21] L.G. Almeida, S.J. Lee, G. Perez, I. Sung, J. Virzi, Top jets at the LHC, Phys. Rev. D 79 (2009) 074012, <http://dx.doi.org/10.1103/PhysRevD.79.074012>, arXiv:0810.0934.
- [22] S.D. Ellis, C.K. Vermilion, J.R. Walsh, Techniques for improved heavy particle searches with jet substructure, Phys. Rev. D 80 (2009) 051501, <http://dx.doi.org/10.1103/PhysRevD.80.051501>, arXiv:0903.5081.
- [23] T. Plehn, G.P. Salam, M. Spannowsky, Fat jets for a light higgs, Phys. Rev. Lett. 104 (2010) 111801, <http://dx.doi.org/10.1103/PhysRevLett.104.111801>, arXiv:0910.5472.
- [24] S.D. Ellis, C.K. Vermilion, J.R. Walsh, Recombination algorithms and jet substructure: pruning as a tool for heavy particle searches, Phys. Rev. D 81 (2010) 094023, <http://dx.doi.org/10.1103/PhysRevD.81.094023>, arXiv:0912.0033.
- [25] T. Plehn, M. Spannowsky, M. Takeuchi, D. Zerwas, Stop reconstruction with tagged tops, J. High Energy Phys. 1010 (2010) 078, [http://dx.doi.org/10.1007/JHEP10\(2010\)078](http://dx.doi.org/10.1007/JHEP10(2010)078), arXiv:1006.2833.
- [26] L.G. Almeida, S.J. Lee, G. Perez, G. Sterman, I. Sung, Template overlap method for massive jets, Phys. Rev. D 82 (2010) 054034, <http://dx.doi.org/10.1103/PhysRevD.82.054034>, arXiv:1006.2035.
- [27] A. Hook, M. Jankowiak, J.G. Wacker, Jet dipolarity: top tagging with color flow, J. High Energy Phys. 1204 (2012) 007, [http://dx.doi.org/10.1007/JHEP04\(2012\)007](http://dx.doi.org/10.1007/JHEP04(2012)007), arXiv:1102.1012.
- [28] M. Jankowiak, A.J. Larkoski, Jet substructure without trees, J. High Energy Phys. 1106 (2011) 057, [http://dx.doi.org/10.1007/JHEP06\(2011\)057](http://dx.doi.org/10.1007/JHEP06(2011)057), arXiv:1104.1646.
- [29] D.E. Soper, M. Spannowsky, Finding top quarks with shower deconstruction, Phys. Rev. D 87 (5) (2013) 054012, <http://dx.doi.org/10.1103/PhysRevD.87.054012>, arXiv:1211.3140.
- [30] A.J. Larkoski, G.P. Salam, J. Thaler, Energy correlation functions for jet substructure, J. High Energy Phys. 1306 (2013) 108, [http://dx.doi.org/10.1007/JHEP06\(2013\)108](http://dx.doi.org/10.1007/JHEP06(2013)108), arXiv:1305.0007.
- [31] S. Schaetzel, M. Spannowsky, Tagging highly boosted top quarks, Phys. Rev. D 89 (2014) 014007, <http://dx.doi.org/10.1103/PhysRevD.89.014007>, arXiv:1308.0540.
- [32] A.J. Larkoski, I. Moulton, D. Neill, Building a better boosted top tagger, Phys. Rev. D 91 (3) (2015) 034035, <http://dx.doi.org/10.1103/PhysRevD.91.034035>, arXiv:1411.0665.
- [33] J.M. Butterworth, A.R. Davison, M. Rubin, G.P. Salam, Jet substructure as a new higgs search channel at the LHC, Phys. Rev. Lett. 100 (2008) 242001, <http://dx.doi.org/10.1103/PhysRevLett.100.242001>, arXiv:0802.2470.
- [34] S. El Hedri, A. Hook, M. Jankowiak, J.G. Wacker, Learning how to count: a high multiplicity search for the LHC, J. High Energy Phys. 1308 (2013) 136, [http://dx.doi.org/10.1007/JHEP08\(2013\)136](http://dx.doi.org/10.1007/JHEP08(2013)136), arXiv:1302.1870.
- [35] D. Krohn, J. Thaler, L.-T. Wang, Jet trimming, J. High Energy Phys. 1002 (2010) 084, [http://dx.doi.org/10.1007/JHEP02\(2010\)084](http://dx.doi.org/10.1007/JHEP02(2010)084), arXiv:0912.1342.
- [36] M. Dasgupta, A. Fregoso, S. Marzani, G.P. Salam, Towards an understanding of jet substructure, J. High Energy Phys. 09 (2013) 029, [http://dx.doi.org/10.1007/JHEP09\(2013\)029](http://dx.doi.org/10.1007/JHEP09(2013)029), arXiv:1307.0007.
- [37] A.J. Larkoski, S. Marzani, G. Soyez, J. Thaler, Soft drop, J. High Energy Phys. 05 (2014) 146, [http://dx.doi.org/10.1007/JHEP05\(2014\)146](http://dx.doi.org/10.1007/JHEP05(2014)146), arXiv:1402.2657.
- [38] G. Soyez, G.P. Salam, J. Kim, S. Dutta, M. Cacciari, Pileup subtraction for jet shapes, Phys. Rev. Lett. 110 (16) (2013) 162001, <http://dx.doi.org/10.1103/PhysRevLett.110.162001>, arXiv:1211.2811.
- [39] M. Cacciari, G.P. Salam, Pileup subtraction using jet areas, Phys. Lett. B 659 (2008) 119–126, <http://dx.doi.org/10.1016/j.physletb.2007.09.077>, arXiv:0707.1378.
- [40] D. Krohn, M.D. Schwartz, M. Low, L.-T. Wang, Jet cleansing: pileup removal at high luminosity, Phys. Rev. D 90 (6) (2014) 065020, <http://dx.doi.org/10.1103/PhysRevD.90.065020>, arXiv:1309.4777.
- [41] M. Cacciari, G.P. Salam, G. Soyez, SoftKiller, a particle-level pileup removal method, Eur. Phys. J. C 75 (2) (2015) 59, <http://dx.doi.org/10.1140/epjc/s10052-015-3267-2>, arXiv:1407.0408.
- [42] D. Bertolini, P. Harris, M. Low, N. Tran, Pileup per particle identification, J. High Energy Phys. 10 (2014) 059, [http://dx.doi.org/10.1007/JHEP10\(2014\)059](http://dx.doi.org/10.1007/JHEP10(2014)059), arXiv:1407.6013.
- [43] A.J. Larkoski, D. Neill, J. Thaler, Jet shapes with the broadening axis, J. High Energy Phys. 04 (2014) 017, [http://dx.doi.org/10.1007/JHEP04\(2014\)017](http://dx.doi.org/10.1007/JHEP04(2014)017), arXiv:1401.2158.
- [44] D. Curtin, R. Essig, B. Shuve, Boosted multijet resonances and new color-flow variables, Phys. Rev. D 88 (2013) 034019, <http://dx.doi.org/10.1103/PhysRevD.88.034019>, arXiv:1210.5523.
- [45] T. Anderson, R. Bahadur, Classification into two multivariate normal distributions with different covariance matrices, Ann. Math. Stat. 33 (1962) 420–431.
- [46] G. Kasieczka, T. Plehn, T. Schell, T. Strebler, G.P. Salam, Resonance searches with an updated top tagger, J. High Energy Phys. 06 (2015) 203, [http://dx.doi.org/10.1007/JHEP06\(2015\)203](http://dx.doi.org/10.1007/JHEP06(2015)203), arXiv:1503.05921.