

Searches for Heavy Resonances with Substructure

Petar Maksimović

The William H. Miller III Department of Physics and Astronomy, The Johns Hopkins University, Baltimore, Maryland, USA; email: petar@jhu.edu

ANNUAL
REVIEWS **CONNECT**

www.annualreviews.org

- Download figures
- Navigate cited references
- Keyword search
- Explore related articles
- Share via email or social media

Annu. Rev. Nucl. Part. Sci. 2022. 72:447–75

First published as a Review in Advance on August 1, 2022

The *Annual Review of Nuclear and Particle Science* is online at nucl.annualreviews.org

<https://doi.org/10.1146/annurev-nucl-102419-055402>

Copyright © 2022 by Annual Reviews. This work is licensed under a Creative Commons Attribution 4.0 International License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited. See credit lines of images or other third-party material in this article for license information.



Keywords

BSM searches, heavy resonances, jet substructure

Abstract

In the past decade, the Large Hadron Collider (LHC) has probed a higher energy scale than ever before. Most models of physics beyond the standard model (BSM) predict the production of new heavy particles; the LHC results have excluded lower masses of such particles. This makes the high-mass regions especially interesting for current and future searches. In most BSM scenarios of interest, the new heavy resonances decay to standard model particles. In a subset of these models, the new particles have large couplings to the top quark, the W and Z bosons, or the Higgs boson. The top quark and W , Z , and Higgs bosons often decay to quarks, giving rise to jets of particles with substructure; event selection based on substructure is used to suppress standard model backgrounds. This review covers the key concepts in experimental searches based on the jet substructure and discusses recent results from the ATLAS and CMS experiments.

Contents

1. INTRODUCTION	448
2. AN ANATOMY OF A HEAVY RESONANCE SEARCH	450
3. JET IDENTIFICATION USING SUBSTRUCTURE	450
3.1. Experiments, Triggers, and Data Sets for Heavy Resonance Searches	450
3.2. Triggers	450
3.3. Jet Reconstruction	451
3.4. Jet Taggers and Their Calibrations	453
3.5. Multivariate Tools for W and Top Jet Tagging	454
3.6. Background Estimation in Searches with Substructure	456
4. EXAMPLES OF RECENT SEARCHES WITH SUBSTRUCTURE	457
4.1. Searches for Resonances with Top Quarks in the Final State	457
4.2. Searches for Diboson Resonances	461
4.3. Di-Higgs and Di-Higgs-Like Resonances	463
4.4. Three-Jet Topologies and Exotic Substructure	467
5. THE FUTURE OF THE SEARCHES WITH SUBSTRUCTURE	470
5.1. Searches with Anomalous Substructure	470
6. CONCLUSIONS	472

1. INTRODUCTION

The standard model (SM) of particle physics allows for an economical formulation of the electroweak symmetry breaking in terms of only one new fundamental degree of freedom—the Higgs boson. However, the SM does not explain the dynamical origin of the symmetry breaking, nor does it explain why the Higgs boson is light. Several realistic models have been built that propose solutions to some of these long-standing problems.

In the past decade, the Large Hadron Collider (LHC) has probed the highest energy scale ever accessed by particle physics experiments. Most models of beyond-the-SM (BSM) physics predict production of new heavy particles, either as a result of the assumption of additional gauge symmetries or as a result of predicted extensions of the Higgs sector. The LHC's high-energy proton–proton collisions make it possible to produce these new states with masses up to one hundred times the electroweak energy scale, enabling the LHC experiments to explore a large variety of the models of BSM physics. The LHC searches have excluded lower masses of such new particles in many popular models. This makes the high-mass regions especially interesting for current and future searches.

In most BSM scenarios of interest, new heavy resonances decay to SM particles. In a subset of these models, the new particles have large couplings to the top quark, the W and Z bosons, or the Higgs boson. The top quark and W , Z , and Higgs bosons further decay to lighter SM particles, which in turn decay to quarks. The showering and hadronization of partons (quarks and gluons) produce jets: collimated sprays of hadrons like pions and kaons.

When the new resonance is heavy, its decay products—the top quark and W , Z , and Higgs bosons—are highly Lorentz-boosted, and the jets formed by the hadronization of quarks from the decay of a top quark or W , Z , or Higgs boson combine in a merged jet. This jet has two important properties: (a) a large jet mass and (b) a substructure—namely, information about the distribution of energy within it and dispersion of soft particles with respect to the lobes of energy

corresponding to the jets formed by quarks from the SM particle's decay. Some heavy SM particles decay to bottom and charm quarks; both kinds result in (*c*) displaced vertices and parton showers of heavy quarks that can be identified with suitable algorithms collectively known as *b* tagging.

The ATLAS and CMS Collaborations share new results on resonance searches several times a year, and by the time this review is published, parts of it may already be obsolete. Thus, the aim of this review is to provide a broader view of the techniques used, the main issues facing these measurements, and the future possibilities and challenges.

Composite Higgs models (1) treat the observed Higgs boson as a composite particle. This approach could solve the electroweak hierarchy problem (also called the Higgs naturalness problem), and the Higgs boson is a pseudo-Nambu-Goldstone boson of approximate global symmetry of the strong sector. In addition to supersymmetry, the composite Higgs models are, at the moment, the only other compelling solution. Like supersymmetry, there are many variants of the composite Higgs models; many of these models are realistic and designed to agree with the current data. The basic predictions of the composite Higgs models are the presence of the heavy partners of the SM quarks and bosons, which are excitations of the SM states: Z' (also ρ^0) and W' (also ρ^+) as partners of the heavy vector bosons W and Z , and T' and B' as partners of the third-generation quarks top and bottom. The key aspect is that the masses of the new partners are TeV-scale, which makes these models experimentally accessible to the LHC.

Several models give rise to heavy resonances that can decay to a $t\bar{t}$ final state. A topcolor-assisted technicolor (TC2) model (2, 3) predicts a spin-1 color-singlet boson. This leptophobic Z boson couples only to first- and third-generation quarks; it is produced mainly by $q\bar{q}$ annihilation, but it can also decay to $t\bar{t}$. The ATLAS experiment also considers simplified models for dark matter interactions, with both an axial-vector mediator and a vector mediator as proposed by the LHC Dark Matter Working Group (4).

Another $t\bar{t}$ benchmark is a spin-1 color-octet boson, the first Kaluza-Klein (KK) excitation of the gluon, g_{KK} , predicted by a Randall-Sundrum (RS) model with the SM fields propagating in the bulk of a single warped extra dimension (5). As predicted in Reference 5, the g_{KK} is created mostly via $q\bar{q}$ annihilation and decays predominantly into a $t\bar{t}$ pair with a branching fraction of $\approx 90\%$. The fourth benchmark is the bulk RS model (6, 7), which inherits the SM fields propagating in the bulk from the original RS model and predicts a spin-2 color-singlet boson. In this model, the first KK excitation of the graviton, G_{KK} , is dominantly produced via gluon fusion. The dimensionless coupling constant k/M_{Pl} controls the production rate and the decay width.

Several benchmark models are also used to interpret searches for the heavy resonances decaying to two bosons. Spin-0 radions and spin-2 gravitons from the RS model of warped extra dimensions are used as prototype resonances of these spins. Spin-1 resonances decaying to VV or VH are usually studied within an effective Lagrangian framework called the heavy vector triplet (HVT) model (8). HVT, a broad phenomenological framework for heavy resonances coupling to bosons and fermions of the SM, is used to interpret resonances decaying to VV and VH final states. The heavy partners Z' and W' of SM vector bosons interact with quarks and the Higgs field with coupling strengths g_q and g_H , respectively. In addition, the coupling to the Higgs field results in the interactions with longitudinally polarized Z and W bosons. Two variants of the HVT model are typically used for result interpretation in both the ATLAS and CMS experiments. HVT Model A uses $g_q = -0.55$ and $g_H = -0.56$ and provides the same phenomenology as weakly coupled models based on an extended gauge symmetry (9). HVT Model B sets $g_q = 0.14$ and $g_H = -2.9$ and corresponds to strongly coupled scenarios like those in composite Higgs models. The CMS experiment has used HVT Model C ($g_V \approx 1$, $c_H \approx 1$, $c_F = 0$), which allows heavy resonance production via the vector boson fusion (VBF) process.

2. AN ANATOMY OF A HEAVY RESONANCE SEARCH

In the past decade, the majority of searches that involved substructure targeted heavy particles decaying into top quarks or W , Z , and Higgs bosons. The first generation of searches focused on $Z' \rightarrow t\bar{t}$ followed by the diboson searches $X \rightarrow VV$ (where V is either a W or a Z boson). Eventually, all combinations were probed, including VH , HH , $t\bar{t}$, tV , tH , and bH .

Each of these particles is reconstructed as a single, merged, massive jet. In this review I use the term large- R jet; the reason should become clear in Section 3.3, in which the reconstruction of such jets in the ATLAS and CMS experiments is discussed. A combination of selection criteria based on the jet mass and a number of substructure variables [including, in recent years, machine learning (ML) discriminants] is applied to suitable large- R jets to identify candidates for boosted top quarks and W , Z , or Higgs bosons. These tools are introduced and discussed in Section 3.4.

The data analysis is conceptually straightforward: To measure the production rate of the BSM particles, or to set an upper limit on their production, one needs data collected with a suitable mix of triggers, a procedure to predict the background yield (briefly surveyed in Section 3.6), and a way of estimating the signal efficiency—the probability that a signal event will pass the selection. Understanding the systematic uncertainty on the signal efficiency usually requires calibration of the jet-tagging efficiency in data, and estimating it is especially challenging when good standard candles are not available. This issue is also discussed in Section 3.4. Calibrating jet taggers for signatures with more complex substructure is one of the outstanding tasks in this area.

Depending on whether the final state involves high- p_T leptons or not, the main backgrounds are typically either $t\bar{t} + \text{jets}$ and $W + \text{jets}$, or quark- or gluon-initiated multijet production (hereafter referred to as QCD multijet background). The background composition usually dictates the type of the background estimate, which, in turn, influences the formulation of the likelihood and the approach used for the signal extraction.

3. JET IDENTIFICATION USING SUBSTRUCTURE

3.1. Experiments, Triggers, and Data Sets for Heavy Resonance Searches

The ATLAS and CMS detectors have been described elsewhere (10, 11). Both detectors provide comprehensive coverage consisting of a pixel detector closest to the beam line, a tracker, electromagnetic and hadronic calorimeters, and muon systems on the outside. Both the ATLAS and CMS experiments deploy two-level trigger systems to save events for offline analysis. While the technological choices made in the two experiments have been quite different, the sensitivity for new heavy particles has, somewhat remarkably, been nearly identical.

3.2. Triggers

The triggers define data sets that can be used in offline data analyses. Searches with substructure, by definition, require the presence of large- R jets, so the typical triggers used are either hadronic, leptonic, or based on p_T^{miss} .

The hadronic triggers require either a large p_T of a single jet or multiple jets with the scalar sum of their p_T above a certain threshold. In some cases there is also a requirement for the presence of a large- R jet with a trimmed jet mass above 30 GeV.

Leptonic triggers require at least one high- p_T lepton, isolated or nonisolated. The isolation strongly suppresses the QCD multijet background from leptons arising from decays of bottom and charm hadrons, decays in flight, and misidentified leptons inside jets. Given lower backgrounds, the p_T thresholds for isolated leptons can be lower than for the nonisolated ones. However, non-isolated triggers are often important because highly boosted objects decaying semileptonically

(like $t \rightarrow \nu \ell b$ or $H \rightarrow WW^* \rightarrow \nu \ell q q'$) produce leptons that are too close to jets from the same decay, and imposing an isolation requirement would result in a substantial loss of the signal efficiency. Leptonic, hadronic, and p_T^{miss} triggers are often combined with each other, and the trigger combinations used in a particular analysis depend on the final state being probed.

3.3. Jet Reconstruction

As inputs to jet reconstruction, the CMS experiment combines tracker, calorimeter, and muon system data into particle flow (PF) candidates (12). In contrast, the ATLAS experiment has generally reconstructed jets solely from calorimeter information and used tracks to supplement it and improve performance. Because of the ATLAS detector's weaker magnetic field and longitudinally segmented calorimeter, it benefits less from PF than the CMS detector does.

Both the ATLAS and CMS experiments use the anti- k_t jet clustering algorithm (13) implemented in the FastJet software package (14). The anti- k_t algorithm produces approximately conical jets, and the distance parameter R roughly corresponds to a radius ΔR of a circle in the η - ϕ plane. Jets initiated by quarks and gluons that are not merged with other jets are reconstructed as small-radius (small- R) jets using $R = 0.4$ in both experiments. Decays of boosted objects that produced merged jets are reconstructed as large-radius (large- R) jets; ATLAS uses $R = 1.0$, whereas CMS uses $R = 0.8$. Because of their larger mass and radius, large- R jets are sometimes also referred to as fat jets. CMS reconstructs both kinds of jets using PF candidates. ATLAS reconstructs small- R jets from topological clusters in the calorimeter (15) and reconstructs large- R jets from track-calorimeter clusters (16) as inputs. While the ATLAS experiment is currently moving to a combination of PF (17) and track-calorimeter clusters called unified flow objects, its analyses featured in this review do not use it yet.

3.3.1. Jet grooming. The mass of a large- R jet from merged decay products of a top quark, from W , Z or Higgs bosons, or from a boosted lighter BSM particle emerging from a decay of a heavier resonance will peak near the particle's true mass. In contrast, the distribution of the jet mass from the QCD multijet backgrounds, as well as other SM backgrounds resulting from accidental merging of small- R jets, is usually smooth in the same region. For this reason, the jet mass is an excellent discriminant between the signal and the smooth SM backgrounds.

QCD radiation adds to the mass distribution of the multiprong signal jets as well as the QCD background. It makes the signal more broad and also increases the tail of the QCD background so that the QCD jets are both more massive and more likely to be reconstructed as large- R jets. Jet grooming is a systematic removal of soft and wide-angle radiation from within a jet to both clean the jet mass and reveal the underlying substructure.

Three grooming algorithms are mainly used. The ATLAS experiment uses trimming (18), whereas the CMS experiment used pruning (19) in Run 1 but has switched to soft drop (20) in Run 2. In trimming, the k_t algorithm is used to recluster the jet constituents into subjets using $R_{\text{sub}} = 0.2$, and then subjets with p_T less than 5% of the p_T parent jet are removed.

In pruning and soft drop, a condition is imposed in each $2 \rightarrow 1$ clustering step by rewinding the jet clustering sequence. The transverse momentum fraction of the softer particle to the merged system, $z = \min(p_{T,1}, p_{T,2})/(p_{T,1} + p_{T,2})$, is a proxy for the scale of the soft radiation, and the angular distance ΔR between the two particles is used for identifying wide-angle radiation. For pruning, in the $2 \rightarrow 1$ clustering step, the softer particle is removed if $z < 0.1$ and $\Delta R < 0.5$. For soft drop, the softer particle is removed if $z < 0.1 \times (\Delta R/R)^\beta$. An important feature of soft drop is that the groomed observables are analytically calculable to high-order resummation accuracy. Most applications of soft drop set $\beta = 0$, which makes it equivalent to an earlier algorithm known as the modified mass drop tagger (21).

In addition to removing soft and wide-angle radiation within the jet, grooming also removes the contribution from particles that originated from initial-state radiation, the underlying event, and the pileup interactions.

3.3.2. Pileup and its mitigation. At the instantaneous luminosity characteristic of Run 2 of the LHC, dozens of additional proton–proton interactions in the same bunch crossing are possible, resulting in a number of additional primary vertices from soft QCD interactions called pileup. Pileup pp collisions produce, on average, a few tens of soft hadrons in their final state, adding hundreds to thousands of soft hadrons to the hard collision we seek to study. The particles emerging from the pileup collisions are interspersed among the particles from the hard collision, increasing the energy of jets and smearing the jet substructure information. Thus, pileup mitigation is a necessary component of most physics at the LHC.

To suppress central jets from pileup interactions, the ATLAS experiment requires that they pass the jet vertex tagger selection (22) if they are in the range of $p_T < 120$ GeV and $|\eta| < 2.5$. In addition, the trimming serves a dual role in the ATLAS experiment; its purpose is also to reduce contributions to the jet transverse momentum from pileup. In the CMS experiment, the pileup contribution to small- R jets is suppressed by removing the charged jet constituents (tracks) consistent with originating from pileup vertices, and an offset correction is applied to adjust for remaining contributions (12, 14). For large- R jets, a separate PUPPI (pileup per particle identification) algorithm (23, 24) further reduces the effect of pileup by rescaling the momentum of each neutral jet constituent according to its probability of originating from the pileup vertex; this probability is estimated using the local density of charged jet constituents (24).

3.3.3. Mass decorrelation. Many searches for heavy resonances rely on the groomed jet mass as one of the discriminants. The regions around the signal peak (sidebands) can serve either as a backbone of the background estimate or as a powerful control region kinematically similar to the signal region that can be used to validate the analysis procedures.

We have seen that in the QCD jets the additional radiation contributes both to jet mass and to the substructure, especially in the case of hard gluons that are not removed by jet grooming. This creates the correlation between jet mass and substructure, and many jet taggers, out of the box, sculpt the jet mass distribution once the selection on the tagger discriminant is applied, making the background broadly peak closer to signal. There are three ways to ameliorate this effect.

1. For each point in jet mass versus jet p_T (or resonance mass) space, a separate tagger requirement is defined. This idea has been pioneered in the Designing Decorrelated Taggers (DDT) approach (25) and has been further generalized into a DDT map in Reference 26.
2. For neural network (NN) taggers, one can use a loss function in training that penalizes deviation from the original mass distribution.
3. If the tagger is a boosted decision tree (BDT) or an NN, signal and background samples that have the same jet mass distribution are used for training, so that the tagger does not learn how to use the jet mass to differentiate the signal from the background. This is usually achieved by generating signal that spans a large mass range and then reweighting either signal or background simulation samples so that they have the same shape.

The first and the third approaches usually have better decorrelation (27). The second approach also results in a loss of performance compared with the non-mass-decorrelated implementation of the tagger.

3.4. Jet Taggers and Their Calibrations

The easiest way to tag a jet with substructure is to consider a jet with a groomed mass in an appropriate mass window and then apply a selection on a substructure discriminant. Several substructure variables rely on the sum of p_T -weighted distances between particles within the large- R jet. The top jet and $H \rightarrow b\bar{b}$ jets also require some form of b tagging. These tools have been used in many analyses and are well understood and robust. I cover them first and then move to the ML taggers that are mostly used in the CMS experiment.

3.4.1. N -subjettiness for top and V tagging. The inclusive jet shape called N -subjettiness (28, 29) starts by assuming the number of subjects, N , reclustering the jet using the exclusive k_t algorithm that returns exactly N subjects, and then calculating the sum of p_T of all constituents weighted by the angular difference with respect to the nearest exclusive subject axis:

$$\tau_N^\beta = \frac{1}{d_0} \sum_k p_{T,k} \min\{R_{1,k}^\beta, R_{2,k}^\beta, \dots, R_{N,k}^\beta\}, \quad 1.$$

where k sums over the jet constituents, $p_{T,k}$ are their transverse momenta, and $R_{i,k}$ are the angular distances with respect to the exclusive subject axis i . The normalization factor is $d_0 = \sum_k p_{T,k} R$, where R is the distance parameter used to originally cluster the jet. The exponent $\beta = 1$ is nearly always used. Jets with $\tau_N \approx 0$ have all their radiation aligned with the candidate subject axes and therefore have N (or fewer) subjects. Jets with $\tau_N \gg 0$ have a large fraction of their energy distributed away from the candidate subject directions and therefore have at least $N + 1$ subjects. The most widely used N -subjettiness variables are the ratio $\tau_{21} = \tau_2/\tau_1$, which is sensitive to two-prong jets like W or Z , and the ratio $\tau_{32} = \tau_3/\tau_2$ (**Figure 1**), which is sensitive to three-prong jets and is often used for top tagging.

3.4.2. Energy correlation functions. Generalized energy correlation functions (30) are based on the energies and pairwise angles of particles within a jet, with $(N + 1)$ -point correlators sensitive to N -prong substructure. However, they do not require the explicit identification of the candidate subject axes. In addition, these correlation functions are sensitive to certain soft and collinear features that are obscured by other methods. The most popular combinations of the energy correlation functions are C_2^β (30), D_2^β (31, 32), and N_2^β (30). The variable D_2 in particular has been widely used in ATLAS analyses to identify boosted two-prong jets. It is defined as $D_2^{(\beta)} = \frac{e_3^{(\beta)}}{(e_2^{(\beta)})^3}$, where the n -point energy correlation functions $e_n^{(\beta)}$ are defined as

$$e_2^{(\beta)} = \frac{1}{p_{T,J}^2} \sum_{1 \leq i \leq j \leq n_J} p_{T,i} p_{T,j} R_{ij}^\beta, \quad 2.$$

$$e_3^{(\beta)} = \frac{1}{p_{T,J}^3} \sum_{1 \leq i \leq j \leq k \leq n_J} p_{T,i} p_{T,j} p_{T,k} R_{ij}^\beta R_{ik}^\beta R_{jk}^\beta, \quad 3.$$

where $p_{T,J}$ is the p_T of the whole jet, $p_{T,k}$ is the p_T of the jet constituent k , and n_J is the number of jet constituents. Usually $\beta = 1$ is used. The D_2 variable exploits the sensitivity of e_2 to radiation about a single direction and of e_3 to radiation about two directions, such as for the two-pronged jets.

3.4.3. b tagging. If b quarks are a part of the decay of the heavy resonance, and their hadronization produces one of the subjects of a large- R jet (e.g., in $t \rightarrow bqq$ decays), then application of b -tagging algorithms to the subjects is used as a part of the large- R jet tagging.




Figure 1

(a) A comparison of the observed data and predicted distributions from simulation of the anti- k_t $R = 1.0$ trimmed jet τ_{32} for the top quark selection in a sample enriched in lepton + jets $t\bar{t}$ events. The fully merged three-prong $t \rightarrow bqq$ events (*white*) peak at lower τ_{32} values than other processes. Panel adapted from Reference 38 (CC BY 4.0). (b) Performance of the ParticleNet algorithm for identifying hadronically decaying $H \rightarrow b\bar{b}$ decays compared with the previously used CMS machine learning algorithm, DeepAK8. The suffix MD indicates the mass-decorrelated version of the discriminant, which is typically used in searches. Panel adapted from Reference 27 (CC BY 4.0). Abbreviations: DDT, Designing Decorrelated Taggers; MD, mass decorrelated.

In the ATLAS experiment, variable- R track jets are clustered with the anti- k_t algorithm, with a configurable radius parameter R ranging from 0.02 to 0.4, inversely proportional to the jet p_{T} . Such track jets, associated with the large- R jet, can be b tagged by two kinds of multivariate discriminants (33, 34). One such algorithm, MV2, is based on a BDT that combines the results of three b -tagging algorithms with varying efficiencies and background rejections; typically the variant MV2c10 is used. The other algorithm, DL1, is implemented as a deep feed-forward NN. The ATLAS experiment uses either algorithm for tagging subjects in top and Higgs jets.

The CMS experiment employs the same principle in some analyses as part of top tagging (usually alongside τ_{32}); in these cases, the PF candidates from the soft drop subjects groomed with soft drop are b tagged using either the DeepCSV algorithm (35) or DeepJet (36, 37). However, for Higgs jet tagging, dedicated algorithms are used.

3.5. Multivariate Tools for W and Top Jet Tagging

The ATLAS experiment has developed multivariate classifiers based on the deep NN (DNN) architecture (38) for large- R jets that contain hadronically decaying W bosons and top quarks. Multiple jet-level discriminants are used as inputs: calibrated jet p_{T} and mass, N -subjettiness, energy correlation functions, splitting scales (39), and the minimum pairwise invariant mass Q_w between the subjects (39). The W jet tagging uses additional variables like Fox-Wolfram moments, planar flow, angularity, and aplanarity (40).

The CMS experiment has also studied multivariate classifiers based on similar variables, but eventually settled on NN taggers that, as inputs, use the jet constituents directly. DeepAK8 is

a family of DNN taggers (41) that are based on a convolutional NN architecture and use both the displaced vertices and individual PF candidates; the mass decorrelation is implemented via a loss function. The ImageTop tagger is a convolutional NN that treats jets as calorimeter images; its performance is on par with that of the DeepAK8 top tagger but exhibits a much better mass decorrelation enforced by training. The CMS experiment employs its own version (42, 43) of the JHU Top Tagger (44) based on Cambridge-Aachen jet clustering (45) and employing additional selection according to the relative kinematics of the three subjets. HOTVR (heavy object tagger with variable R) (46) is also used in some searches, where the principle behind the variable- R jet clustering is applied to the whole event to identify large- R jets. For HOTVR, jets as large as $R = 1.5$ are reconstructed, but the jet area shrinks inversely with p_T , thus keeping the same signal efficiency over a large range of top quark momenta. The HOTVR-based top tagger also relies on the kinematics of the subjets for further selection.

3.5.1. Higgs jet tagging. Development of Higgs jet tagging began after the seminal work by Butterworth et al. (47), which inaugurated the field of jet substructure in collider physics.

The approach most commonly used in the ATLAS experiment to tag boosted $H \rightarrow b\bar{b}$ decays relies on b tagging variable- R track jets associated with the large- R jet. However, a new algorithm has recently been developed (48) that uses the kinematics of the jet constituents in the center-of-mass (CM) frame of the large- R jet, since there the two b quarks of a two-body $H \rightarrow b\bar{b}$ decay can be easily separated into a back-to-back topology. The topological clusters of the large- R jet (Section 3.3) and the tracks associated with the jet are boosted to the jet's CM frame. There, the topological clusters are reclustered to form exactly two subjets using the EEkT jet algorithm (49). Tracks are associated with the CM frame subjets, and then both are boosted back to the laboratory frame, and the standard ATLAS b -tagging algorithm (MV2c10) is used to identify Higgs jets with two b -tagged subjets.

The CMS experiment uses three Higgs jet-tagging algorithms. The first one is the double- b tagger, a BDT trained on high-level features that are sensitive to both displaced vertices and the properties of B hadron shower and fragmentation. It has an excellent decorrelation in both jet mass and jet p_T , which was enforced by training. The next $H \rightarrow b\bar{b}$ tagger is the DeepAK8-MD Hbb tagger. The DeepAK8-MD Hbb tagger is a mass-decorrelated tagger targeting $H \rightarrow b\bar{b}$. For a given signal efficiency, typical of CMS searches, this tagger's rejection rate is better than that of the double- b tagger by about a factor of two. Lastly, a new DNN platform called ParticleNet (50), based on the graph NN architecture treating each jet as an unordered set of particles in space (particle cloud), has superseded the DeepAK8 Hbb tagger since it has about a factor of two better background suppression for signal efficiency of $\approx 50\%$. It is based on permutation-invariant graph NNs, and it uses the jet substructure and the flavor content simultaneously.

3.5.2. Calibrating taggers. Jet taggers based on substructure can exploit features that are hard to accurately model, and, therefore, the uncertainty on the signal-tagging efficiency is one of the most important experimental concerns when jet substructure is used in searches for BSM physics. In the past, some effort has been made to use theoretically stable observables like the soft drop mass and energy correlation fractions. However, any tagger can in principle be a black box as long as its signal efficiency is measured in data.

In a sample of $t\bar{t}$ -enriched lepton + jets events, the decay of two top quarks results in a final state with two b quarks and two W bosons; one W boson decays leptonically and the other decays to hadrons. This topology provides a relatively pure source of W jets in data and, after background subtraction in a likelihood fit, is used to compare the efficiencies of W tagging in data and in simulation. This correction factor is also used for Z tagging and, with additional systematics, for the $H \rightarrow b\bar{b}$ tagging.

The kinematic tail of the SM production of $t\bar{t}$ pairs results in a sizable number of boosted top quarks, which can be used for calibrating the boosted top taggers. However, care needs to be taken if the semileptonic sample is a part of the measurement (e.g., for $Z' \rightarrow t\bar{t}$) since then the measurement of the signal efficiency is done simultaneously with the extraction of the signal yield.

Unfortunately, for the various Higgs taggers there are no standard candles, and the situation is more complicated. DeepAK8 and ParticleNet $H \rightarrow b\bar{b}$ taggers are calibrated in data using the merged $g \rightarrow b\bar{b}$ events as a proxy for $H \rightarrow b\bar{b}$ jets. A dedicated BDT has been developed to select $g \rightarrow b\bar{b}$ events that resemble $H \rightarrow b\bar{b}$ jets. This BDT also provides a handle on the level of similarity between the proxy jets and the signal jets; varying it allows an estimate of an uncertainty on the signal efficiency. The distributions of the secondary vertex mass in pass and fail events are fitted simultaneously to extract the signal efficiency. The $H \rightarrow WW^* \rightarrow 4q$ tagger, used in the triboson search (Section 4.4.2), is calibrated on the sample of four-prong boosted top jets with an additional radiated gluon.

3.6. Background Estimation in Searches with Substructure

The background estimation methods used in the searches with substructure are broadly similar to procedures used elsewhere in collider physics except that the addition of substructure-related variables provides new handles on SM backgrounds, allowing for the definition of control regions to estimate or constrain them. In most cases, the dominant backgrounds are $t\bar{t} + \text{jets}$, as estimated from simulation-assisted approaches, and the QCD production of quark- and gluon-initiated events. The latter is estimated from data given that its modeling is not entirely reliable, especially in the kinematic regime probed by heavy resonances. While there are many procedures for data-driven background estimation, they generally fall into two types: a background estimate based on a transfer function (TF), and a bump hunt.

3.6.1. Bump hunt and its derivatives. A resonance has a pole mass, and thus the reconstructed mass usually has a peak near it, typically with a longer tail toward the lower masses. In contrast, SM backgrounds do not have a preferred mass value, and the distribution of reconstructed masses is smooth. This is the essence of the bump hunt approach to the background prediction and the signal extraction: The signal is a peak on top of a smoothly varying background. As a technique, the bump hunt is ubiquitous in particle and nuclear physics.

Most often, the bump hunt is implemented by analytically parameterizing the background probability density function (PDF). This approach is easy to set up and use, although it requires special handling, in particular when determining the optimal number of parameters (via an F-test) and when studying sources of systematic uncertainty, effectively marginalizing over different analytical functions and possibly creating a spurious signal shape (Section 4.2.1) by an undulating background shape.

A special case of the bump hunt is when the shapes of the background components can be described reasonably well by simulation, but the normalization is only loosely constrained. The best example of this is the $Z' \rightarrow t\bar{t}$ search in a lepton + jets sample: The background is dominated by $t\bar{t} + \text{jets}$, but the signal region itself is the main source of our knowledge of $t\bar{t} + \text{jets}$ in the high- $m_{t\bar{t}}$ regime. In this case, the background is modeled by templates obtained from the simulation; however, the templates are allowed to vary, guided by a number of nuisance parameters. The background normalization itself is determined in the fit from the sidebands around the hypothesized signal peak.

In recent years, several analyses highlighted in this review have employed the bump hunt in two and three dimensions (Section 4.2.2). These analyses either use distributions of simulated events to motivate analytical background shapes or use the simulation directly by building smooth

templates from smeared truth information. The best fit to data with multidimensional templates then amounts to tuning the Monte Carlo generators but doing it in situ. The multidimensional bump hunt techniques often result in reduced background uncertainty and improved signal significance, at the cost of the substantial human effort required to set them up properly.

3.6.2. Background estimates based on a transfer function. The other class of background estimation procedures is based on the concept of a TF. The most basic variant is the ABCD method; it requires a pair of variables for which the joint 2-dimensional PDF can be written as a product of two 1-dimensional PDFs. Then the number of events in region A can be written as $N_A = N_B(N_C/N_D)$, where N_C/N_D is essentially the TF that is measured in regions C and D, but it allows the prediction of the yield in region A based on the yield from region B.

Many variants of the TF-based background estimate are in use. The plain ABCD approach can be evaluated directly in each bin of the invariant mass distribution. If the TF is based on jet substructure, in some cases it needs to be parameterized as a function of jet p_T and η . The TF can also be obtained from simulation (known as the α method).

A special extension of the ABCD approach is the alphabet procedure, where the TF measured in regions C and D varies as a function of the search variable. For example, a correlation between the substructure discriminants like τ_{32} and the jet mass arises from the gluon radiation that contributes to both. In the alphabet procedure, the ratio of the events that pass and fail the substructure selection is measured in the mass sidebands of the jet mass (of a top or Higgs jet) and is applied to the failing events in the jet mass signal region to predict the passing events in the signal region. Mass decorrelation of the jet taggers simplifies this dependence, both reducing the statistical uncertainty on the background estimate and making the alphabet procedure particularly suitable for these kinds of taggers.

4. EXAMPLES OF RECENT SEARCHES WITH SUBSTRUCTURE

This section considers concrete examples of the experimental principles for heavy resonance searches presented above, sampling different models, final states, jet-tagging techniques, and background estimation procedures.

4.1. Searches for Resonances with Top Quarks in the Final State

Merged top quark jets have rich substructure, and the first jet taggers were employed in the resonance searches involving at least one top jet. Examples of such searches are given below.

4.1.1. $Z' \rightarrow t\bar{t}$. Many BSM models prominently involve couplings of top quarks to new states, and thus they can be used as probes for new phenomena at the TeV scale. Models such as the two-Higgs-doublet model (2HDM) (51), the TC2 model (2, 3, 52), and RS models of warped extra dimensions (5, 53) all provide heavy states decaying to a $t\bar{t}$ pair. Experimentally, the limits are set either on a generic narrow Z' resonance or on KK gluon or KK graviton states from the RS model. The $Z' \rightarrow t\bar{t}$ final state was targeted from the very beginning of the LHC data analysis, and these searches were the first heavy resonance searches using jet substructure that were published with the LHC data (54, 55). Since then, both experiments have put out several updates with increased beam energy and integrated luminosity.

Three $t\bar{t}$ final states are of interest: dileptonic, where both top quarks decay semileptonically; semileptonic, where one top quark decays hadronically and the other to $\nu\ell b$; and hadronic (or all-hadronic or all-jets), where both top quarks decay to quarks. The latter two channels drive the sensitivity at higher masses; they have a nearly equal share of the $t\bar{t}$ branching ratio. Because

of the presence of a high- p_T lepton, the semileptonic channel has a low contribution from QCD multijets, and the background is dominated by the SM $t\bar{t}$ and $W + \text{jets}$ productions. Searches in this channel generally use jet substructure only for the event categorization. Background estimates for the semileptonic channel are based on templates from simulation, morphing in the fit (Section 3.6.1). Somewhat surprisingly, after additional selection based on substructure discriminants, the hadronic channel reaches the sensitivity of the semileptonic one. Historically, this fact has notably increased the interest in the use of the jet substructure in searches for new BSM states.

While the ATLAS Collaboration and the CMS Collaboration both published searches based on the 2016 data, ATLAS has a result for the full Run 2 in the hadronic final state (56) that provides an excellent example of how the searches for heavy resonances are performed. This analysis uses two large- R jets, identified with a combination of the DNN top tagger (Section 3.5) and the variable- R track jets associated with the large- R jet that are b tagged using the DL1 algorithm. Events containing charged leptons (electrons or muons) are removed to keep this search statistically disjunct from the $t\bar{t}$ search in the lepton + jets channel, thus allowing for a future combination. The two leading large- R jets need to be back-to-back in the transverse plane ($|\Delta\phi| > 1.6$), and the difference of their rapidities, $|\Delta y|$, must be less than 1.8; because of the large resonance mass, this cut is efficient for the signal but suppresses the QCD multijet background, dominated by processes with a t -channel gluon exchange. Similar versions of these two cuts are applied in nearly all searches involving two jets described in this review. The two large- R jets are required to be top tagged; this is sufficient since the top tagging has $\approx 80\%$ efficiency for the $Z'_{\text{TC2}} \rightarrow t\bar{t}$ signal over the whole mass range. The events are then divided into two signal regions, SR1 b and SR2 b , depending on whether one or both large- R jets are associated with a b -tagged variable- R jet. The product of geometric acceptance and selection efficiency, for the union of SR1 b and SR2 b , exceeds 10% up to 6 TeV.

The dominant QCD background is predicted from data following the bump hunt approach, where the background is parameterized with a smoothly falling analytical function. However, the uncertainty in the background arises from the choice of functional form and fit range. To choose the appropriate functional form, QCD background is also estimated by multiplying the distribution of $m_{t\bar{t}}$ from the CR where both large- R jets fail the b tagging by the TF chosen to model the b -tagging false positives in an appropriate signal region, measured in events where one of the large- R jets fails the top tagging. Then, the resulting prediction of the $m_{t\bar{t}}$ distribution is extensively studied with different functional forms, and a three-parameter function $p_0(1-x)^{p_1}x^{p_2+p_3}\log(x)$ is chosen, where $x = m_{t\bar{t}}/\sqrt{s}$ and p_i are free parameters. This form is used to fit data in both SR1 b and SR2 b to estimate the background in each. An example of the fit to data distribution of $m_{t\bar{t}}$ in SR2 b is shown in **Figure 2**. An additional uncertainty in the background modeling due to a spurious signal can arise as a bias in the signal estimate obtained from a signal + background fit to the $m_{t\bar{t}}$ distribution under the background-only hypothesis. This uncertainty is explicitly included in the likelihood fit to data. The upper limits on $\sigma B(Z' \rightarrow t\bar{t})$ are provided only up to 5 TeV for the Z'_{TC2} signal mass because of the large spurious-signal uncertainty (exceeding 200%) at masses beyond 5 TeV, making the limit calculation unreliable beyond ≈ 5.2 TeV. The upper limits reach below ≈ 3 fb and result in the exclusion of Z'_{TC2} masses up to 3.9 for decay width of 1% and 4.7 TeV for decay width 3%. The sensitivity of this search is limited by the statistical uncertainty of the background estimation up to ≈ 4.5 TeV, where the systematic uncertainty due to the spurious signal begins to dominate.

The CMS search (57), based on 2016 data, considers all three exclusive final states. In the dilepton and the semileptonic channels, the leptons are required to be nonisolated but pass additional selection; the backgrounds are obtained from simulation templates fitted to data. In the hadronic channel, the CMS top tagger is used in combination with the subset b tagging and τ_{32} to define




Figure 2

(a) Observed reconstructed $m_{t\bar{t}}$ distribution in data for the region SR2b with the background fit overlaid. The predicted Z'_{TC2} signals with masses of 2 and 4 TeV (scaled up by a factor of 5) are superimposed on the background prediction. Panel adapted from Reference 56 (CC BY 4.0). (b) Observed limits on the heavy vector triplet model at 95% CL in the g_F versus g_H plane for W' resonance masses of 2, 3, and 4 TeV. The circles indicate the coupling values for Models A and B, and the gray region corresponds to the area of phase-space where the decay width of the resonance is no longer negligible and the signal m_{WH} shape is no longer expected to be dominated by the experimental resolution. Panel adapted from Reference 67 (CC BY 4.0).

signal regions of varying purity; the background is estimated from the preselection control region and weighted by a TF. Upper limits on $\sigma\mathcal{B}(Z' \rightarrow t\bar{t})$ are derived for a leptophobic topcolor Z' resonance with widths of 1, 10, and 30% relative to its mass, and masses up to 3.80, 5.25, and 6.65 TeV are respectively excluded. KK gluons in the RS model are excluded up to 4.55 TeV. At the highest resonance masses, the larger-width searches are limited not by statistics but by the fact that the signal shape at such high masses ceases to be a resonance: The falling parton density function enhances the low-mass tail of the relativistic Breit-Wigner distribution, producing a shape where the bulk of the signal yield is in the tail and not localized in a peak anymore.

4.1.2. $W' \rightarrow t\bar{b}$. The search for $W' \rightarrow t\bar{b}$ is in many ways similar to that for $Z' \rightarrow t\bar{t}$: The decay of the top quark can be either semileptonic or hadronic. For high-mass W' decays, the top is boosted and hadronic decays are reconstructed with top jet tagging, and the small- R jet recoiling against it is b tagged (103, 104). However, in this case there is only one top quark, so the semileptonic channel has an edge. In the hadronic channel, the contribution from $t\bar{t} + \text{jets}$ is suppressed by requiring the small- R jet to have a low jet mass; inverting this selection allows for a definition of a good control region enriched in $t\bar{t}$. The right-handed W'_R does not mix with the SM single top production, so it is a straightforward bump hunt on top of a smoothly falling background. The left-handed W'_L does, and it also results in weaker limits at higher masses.

4.1.3. b^* , B' , and T' searches. Most searches for the vector-like quarks (VLQs) T' and B' involve pair production since the QCD production cross section dominates at medium masses. However, T' and B' can also be singly produced with an accompanying b or t quark; this production

mode is of interest when searching for heavier VLQ masses since only one heavy state is produced. A single VLQ is a heavy resonance that decays to boosted top/ $W/Z/H$ jets; these searches are conceptually similar to searches for Z' and W' bosons and diboson searches described below: Two boosted objects are nearly back-to-back in the transverse plane, and they are reconstructed using the usual jet substructure techniques. As examples, we may consider a search for a single B' from the ATLAS experiment and searches for a single T' and an excited b^* quark from the CMS experiment.

4.1.3.1. ATLAS $B' \rightarrow bH$ search. The ATLAS experiment searched for $B' \rightarrow bH$ (58) where the b quark from B' decay is reconstructed as a small- R -tagged jet and the boosted $H \rightarrow b\bar{b}$ candidate is reconstructed as a large- R jet. The events are categorized depending on whether one or two variable- R track jets matched to the large- R jet are b tagged by a version of the DL1 algorithm (34) specifically trained for these kinds of jets. The jet mass of the Higgs candidate jet must be in the range of 105 to 135 GeV. A kinematic cut on the ratio between the Higgs candidate p_T and the reconstructed VLB invariant mass, $p_T^H/m_B > 0.4$, is sensitive to whether the QCD background event is produced as a result of an s - or t -channel process.

The VLQ mass, m_B , is used as the search variable. The dominant QCD background is predicted using a modified ABCD-like procedure, where the events in the signal region, with present forward (high- $|\eta|$) jets and a b -tagged small- R jet from the B' decay, are estimated using the events without forward jets and the events in which the jet from B' decay has not been b tagged.

4.1.3.2. CMS $T' \rightarrow tZ, Z \rightarrow \nu\bar{\nu}$ search. The CMS experiment has searched for $T' \rightarrow tZ$ in the final state with a boosted hadronically decaying top quark and with Z decaying to neutrinos (59). This channel is quite sensitive since a large cut on p_T^{miss} is effective in suppressing the backgrounds from both QCD multijet and also $t\bar{t} + \text{jets}$: The remaining dominant SM background is $Z + b\bar{b}$, where Z decays to neutrinos. However, it is still hard for the b quarks that recoil against the Z boson to be aligned with enough additional gluon radiation to fake a hadronic top quark decay. This is the most sensitive search in the tZ channel.

4.1.3.3. CMS search for an excited bottom quark. An excited bottom quark, b^* , can be produced by an interaction of a sea b quark and a gluon, and thus with a higher cross section than single VLQs. However, in the model of interest, b^* dominantly decays to a top quark and a W boson. The CMS experiment has two analyses. The hadronic channel (60) requires two large- R jets; τ_{32} and subjet b tagging are used for the top jet identification, and τ_{21} is used for tagging the hadronic W jet on the other side. This analysis is the first to apply the alphabet background prediction technique in two dimensions, as a function of the top jet mass and the invariant mass of the tW candidate. The two-dimensional distribution of the events that fail the top tagging is multiplied by a parametric two-dimensional fail-to-pass TF to predict the backgrounds that pass the top tagging. The fit parameters are obtained in situ from the fit to data, thereby providing an m_{tW} -dependent interpolation through the top mass signal region.

The CMS detector also has a semileptonic $b^* \rightarrow tW$ channel (61) in which the W boson decays leptonically to an electron or a muon and missing momentum. This signature results in striking events where a hadronic top jet is balanced by a single, energetic, very isolated lepton with no additional activity nearby. The top jet is reconstructed using the HOTVR algorithm (Section 3.5), and one of the HOTVR subjets must be b tagged. The semileptonic channel is as sensitive as the hadronic channel since the isolation requirement eliminates nonprompt leptons from heavy flavor decays from the QCD multijet background, and the veto on the second b -tagged jet suppresses $t\bar{t} + \text{jets}$. This leaves $W + b\bar{b}$ as the leading background, which is reduced by the selections on the jet mass and substructure.

4.2. Searches for Diboson Resonances

Heavy resonances decaying to a pair of SM bosons— W , Z , H , or γ —are well motivated in many BSM models and represent a large subset of the searches with substructure. Since the hadronic W and Z decays to light quarks have nearly identical substructure and mostly overlapping mass windows, they are often collectively referred to as V to highlight that both final states are included. These searches sample many experimental techniques, and several are examined here in more detail.

4.2.1. $\gamma + V/H$ resonances. As the first example of a diboson resonance search, we may consider the case in which one of the bosons is a photon and the other is a W , Z , or Higgs boson—a case that naturally arises in several models (8, 62). Final states with an energetic, isolated photon are attractive because this selection is efficient for the γV signal but effectively suppresses the otherwise dominant SM background sources like QCD multijet, $t\bar{t} + \text{jets}$, and $V + \text{jets}$ production, leaving $\gamma + \text{jets}$ as the leading background component. For $m_X \gtrsim 1$ TeV, decays to γV result in a boosted V boson. In the case of hadronic $V \rightarrow q\bar{q}'$ decays, jet substructure can be used to suppress the combinatorial background from other jets in a $\gamma + \text{jets}$ event. Both γ and the V jet are required to be in the central η region, as that favors the signal.

The CMS search in this channel looked for $W' \rightarrow W\gamma$ (63), where $W \rightarrow q\bar{q}'$ is identified using the τ_{21} variable. The background, dominated by $\gamma + \text{jets}$, is modeled by a smoothly falling analytical function.

4.2.1.1. ATLAS $V\gamma$ search. The ATLAS analysis (64) in this channel is more elaborate, and it is a good example of an application of the bump hunt approach. The search is optimized for spin-0, spin-1, and spin-2 states. The size of the V jet mass window increases from about 20 to 50 GeV as p_T increases from 500 to 2,500 GeV. Two analyses are performed, one for $W\gamma$ and the other for $Z\gamma$ final states. Both require the V jet mass to be in the W or Z signal mass window and classify the events into multiple categories based on the D_2 variable (Section 3.4.2), which is sensitive to the two-prong substructure. The $Z\gamma$ selection also has a high purity category for $Z \rightarrow b\bar{b}$ events with two b -tagged track jets matched to the V jet. The total signal efficiency ranges from about 20% at low masses to as high as 60% at 6.8 TeV. This allows an improved search sensitivity since the SM backgrounds drop as a function of the invariant mass of the $V\gamma$ system ($m_{J\gamma}$), while the signal efficiency is increased at higher masses where it is at a premium.

The signal is modeled as a peak in the $m_{J\gamma}$ distribution, on top of a smoothly falling background, parameterized with an analytical function flexible enough to accommodate the shape in each category. The function chosen to model the background is the same as in the $Z' \rightarrow t\bar{t}$ search. The likelihood fit also includes spurious signal as a source of possible uncertainty on the signal yield. Events in which the photon is in the forward pseudorapidity region are used to confirm that the chosen functional form is flexible enough to model the $m_{J\gamma}$ distribution in data.

The upper limits on $\sigma(pp \rightarrow X) \times \mathcal{B}(X \rightarrow V\gamma)$ as a function of m_X are evaluated for spin-0 and spin-2 $gg \rightarrow X^0 \rightarrow Z\gamma$, spin-2 $q\bar{q} \rightarrow X^0 \rightarrow Z\gamma$, and spin-1 $q\bar{q} \rightarrow X^\pm \rightarrow W^\pm\gamma$. All limits are similar and range from 10 fb at 1 TeV down to about 0.5 fb at 7 TeV.

4.2.1.2. ATLAS $H\gamma$ search. The ATLAS experiment also searched for a resonance in the $H\gamma$ final state (65), employing the novel CM Higgs tagger (Section 3.5.1) for the identification of boosted $H \rightarrow b\bar{b}$ decays. The Higgs jet mass window is optimized separately for each Z' hypothesis, and the width of the Higgs mass window increases from 30 GeV at $p_T = 0.5$ TeV to 70 GeV at $p_T = 2$ TeV. The signal extraction also follows the bump hunt approach with an analytically

parameterized background PDF. Upper limits on $\sigma(pp \rightarrow X) \times \mathcal{B}(X \rightarrow H\gamma)$ for narrow spin-1 resonances decrease from 11.6 fb at 0.7 TeV to 0.11 fb at 4 TeV.

4.2.1.3. CMS $H\gamma$ search. The CMS experiment also has dedicated searches for $H\gamma$ resonance (66) based on the 2016 data. The $H \rightarrow b\bar{b}$ decays reconstructed as large- R jets are identified using the double- b Hbb tagger (Section 3.5.1), although the analysis also includes an untagged category. This category retains optimal sensitivity to resonances above 2 TeV: In this regime, the efficiency of the double- b tagger deteriorates as it becomes harder to resolve the tracks originating from secondary vertices; at the same time, the background is small enough that even untagged events contribute to sensitivity. Like in the ATLAS analysis, the background estimate is based on an analytically parameterized function, although a different functional form is used. Upper limits on the production cross section of $H\gamma$ resonances ranging from 25 to 0.4 fb are set as a function of the resonance mass in the range of 720 to 3,250 GeV.

4.2.2. VV and VH resonances. Resonances decaying to two heavy bosons, W , Z , or H , have seen some of the earliest use of the substructure. Over the past decade, all possible decays have been covered, from fully hadronic VV and VH searches to semileptonic ones (with $W \rightarrow \nu\ell$, $Z \rightarrow \ell^+\ell^-$, and $Z \rightarrow \nu\bar{\nu}$, where ℓ is an electron or a muon). Both $H \rightarrow b\bar{b}$ and $H \rightarrow WW^*$ have been used, with the latter also including the final states with leptons. The following subsections take a closer look at the WH analysis from the ATLAS experiment and a combined WV and WH search from the CMS experiment since they provide examples of very different technological approaches, even when the final states are the same.

4.2.2.1. ATLAS WH search. The ATLAS experiment has performed a search for a W' decaying into a W boson and a 125-GeV Higgs boson H in the $\ell^\pm \nu b\bar{b}$ final state (67). The data set is based on leptonic triggers using both isolated and nonisolated leptons.

The $W \rightarrow \nu\ell$ candidate is reconstructed from the lepton and p_T^{miss} by imposing a W boson mass constraint on the lepton–neutrino system and solving the quadratic equation for the z component of the neutrino momentum. To reconstruct the merged $H \rightarrow b\bar{b}$ decay, two leading variable- R track jets associated with the large- R jet are considered for b tagging, and both $1b$ -tag and $2b$ -tag categories are used in the analysis. The $1b$ -tag category helps with the signal efficiency at large WH invariant masses, but it has larger background levels and a different background composition. Events with a b -tagged variable- R jet outside the large- R jet are vetoed to suppress the $t\bar{t} + \text{jets}$ background. The efficiency of the merged category increases with the mass and dominates for masses above approximately 1.4 TeV. For the merged selection in the $1b$ -tag category, the efficiency plateaus around 20%, and the total reconstruction efficiency from all categories is around 30%.

The background composition in the signal region is dominated by $W + \text{jets}$ and $t\bar{t}$ with a nonnegligible fraction of $Z + \text{jets}$ events passing the selections. The simulated $W + \text{jets}$ events are split into multiple components: $W + \text{hf}$ (dominates $2b$ tags) and $W + \text{hl}$ with one heavy (b or c) and one light jet. The $t\bar{t}$ and single t processes are combined into one component (top background). The normalizations of the top, $W + \text{hf}$, and $W + \text{hl}$ backgrounds are determined in the fit to the m_J sidebands in data. The 95% CL upper limits between 1.3 pb and 0.56 fb are placed on $\sigma \times BR(W' \rightarrow WH)$ in HVT models. Limits as a function of g_F and g_H couplings are shown in **Figure 2**.

4.2.2.2. CMS WV and WH search, including vector boson fusion production. A recent CMS report (68) combines the searches for $X \rightarrow WV$ and $X \rightarrow WH$ resonances since both $V \rightarrow qq'$ and

$H \rightarrow b\bar{b}$ decays are considered as separate event categories. The lepton and p_T^{miss} selection, as well as the kinematic reconstruction of the $W \rightarrow \nu\ell$ decay, proceeds similarly to the ATLAS analysis.

The main difference is in the deployment of substructure. To reconstruct two-prong $V \rightarrow qq'$ decays, this search uses τ_{21}^{DDT} , a mass-decorrelated variant of the τ_{21} variable (Section 3.4.1), because the mass of the large- R V or H jet is one of the variables used in the two-dimensional fit. For the $H \rightarrow b\bar{b}$ category, the double- b tagger (Section 3.5.1) is used for H or Z bosons decaying to $b\bar{b}$. Events with b -tagged small- R jets are enriched in $t\bar{t} + \text{jets}$; they are removed, and this category provides a control region used to constrain the $t\bar{t}$ component in the fit. A VBF-tagging criterion is defined as $m_{jj} > 500$ GeV and $|\Delta\eta_{jj}| > 4$, where m_{jj} and $|\Delta\eta_{jj}|$ are the invariant mass and pseudorapidity separation of the two highest- p_T small- R jets. The use of a large window for m_J allows the selection of background events containing V jets as well as top quark jet candidates while retaining sizable sidebands to constrain shapes and normalizations. The events are also divided according to rapidity, y : The low Δy region corresponds to a difference in rapidity between the reconstructed bosons of $|\Delta y| \leq 1$, and the high Δy region corresponds to $|\Delta y| > 1$. The overall signal selection efficiency times acceptance ranges from ≈ 20 to 80%, depending on the benchmark model and increasing with resonance mass.

The signal extraction, along with the background estimation, is obtained by a simultaneous maximum likelihood fit to the two-dimensional m_{WV} versus m_J data distributions in the 24 search categories. The templates for the signal and background processes are constructed from simulation. Analytical shapes are used to model the signal, while binned templates are used for background components. Particular care is devoted to constructing smooth background templates, modifying the strategy to accommodate the larger 2D signal region and the fact that new categorization criteria such as VBF tagging and double- b tagging result in low statistics in some categories.

4.3. Di-Higgs and Di-Higgs-Like Resonances

Searches for HH resonances in the high- m_{HH} regime are well motivated not only because of the sensitivity to radion or bulk gravitons but also because the tail of this distribution is sensitive to BSM contributions and is even able to aid the nonresonant search for the SM HH production, which is one of the main goals of Runs 3 and 4 of the LHC.

For these reasons, nearly all permutations of the Higgs boson's decays have been explored, especially at lower HH invariant masses. In the boosted regime, the backgrounds are lower, and thus all resonant searches consider one Higgs boson that decays to a $b\bar{b}$ pair. They are classified as follows according to the decay of the other Higgs boson.

- $X \rightarrow HH \rightarrow 4b$: The other Higgs boson decays to a $b\bar{b}$ pair. This channel has the highest overall branching ratio but also the largest background from the QCD multijet production of the heavy flavor.
- $X \rightarrow HH \rightarrow bbWW^*$: The other Higgs boson decays to WW^* , which further decays either to two leptons, or to a lepton and a boosted quark pair. In this case the dominant backgrounds are $t\bar{t} + \text{jets}$ and $W + \text{jets}$.
- $X \rightarrow HH \rightarrow bb\tau\tau$: The other Higgs boson decays to a $\tau^+\tau^-$ pair. The overall branching ratio is the lowest of the three, but the backgrounds are quite suppressed.

4.3.1. $HH \rightarrow bb\tau\tau$. The ATLAS Collaboration has published a search for $X \rightarrow HH \rightarrow bb\tau\tau$ (69). The key ingredient is the new di- τ tagger that can identify hadronically decaying $\tau^+\tau^-$ pairs with a large Lorentz boost; a new approach was necessary since, for $m_X > 2$ TeV, more than 50% of the $\tau^+\tau^-$ pairs have $\Delta R(\tau^+, \tau^-) < 0.4$.




Figure 3

(a) Distributions of the visible mass of the di- τ object in the region with no b tagging in the $HH \rightarrow b\bar{b}\tau\tau$ analysis. All simulated events containing a generator-level $\tau^+\tau^-$ pair matched to a simulated di- τ object are referred to as true di- τ . Panel adapted from Reference 69 (CC BY 4.0). (b) The fit result compared with data projected onto $m_{b\bar{b}}$ for the single lepton channel in the $X \rightarrow HH \rightarrow b\bar{b}WW^*$ search. Panel adapted from Reference 71 (CC BY 4.0). Abbreviations: LP, low purity; SF, scale factor.

A BDT discriminant is built¹ using information from the two leading subjects: the shapes of clusters in the calorimeter, and tracks and vertices from the track jets matched to the two subjects. The tracks found in the isolation region (namely, the area of the large- R jet excluding the di- τ subjects) are also used. Hadronic decays of τ leptons result in a neutrino and one or three pions, thus producing one or three isolated tracks that carry most of the momentum of their subject. In contrast, the jets from the QCD background have a larger fraction of energy in the isolation region, fewer collimated tracks, and more total tracks, each of which carries a smaller fraction of the jets' transverse momentum. In training, the p_T spectra of the di- τ system of both signal and background are reweighted so that they are flat, in order to reduce the dependence on the p_T (see Section 3.3.3).

Figure 3 shows the distribution of visible di- τ mass, calculated from the observed objects, in preselected events. A broad but clear $Z \rightarrow \tau^+\tau^-$ peak demonstrates the success of this approach.² The background estimate is based on the misidentification rate of the di- τ tagger, measured in a large multijet sample. For the HH resonance, a heavy, narrow scalar resonance produced via gluon–gluon fusion is used. For $1.2 < m_X < 3$ TeV, the observed upper limits lie between 94 and 28 fb. Below 2 TeV, the limits deteriorate mostly because of the uncertainties in the di- τ reconstruction.

4.3.2. $HH \rightarrow b\bar{b}WW^*$. The CMS experiment has a search for $X \rightarrow HH$ where one H boson decays to a $b\bar{b}$ pair, and the other to a WW^* , with at least one W boson subsequently decaying

¹A similar algorithm was implemented by the CMS Collaboration (70).

²Note that the peak is shifted down in mass because the plotted quantity is the invariant mass of the observed objects.

leptonically (71). The $H \rightarrow b\bar{b}$ decay is reconstructed as a single large- R jet and is tagged with the DeepAK8 Hbb mass-decorrelated tagger (cf. Section 3.5.1). In the single lepton category, one lepton that passes the mini-isolation criterion is required. The remaining two quarks form a large- R jet with a two-prong substructure. The dilepton category requires two opposite-sign leptons and missing momentum in the same hemisphere. This category is also sensitive to $H \rightarrow \tau\tau \rightarrow \ell\nu\ell\nu\nu$ decays.

The signal is extracted using a two-dimensional maximum likelihood fit of the $H \rightarrow b\bar{b}$ jet mass and HH invariant mass distributions, using the same approach employed in the CMS WH resonance search (Section 4.2.2.2). The projection onto the H jet mass of 1 of the 12 signal categories is shown in **Figure 3**. Model-independent exclusion limits are evaluated for spin-0 and spin-2 massive bosons decaying to HH . The results are interpreted in the context of radion and bulk graviton production in models with a warped extra spatial dimension.

4.3.3. $HH \rightarrow 4b$. Both the ATLAS and CMS experiments have results in the channel with two large- R jets, where each jet encapsulates a merged $H \rightarrow b\bar{b}$ decay. The kinematic event preselection is nearly the same: It requires two large- R jets above $p_T > 450$ GeV (due to the trigger), and a $|\Delta\eta|$ between them that is less than 1.3, in order to suppress the QCD multijet background. The main backgrounds are QCD and $t\bar{t}$ + jets, although their proportions (and thus importance) differ. The basic idea for the background estimates is nevertheless the same: The QCD multijet background is estimated from data by deriving a TF from the events that fail the tagging and those that pass in the control region defined by the jet mass of a Higgs candidate. This TF is used to reweight the events that fail the tagging in the Higgs jet mass signal region. The $t\bar{t}$ + jets component is obtained from simulation corrected from data control regions.

4.3.3.1. ATLAS $HH \rightarrow 4b$ search. The ATLAS search for HH resonances in the $4b$ final state (72) contains both a resolved channel and a boosted channel. The boosted H jet candidates are identified with b tagging of the variable- R track jets associated with a large- R jet. At moderate Higgs boosts, the two b hadrons are reconstructed separately but still within the same large- R jet. At large boosts, the b quarks are close enough that their decay products are reconstructed as a single track jet. To maximize the efficiency of reconstructing $H \rightarrow b\bar{b}$ decays over a large range of HH invariant masses, three event categories are defined: $4b$, with two b -tagged track jets associated with each H jet; $3b$, where one H jet has only one b -tagged track jet; and $2b$, where there is exactly one b -tagged track jet per H candidate. The signal efficiency in the $4b$ category peaks around 1.5 TeV. Above that jet p_T , the track jets begin to merge, so for higher m_{HH} , the $3b$ and $2b$ categories are more efficient. At very large boosts, most signal events contain only one track jet per H jet candidate, so the otherwise less performant $2b$ category is used only for signal hypotheses with $m_X \geq 2$ TeV.

The dominant background, the QCD multijet production, is estimated from data. Three additional low-tag categories are defined by inverting the requirements on the b tagging in one of the two Higgs candidates. One low-tag category is defined for each signal category. The events in the low-tag category are reweighted by a TF to produce the predicted QCD multijet distribution of the HH invariant mass in the corresponding signal category. The derivation of the TF, as well as the normalizations of the QCD multijet and $t\bar{t}$ + jets backgrounds, is done in the control region. The signal region, the validation region, and the control region are defined in the two-dimensional plane of the invariant masses of two large- R jets, as shown in **Figure 4a**. The subdominant contribution from $t\bar{t}$ + jets is derived from simulation. The normalization of the QCD multijet and $t\bar{t}$ + jets components in each category is obtained from a fit to the distribution of the leading H candidate jet mass in the control region. Finally, the HH invariant mass




Figure 4

(a) Kinematic region definitions superimposed on the data in the $2b\text{-}1f$ category of the ATLAS $HH \rightarrow 4b$ search. H_1 and H_2 are the reconstructed Higgs boson candidates. The CR is shifted to higher masses relative to the SR and VR to maximize the number of selected events while avoiding the overwhelming QCD multijet background at lower jet mass values. Panel adapted from Reference 72 (CC BY 4.0). (b) The distributions of the H and the Y candidate jets' ParticleNet scores for the signal with $m_X = 1,600$ GeV and $m_Y = 90$ GeV (filled squares) and multijet background (open circles), in the CMS $X \rightarrow HY \rightarrow 4b$ search. The grid lines show the different event categories defined using the ParticleNet scores of the two jets. Panel adapted from Reference 74 (CC BY 4.0). Abbreviations: CR, control region; SB, sideband; SR, signal region; VR, validation region; VS, validation signal.

distribution of the reweighted and normalized events from the low-tag categories is smoothed by a parametric dijet function. The resolved and boosted channels are fitted together, and upper limits on $\sigma(X \rightarrow HH)$ vary from $\approx 1,000$ fb at 250 GeV to about 2 fb at 3 TeV. The bulk RS model is excluded up to ≈ 1.4 TeV.

4.3.3.2. CMS $X \rightarrow HH$ and $X \rightarrow HY \rightarrow 4b$ and $X \rightarrow YY \rightarrow 4b$ searches. The CMS experiment performed the same search (73) employing the DeepAK8 $H \rightarrow b\bar{b}$ mass-decorrelated tagger (Section 3.5.1) for both Higgs jets, and results were on par with those of the ATLAS experiment. However, the CMS analysis has been surpassed by an extended search that is also looking for the resonant production of a new massive scalar X decaying into a new scalar Y and the SM Higgs boson, with both Y and H subsequently decaying into a $b\bar{b}$ pair (74), covering $0.9 < m_X < 4$ TeV and $60 < m_Y < 600$ GeV. The $Y \rightarrow b\bar{b}$ decay is also reconstructed as a single large- R jet, and this selection is efficient only for large Y boosts, translating to an m_X/m_Y ratio larger than about seven.

Both jets are reconstructed using the ParticleNet Hbb tagger (cf. Section 3.5.1), and it is the first application of this tagger in a search for a heavy resonance. For the same signal efficiency, this tagger has twice the QCD background suppression rate compared with the DeepAK8 Hbb tagger. The $X \rightarrow HY$ events are reconstructed from two large- R jets required to pass the ParticleNet Hbb tagger; one jet must be compatible with the Higgs boson's mass, and the other jet is assumed to be Y . The two jets are used to reconstruct X . The QCD multijet background is derived from a two-dimensional pass/fail TF derived in situ; however, the requirement for one jet to be close to the Higgs mass necessitates building the TF from a control region and then fitting a correction

on top of it. **Figure 4b** illustrates both an excellent separation of signal events from the QCD multijet events and the categorization of events into two pass regions (SR1 and SR2) and the fail regions from which they are estimated (SB1 and SB2). The $t\bar{t}$ + jets component is estimated from simulation templates that also morph in the fit.

The results are interpreted as scalar resonances predicted in the next-to-minimal supersymmetric SM and the two-real-scalar-singlet extension of the 2HDM. Upper limits are placed on the production cross section as a function of the masses of X and Y . This is the first search for this process using Lorentz-boosted event topologies, and it significantly extends the sensitivity to these models. For $m_Y = 125$ GeV, it also provides significantly stronger limits for the $X \rightarrow HH \rightarrow 4b$ process.

4.3.4. $X \rightarrow \phi\phi \rightarrow 4b$. The last example of a search in the $4b$ final state covers the decay of a new resonance to a pair of Higgs-like scalars, ϕ , each further decaying to a $b\bar{b}$ pair. Such a situation arises in many BSM scenarios—for instance, when there is a spontaneously broken additional approximate global symmetry (75–80). If $m_X > 2m_\phi$, then $X \rightarrow \phi\phi$ is the dominant decay of X , while ϕ couples to fermions similarly to the SM Higgs boson. The search performed in the CMS experiment (81) focuses on $m_\phi < m_H$ and uses the double- b Higgs jet tagger (Section 3.5.1). The central idea is to switch from using the jet masses of the two jets to the average mass between the two jets and the absolute value of their difference, which should be near zero for the signal. The search is performed in m_X versus the average jet mass. The background estimate is also based on a pass/fail TF; however, it is measured in the large $|\Delta\eta|$ region, which is signal-depleted.

Model-specific exclusion limits on the production cross section of X are set in the m_X versus m_ϕ plane. The branching ratios of $X \rightarrow \phi\phi$ and $\phi \rightarrow b\bar{b}$ are assumed to be 100%. The limits range from 30 fb at $m_X = 1$ TeV, down to 1 fb at $m_X = 3$ TeV. These are the first such limits on this process.

4.4. Three-Jet Topologies and Exotic Substructure

While most heavy resonance searches assume two-body decays with two- or three-prong substructure, the next generation of searches involves either more jets in the final state or possibly yet-unexplored jet substructure. Some of these searches are described in the following section.

4.4.1. W' decaying to a vector-like quark and a third-generation quark. In a generic model with heavy partners of the Z and W bosons and of the top and bottom quarks, it is usually assumed that only one particle will be produced. However, if a heavy resonance is produced and decays to a VLQ and another quark, the final state is substantially modified and a dedicated search is warranted. The CMS experiment looked for both W' decays either to $T'b$, followed by $T' \rightarrow tH$ or tZ , or to $B't$, followed by $B' \rightarrow bH$ or bZ (82). In both cases, the final state is tbH or tbZ . Given that VLQs have been ruled out below about 1.2 TeV, for most of the allowed parameter space the signal events produce a Mercedes-sign topology with one boosted top jet, one boosted H or Z jet (decaying to $b\bar{b}$), and an energetic b -tagged small- R jet. This analysis uses the ImageTop tagger (Section 3.5), and the search variable is the invariant mass of the three candidate jets. The background estimation is based on the ABCD method in the plane defined by the discriminants of the ImageTop and the Higgs/ Z taggers. This is the first search to W' in this channel using the full Run 2 data set, and W' with these decays is excluded up to 3.2 TeV.

4.4.2. VWW resonances. The CMS experiment searched for models with a massive KK excitation of W and Z bosons, in the RS1 warped extra dimensions framework where only the electroweak fields propagate in the bulk. In this case, the radion preferentially decays to a pair of

electroweak bosons, WW , ZZ , $W\gamma$, or $\gamma\gamma$. The first two radion decays were considered, resulting in the triboson states WWW and ZWW .

Two analyses are performed: The semileptonic channel (83) considers the case in which one W decays leptonically; in the hadronic channel (84), all W bosons decay to quarks. Depending on the ratio of the masses of the W_{KK} and the radion, the decay of the latter is either resolved into two separately reconstructed W bosons or merged into one large- R jet containing either four quarks or two quarks and a lepton. In both searches, the events are divided into categories based on the kinematics of the event—namely, which W boson decayed leptonically, and whether the radion is merged or resolved. A special procedure was used to calibrate the tagging discriminants since their whole distributions were used. The tagger targeting radion jets with four merged quarks, and thus four-prong substructure, was calibrated using the top decays with additional gluon jets. The semileptonic and hadronic channels are combined at the likelihood level, and the limits on the cross section are evaluated in the W_{KK} versus radion mass two-dimensional plane. They are shown in **Figure 5a** and are the most stringent limits on the triboson resonances to date, with W_{KK} being excluded up to masses of ≈ 3.7 TeV.

4.4.3. Trigluon resonances. In the default two-brane warped dimension model where all fields propagate in the bulk, the radion ϕ dominantly couples to gluons. In this case, the dominant production mode is of a KK excitation of the gluon (KKg). KKg decays to a gluon and a radion, which subsequently preferentially decays to a pair of gluons. In this three-boson resonance, all three bosons are gluons. The CMS experiment has searched for this signature (85) as well but has focused only on the case where $m_{KKg} \gg m_\phi$ and the radion is boosted, resulting in the merging of the gluon pair into one large- R jet with a two-prong substructure. This jet is identified using the τ_{21} variable, significantly improving the sensitivity compared with the simple dijet search.

However, the τ_{21} variable is not a perfect discriminant, and when applied to a signal event, it sometimes misidentifies the radion jet. For this reason a sliding mass window is applied to either jet in the event. So, for every pair of resonance masses (m_{KKg}, m_ϕ), a new selection on the jet mass is applied, and an m_{JJ} distribution is produced and fitted with a smoothly falling analytical function. Limits on the $\sigma_{KKg} \times \mathcal{B}(\phi \rightarrow gg)$ are evaluated as a function of the KKg and radion masses. This search excluded at the 95% CL a KK gluon with a mass of 4.2 TeV for a radion mass of 0.42 TeV, and a radion with a mass of 0.74 TeV for a KK gluon mass of 3.7 TeV.

4.4.4. W'_R decaying to a heavy neutrino and a lepton. A new approach to event selection based on the substructure was employed in the CMS search (86) for the right-handed heavy partner of the W boson (W_R) decaying into a lepton and a heavy neutrino, N , in a final state consisting of two same-flavor leptons (e or μ) and two quarks. Such W_R arise in left–right symmetric models (87–90) that extend the electroweak sector of the SM by a right-handed $SU(2)$ group. These models predict a W_R that couples to the right-handed fermions, and they also explain small neutrino masses in the SM via the see-saw mechanism (91–93) by adding heavy right-handed neutrinos. The heavy neutrinos decay to a lepton and two quarks. Depending on the ratio of the masses of W_R and N , the decay of N can be resolved or merged; this review focuses only on the latter.

Experimentally, identifying prompt leptons in a hadronic environment is challenging because of a large QCD production of heavy flavor, which results in nonisolated leptons. The usual solution is to select isolated leptons, where the p_T of the lepton is compared with the $\sum p_T$ of the particles in a cone around it. The cone is either fixed, or it can also shrink with the p_T of the lepton [the so-called mini-isolation (94)]. However, neither selection is efficient at very high boosts. The lepton subjet fraction (LSF) has been proposed (95) as a boost-invariant replacement for the isolation. LSF is calculated like a traditional isolation, but the effective cone size is determined




Figure 5

(a) The exclusion limit on the production cross section of W_{KK} obtained from the combination of the all-hadronic and single-lepton searches. The blue dashed lines are the boundaries between different merged and resolved decay topologies. Panel adapted from Reference 84 (CC BY 4.0). (b) The exclusion limit on the product of the production cross sections and the branching fractions of a right-handed W_R boson divided by the theory expectation for a coupling constant of the W_R equal to the standard model coupling, for the muon channel. The previous search is shown in magenta, and the biggest improvement can be seen in the $m_N < 0.5$ TeV region, where the new boosted category greatly improves the sensitivity. Panel adapted from Reference 86 (CC BY 4.0).

by the exclusive k_t clustering into n subjets. The clustering includes leptons, and locally isolated leptons dominate their own subjet. The LSF variable is thus defined as $\text{LSF}_n = \frac{p_{\text{T}}^{\ell}}{p_{\text{T}}^{\text{subjet}}}$, where n is the predetermined number of subjets and is chosen according to the signal topology. For the $N \rightarrow \ell q q'$ jet there are three subjets, and the selection is based on LSF_3 . After the full event selection, the LSF_3 distribution for the $N \rightarrow \ell q q'$ signal, as well as $t\bar{t} + \text{jets}$ with $t \rightarrow \nu \ell b$ decays, peaks at 1. Other backgrounds—most notably $W + \text{jets}$ —have a broad shape peaking at lower values since the presence of an additional high- p_{T} energetic lepton will have already eliminated the majority of the QCD multijet background. A requirement on LSF_3 is beneficial to suppress the otherwise dominant $W + \text{jets}$ component, and it reduces the uncertainty on the background estimation caused by the lack of a suitable control region to estimate this component.

The background shapes are derived from simulation, but their normalizations are obtained from a likelihood fit to the data distribution of the invariant mass of the lepton and the heavy neutrino jet. The normalization of the Drell-Yan process is unconstrained in data; however, the two W bosons in $t\bar{t} + \text{jets}$ and tW events decay independently and are constrained in the likelihood by the event yield in the $e\mu$ control region. The limits for the $\mu\mu$ channel are shown in **Figure 5b**; the limits in the ee channel are similar with the same broad features. The boosted channel (blue curve in **Figure 5b**) clearly dominates the case where the ratio of $m_{W_{\text{R}}}$ and m_N masses is large—in the lower right corner of the two-dimensional limit plot. For $m_N = m_{W_{\text{R}}}/2$ ($m_N = 200$), the mass of the W_{R} is excluded at the 95% CL up to 4.7 (4.8) TeV and 5.0 (5.4) TeV for the electron and muon channels, respectively. This analysis provides the most stringent limits on the W_{R} mass to date.

5. THE FUTURE OF THE SEARCHES WITH SUBSTRUCTURE

The LHC experiments have collected only 3–4% of the total LHC data set, so the above searches will continue to be updated throughout LHC Runs 3, 4, and beyond. The most obvious way to improve them involves building more powerful taggers for top quarks and W , Z , and Higgs bosons based on ML technology. The second generation of these ML jet taggers is already on the market (Section 3.4).

In addition, the LHC experiments are beginning to broaden the set of models considered in these searches. Models that give rise to exotic jet substructure are particularly exciting. One example is a jet from a boosted radion that decays to a WW pair with each W subsequently decaying to two quarks (Section 4.4.2). Another example is the heavy neutrino (or the R -parity-violating neutralino) decaying to $\ell q q'$ that produces a jet with a nonisolated lepton (Section 4.4.4). These two recent results represent the first generation of searches that involve BSM jets with new kinds of substructure.

However, numerous models predict jets with even more interesting substructure. For example, a boosted H^+ decaying to $t\bar{b}$ would result in a four-prong jet with two b -tagged subjets, three levels of mass hierarchy, and a rather complicated color flow. A boosted heavy Higgs boson decaying to a $t\bar{t}$ pair would produce either a six-prong jet with two b tags or a five-prong jet where one of the subjets is dominated by a locally isolated lepton. It is not hard to find examples with even more complex jet substructure. What must be emphasized is that the existing searches that fill much of this review are generally quite insensitive to these decays, and New Physics could still be hiding in plain sight.

5.1. Searches with Anomalous Substructure

I conclude this review by considering a new kind of search for heavy resonances that is expected to be much more sensitive to jets with unusual substructure. A new generation of ML algorithms

focused on anomaly detection is well suited for new heavy resonances, particularly for new heavy states in models underrepresented in the current search portfolios of the LHC experiments. These are sophisticated ML tools that use unsupervised or weakly supervised training on data; they essentially bundle the training of the new exotic taggers in situ with the background estimate in a self-consistent and largely automated way. Once operational, this kind of tool could become a one-stop shop for new heavy resonance searches with jet substructure.

5.1.1. A generic $A \rightarrow BC$ search from the ATLAS experiment. The ATLAS Collaboration deserves kudos for performing and publishing the first heavy resonance search using anomaly detection (96). This analysis has no specific signal model hypothesis; instead, it is essentially a three-dimensional search, $A \rightarrow BC$, for a new heavy resonance A with its mass in the TeV range, and BSM particles B and C with masses of several hundred GeV. Given the m_B/m_A and m_C/m_A ratios, B and C are highly Lorentz-boosted; they are assumed to decay hadronically and are reconstructed as large- R jets, and the search targets dijet topology.

Potential signals can be enhanced by classifiers trained on data using weakly supervised learning, and the features used for ML are the masses of the two jets. The search is based on the full Run 2 data set, and the dijet invariant mass spectrum covers the range from 1.8 to 8.2 TeV. Cross-section limits for narrow-width A , B , and C particles vary with their masses. In some parts of the $m_A:m_B:m_C$ space, the obtained limits are up to 10 times more sensitive compared with the inclusive dijet resonance search.

This approach is complementary to the dedicated searches in the two-large- R -jet topology discussed earlier. And even though this analysis does not use any of the jet substructure information in the training of the ML classifier, it is the first step toward the future resonance searches that will, and for that reason it has been included in this review.

5.1.2. Proposed anomaly detection methods for resonance searches. The nagging worry of the LHC physics program is “What if the New Physics is hiding in the data but we have not searched for it in the right places?” The goal of anomaly detection is to find unanticipated BSM physics by learning directly from data, thus reducing a priori bias as much as possible. The number of proposed methods is large, and here I mention only the few suitable for the heavy resonance searches with substructure.

An extension of the bump hunt dijet search using the classification without labels (CWoLa) (97) (also known as CWoLa hunting) uses NNs to identify differences between the signal region of the resonance mass spectrum and the sideband regions that surround it. The sidebands are used for supervised learning, which is then applied to the signal region. CWoLa hunting can be combined with the simulation-assisted likelihood-free anomaly detection (SALAD) (98), which uses the simulation as the reference, but a parameterized reweighting NN model is trained in the sidebands (99). This allows the CWoLa classifier to ignore the information correlated with the resonance mass by relying on SALAD to interpolate into the signal region.

In contrast with CWoLa, Tag-and-Train (100) assumes that the BSM physics will produce two anomalous objects in the event, and it assumes that those two objects are simultaneously anomalous (although not necessarily of the same kind). It uses autoencoders to prefilter events, trains two separate classifiers on the data selected by one object, and then applies each classifier to the other object. After a few iterations, the significance of the signal can be considerably improved with respect to the autoencoder alone.

Two other methods use NNs to improve and automate the background estimation. Anomaly detection with density estimation (ANODE) (101) interpolates the probability density from the sidebands into the signal region to estimate the background and then constructs a likelihood

ratio of data versus background, which is broadly sensitive to overdensities in the data that could be due to localized anomalies. The background is directly obtained from the learned densities. ANODE can enhance the significance of the bump hunt by up to sevenfold. The ABCD background estimate can also be improved with ML (102) in a procedure where the two independent classifiers used in the ABCD methods are designed using ML techniques. The state-of-the-art decorrelation methods are used to construct powerful yet independent discriminators. This approach significantly improves performance in terms of the closure of the background estimate, background rejection, and signal contamination in the control regions of the ABCD plane.

6. CONCLUSIONS

The field of searches for heavy resonances using substructure is vibrant and dynamic. Numerous analyses have been performed by the ATLAS and CMS Collaborations in the past decade; however, this review focuses only on the recent results and experimental techniques, favoring the state of the art over the historical perspective. Both experiments are developing new analysis techniques and approaches, including the new taggers. Jet substructure is a fertile ground for the use of ML, and its use will increase beyond the jet taggers into the overall analysis design. We should expect a host of new results, including models that have not been probed so far, in Runs 3 and 4 of the LHC.

SUMMARY POINTS

1. As the lighter masses of new particles in popular models of physics beyond the standard model have been excluded, the usefulness of jet substructure has dramatically increased. It is now one of the essential tools for searches at the Large Hadron Collider (LHC).
2. The ATLAS and CMS experiments have searched for a number of two-prong resonances, including $Z' \rightarrow t\bar{t}$, $W' \rightarrow tb$, diboson, and di-Higgs boson resonances.
3. Recently, new topologies (like triboson resonances) and new signatures (like radion four-prong jets or nonisolated leptons from heavy neutrino decays) have also been searched for.
4. The second generation of the machine learning (ML)-based jet taggers is now being deployed.

FUTURE ISSUES

1. Jet taggers for top quark and W and Z boson decays have natural standard candles, and their efficiencies can be measured in data. Higgs jet taggers can also be calibrated in data, but with some effort. However, jet taggers trained to identify complex cascade decays that are merged within a single jet cannot be directly calibrated in data, and new approaches will be needed.
2. Understanding and improving the quality of the Monte Carlo simulation of the jet shower may eventually be a part of the solution to the above issue. In addition, making the simulation of the jet shower more realistic will enable the development of even better ML taggers using deep neural networks.

3. The ATLAS and CMS experiments have begun exploring searches for heavy resonances using anomalous jet substructure, and this trend will continue in Run 3 of the LHC. ML will also be used more broadly to select the search variables, design new classifiers, and use them in background estimates.

DISCLOSURE STATEMENT

The author is not aware of any affiliations, memberships, funding, or financial holdings that might be perceived as affecting the objectivity of this review.

ACKNOWLEDGMENTS

The author would like to thank Arnaud Ferrari for the discussion of ATLAS results, and also Amitav Mitra and Maryam Esmat for providing feedback on a draft of this manuscript. This work was supported by the National Science Foundation (NSF) under grant PHY-2012584.

LITERATURE CITED

1. Dugan MJ, Georgi H, Kaplan DB. *Nucl. Phys. B* 254:299 (1985)
2. Hill CT, Parke SJ. *Phys. Rev. D* 49:4454 (1994)
3. Hill CT. *Phys. Lett. B* 345:483 (1995)
4. Albert A, et al. *Phys. Dark Univ.* 26:100377 (2019)
5. Lillie B, Randall L, Wang LT. *J. High Energy Phys.* 0709:074 (2007)
6. Agashe K, Davoudiasl H, Perez G, Soni A. *Phys. Rev. D* 76:036006 (2007)
7. Fitzpatrick AL, Kaplan J, Randall L, Wang LT. *J. High Energy Phys.* 0709:013 (2007)
8. Pappadopulo D, Thamm A, Torre R, Wulzer A. *J. High Energy Phys.* 1409:60 (2014)
9. Barger VD, Keung WY, Ma E. *Phys. Rev. D* 22:727 (1980)
10. Aad G, et al. *J. Instrum.* 3:S08003 (2008)
11. Chatrchyan S, et al. *J. Instrum.* 3:S08004 (2008)
12. Sirunyan AM, et al. *J. Instrum.* 12:P10003 (2017)
13. Cacciari M, Salam GP, Soyez G. *J. High Energy Phys.* 0804:063 (2008)
14. Cacciari M, Salam GP, Soyez G. *Eur. Phys. J. C* 72:1896 (2012)
15. Aad G, et al. *Eur. Phys. J. C* 77:490 (2017)
16. ATLAS Collab. *Improving jet substructure performance in ATLAS using Track-CaloClusters*. Rep. ATL-PHYS-PUB-2017-015, CERN, Geneva (2017)
17. Aaboud M, et al. *Eur. Phys. J. C* 77(7):466 (2017)
18. Krohn D, Thaler J, Wang LT. *J. High Energy Phys.* 1002:84 (2010)
19. Ellis SD, Vermilion CK, Walsh JR. *Phys. Rev. D* 81:094023 (2010)
20. Larkoski AJ, Marzani S, Soyez G, Thaler J. *J. High Energy Phys.* 1405:146 (2014)
21. Dasgupta M, Fregoso A, Marzani S, Salam GP. *J. High Energy Phys.* 1309:29 (2013)
22. Aad G, et al. *Eur. Phys. J. C* 76(11):581 (2016)
23. Sirunyan AM, et al. *J. Instrum.* 15(09):P09018 (2020)
24. Bertolini D, Harris P, Low M, Tran N. *J. High Energy Phys.* 1410:59 (2014)
25. Dolen J, et al. *J. High Energy Phys.* 1605:156 (2016)
26. Sirunyan AM, et al. *J. High Energy Phys.* 1801:97 (2018)
27. CMS Collab. *Identification of highly Lorentz-boosted heavy particles using graph neural networks and new mass decorrelation techniques*. Rep. CMS-DP-2020-002/CERN-CMS-DP-2020-002, CERN, Geneva (2020)
28. Thaler J, Van Tilburg K. *J. High Energy Phys.* 1103:15 (2011)
29. Thaler J, Van Tilburg K. *J. High Energy Phys.* 1202:93 (2012)
30. Larkoski AJ, Salam GP, Thaler J. *J. High Energy Phys.* 1306:108 (2013)

31. Larkoski AJ, Moulton I, Neill D. *J. High Energy Phys.* 1412:9 (2014)
32. Larkoski AJ, Moulton I, Neill D. *J. High Energy Phys.* 1605:117 (2016)
33. ATLAS Collab. *Optimisation and performance studies of the ATLAS b-tagging algorithms for the 2017–18 LHC run.* Rep. ATL-PHYS-PUB-2017-013, CERN, Geneva (2017)
34. Aad G, et al. *Eur. Phys. J. C* 79(11):970 (2019)
35. Sirunyan AM, et al. *J. Instrum.* 13:P05011 (2018)
36. Bols E, et al. *J. Instrum.* 15:P12012 (2020)
37. CMS Collab. *Performance of the DeepJet b tagging algorithm using 41.9/fb of data from proton-proton collisions at 13 TeV with Phase 1 CMS detector.* CMS Detect. Perform. Note CMS-DP-2018-058, CERN, Geneva (2018)
38. Aaboud M, et al. *Eur. Phys. J. C* 79(5):375 (2019)
39. Thaler J, Wang LT. *J. High Energy Phys.* 0807:092 (2008)
40. Chen C. *Phys. Rev. D* 85:034007 (2012)
41. Sirunyan AM, et al. *J. Instrum.* 15(06):P06005 (2020)
42. CMS Collab. *A Cambridge-Aachen (C-A) based jet algorithm for boosted top-jet tagging.* Rep. CMS-PAS-JME-09-001, CERN, Geneva (2009)
43. CMS Collab. *Boosted top jet tagging at CMS.* Rep. CMS-PAS-JME-13-007, CERN, Geneva (2014)
44. Kaplan DE, Rehermann K, Schwartz MD, Tweedie B. *Phys. Rev. Lett.* 101:142001 (2008)
45. Dokshitzer YL, Leder GD, Moretti S, Webber BR. *J. High Energy Phys.* 9708:001 (1997)
46. Lapsien T, Kogler R, Haller J. *Eur. Phys. J. C* 76(11):600 (2016)
47. Butterworth JM, Davison AR, Rubin M, Salam GP. *Phys. Rev. Lett.* 100:242001 (2008)
48. Chen C. *Phys. Rev. D* 92(9):093010 (2015)
49. Catani S, et al. *Phys. Lett. B* 269:432 (1991)
50. Qu H, Gouskos L. *Phys. Rev. D* 101(5):056019 (2020)
51. Branco GC, et al. *Phys. Rep.* 516:1 (2012)
52. Harris RM, Jain S. *Eur. Phys. J. C* 72:2072 (2012)
53. Randall L, Sundrum R. *Phys. Rev. Lett.* 83:3370 (1999)
54. Chatrchyan S, et al. *J. High Energy Phys.* 1209:29 (2012). Erratum. *J. High Energy Phys.* 1403:132 (2014)
55. Aad G, et al. *J. High Energy Phys.* 1209:41 (2012)
56. Aad G, et al. *J. High Energy Phys.* 2010:61 (2020)
57. Sirunyan AM, et al. *J. High Energy Phys.* 1904:31 (2019)
58. ATLAS Collab. *Search for single vector-like B quark production and decay via $B \rightarrow bH(b\bar{b})$ in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector.* Rep. ATLAS-CONF-2021-018, CERN, Geneva. <https://atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/CONFNOTES/ATLAS-CONF-2021-018/> (2021)
59. Tumasyan A, et al. (CMS Collab.) *J. High Energy Phys.* 2205:93 (2022)
60. Sirunyan AM, et al. *J. High Energy Phys.* 2112:106 (2021)
61. Tumasyan A, et al. (CMS Collab.) *J. High Energy Phys.* 2204:48 (2022)
62. Artoisenet P, et al. *J. High Energy Phys.* 1311:43 (2013)
63. Tumasyan A, et al. (CMS Collab.) *Phys. Lett. B* 826:136888 (2022)
64. ATLAS Collab. *Search for high-mass $W\gamma$ and $Z\gamma$ resonances in the hadronic final state using 139 fb⁻¹ of pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS Detector.* Rep. ATLAS-CONF-2021-041, CERN, Geneva. <https://atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/CONFNOTES/ATLAS-CONF-2021-041> (2021)
65. Aad G, et al. *Phys. Rev. Lett.* 125:251802 (2020)
66. Sirunyan AM, et al. *Phys. Rev. Lett.* 122(8):081804 (2019)
67. ATLAS Collab. *Search for heavy resonances decaying into a W boson and a Higgs boson in final states with leptons and b-jets in 139 fb⁻¹ of pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector.* Rep. ATLAS-CONF-2021-026, CERN, Geneva. <https://atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/CONFNOTES/ATLAS-CONF-2021-026> (2021)
68. Tumasyan A, et al. (CMS Collab.) *Phys. Rev. D* 105:032008 (2022)
69. Aad G, et al. *J. High Energy Phys.* 2011:163 (2020)
70. Sirunyan AM, et al. *J. Instrum.* 13(10):P10005 (2018)
71. Tumasyan A, et al. *J. High Energy Phys.* 2205:5 (2022)

72. Aad G, et al. (ATLAS Collab.) *Phys. Rev. D* 105:092002 (2022)
73. CMS Collab. *Search for resonant Higgs boson pair production in four b quark final state using large-area jets in proton-proton collisions at $\sqrt{s} = 13$ TeV*. Rep. CMS-PAS-B2G-20-004, CERN, Geneva (2021)
74. CMS Collab. *Search for a massive scalar resonance decaying to a light scalar and a Higgs boson in the four b quark final state with boosted topology*. Rep. CMS-PAS-B2G-21-003, CERN, Geneva (2021)
75. Dorsch GC, Huber SJ, Mimasu K, No JM. *Phys. Rev. D* 93(11):115033 (2016)
76. Kling F, No JM, Su S. *J. High Energy Phys.* 1609:93 (2016)
77. Baum S, Shah NR. *J. High Energy Phys.* 1812:44 (2018)
78. Ellwanger U, Rodríguez-Vázquez M. *J. High Energy Phys.* 1711:8 (2017)
79. Baum S, Shah NR, Freese K. *J. High Energy Phys.* 1904:11 (2019)
80. Robens T, Stefaniak T, Wittbrodt J. *Eur. Phys. J. C* 80(2):151 (2020)
81. CMS Collab. *Search for new particles in an extended Higgs sector in the four b quark final state at $\sqrt{s} = 13$ TeV*. Rep. CMS-PAS-B2G-20-003, CERN, Geneva (2021)
82. CMS Collab. *Search for W' decaying to a vector-like quark and a top or bottom quark in the all-jets final state*. Rep. CMS-PAS-B2G-20-002, CERN, Geneva (2021)
83. CMS Collab. *Search for resonances decaying to three W bosons in proton-proton collisions at $\sqrt{s} = 13$ TeV*. Rep. CMS-B2G-20-001, CERN, Geneva (2022)
84. Tumasyan A, et al. (CMS Collab.) *Phys. Rev. D* 129:021802 (2022)
85. Tumasyan A, et al. (CMS Collab.) *Search for high-mass resonances decaying to a jet and a Lorentz-boosted resonance in proton-proton collisions at $\sqrt{s} = 13$ TeV*. Rep. CMS-EXO-20-007/CERN-EP-2021-238, CERN, Geneva (2022)
86. Tumasyan A, et al. (CMS Collab.) *J. High Energy Phys.* 2204:47 (2022)
87. Pati JC, Salam A. *Phys. Rev. D* 10(1):275 (1974)
88. Mohapatra RN, Pati JC. *Phys. Rev. D* 11(9):2558 (1975)
89. Senjanovic G, Mohapatra RN. *Phys. Rev. D* 12(5):1502 (1975)
90. Keung WY, Senjanović G. *Phys. Rev. Lett.* 50(19):1427 (1983)
91. Mohapatra RN, Senjanović G. *Phys. Rev. Lett.* 44(14):912 (1980)
92. Das A, Dev PSB, Mohapatra RN. *Phys. Rev. D* 97(1):015018 (2018)
93. Gell-Mann M, Ramond P, Slansky R. *Conf. Proc. C* 790927:315 (1979)
94. Rehermann K, Tweedie B. *J. High Energy Phys.* 1103:59 (2011)
95. Brust C, et al. *J. High Energy Phys.* 1504:79 (2015)
96. Aad G, et al. *Phys. Rev. Lett.* 125(13):131801 (2020)
97. Collins JH, Howe K, Nachman B. *Phys. Rev. Lett.* 121(24):241803 (2018)
98. Andreassen A, Nachman B, Shih D. *Phys. Rev. D* 101(9):095004 (2020)
99. Benkendorfer K, Pottier LL, Nachman B. *Phys. Rev. D* 104(3):035003 (2021)
100. Amram O, Suarez CM. *J. High Energy Phys.* 2101:153 (2021)
101. Nachman B, Shih D. *Phys. Rev. D* 101:075042 (2020)
102. Kasieczka G, Nachman B, Schwartz MD, Shih D. *Phys. Rev. D* 103(3):035021 (2021)
103. Aaboud M, et al. (ATLAS Collab.) *Phys. Lett. B* 781:327 (2018)
104. Sirunyan AM, et al. (CMS Collab.) *Phys. Lett. B* 820:136535 (2021)



Contents

The Road to Precision Cosmology <i>Michael S. Turner</i>	1
<i>B</i> Flavor Anomalies: 2021 Theoretical Status Report <i>David London and Joaquim Matias</i>	37
Testing Lepton Flavor Universality with Pion, Kaon, Tau, and Beta Decays <i>Douglas Bryman, Vincenzo Cirigliano, Andreas Crivellin, and Gianluca Inguglia</i>	69
Something Can Come of Nothing: Surface Approaches to Quantum Fluctuations and the Casimir Force <i>Giuseppe Bimonte, Thorsten Emig, Noah Graham, and Mebran Kardar</i>	93
Exotic Higgs Decays <i>María Cepeda, Stefania Gori, Verena Ingrid Martinez Outschoorn, and Jessie Shelton</i>	119
Fundamental Neutron Physics at Spallation Sources <i>Nadia Fomin, Jason Fry, Robert W. Pattie Jr., and Geoffrey L. Greene</i>	151
Exploring Stars in Underground Laboratories: Challenges and Solutions <i>Marialuisa Aliotta, Axel Boeltzig, Rosanna Depalo, and György Gyürky</i>	177
Status of Lattice QCD Determination of Nucleon Form Factors and Their Relevance for the Few-GeV Neutrino Program <i>Aaron S. Meyer, André Walker-Loud, and Callum Wilkinson</i>	205
Precision QCD Physics at the LHC <i>Thomas Gebrmann and Bogdan Malaescu</i>	233
Probing the Neutrino-Mass Scale with the KATRIN Experiment <i>Alexey Lokhov, Susanne Mertens, Diana S. Parno, Magnus Schlösser, and Kathrin Valerius</i>	259
Electroweak Penguin Decays of <i>b</i> -Flavored Hadrons <i>Ulrik Egede, Shobei Nishida, Mitesh Patel, and Marie-Hélène Schune</i>	283
Progress in Understanding Short-Range Structure in Nuclei: An Experimental Perspective <i>John Arrington, Nadia Fomin, and Axel Schmidt</i>	307

Short-Lived Nuclides in the Early Solar System: Abundances, Origins, and Applications <i>Andrew M. Davis</i>	339
High-Energy Extragalactic Neutrino Astrophysics <i>Naoko Kurahashi, Kohta Murase, and Marcos Santander</i>	365
The Proton Structure in and out of Muonic Hydrogen <i>Aldo Antognini, Franziska Hagelstein, and Vladimir Pascalutsa</i>	389
Novel Quantum Sensors for Light Dark Matter and Neutrino Detection <i>Sunil R. Gokhale and Eneali Figueroa-Feliciano</i>	419
Searches for Heavy Resonances with Substructure <i>Petar Maksimović</i>	447

Errata

An online log of corrections to *Annual Review of Nuclear and Particle Science* articles may be found at <http://www.annualreviews.org/errata/nucl>