



Tagging the Higgs boson decay to bottom quarks with colour-sensitive observables and the Lund jet plane

Luca Cavallini^{1,a}, Andrea Coccaro^{2,b}, Charanjit K. Khosa^{2,3,4,c}, Giulia Manco^{1,5,d}, Simone Marzani^{2,4,e},
Fabrizio Parodi^{2,4,f}, Daniela Rebuzzi^{1,5,g}, Alberto Rescia^{1,h}, Giovanni Stagnitto^{6,i} 

¹ Dipartimento di Fisica, Università di Pavia, Via A. Bassi 6, 27100 Pavia, Italy

² INFN, Sezione di Genova, Via Dodecaneso 33, 16146 Genova, Italy

³ H.H. Wills Physics Laboratory, University of Bristol, Tyndall Avenue, Bristol BS8 1TL, UK

⁴ Dipartimento di Fisica, Università di Genova, Via Dodecaneso 33, 16146 Genova, Italy

⁵ INFN, Sezione di Pavia, Università di Pavia, Via A. Bassi 6, 27100 Pavia, Italy

⁶ Physik-Institut, Universität Zürich, Winterthurerstrasse 190, 8057 Zurich, Switzerland

Received: 29 December 2021 / Accepted: 17 May 2022 / Published online: 28 May 2022
© The Author(s) 2022

Abstract We study the problem of distinguishing b -jets stemming from the decay of a colour singlet, such as the Higgs boson, from those originating from the abundant QCD background. In particular, as a case study, we focus on associate production of a vector boson and a Higgs boson decaying into a pair of b -jets, which has been recently observed at the LHC. We consider the combination of several theory-driven observables proposed in the literature, together with Lund jet plane images, in order to design an original Hbb tagger. The observables are combined by means of standard machine learning algorithms, which are trained on events obtained with fast detector simulation techniques. We find that the combination of high-level single-variable observables with the Lund jet plane provides an excellent discrimination performance. We also study the dependence of the tagger on the invariant mass of the decaying particles, in order to assess the extension to a generic Xbb tagger.

1 Introduction

Since the discovery of the Higgs boson at the Large Hadron Collider (LHC) by the ATLAS [1] and CMS [2] experiments in 2012, our understanding of the properties of this particle has progressively evolved. In addition to the “golden” decay modes of the Higgs boson, $H \rightarrow \gamma\gamma$ and $H \rightarrow 4l$, in the past few years other decay channels have been observed, usually in association with particular production modes. For instance, due to its large branching ratio, the $H \rightarrow b\bar{b}$ decay plays a central role in studies that aim at probing the structure of the Higgs couplings to the fermions. In this regard, one of the most interesting processes is the associated production of a Higgs boson H and a vector boson V (W or Z), with the vector boson decaying leptonically and the Higgs boson decaying hadronically into a pair of b -quarks, $V(l\bar{l})H(b\bar{b})$: the decay products of the vector boson provide us with a clean experimental signature, as well as a recoil system for the Higgs particle. Both ATLAS and CMS experiments have reported the observation of the $H \rightarrow b\bar{b}$ decay and of the VH production mode [3,4] and the ATLAS experiment reported the first cross-section measurements [5–8] targeting different regimes of reconstructed transverse momenta of the vector boson and in fiducial volumes, as defined by the simplified template cross-section framework [9]. The experimental focus is therefore shifting towards precision measurements of the kinematics of the $H \rightarrow b\bar{b}$ decay channel and, as suggested in Ref. [10], additional differential information, and hence discrimination of this process against sources of backgrounds, is crucial for the sensitivity to beyond-the-Standard-Model operators. After the fragmentation and hadronisation process, the hard b -quarks produced by the Higgs boson

^a e-mail: luca.cavallini02@universitadipavia.it

^b e-mail: andrea.coccaro@ge.infn.it

^c e-mail: charanjit.kaur@bristol.ac.uk

^d e-mail: giulia.manco@pv.infn.it

^e e-mail: simone.marzani@ge.infn.it

^f e-mail: fabrizio.parodi@ge.infn.it

^g e-mail: daniela.rebuzzi@unipv.it

^h e-mail: alberto.rescia@desy.de

ⁱ e-mail: giovanni.stagnitto@physik.uzh.ch (corresponding author)

decay are usually detected as two separate b -jets [11,12]. In simulation, a b -jet is defined by a suitable particle-level observable, based on the angular distance of B -hadrons with respect to the jet axis, or by ghost association [13,14]. On real data, b -jets are identified by means of dedicated b -tagging algorithms.

In order to make the most out of the large set of accumulated data, strategies to better discriminate the $H(b\bar{b})$ process over the large QCD background (where the pair of b -quarks is produced by pure strong interaction, mostly by $g(b\bar{b})$ collinear splitting) are being actively developed. The signal/background discrimination is especially compelling in the boosted regime, when the transverse momentum of the jets is much greater than their invariant mass: in such a situation, the b -quarks may be close in angle, and hence reconstructed as a single jet. Since the seminal work of Ref. [15], several jet substructure techniques – which aim to improve the discrimination performance by finding hard prongs inside a large radius jet – have been designed, tested and implemented in the analyses by the experimental collaborations (see, for instance Ref. [16] and references therein). Broadly speaking, two main strategies exist: design high-level theoretically-motivated observables, sensitive to particular features of the signal distribution, which can be measured on data and used as single-variable discriminant; or produce some low-level representation of the jets (list of particles, calorimetric images,...), to be used as input for machine learning (ML) techniques. We refer the reader to the recent literature about ML based approaches for $H \rightarrow b\bar{b}$ tagging [17–27], which is continuously being updated in Ref. [28].

Our case of interest is particularly challenging because both signal and background feature a similar flavoured two-prong structure. However, the two processes have a different behaviour with respect to the QCD radiation pattern. Namely, in the signal case, the b -jets originate from the decay of a colour singlet, and thus radiation will be mostly contained within the two b -quark system. Instead, in the background case, we expect QCD radiation to be more diffuse, due to colour connections with the rest of the event. Therefore, we would like to exploit observables that are particularly sensitive to the colour flows in the event. In this paper we select several theory-driven single-valued observables (see Sect. 2) and a theoretically-motivated representation of a jet (specifically, the Lund jet plane [29]) and build a combined Hbb tagger, with the aim to exploit the best of both strategies. Such a combination is performed by means of standard ML algorithms, for different input choices. We use boosted decision trees (BDT) for single-valued observables and convolutional neural networks (CNN) for Lund jet images. BDTs have been part of HEP analyses for a long time [30]. CNNs are showing promising potential for image based data sets for various applications in HEP see e.g. [31–33]. Moreover,

in our analysis we account for the experimental detection and reconstruction of the physical quantities of interest by performing a fast detector simulation on the generated Monte Carlo events, and we assess the impact of these so-called detector effects on individual variable distributions and on the overall performance of the tagger. Finally, in an ideal scenario we would like to apply the same tagger for the decay products of a generic colour singlet X , without any prior knowledge on the value of its mass, so as to design a global Xbb tagger. In this view, it would be desirable to keep the tagger uncorrelated with the invariant mass of the b -quark pair, in order to ease its calibration on Z +jet events and to simplify the determination of the non-resonant background shape in data-driven approaches.

The paper is organised as follows. In Sect. 2 we briefly introduce the colour sensitive observables under study and the Lund jet plane. In Sect. 3 we discuss the event generation set-up and the selection cuts adopted in our analysis. In Sect. 4 we study the individual distributions of the colour sensitive observables and the output of the Lund jet plane CNN for the signal and the background processes, before and after detector simulation. Moreover, we assess the discrimination performance of the combination of several observables, by also including the Lund jet plane CNN output as an additional input. In Sect. 5 we discuss the BDT dependence on the invariant mass of the large radius jet used in the analysis (see Sect. 3), to determine the mass bias of our Hbb tagger. Finally, in Sect. 6 we draw our conclusions.

2 Description of the observables

We make use of high-level colour sensitive variables introduced in the literature in the past few years: jet pull and its projections, namely the pull angle, θ_p [34,35], and the parallel and perpendicular components of the pull vector, t_{\parallel} and t_{\perp} [36,37]; the colour ring \mathcal{O} [38]; D_2 [39,40] and the Lund jet plane [29]. We limit ourselves to a brief introduction of the relevant variables, and we refer the interested reader to the original papers.

2.1 Jet pull

Let us consider a hard jet J_a . The *pull vector* \mathbf{t} is the jet shape observable that is defined as

$$\mathbf{t} = \frac{1}{p_{T_a}} \sum_{i \in J_a} p_{T_i} |\mathbf{r}_i|^2 \hat{\mathbf{r}}_i, \quad (1)$$

where p_{T_a} is the transverse momentum of the jet, and the sum runs over all the jet constituents. y and ϕ represent rapidity and azimuthal angle, and \mathbf{r}_i is the distance vector between

the jet axis and its i -th constituent in the y - ϕ plane

$$\mathbf{r}_i = (y_i - y_a, \phi_i - \phi_a). \tag{2}$$

The pull vector is sensitive to the different colour connections of the entire event in which the jet is formed. If we consider events with two hard jets (or subjets) J_a and J_b that originate from the decay of a colour singlet, additional QCD radiation tends to be emitted between the two jets, causing the pull vector of J_a to point in the direction of J_b and vice-versa. If instead the two jets originate from the decay of a colour octet, such as the gluon, then the pull vectors will instead tend to point in different directions.

In order to make these considerations more quantitative, we can introduce suitable projections of the pull vector \mathbf{t} along two directions: the one given by the unit vector which points from the centre of J_a to the centre of J_b

$$t_{\parallel} = \mathbf{t} \cdot \hat{n}_{\parallel}, \quad \text{with } \hat{n}_{\parallel} = \frac{1}{\sqrt{\Delta y^2 + \Delta\phi^2}} (\Delta y, \Delta\phi), \tag{3}$$

and the other generated by the unit vector perpendicular to \hat{n}_{\parallel}

$$t_{\perp} = \mathbf{t} \cdot \hat{n}_{\perp}, \quad \text{with } \hat{n}_{\perp} = \frac{1}{\sqrt{\Delta y^2 + \Delta\phi^2}} (-\Delta\phi, \Delta y), \tag{4}$$

where in the above equations we have introduced $\Delta y = y_a - y_b$ and $\Delta\phi = \phi_a - \phi_b$. We also consider θ_p , known as the pull angle, defined as

$$\theta_p = \arccos \frac{t_{\parallel}}{|\mathbf{t}|}. \tag{5}$$

Of all the variables built out of the jet pull vector, the pull angle has been shown to be one of the most effective discriminants of the two different colour configurations [34]. However, the comparison between experimental measurements of the pull angle and theoretical calculations has shown that this observable is not under good theoretical control [41–43]. This problem can be traced back to the fact that θ_p is not infra-red and collinear (IRC) safe, but only Sudakov safe [35] (for discussions about Sudakov safety, see Refs. [44–46]). Instead, t_{\parallel} and t_{\perp} are IRC safe observables, so it is interesting to assess whether individually or in some combination t_{\parallel} and t_{\perp} possess the same discriminating power of its non IRC safe counterpart. It is also interesting to study how the discriminating power of the pull vector variables is affected when combining the variables for both jets J_a and J_b . Thus, we will include all the three pull vector variables, for both jets J_a and J_b , as input to the ML algorithms. Furthermore, although θ_p is not IRC safe, it is also included.

Finally, we note that there is potential overlap in including both t_{\parallel} , t_{\perp} and θ_p because the jet pull is only a two-

component vector. However, this is not an issue, since the machine learning algorithms are trained to be robust against interdependence between input variables.

2.2 Jet colour ring

The jet colour ring was introduced in Ref. [38], as an observable that is provably optimal, in certain kinematic limits. The starting point of its construction is the observation that, according to the Neyman–Pearson lemma, the ratio of the matrix elements squared for the signal and the background process should be monotonic to the optimal single-variable discriminant [47]. When considering a decay of a colour singlet as signal and a colour octet as background, with a subsequent gluon emission in the boosted regime, and working in the soft-collinear limit approximation, the ratio simplifies to

$$\frac{|\mathcal{M}_S|^2}{|\mathcal{M}_B|^2} \simeq \frac{\theta_{ak}^2 + \theta_{bk}^2}{\theta_{ab}^2}, \tag{6}$$

where the indices a and b refer to the hard partons, the index k to an additional (gluon) emission, and θ_{ij} 's are the angles between them. The above considerations lead to the definition of the jet colour ring

$$\mathcal{O} = \frac{\Delta_{ak}^2 + \Delta_{bk}^2}{\Delta_{ab}^2}, \tag{7}$$

where now Δ_{ij} are distances between jets (or subjets) in the azimuth-rapidity plane. The observable name originates from its geometric interpretation: radiation from colour singlets will tend to fall between the two jets, leading to values of $\mathcal{O} < 1$, while in the case of colour octets, we will tend to have $\mathcal{O} > 1$.

2.3 D_2

The variable D_2 [39] is defined as the ratio of two normalized N -point energy correlation functions (ECFs) [48], e_k^β :

$$D_2^{(\beta)} = \frac{e_3^{(\beta)}}{(e_2^{(\beta)})^3}, \tag{8}$$

where β is a parameter which we have set to $\beta = 2$. The variable is usually calculated on a large radius jet, and is useful to discriminate 2-prong jets from 1-prong jets. Furthermore, because of its sensitivity to soft radiation at wide angles, D_2 also probes colour correlations and it is therefore useful to disentangle different colour configurations. However, we note that D_2 retains a correlation with the mass of the large radius jet; this may be a problem when designing a tagger

free of any mass bias, see for instance [49]. We will come back to this aspect in Sect. 5.

2.4 Lund jet plane

The Lund jet plane is a theory-inspired representation of a jet [29]. It is formed parsing backwards the Cambridge-Aachen (C/A) [50, 51] clustering history of the jet. The procedure starts by undoing the final clustering step and by recording the kinematics of the splitting. The primary Lund jet plane is obtained by iterating the above procedure, always following the hardest branch in each splitting and recording the azimuth-rapidity separation of the branches involved in the splitting and the relative transverse momentum of the emission. The Lund jet plane has been exploited in the context of vector boson [29], top [52] and Higgs [27] tagging. Furthermore, it has been successfully measured at ATLAS [53] and first-principle theoretical predictions have been performed in Ref. [54].

3 Event simulation and selection

We generate about 300k events for the $pp \rightarrow H(b\bar{b})Z(\nu_\ell\bar{\nu}_\ell)$ signal and 4M events for the $pp \rightarrow b\bar{b}\nu_\ell\bar{\nu}_\ell$ background processes, so as to have about 50k events remaining after all analysis cuts, in accordance with Table 1, as will be detailed below. In case of signal, $b\bar{b}$ pair is produced from the decay of the Higgs boson, while it comes mainly from QCD interaction in case of background. Fig. 1a, b show representative Feynman diagrams for the signal and the background processes, respectively.

We generate hard events using MG5_AMC@NLO v2.8.3.2 [55], by imposing a 200 GeV cut on the p_T for the neutrino pair in the final state. This is done to ensure that the events generated are firmly in the boosted regime. These parton-level events are subsequently showered in PYTHIA v8.305 [56], including MPI and underlying events, to produce particle-level events.

Finally, rather than simulating an entire detector, DELPHES v3.5.0 is used to perform a fast detector simulation [57, 58]. This allows us to understand how the discrimination power could be affected by real-life detector effects, without having to run a computationally expensive full simulation, which in addition is strongly detector dependent. From DELPHES, we

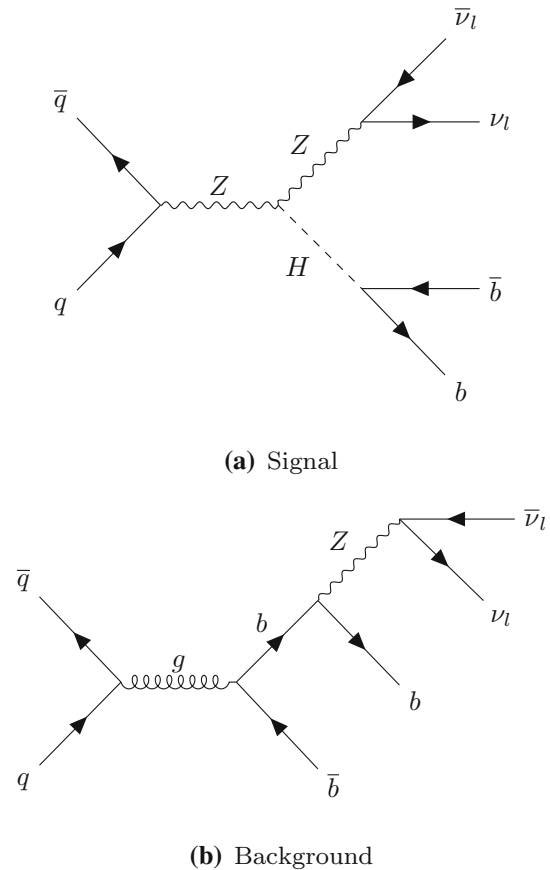


Fig. 1 Representative Feynman diagrams

extract both the Monte Carlo truth of the event, containing the particle-level information, e.g. the same one would get from a perfectly efficient detector (henceforth referred to as *truth*), as well as the reconstructed events including the detector effects (henceforth referred to as *reco*). The DELPHES simulation is run using the ATLAS card with minor modifications, described below, to fit our needs. For the truth case, we consider all visible, stable particles with $p_T > 0.5$ GeV. Instead, in the reco case, the jets are built using the simulated calorimeter towers and tracks. All electromagnetic calorimeter towers with energy $E > 0.5$ GeV and significance $S > 2.0$ and all hadronic calorimeter towers with energy $E > 1.0$ GeV and significance $S > 2.0$ are considered. Tracks are required to have $p_T > 0.5$ GeV. DELPHES uses FASTJET v3.3.4 [59] to perform the jet clustering.

At this point the analysis is the same in both the truth case and the reco case. First, the constituents are clustered into jets with radius $R = 1.0$ using the anti- k_T algorithm [60]. For each event, we choose the jet with the highest p_T as the *large radius jet*. We only accept the event if the large jet has $p_T > 250$ GeV and $|y| < 1.5$, because of the tracking detector acceptance.

Table 1 The efficiencies of the analysis after cuts are applied

	Truth (%)	Reco (%)
Signal	20	17
Background	1.6	1.3

We also cluster the constituents into smaller jets with radius $R = 0.2$. We identify those jets having $\Delta R < 0.8$ from the large jet, and call these subjets. We then proceed to identify the b -subjets which originate from the b -partons, through a process known as b -labelling. We do this by first identifying the b -partons originating from the hard scattering in the event record, requiring a minimum p_T of 5.0 GeV. For each b -parton, we compute the distance between each b -parton and subjet; the subjet which is closest to the b -parton, provided that the distance is below 0.2, is labelled as b -subjet. The association between a b -parton and b -subjet is unique.

For the event to be accepted, we require two b -labelled subjets with $p_T > 10$ GeV. The pull variables are calculated on these two b -labelled subjets, and D_2 is calculated on the large jet. For the colour ring to be defined, there must also be a third non- b -subjet within $\Delta R = 0.8$ from the large jet. In a majority of cases, this third jet is not present. To avoid discarding too many events, in these cases we assign a default value of $\mathcal{O} = -1$ to the colour ring. This allows for higher statistics, but also provides useful information to the machine learning algorithms.

For the Lund jet plane, we consider large radius jet constituents and re-cluster them using C/A algorithm. Considering the declustering history of this jet, we get the primary Lund jet plane. Considering 25×25 pixels for each image, we put 1 or 0 in a pixel depending on if $(\ln 1/\Delta, \ln k_T)$ value of the splitting falls in that pixel or not. Our implementation of the Lund jet plane is based on the one present in FASTJET-CONTRIB [61] repository.

Table 1 shows the percentage of events which pass the selections in all cases considered. The selection discards more background than signal events in both the truth and reco cases. The most important cut is the p_T cut on the large radius jet, accounting for 60% of discarded events. The second most important cut is the rapidity cut on the large radius jet which rejects 10% of the events.

4 Discrimination performance

After event selection, we are ready to evaluate observables on the selected events. We first show in Fig. 2 the normalised distributions for the eight colour sensitive (CS) observables introduced in Sect. 2, both for signal and background, and at the truth and reco level. By just looking at the plots, we can appreciate the strong discrimination power of the \mathcal{O} and D_2 , which is retained at the reco level. Instead, the observables related to the jet pull vector are more affected by detector effects: there is a visible difference in $t_{\parallel i}$, $t_{\perp i}$ and θ_{pi} , both for the leading jet a and the sub-leading jet b , between the truth and the reco cases. In particular, the pull angle observables θ_{pa} and θ_{pb} seem to be good discriminants at the truth level,

but detector effects noticeably flatten the signal distribution, hence leading to a worsening of the discrimination power.

We then show in Fig. 3 the averaged Lund images for the signal and background process, in the truth and reco cases respectively. We note that detector effects lead to an overall decrease of the image quality, in the sense that the distinctive features of the truth case are still present, but in the reco case there is additional radiation for middle values of Δ and k_T for both the signal and the background events. However, the high density patch at large Δ and high k_T in the case of the signal is still clearly visible by eye also in the reco case.

After having determined the distributions of the CS observables and the Lund jet images, we now use these as inputs to ML algorithms in order to build combined classifiers. Specifically, we train a BDT¹ on the CS observables, whereas Lund images are classified using a CNN. More details about these methods and architectures are provided in Appendix A. The output distribution of the CNN Lund jet plane classifier (LP_{CNN}) is shown in Fig. 2. In the following, we consider also the combination of (some of) the CS and LP_{CNN} observables, in order to improve the total discrimination power; in such cases, we adopt a two-step procedure, by using the output of the CNN Lund jet plane classifier as an additional input to the BDT.

In Fig. 4, we show the receiver operating characteristic (ROC) curves for several combinations of observables. Namely, we consider all the colour sensitive observables (CS) or just the D_2 and the colour ring (D_2 +CR), combined through a BDT; the LP_{CNN} ; the combination of all the CS observables with the LP_{CNN} (CS+ LP_{CNN}), by means of the two-step procedure explained above. For each curve in Fig. 4, we report the value of the area under the ROC curve (AUC) in Table 2, both for the truth and the reco cases. A perfect classifier would have $AUC = 1$, whereas a random classifier is associated with $AUC = 1/2$.

As expected, we first observe a worsening of the discrimination power by moving from the truth case to the reco case, as can also be seen by comparing the value in the left and in the right column of Table 2. However, the performance after taking into account detector effects is still good for most of the combinations, close to 0.85 for the CS+ LP_{CNN} combination. Furthermore, we see that most of the discriminating power of the set of CS observables is actually coming from the combination of D_2 +CR alone. This is in agreement with what observed at the level of distributions at the beginning of Sect. 4: the jet pull observables, including the pull angle, seem not to add any additional information useful for classification. At reco level, their discrimination power is almost unnoticeable. Moving to the combinations involving

¹ We have also tried a neural network, which returns similar results in all the cases analysed. Hence, we only report results obtained with the BDT.

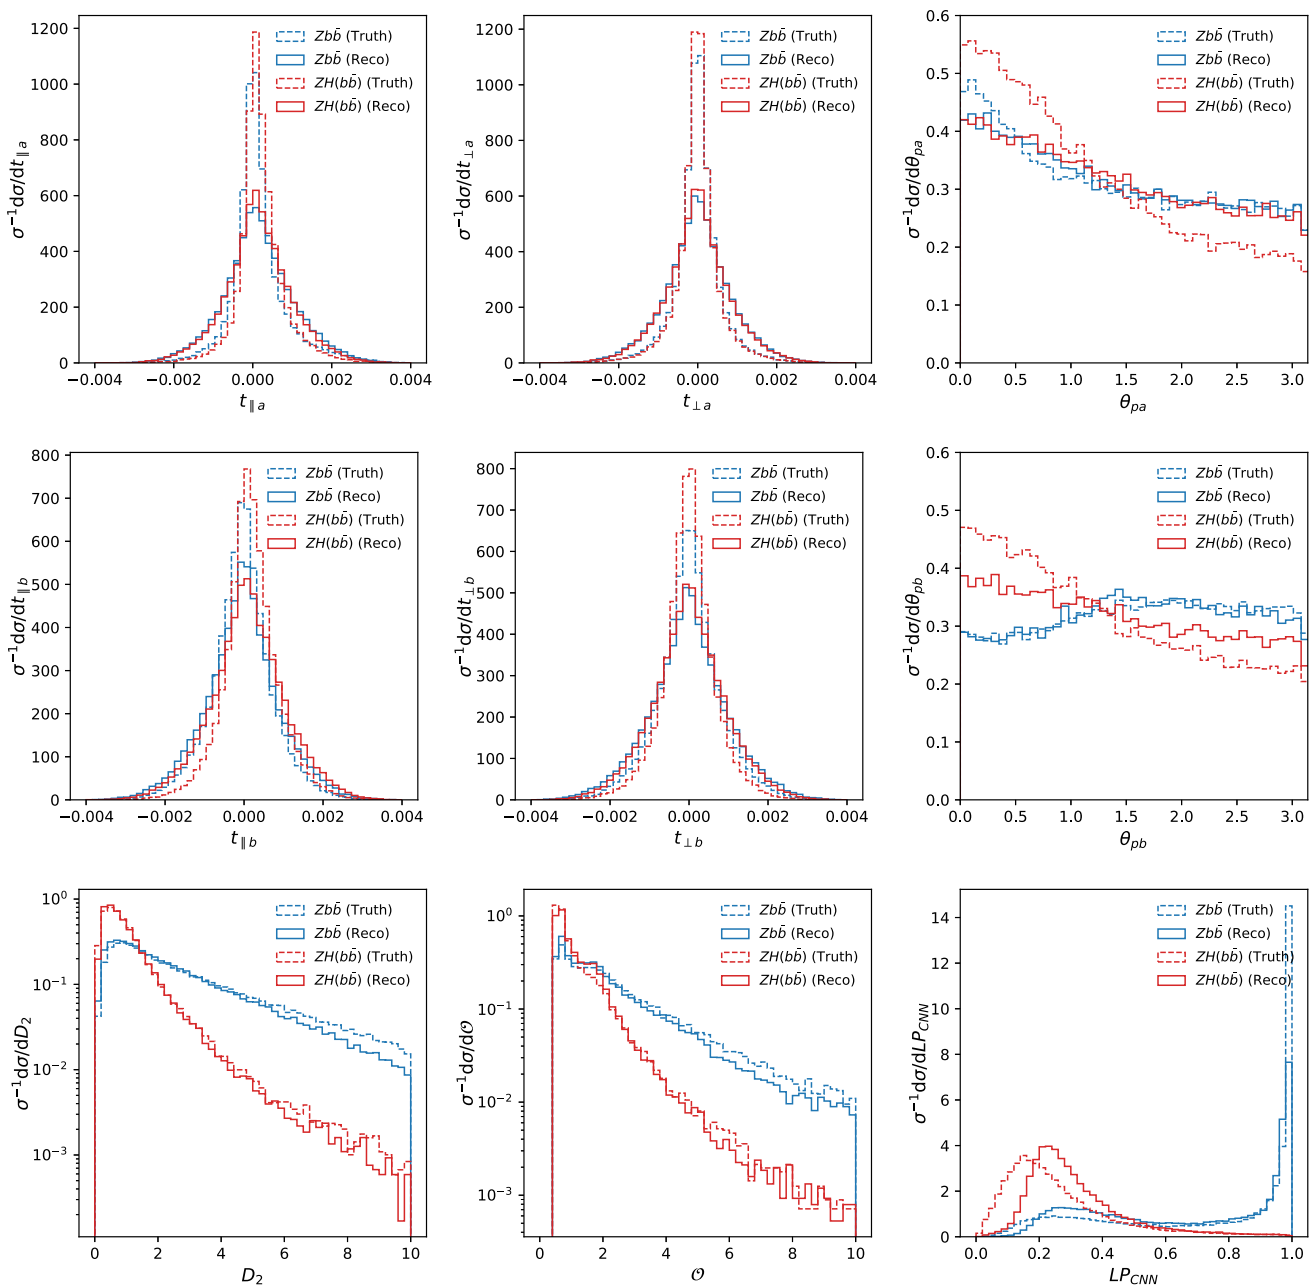


Fig. 2 The distributions of the colour sensitive variables and Lund jet plane after the selection cuts. Signal ($ZH(bb)$) and background (Zbb) distributions are shown in red and blue, respectively, for both truth and reco cases

the Lund jet plane, we observe that the Lund jet plane alone is performing better than the whole set of CS observables, especially in the region of high signal efficiencies. When we combine LP_{CNN} with the CS observables, we see a noticeable improvement of the overall classification power, with a value of AUC equal to 0.893 in the truth case and 0.846 in the reco case.

Finally, in Table 3 we rank the variables based on their importance in the BDT, both in the truth and reco case. The ranking presented here also includes the output of the Lund

jet plane CNN as an additional input. LP_{CNN} is the most discriminating variable, both in the truth and in the reco case. It is followed in order by D_2 and the colour ring \mathcal{O} . The jet pull variables are all of similar importance, at the bottom of the ranking score. For the reconstructed case, \mathcal{O} gains additional importance with respect to the pull variables.

We end this Section with a comment about the usage of the pull angle. As already mentioned in Sect. 2, the pull angle θ_p is only Sudakov safe. One may wonder what is the effect of keeping only its (IRC-safe) projections t_{\parallel} and t_{\perp} , instead




Fig. 3 Averaged primary Lund jet plane images for $ZH(bb)$ and Zbb in the truth and reco case

of using all the three observables t_{\parallel} , t_{\perp} and θ_p as input in the BDT, as we have done. Unsurprisingly, by looking at the correlation matrix in the BDT, we observe a strong correlation between these variables. Given that the variables derived from the jet pull vector(s) have a small influence on the overall performance, a variant of the tagger could be conceived, with comparable performance, by dropping θ_p among the list of inputs to the BDT.

5 Invariant mass dependence

Since our goal is to develop a tagger purely sensitive to the colour configuration of the decaying particle, in order to be applied to other contexts (such as Z or W boson hadronic

decays), ideally our procedure should be insensitive to the invariant mass of decaying system, specifically to the invariant mass of the pair of b -jets.

In Fig. 5 we show the distribution of the invariant mass measured on the whole set of background events and on three subset of events, each of the same size, corresponding to signal-enriched, intermediate and background-enriched regions. These regions are defined by means of a set of cuts on different discriminant variables: D_2 alone (Fig. 5a), combined BDT with CS but D_2 (Fig. 5b), Lund plane CNN (Fig. 5c). We show results for the reco case only, since the ones in the truth case are similar. In an ideal scenario, a cut on the discriminating variable should not also concomitantly imply a cut on the invariant mass of the system, hence the

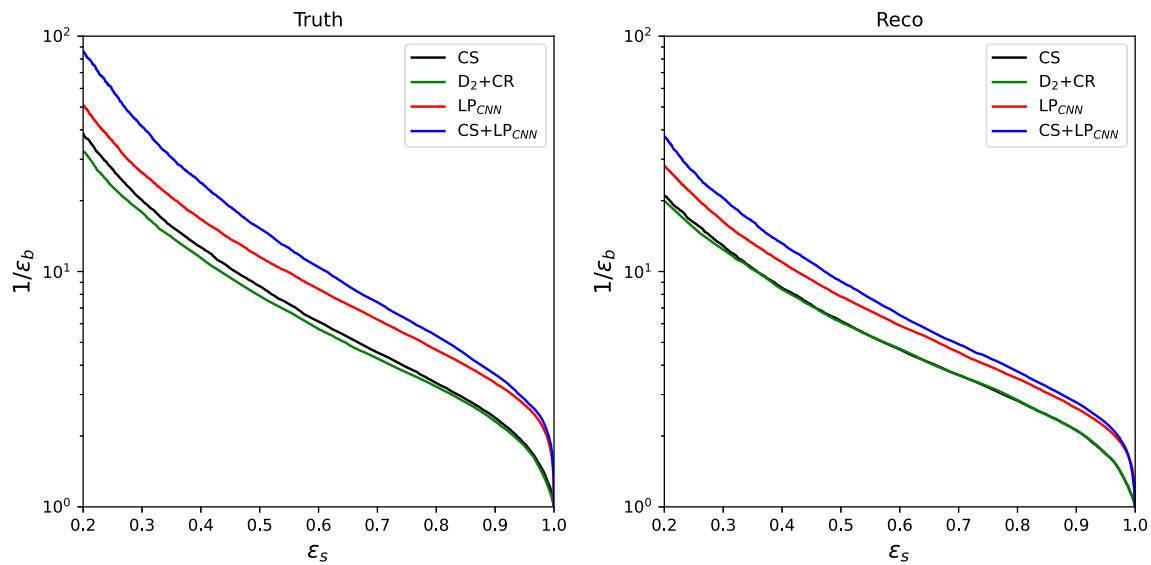


Fig. 4 The ROC curves showing background rejection as a function of signal efficiency for the truth (left) and reco case (right) for CS variables, LP_{CNN} and the combined cases

Table 2 Area under the ROC curves for different combination of observables

	AUC-test sample	
	Truth	Reco
CS observables	0.826	0.788
D_2 +CR	0.817	0.787
LP _{CNN}	0.876	0.828
CS + LP _{CNN}	0.893	0.846

Table 3 BDT observable ranking for the truth and reco cases

Observable ranking				
Rank	Truth		Reco	
	Obs.	Importance	Obs.	Importance
1	LP _{CNN}	6.6×10^{-1}	LP _{CNN}	4.8×10^{-1}
2	D_2	1.4×10^{-1}	\mathcal{O}	1.0×10^{-1}
3	\mathcal{O}	5.7×10^{-2}	D_2	9.3×10^{-2}
4	θ_{pb}	3.0×10^{-2}	θ_{pb}	7.0×10^{-2}
5	θ_{pa}	2.9×10^{-2}	θ_{pa}	6.5×10^{-2}
6	$t_{\parallel b}$	2.6×10^{-2}	$t_{\perp b}$	6.0×10^{-2}
7	$t_{\parallel a}$	2.4×10^{-2}	$t_{\parallel a}$	4.5×10^{-2}
8	$t_{\perp b}$	1.9×10^{-2}	$t_{\perp a}$	4.3×10^{-2}
9	$t_{\perp a}$	1.0×10^{-3}	$t_{\parallel a}$	3.3×10^{-2}

curves for the three regions should overlap, and agree with the curve without any cut.

We find that the sensitivity to the invariant mass is introduced mainly through the D_2 observable, which is highly correlated to the value on the mass, as it is clear from Fig. 5a.

Such a correlation has been already investigated in the literature [49,62]. By removing D_2 from the CS input variables of the BDT, the mass bias is greatly reduced, as can be observed by comparing Fig. 5a and Fig. 5b. However, given the fact that the D_2 observable is ranked as one of the most important (see Table 3), the removal of this variable comes at the price of losing a good part of efficiency.

Finally, it is interesting to study whether the LP_{CNN} alone retains or not a dependence on the invariant mass of the decaying system. This is shown in Fig. 5c. Unfortunately, the output of the LP_{CNN} appears to be notably correlated to the invariant mass of the pair of b -jets. In order to better understand this behaviour, by looking at the correlation matrix of the BDT with the CS variables and the LP_{CNN} as input, we note that LP_{CNN} and D_2 are largely correlated. Hence, the behaviour we observe in Fig. 5c for LP_{CNN} can be related to the known behaviour of D_2 of Fig. 5a. Given the fact that LP_{CNN} is our best discriminating variable, considering removing it from the combination comes at the price of losing a good part of discrimination power.

6 Conclusions

In this paper, we have investigated the problem of distinguishing the b -jets originating from the decay of a colour singlet from those originating from the pure QCD background (mostly through $g \rightarrow b\bar{b}$ collinear splitting). We have focused on the signal process $pp \rightarrow H(b\bar{b})Z(\nu_\ell\bar{\nu}_\ell)$, but we are confident that our strategy is valid in a more general context. Specifically, we have trained a BDT architecture on eight high-level, colour sensitive observables, in order to

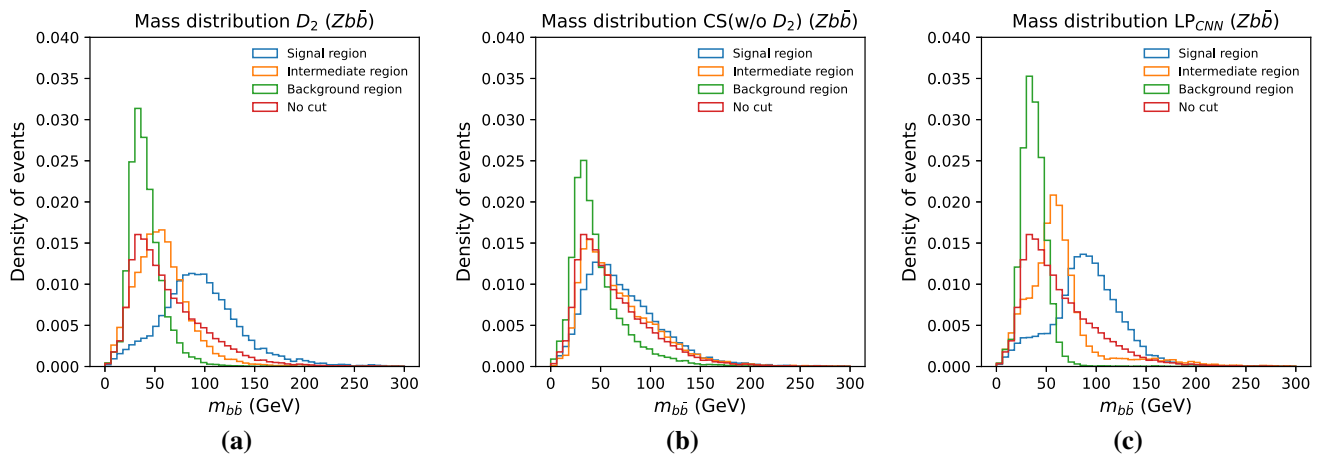


Fig. 5 The distribution of the invariant mass of the b -jets pair in the reco background sample for **a** different cuts on D_2 , **b** different ranges of the BDT output for the CS without the inclusion of D_2 , **c** and for different cuts on Lund plane CNN output

develop a combined colour tagger. We have also explored the discrimination performance of a CNN architecture trained on Lund jet plane images. Finally, we have combined the high-level observables and the output of the Lund jet plane CNN in a common BDT architecture. We have also performed a fast detector simulation in order to better assess the experimental feasibility of this tagging strategy. Namely, we have compared individual distributions and the final performance of the tagger before and after the inclusion of detector effects.

We have found a good discrimination power for our combination of colour sensitive observables with the output of the Lund jet plane CNN (AUC = 0.893), slightly deteriorated when including detector effects (AUC = 0.846). The Lund jet plane alone has been proven to be a powerful Hbb tagger even in presence of detector effect, thus extending the results of Ref. [27], and when combined with theoretically motivated single-variable observables, such as D_2 or the colour ring, the overall performance appreciably improves. In the end, we have shown that our tagger, which is a combination of several theory-driven single-variable observables with a representation of radiation pattern within a jet, is not only effective in theory, but also shows promising prospects for application to experimental analyses.

We have also studied to what extent the tagger is sensitive to the mass of the decaying particle. In the case of the colour sensitive observables, the mass bias has been shown to be ultimately due to the D_2 input variable, as already known from the literature. However, we have also found that the Lund plane CNN retains a large mass bias, and to the best of our knowledge this has not been pointed out in the literature so far. The elimination of these variables come at the cost of classification efficiency, especially in the case of the Lund plane CNN. Further studies are needed in order to understand how to remove such a mass bias. For instance, one could

plan to explore techniques similar to the ones presented in Refs. [49,62], based on a debiasing *a posteriori*.

Even if a fast detector simulation offers a good starting point in order to assess the feasibility of a tagging strategy, in the end a full detector simulation within a more defined experimental context would be required. Such a more realistic scenario would also entail the inclusion of the efficiency of real-life existing b -tagging algorithms used by the experimental collaborations, whereas in this paper we have assumed a b -tagging algorithm with 100% efficiency. Further, although the main scope of this study is to assess the performance of relatively simple observables and ML architectures for Hbb tagging, it is interesting to investigate the importance of physics information beyond the primary Lund plane, which can be exploited by including planes originated by secondary splittings. In this context, we have made a preliminary investigation using the LundNet-5 model of Ref. [52], which employs graph neural networks, and obtaining an improved performance corresponding to AUC = 0.925 and AUC = 0.894, for truth and reco, respectively.

We leave the implementation of these suggestions, as well as the opportunity to consider new high-level observables as input for our tagger, for future studies.

Acknowledgements The work of CK and SM is supported by Università di Genova under the curiosity-driven grant “Using jets to challenge the Standard Model of particle physics” and by the Italian Ministry of Research (MUR) under Grant PRIN 20172LNEEZ and under the project “Dipartimenti di Eccellenza”. AR acknowledges support from DESY (Hamburg, Germany), a member of the Helmholtz Association HGF and from the University of Zürich.

Data Availability Statement This manuscript has no associated data or the data will not be deposited. [Authors’ comment: For this publication only computer-generated pseudo data were used. These have been obtained with either publicly available Monte Carlo event generators or with private code, which is made available by the Authors upon request.]

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>. Funded by SCOAP³.

Appendix A: Details about ML algorithms

In this Appendix, we report some details about BDT parameters and CNN architecture adopted in the analysis.

A.1 BDT parameters

The boosted decision tree (BDT) has been implemented using the ROOT Toolkit for Multivariate Analysis (TMVA) library [63]. The BDT is made up of 50 trees, with a maximum depth of 5 and minimum node size set to 2.5% of the total number of events (Table 4). The Gini index is used as the optimisation criterion. AdaBoost has been chosen as the boosting model and the number of cuts is set to 80. We use a 50/50 train/test sample, and there is no downsampling.

A.2 CNN architecture

We used CNN (implemented using Keras [64]) for the Lund jet images data set. Balanced data set for the binary classification is used; 70% for the training, 15% for the validation and 15% for the testing. We tried several models with different combinations of hyper-parameters and the best architecture is described in the Table 5. Here N_i , $i = 1 \dots 4$ denotes the number of filters in the corresponding convolutional layer. Filter size is 3×3 in all the convolutional layers. Activation function 'relu' and 'softmax' are used for the intermediate

Table 4 BDT parameters used for the truth and reco data set

Parameters	Value
No. of trees	100
Max depth	3
MinNodeSize	2.5%
Boost type	AdaBoost
Train/test	50/50
No. of cuts	200
Downsampling	No

Table 5 CNN architectures used for truth and reco cases

Parameter	Value
N_1 Conv2D	30
N_2 Conv2D	30
Dropouts	-(0.3)
N_3 Conv2D	30
Dropouts	-(0.3)
N_4 Conv2D	10
Dropouts	-(0.1)
Flat layer	150
Epochs	30
Batch size	800

and last layer, respectively. For the optimisation, 'adam' optimiser is used. We used pooling (MaxPooling) operation after the second and fourth convolutional layer for image down-sizing. Note that CNN architecture used is same for truth and reco data except the dropouts. The numbers mentioned in the brackets are the dropouts strength used for the reco data set.

References

1. ATLAS Collaboration, Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC. Phys. Lett. B **716**, 1 (2012). <https://doi.org/10.1016/j.physletb.2012.08.020>. arXiv:1207.7214
2. CMS Collaboration, Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC. Phys. Lett. B **716**, 30 (2012). <https://doi.org/10.1016/j.physletb.2012.08.021>. arXiv:1207.7235
3. ATLAS Collaboration, Observation of $H \rightarrow b\bar{b}$ decays and VH production with the ATLAS detector. Phys. Lett. B **786**, 59 (2018). <https://doi.org/10.1016/j.physletb.2018.09.013>. arXiv:1808.08238
4. CMS Collaboration, Observation of Higgs boson decay to bottom quarks. Phys. Rev. Lett. **121**, 121801 (2018). <https://doi.org/10.1103/PhysRevLett.121.121801>. arXiv:1808.08242
5. ATLAS Collaboration, Measurement of VH , $H \rightarrow b\bar{b}$ production as a function of the vector-boson transverse momentum in 13 TeV pp collisions with the ATLAS detector. JHEP **05**, 141 (2019). [https://doi.org/10.1007/JHEP05\(2019\)141](https://doi.org/10.1007/JHEP05(2019)141). arXiv:1903.04618
6. ATLAS Collaboration, Measurement of VH , $H \rightarrow b\bar{b}$ production as a function of the vector-boson transverse momentum in 13 TeV pp collisions with the ATLAS detector. JHEP **05**, 141 (2019). [https://doi.org/10.1007/JHEP05\(2019\)141](https://doi.org/10.1007/JHEP05(2019)141). arXiv:1903.04618
7. ATLAS Collaboration, Measurements of WH and ZH production in the $H \rightarrow b\bar{b}$ decay channel in pp collisions at 13 TeV with the ATLAS detector. Eur. Phys. J. C **81**, 178 (2021). <https://doi.org/10.1140/epjc/s10052-020-08677-2>. arXiv:2007.02873
8. ATLAS Collaboration, Measurement of the associated production of a Higgs boson decaying into b -quarks with a vector boson at high transverse momentum in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector. Phys. Lett. B **816**, 136204 (2021). <https://doi.org/10.1016/j.physletb.2021.136204>. arXiv:2008.02508
9. N. Berger et al., Simplified template cross sections—stage 1.1. arXiv:1906.02754

10. J. Brehmer, S. Dawson, S. Homiller, F. Kling, T. Plehn, Benchmarking simplified template cross sections in WH production. *JHEP* **11**, 034 (2019). [https://doi.org/10.1007/JHEP11\(2019\)034](https://doi.org/10.1007/JHEP11(2019)034) arXiv:1908.06980
11. ATLAS Collaboration, ATLAS b-jet identification performance and efficiency measurement with $t\bar{t}$ events in pp collisions at $\sqrt{s} = 13$ TeV. *Eur. Phys. J. C* **79**, 970 (2019). <https://doi.org/10.1140/epjc/s10052-019-7450-8> arXiv:1907.05120
12. CMS Collaboration, Identification of heavy-flavour jets with the CMS detector in pp collisions at 13 TeV. *JINST* **13**, P05011 (2018). <https://doi.org/10.1088/1748-0221/13/05/P05011> arXiv:1712.07158
13. M. Cacciari, G.P. Salam, G. Soyez, The catchment area of jets. *JHEP* **04**, 005 (2008). <https://doi.org/10.1088/1126-6708/2008/04/005> arXiv:0802.1188
14. ATLAS Collaboration, Performance of jet substructure techniques for large- R jets in proton-proton collisions at $\sqrt{s} = 7$ TeV using the ATLAS detector. *JHEP* **09**, 076 (2013). [https://doi.org/10.1007/JHEP09\(2013\)076](https://doi.org/10.1007/JHEP09(2013)076) arXiv:1306.4945
15. J.M. Butterworth, A.R. Davison, M. Rubin, G.P. Salam, Jet substructure as a new Higgs search channel at the LHC. *Phys. Rev. Lett.* **100**, 242001 (2008). <https://doi.org/10.1103/PhysRevLett.100.242001> arXiv:0802.2470
16. S. Marzani, G. Soyez, M. Spannowsky, Looking inside jets: an introduction to jet substructure and boosted-object phenomenology, vol. 958, Springer (2019). <https://doi.org/10.1007/978-3-030-15709-8> arXiv:1901.10342
17. K. Datta, A. Larkoski, B. Nachman, Automating the construction of jet observables with machine learning. *Phys. Rev. D* **100**, 095016 (2019). <https://doi.org/10.1103/PhysRevD.100.095016> arXiv:1902.07180
18. J. Lin, M. Freytsis, I. Moul, B. Nachman, Boosting $H \rightarrow b\bar{b}$ with Machine Learning. *JHEP* **10**, 101 (2018). [https://doi.org/10.1007/JHEP10\(2018\)101](https://doi.org/10.1007/JHEP10(2018)101) arXiv:1807.10768
19. E.A. Moreno, T.Q. Nguyen, J.-R. Vlimant, O. Cerri, H.B. Newman, A. Perival et al., Interaction networks for the identification of boosted $H \rightarrow b\bar{b}$ decays. *Phys. Rev. D* **102**, 012010 (2020). <https://doi.org/10.1103/PhysRevD.102.012010> arXiv:1909.12285
20. A. Chakraborty, S.H. Lim, M.M. Nojiri, Interpretable deep learning for two-prong jet classification with jet spectra. *JHEP* **19**, 135 (2020). [https://doi.org/10.1007/JHEP07\(2019\)135](https://doi.org/10.1007/JHEP07(2019)135) arXiv:1904.02092
21. CMS Collaboration, Identification of heavy, energetic, hadronically decaying particles using machine-learning techniques. *JINST* **15**, P06005 (2020). <https://doi.org/10.1088/1748-0221/15/06/P06005> arXiv:2004.08262
22. Y.-L. Chung, S.-C. Hsu, B. Nachman, Disentangling boosted Higgs boson production modes with machine learning. arXiv:2009.05930
23. B. Tannenwald, C. Neu, A. Li, G. Buehlmann, A. Cuddeback, L. Hatfield et al., Benchmarking machine learning techniques with di-higgs production at the LHC. arXiv:2009.06754
24. J. Guo, J. Li, T. Li, The boosted higgs jet reconstruction via graph neural network. arXiv:2010.05464
25. M. Abbas, A. Khan, A.S. Qureshi, M.W. Khan, Extracting signals of Higgs boson from background noise using deep neural networks. arXiv:2010.08201
26. C. Jang, S.-K. Ko, Y.-K. Noh, J. Choi, J. Lim, T.J. Kim, Learning to increase matching efficiency in identifying additional b-jets in the $t\bar{t}b\bar{b}$ process. arXiv:2103.09129
27. C.K. Khosa, S. Marzani, Higgs boson tagging with the Lund jet plane. *Phys. Rev. D* **104**, 055043 (2021). <https://doi.org/10.1103/PhysRevD.104.055043> arXiv:2105.03989
28. HEP ML Community, A living review of machine learning for particle physics. <https://iml-wg.github.io/HEPML-LivingReview/>
29. F.A. Dreyer, G.P. Salam, G. Soyez, The Lund jet plane. *JHEP* **12**, 064 (2018). [https://doi.org/10.1007/JHEP12\(2018\)064](https://doi.org/10.1007/JHEP12(2018)064) arXiv:1807.04758
30. H.-J. Yang, B.P. Roe, J. Zhu, Studies of boosted decision trees for MiniBooNE particle identification. *Nucl. Instrum. Methods A* **555**, 370 (2005). <https://doi.org/10.1016/j.nima.2005.09.022> arXiv:physics/0508045
31. C.K. Khosa, V. Sanz, Anomaly awareness. arXiv:2007.14462
32. C.K. Khosa, L. Mars, J. Richards, V. Sanz, Convolutional neural networks for direct detection of dark matter. *J. Phys. G* **47**, 095201 (2020). <https://doi.org/10.1088/1361-6471/ab8e94> arXiv:1911.09210
33. C.K. Khosa, V. Sanz, M. Soughton, Using machine learning to disentangle LHC signatures of Dark Matter candidates. *SciPost Phys.* **10**, 151 (2021). <https://doi.org/10.21468/SciPostPhys.10.6.151> arXiv:1910.06058
34. J. Gallicchio, M.D. Schwartz, Seeing in color: jet superstructure. *Phys. Rev. Lett.* **105**, 022001 (2010). <https://doi.org/10.1103/PhysRevLett.105.022001> arXiv:1001.5027
35. A.J. Larkoski, S. Marzani, C. Wu, Theory predictions for the pull angle. *Phys. Rev. D* **99**, 091502 (2019). <https://doi.org/10.1103/PhysRevD.99.091502> arXiv:1903.02275
36. Y. Bao, A.J. Larkoski, Calculating pull for non-singlet jets. *JHEP* **12**, 035 (2019). [https://doi.org/10.1007/JHEP12\(2019\)035](https://doi.org/10.1007/JHEP12(2019)035) arXiv:1910.02085
37. A. Larkoski, S. Marzani, C. Wu, Safe use of jet pull. *JHEP* **01**, 104 (2020). [https://doi.org/10.1007/JHEP01\(2020\)104](https://doi.org/10.1007/JHEP01(2020)104) arXiv:1911.05090
38. A. Buckley, G. Callea, A.J. Larkoski, S. Marzani, An optimal observable for color singlet identification. *SciPost Phys.* **9**, 026 (2020). <https://doi.org/10.21468/SciPostPhys.9.2.026> arXiv:2006.10480
39. A.J. Larkoski, I. Moul, D. Neill, Power counting to better jet observables. *JHEP* **12**, 009 (2014). [https://doi.org/10.1007/JHEP12\(2014\)009](https://doi.org/10.1007/JHEP12(2014)009) arXiv:1409.6298
40. I. Moul, L. Necib, J. Thaler, New angles on energy correlation functions. *JHEP* **12**, 153 (2016). [https://doi.org/10.1007/JHEP12\(2016\)153](https://doi.org/10.1007/JHEP12(2016)153) arXiv:1609.07483
41. D0 Collaboration, Measurement of color flow in $t\bar{t}$ events from $p\bar{p}$ collisions at $\sqrt{s} = 1.96$ TeV. *Phys. Rev. D* **83**, 092002 (2011). <https://doi.org/10.1103/PhysRevD.83.092002> arXiv:1101.0648
42. ATLAS Collaboration, Measurement of colour flow with the jet pull angle in $t\bar{t}$ events using the ATLAS detector at $\sqrt{s} = 8$ TeV. *Phys. Lett. B* **750**, 475 (2015). <https://doi.org/10.1016/j.physletb.2015.09.051> arXiv:1506.05629
43. ATLAS Collaboration, Measurement of colour flow using jet-pull observables in $t\bar{t}$ events with the ATLAS experiment at $\sqrt{s} = 13$ TeV. *Eur. Phys. J. C* **78**, 847 (2018). <https://doi.org/10.1140/epjc/s10052-018-6290-2> arXiv:1805.02935
44. A.J. Larkoski, J. Thaler, Unsafe but calculable: ratios of angularities in perturbative QCD. *JHEP* **09**, 137 (2013). [https://doi.org/10.1007/JHEP09\(2013\)137](https://doi.org/10.1007/JHEP09(2013)137) arXiv:1307.1699
45. A.J. Larkoski, S. Marzani, G. Soyez, J. Thaler, Soft drop. *JHEP* **05**, 146 (2014). [https://doi.org/10.1007/JHEP05\(2014\)146](https://doi.org/10.1007/JHEP05(2014)146) arXiv:1402.2657
46. A.J. Larkoski, S. Marzani, J. Thaler, Sudakov safety in perturbative QCD. *Phys. Rev. D* **91**, 111501 (2015). <https://doi.org/10.1103/PhysRevD.91.111501> arXiv:1502.01719
47. J. Neyman, E.S. Pearson, On the problem of the most efficient tests of statistical hypotheses. *Phil. Trans. R. Soc. Lond. A* **231**, 289 (1933). <https://doi.org/10.1098/rsta.1933.0009>
48. A.J. Larkoski, G.P. Salam, J. Thaler, Energy correlation functions for jet substructure. *JHEP* **06**, 108 (2013). [https://doi.org/10.1007/JHEP06\(2013\)108](https://doi.org/10.1007/JHEP06(2013)108) arXiv:1305.0007
49. J. Dolen, P. Harris, S. Marzani, S. Rappoccio, N. Tran, Thinking outside the ROCs: designing decorrelated taggers (DDT) for

- jet substructure. *JHEP* **05**, 156 (2016). [https://doi.org/10.1007/JHEP05\(2016\)156](https://doi.org/10.1007/JHEP05(2016)156) arXiv:1603.00027
50. Y.L. Dokshitzer, G.D. Leder, S. Moretti, B.R. Webber, Better jet clustering algorithms. *JHEP* **08**, 001 (1997). <https://doi.org/10.1088/1126-6708/1997/08/001> arXiv:hep-ph/9707323
51. M. Wobisch, T. Wengler, Hadronization corrections to jet cross-sections in deep inelastic scattering, in Workshop on Monte Carlo Generators for HERA Physics (Plenary Starting Meeting), pp. 270–279, 4, (1998). arXiv:hep-ph/9907280
52. F.A. Dreyer, H. Qu, Jet tagging in the Lund plane with graph networks. *JHEP* **03**, 052 (2021). [https://doi.org/10.1007/JHEP03\(2021\)052](https://doi.org/10.1007/JHEP03(2021)052) arXiv:2012.08526
53. ATLAS Collaboration, Measurement of the Lund jet plane using charged particles in 13 TeV proton-proton collisions with the ATLAS detector. *Phys. Rev. Lett.* **124**, 222002 (2020). <https://doi.org/10.1103/PhysRevLett.124.222002>. arXiv:2004.03540
54. A. Lifson, G.P. Salam, G. Soyez, Calculating the primary Lund jet plane density. *JHEP* **10**, 170 (2020). [https://doi.org/10.1007/JHEP10\(2020\)170](https://doi.org/10.1007/JHEP10(2020)170) arXiv:2007.06578
55. J. Alwall, R. Frederix, S. Frixione, V. Hirschi, F. Maltoni, O. Matelaer et al., The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations. *JHEP* **07**, 079 (2014). [https://doi.org/10.1007/JHEP07\(2014\)079](https://doi.org/10.1007/JHEP07(2014)079) arXiv:1405.0301
56. T. Sjöstrand, S. Ask, J.R. Christiansen, R. Corke, N. Desai, P. Ilten et al., An introduction to PYTHIA 8.2. *Comput. Phys. Commun.* **191**, 159 (2015). <https://doi.org/10.1016/j.cpc.2015.01.024> arXiv:1410.3012
57. S. Ovin, X. Rouby, V. Lemaitre, DELPHES, a framework for fast simulation of a generic collider experiment. arXiv:0903.2225
58. DELPHES 3 Collaboration, DELPHES 3, A modular framework for fast simulation of a generic collider experiment. *JHEP* **02**, 057 (2014). [https://doi.org/10.1007/JHEP02\(2014\)057](https://doi.org/10.1007/JHEP02(2014)057). arXiv:1307.6346
59. M. Cacciari, G.P. Salam, G. Soyez, FastJet user manual. *Eur. Phys. J. C* **72**, 1896 (2012). <https://doi.org/10.1140/epjc/s10052-012-1896-2> arXiv:1111.6097
60. M. Cacciari, G.P. Salam, G. Soyez, The anti- k_t jet clustering algorithm. *JHEP* **04**, 063 (2008). <https://doi.org/10.1088/1126-6708/2008/04/063> arXiv:0802.1189
61. fastjet-contrib. <https://fastjet.hepforge.org/contrib/>
62. I. Moult, B. Nachman, D. Neill, Convolved substructure: analytically decorrelating jet substructure observables. *JHEP* **05**, 002 (2018). [https://doi.org/10.1007/JHEP05\(2018\)002](https://doi.org/10.1007/JHEP05(2018)002) arXiv:1710.06859
63. A. Hocker et al., TMVA—toolkit for multivariate data analysis. arXiv:physics/0703039
64. F. Chollet et al., Keras (2015). <https://keras.io>