

Boosting $H \rightarrow b\bar{b}$ with machine learning

Joshua Lin,^{a,b} Marat Freytsis,^c Ian Mould^{d,e} and Benjamin Nachman^a

^a*Physics Division, Lawrence Berkeley National Laboratory,
Berkeley, CA 94720, U.S.A.*

^b*Department of Physics, University of California,
Berkeley, CA 94720, U.S.A.*

^c*Institute of Theoretical Science, University of Oregon,
Eugene, OR 97403, U.S.A.*

^d*Berkeley Center for Theoretical Physics, University of California,
Berkeley, CA 94720, U.S.A.*

^e*Theoretical Physics Group, Lawrence Berkeley National Laboratory,
Berkeley, CA 94720, U.S.A.*

E-mail: joshua.z.lin@berkeley.edu, freytsis@uoregon.edu,
ianmould@lbl.gov, bpnachman@lbl.gov

ABSTRACT: High- p_T Higgs production at hadron colliders provides a direct probe of the internal structure of the $gg \rightarrow H$ loop with the $H \rightarrow b\bar{b}$ decay offering the most statistics due to the large branching ratio. Despite the overwhelming QCD background, recent advances in jet substructure have put the observation of the $gg \rightarrow H \rightarrow b\bar{b}$ channel at the LHC within the realm of possibility. In order to enhance the sensitivity to this process, we develop a two-stream convolutional neural network, with one stream acting on jet information and one using global event properties. The neural network significantly increases the discovery potential of a Higgs signal, both for high- p_T Standard Model production as well for possible beyond the Standard Model contributions. Unlike most studies for boosted hadronically decaying massive particles, the boosted Higgs search is unique because double b -tagging rejects nearly all background processes that do not have two hard prongs. In this context — which goes beyond state-of-the-art two-prong tagging — the network is studied to identify the origin of the additional information leading to the increased significance. The procedures described here are also applicable to related final states where they can be used to identify additional sources of discrimination power that are not being exploited by current techniques.

KEYWORDS: Jets

ARXIV EPRINT: [1807.10768](https://arxiv.org/abs/1807.10768)

Contents

1	Introduction	1
2	Machine learning architecture	3
2.1	Neural network architecture	3
2.2	Inputs and preprocessing	4
3	Boosting standard model Higgs tagging	5
3.1	Simulation setup and validation	5
3.2	Machine learning results	7
3.3	What is the neural network learning?	8
3.3.1	Jet substructure	9
3.3.2	Global event	12
3.4	Recommendations for future searches	12
4	High-p_T Higgs for BSM physics	13
4.1	Simulation	15
4.2	Results and interpretation	15
5	Conclusions	17

1 Introduction

Even though it has been over five years since the discovery of the Higgs boson [1, 2], the final state with the largest branching ratio ($\mathcal{B}(H \rightarrow b\bar{b}) \approx 58\%$ [3]) has not been probed with great precision. This process is difficult to measure due to large and nearly irreducible background processes — only recently has $(V)H \rightarrow b\bar{b}$ been confirmed [4–6]. However, many interesting and largely untested features of high- p_T Higgs production [7–11] are challenging to probe with cleaner final states such as $H \rightarrow \gamma\gamma$ or $H \rightarrow ZZ^* \rightarrow 4\ell$ due to their low branching ratios.¹ Boosted $H \rightarrow b\bar{b}$ provides access to the highest p_T Higgs bosons at the Large Hadron Collider (LHC); if they can be measured with good precision, a door leading beyond the Standard Model (BSM) could be opened.

Major advances in the use of jet substructure and machine learning techniques have revolutionized the ability to look for hadronic signals in increasingly extreme regions of phase space.² Most analyses that exploit the hadronic decays of boosted heavy particles have so far used modern tools to purify the event selection but not to directly identify the

¹Studies of the recently observed $t\bar{t}H$ production [12, 13] provide complementary information on the underlying physics of $pp \rightarrow H$ production.

²For a review of recent theoretical and experimental progress, see refs. [14, 15].

main objects of interest. However, pioneering work by the ATLAS and CMS collaborations have used these techniques to directly measure boosted particle cross-sections [16–19] and to search directly for BSM particles [20–22]. In particular, the CMS collaboration has used single jets to search for the boosted $H \rightarrow b\bar{b}$ decay [19], the first experimental result on the subject since the idea was originally proposed in ref. [23] (albeit exclusively in the VH channel), as well as combining this analysis with other differential low- p_T data [24]. One reason for the long delay between conception and practical results was the development of advanced techniques for grooming [25–27], 2-prong tagging [28–31], jet four-vector calibrations [32, 33], and boosted b -tagging [34–41].

The presence of multiple nearby boosted b quarks sets boosted Higgs identification apart from other boosted massive particle classification. This is because requiring two b -tagged subjets inside a larger jet necessarily requires that the parent jet has a two-prong structure. For boosted massive object identification, most of the jet substructure community has focused on n -prong taggers [28–31, 42–49], which are not optimized for cases where n -prongs are already present. By probing the full radiation pattern inside boosted boson decays, ref. [50] showed that there is information beyond traditional n -prong tagging and even beyond traditional color flow observables [51]. This was also explored in the context of boosted Higgs boson decays in ref. [49], which identified simple observables that capture the additional information. However, neither these studies nor the more recent ref. [52], which considered generic quark and gluon jets as a background to $H \rightarrow b\bar{b}$, were explicitly predicated on subjet tagging as a baseline and did not probe global information beyond jet substructure. Until the present work, the full potential of information beyond n -prong taggers has not been demonstrated for concrete observables such as cross-sections or BSM coupling limits.³

Modern machine learning (ML) tools have shown great promise for using low-level [39, 46, 49, 50, 52, 56–86] and high-level [87–90] information to classify hadronic final states at the LHC. These techniques must be adapted to cope with significant sparsity, large dynamic ranges, multi-channel inputs and data that has no unique representation. Similar techniques have been demonstrated for full event classification with low-level [63, 82] and high-level inputs [91–93]. In addition to the challenges related to the structure of the data, one of the key challenges for applying state-of-the-art techniques in practice is the need for a background estimation method. As in ref. [77], boosted $H \rightarrow b\bar{b}$ has a natural background estimation technique by using the localization of the Higgs boson in the jet

³The latest CMS $b\bar{b}$ and $c\bar{c}$ tagging techniques use machine learning approaches with a large number of particle- and vertex-level inputs [53]. These approaches could learn information beyond n -prong tagging. However, the background source for training is generic quark and gluon jets, not $g \rightarrow b\bar{b}$. The working point used in the boosted H search operates at the 1% mis-tag rate, while the rate of $g \rightarrow b\bar{b}$ is comparable to or lower than this value (see e.g., ref. [54]) so most of the tagger’s effort must go to reducing the large non- $g \rightarrow b\bar{b}$ background. Performance studies specifically with $g \rightarrow b\bar{b}$ as the background show that the tagger reduces the $g \rightarrow b\bar{b}$ background $3\times$ more than the signal [55]. The equivalent performance shown later in this paper (figure 3) corresponds to about $16\times$ more $g \rightarrow b\bar{b}$ than signal. These numbers are not directly comparable because the latter is also after mass and two-prong tagging requirements (and is thus conservative). Therefore, the techniques presented in this paper are using more information, but further studies are required to understand how much more and what type of information is being used.

mass distribution. For this reason, the algorithms presented here may already be useful to enhance existing analysis efforts.

In this paper we use deep neural networks to examine the potential of using all the available information in boosted Higgs events. We use a two stream convolutional neural network to combine jet substructure information with global event information, finding significant gains coming from both components demonstrating that the search can be greatly improved. Furthermore, we are able to identify the dominant source of jet substructure discrimination in terms of a simple observable.

This paper is organized as follows. section 2 details the ML setup, including the preprocessing and architecture of the neural network. The neural network is then applied to the SM search for boosted $H \rightarrow b\bar{b}$ in section 3. Physics beyond the SM could introduce p_T -dependent effects that are enhanced for boosted Higgs bosons. Implications for the NN classifier on BSM physics are described in section 4. The paper concludes in section 5.

2 Machine learning architecture

This section describes our machine learning setup, with a focus on the neural network architecture and preprocessing.

2.1 Neural network architecture

Our neural network architecture is driven both by physics goals as well as the desire to extract the maximal amount of information from the event. For the boosted $H \rightarrow b\bar{b}$ topology, there are two physically distinct components to the events: the substructure of the hardest jet and the global event structure. Due to the color singlet scalar nature of the Higgs, the radiation pattern within and around the $b\bar{b}$ jet is expected to differ from $g \rightarrow b\bar{b}$ jets. Different production mechanisms can also result in different numbers and orientations of jets in the events. All of these aspects are investigated.

To incorporate both local and global information, a two-stream neural network is constructed. One stream acts on the full event information and the other acts on the image of the Higgs candidate jet. The two streams are then combined. This setup can be used to separately assess how much discrimination power can be obtained from the substructure and the global event separately, as well as in combination. A schematic of this two-stream architecture is shown in figure 1.

In order to account for the compact nature of the detector in the ϕ direction, we use padding layers that take the leftmost few columns and append them to the right before each convolution (for all convolutional layers), effectively performing convolutions over the cylinder rather than over a square. Further details related to the image (pre)processing are discussed in section 2.2.

The details of the convolution and pooling layers of each stream are as follows. Each convolutional filter is 5×5 , and the pooling layers are 2×2 , with rectified linear unit (ReLU) activations, and stride length of 1. The first convolutional layer in each stream has 32 filters, and the second convolutional layer in each stream has 64 filters. The dense layer at the end of each stream has 300 neurons each. Finally, the two dense layers from

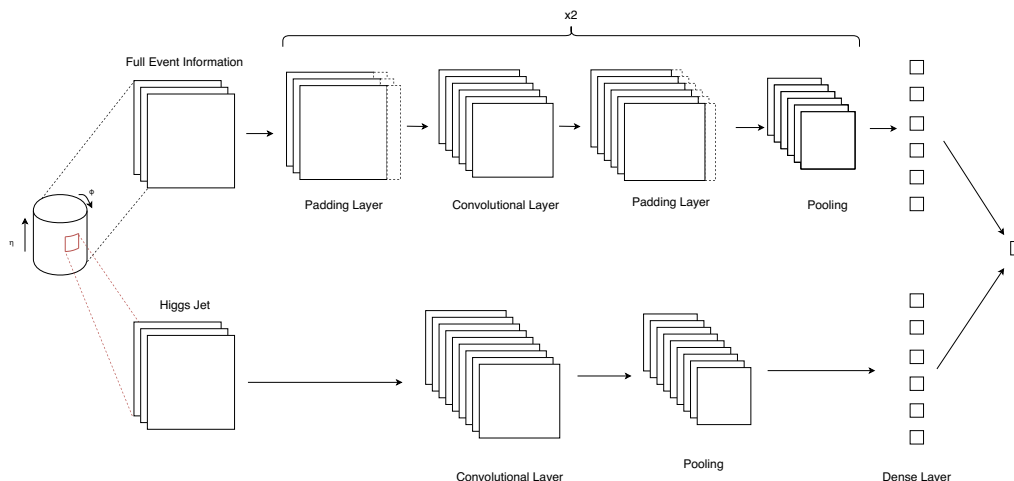


Figure 1. A schematic of the two-stream CNN used in this study. The first stream uses the full event information, while the second stream uses the jet substructure information. More details on the architecture are provided in the text.

each stream are fully connected to an output layer of one neuron with sigmoid activation. In total this gives 2.6 million trainable parameters in the network. We used the AdaDelta optimizer [94], with binary cross entropy as our loss function, and used the relatively simple Early Stopping method as a regularization technique, stopping when the significance improvement of the Higgs measurement at $p_T^{\min} = 450$ GeV stopped improving (with a patience of 2 epochs). We arrived at this final model after testing the performance (measured by the significance improvement of the Higgs measurement at $p_T^{\min} = 450$ GeV) using different optimizers (AdaDelta [94], AdaGrad [95], Adam [96]), different activation functions (mainly testing ReLU against leaky ReLU), and regularization (dropout [97] vs. Early Stopping). Our training was performed using the KERAS [98] Python neural network library with TENSORFLOW [99] backend, on Nvidia GeForce 1080 Ti GPUs.

2.2 Inputs and preprocessing

The inputs to our neural network are jet images [56]. For each event, an image is created for each stream: one image is the full event image and the other is the image of the hardest jet (that has been double b -tagged). Both images are 40×40 pixels. For the jet image, the range (in η - ϕ space) is $2R \times 2R$ where $R = 0.8$ is the radius of the jet. The full event image covers effectively the entire η - ϕ cylinder ($|\eta| < 5$). Inspired by ref. [60], both the jet and event image have three channels analogous to the RGB channels of a color image. The pixel intensity for the three channels correspond to the sum of the charged particle p_T , the sum of the neutral particle p_T , and the number of charged particles. As the neutral particle p_T is particularly sensitive to pileup, additional studies without this channel are included in the results.

To ensure that the neural network is not learning spacetime symmetries, and to reduce the size of the input streams, the jet images are preprocessed in a similar way to previous studies, see e.g., refs. [50, 60]. In particular, all of the images are normalized (sum of

intensities is unity) and standardized (zero-centered and divided by standard deviations). Prior to these steps, the jet images are also rotated so that the two subjets are aligned along the same axis in every image [50, 56]. Details about the subjet identification and b -tagging are discussed in section 3.1.

3 Boosting standard model Higgs tagging

This section studies the neural network performance in the context of improving the significance for the Standard Model boosted $H \rightarrow b\bar{b}$ search.

3.1 Simulation setup and validation

Simulated pp collisions at $\sqrt{s} = 13$ TeV are generated using MADGRAPH5_AMC@NLO 2.6.2 [100] for the hard processes and showered with PYTHIA 8.226 [101].⁴ Background events are generated using two, three and four jet events ($pp > jj$, $pp > jjj$ and $pp > jjjj$) matched using the MLM approach [104]. In order to include finite top mass effects (and BSM contributions in section 4) signal events are generated at one-loop order ($pp > Hj$ [QCD] and $pp > Hj j$ [QCD]), which in this case corresponds to the leading contribution. The overlap between the real emission from the matrix element and the parton shower is also accounted for using the MLM algorithm. Higher order amplitudes for the signal process with full mass dependence are now becoming available [105–111]. These updates could slightly modify the numerical results, but should not change the conclusions and would primarily effect the overall rate and not the features exploited by the machine learning approach, which are primarily associated with the radiation pattern in the jet and the global event. Furthermore, in these studies, higher loop finite top mass effects are found to be flat at high p_T and therefore do not significantly modify the shape of the p_T spectrum.

Events are clustered and analyzed using FASTJET 3.2.1 [112] and the FASTJET CONTRIBUTION extensions. Following a CMS-like analysis [19], jets are clustered with $R = 0.8$ anti- k_t jets [113], which are groomed with the soft drop algorithm [27] with $\beta = 0$ and $z_{\text{cut}} = 0.1$.⁵ Candidate Higgs jets are required to have transverse momentum $p_T > 450$ GeV, and satisfy a double b -tag. In general, b -tagging performance depends heavily on the exact experimental implementation and is detector specific. In this analysis, we use an approach similar to subjet b -tagging in ATLAS⁶ [114] which can be mimicked at particle-level while assuming 100% b -tagging efficiency and infinitely good rejection. This introduces an $O(1)$ correction

⁴In our study, we only use events showered by PYTHIA 8.226 [101]. Although it is true that by switching to a different showering program such as Herwig [102] or Sherpa [103] the quantitative improvement shown in section 3.2 may change, the qualitative gains from the Machine Learning should be robust and motivates a complete analysis to be undertaken by the LHC collaborations. Previous studies with deep learning on images have observed that raw classifier performance varies when different generators are used for testing but give nearly the same result when the testing models are the same (even if the training ones are different). [60]

⁵With this choice of β , the algorithm is the same as the precursor modified mass drop tagger [25] approach. It is called soft drop (SD) in this paper in order to avoid confusion, as this is the name used in the CMS analysis.

⁶The CMS approach does not explicitly reconstruct subjets. See section 5 for additional discussion on this point.

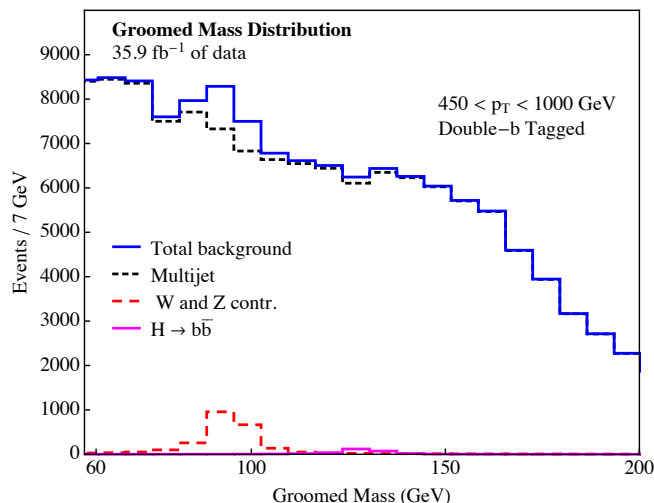


Figure 2. Groomed mass distribution for signal and background (including leading order vector boson contributions). The small Higgs boson signal is visible near $m_{SD} \sim 125$ GeV. Note that the simulation statistical uncertainty is comparable to the data statistical uncertainty. While not ideal, this does not qualitatively change the conclusions.

to the cross-section, but does not qualitatively change the conclusions. The subjects of the large- R jets are ghost-associated [115] $R = 0.2$ anti- k_t jets. Such jets are declared b -tagged if they have a ghost-associated B hadron with $p_T > 5$ GeV. In addition to b -tagging, the leading double b -tagged jet (the Higgs candidate) is required to have $-6.0 < \rho < -2.1$ ($\rho = \log(m_{SD}^2/p_T^2)$). This is chosen following ref. [19] to avoid the deeply nonperturbative region as well as finite cone limitations in the jet clustering, although no re-optimization of this range was performed. Finally, the two-prong observable N_2 [29] is required to be ≤ 0.4 . There is little dependence on the exact N_2 requirement, likely in part because of the two (b -tagged) subjet requirement. figure 2 shows the m_{SD} distribution after applying the above selections. The overall rate, relative rates between processes, and general trends agree with the CMS analysis in ref. [19].

Since the goal of this paper is to emphasize the possible gains for this search using ML, we have made a number of simplifying assumptions, and therefore the exact reproduction of the CMS analysis is not our primary concern. We believe that none of these assumptions significantly change our quantitative conclusions, but they should be revisited with the full analyses in ATLAS and CMS. In particular, the $t\bar{t}$ background is ignored, the background fit is simplified, experimental effects relating to track reconstruction and b -tagging are ignored, and as mentioned above, the Higgs cross-section is only computed at NLO. The top background is small but comparable to the Higgs signal, and since we have consistently ignored it for the pseudo-data and background, any residual contribution is a subleading effect from modeling uncertainty. Tracks reconstructed by ATLAS and CMS are excellent proxies for charged particles, though there are percent-level differences resulting from material interactions and pattern recognition ambiguities. These effects, as well as pileup, can slightly degrade b -tagging performance [39, 84]. Once again, this is

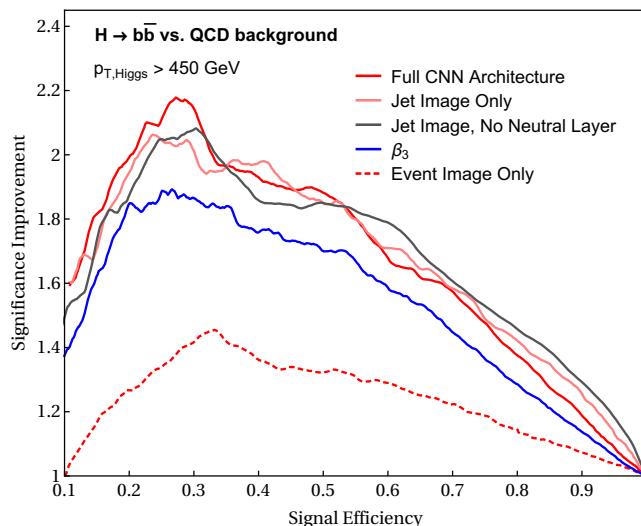


Figure 3. Significance improvement characteristic (SIC) curve for various neural network setups as well as for the single observable β_3 proposed in ref. [49].

important to account for when setting a precise limit, but would not change the relative gains presented here.

3.2 Machine learning results

Having validated the setup specified in section 3.1 against the public CMS results [19], the simulated events are now used as input to our two stream convolutional neural network to identify whether additional discrimination power can be obtained from the jet substructure, jet superstructure, and other global event properties.

Network training proceeds with 50 000 signal and background events passing the selection criteria from section 3.1. The training-validation-test split that we used was 50%–25%–25%. There is no requirement on the jet mass, as the entire spectrum is used to evaluate the significance. In practice, this could make traditional data-driven background estimation techniques more complex to use, though there have also been many techniques proposed to preserve the mass distribution [116–119].

The neural network performance is quantified using the significance improvement characteristic (SIC) curve. Such a curve is approximately equal to $\epsilon/\sqrt{\epsilon_b}$ and quantifies the gain in significance over the baseline selection. Following a CMS-like analysis [19], the full significance is calculated using a binned likelihood fit treating the bin counts as Poisson-distributed random variables. This procedure assumes that the results are dominated by statistical uncertainties, which will always be true for the highest p_T bins. Data statistical uncertainties account for over half of the total uncertainty in ref. [19], so this is a valid approximation. There is no fit to determine the background shape, which is taken directly from the simulation. Once again, this is valid in the statistics limited regime.

The binned likelihood fit is performed in the mass range from 50 GeV to 197 GeV and using bins of width 7 GeV. (The CMS analysis performs the same fit in the 40 GeV to 201 GeV with the same binning [19]). The corresponding SIC curve is shown in figure 3.

A maximum significance gain of about 2.2 is achieved with a signal efficiency of about 25%. This means that if the significance with the nominal selection was 1 for a given dataset size, then after the application of the neural network, the new significance would be 2.2. The maximum significance from the event stream only is about 1.4 while for the same value for the jet stream only is about 2. This indicates that the jet information is much more important than the global information, though a significance gain of 1.4 is still important. Since pileup is not included in the simulation, it is important to show that the performance is similar when pileup sensitive inputs are removed. The “no neutral layer” curve in figure 3 shows that the peak performance is robust and even better than the full network at high significance. Intuitively, a network with more information should not be able to do worse, though in practice, this could occur due to weight sharing or from too few training examples. For reference, the β_3 observable proposed in ref. [49],

$$\beta_3 = \frac{(\tau_1^{(0.5)})^2 (\tau_2^{(1)})^{0.5}}{\tau_2^{(2)}}, \tag{3.1}$$

where $\tau_n^{(j)}$ is the n -subjettiness observable [28, 42] with angular exponent j , is also shown for comparison in figure 3. This single observable captures a significant fraction of the total significance improvement, but there is still more information available from the full two stream setup to boost the significance further. A further investigation into the information learned by the network is described in section 3.3.1.

To understand the impact of a gain of 2.2 in the SIC, the expected significance for the SM $H \rightarrow b\bar{b}$ search is plotted as a function of integrated luminosity in figure 4. A center of mass energy of $\sqrt{s} = 13$ TeV is assumed through the end of LHC Run 3, which corresponds to about 300 fb^{-1} . The curves follow the statical scaling of $\sqrt{\int L dt}$, where L is the instantaneous luminosity. The current CMS result reported an observed (expected) significance of 1.5 (0.7) [19]. As anticipated from the agreement with the mass distribution (figure 2), the significance calculated using the simulation reported in section 3.1 is very similar at 1.227. Without machine learning, “evidence” (3σ) may only be achieved after the full LHC dataset (up to 2023) and “observation” (5σ) may be possible only with the HL-LHC. In contrast, with the application of the neural network, evidence may be achievable with the full Run 2 (2015–2018) dataset (about 150 fb^{-1}) and observation may be possible well before the end of the LHC. This represents one of the main results of this paper, and emphasizes the possible gains to be had with ML.

3.3 What is the neural network learning?

With a significant improvement from the neural network, it is interesting to investigate in more detail what information the machine is exploiting beyond the existing search. This section follows some of the procedures for such a study described in ref. [50].

First, figure 5 shows the (first layer) convolutional filters from both streams of the network. Since both streams are actually three-channel images, there are three sets of filters for each case. While it is difficult to immediately recognize what the network is learning from these filters, there are some hints upon careful inspection. In particular, the

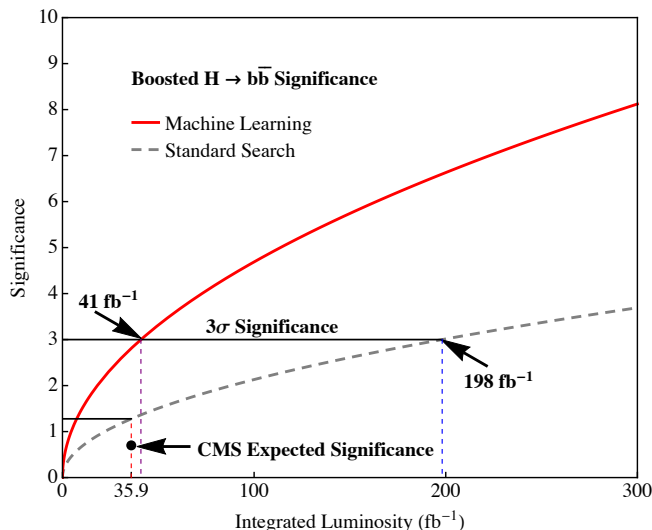


Figure 4. The expected significance as a function of the integrated luminosity for the analysis with and without machine learning. The vertical red dashed line corresponds to the dataset size from the current CMS result [19] while the purple and blue dashed lines indicate the sizes required to reach 3σ with and without machine learning, respectively. The full Run 2 (2015–2018) dataset will be about 150 fb^{-1} and the full LHC dataset (up to 2023), prior to the HL-LHC, will be about 300 fb^{-1} .

event images have a small number of “hot spots.” This may indicate that the network is learning to compute distances between prongs within jets and between jets. In contrast, the jet image filters have many active pixels with complex shapes. These filters are too small to span the typical subjet distance and so may be identifying the pattern of radiation between or around subjets. The following sections examine the two streams of the network in more detail.

3.3.1 Jet substructure

As emphasized earlier, the $H \rightarrow b\bar{b}$ search is different from other boosted hadronically decaying massive boson studies because the application of double b -tagging already enforces a two-prong topology. Therefore, two-prong tagging is not as useful. Studies to further optimize the event selection with N_2 confirm this expectation — little significance gain is possible using only this state-of-the-art two-prong tagging technique (see also ref. [37]). One of the attractive features of jet images is that they can be directly inspected to visualize the information content. For example, figure 6 shows the average of the 100 most signal-like and most background-like jets, according to the neural network. The two-prong structure of both signal and background is clear in all three channels. The main difference between $gg \rightarrow b\bar{b}$ and $H \rightarrow b\bar{b}$ is the orientation of the radiation between and around the two prongs. As expected due to the different color structure, the radiation pattern around the two prongs is more spread out for the gluon case. figure 7 shows additional images that are split by their value of β_3 . It is clear from the images that low β_3 values (background-like) pick out subjets with a broader radiation patterns compared with high β_3 (signal-like)




Figure 5. The 32 filters from the first layer of the total event CNN in (a) and the jet substructure CNN in (b). The top row filters correspond to the charged p_T layer, the second row shows the neutral p_T layer and the bottom row is for the charged-particle multiplicity channel.

images. However, the top plot of figure 7 clearly indicates that β_3 is not the same as the neural network, so there is additional information to learn. figure 8 tries to visualize the additional information. The distribution of β_3 in the signal is reweighted to be the same as the background so that β_3 by itself is not useful for discrimination. The average images for signal and background look very similar by eye, but the difference of the average images reveals interesting structure. These structures still show an enhanced radiation pattern around the subjets for the background relative to the signal — there is thus more color flow information available to learn than is captured by β_3 alone.

Drilling down into the information content of the jet images in more detail, perhaps using more of the techniques from ref. [50] and understanding to what extent β_3 captures color flow and other effects is of great interest for future studies.⁷

⁷Explicit color flow tagging observables have been proposed in the past, such as the jet pull [51]. However, these observables are not powerful tools for tagging [50, 120], though they are very useful for precision studies [121, 122]. We have checked and the SIC for jet pull does not significantly exceed unity.

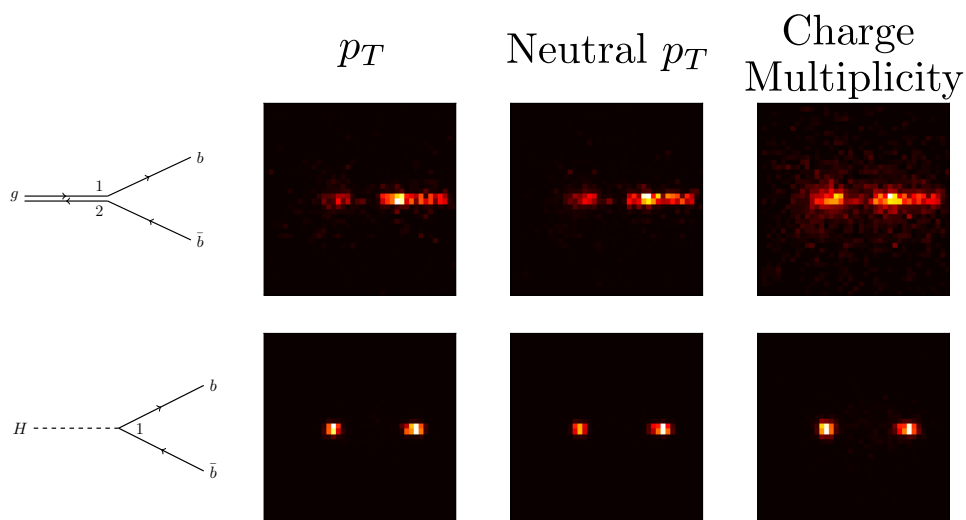


Figure 6. Average jet images for the 100 most background like (top) and signal like (bottom) jets. The jet images are weighted by the p_T in the first column, the neutral p_T in the second column, and the charge multiplicity in the third column. Due to the different color flows, the signal like ($H \rightarrow b\bar{b}$) jets have a more contained color flow pattern.

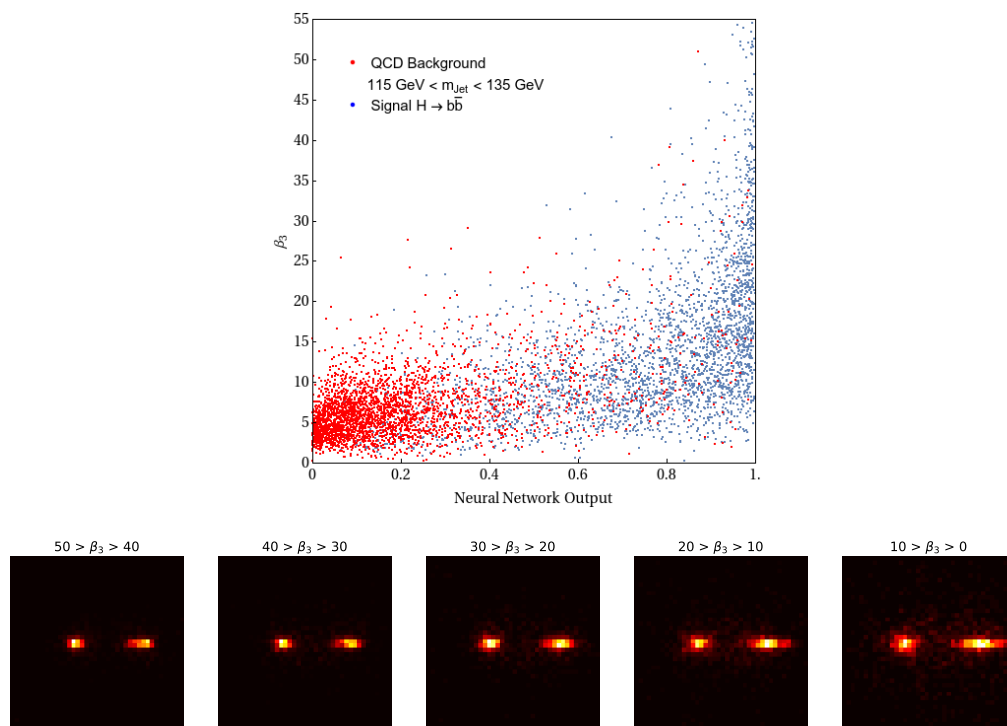


Figure 7. Top: the joint distribution of β_3 and the neural network. Bottom: the average of many signal images in decreasing β_3 intervals.

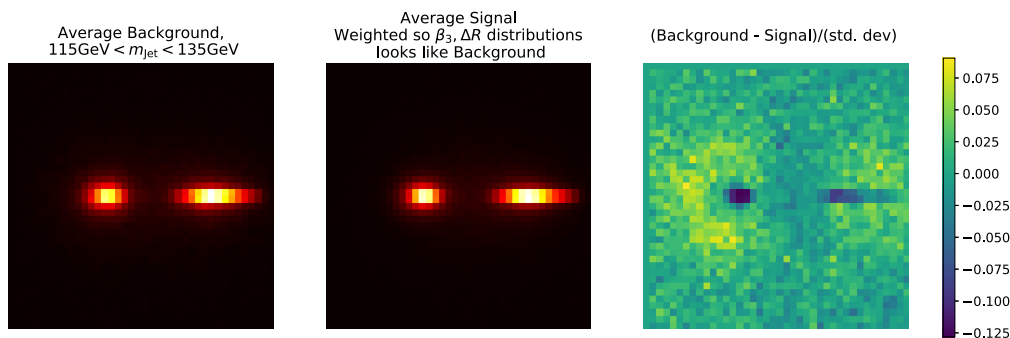


Figure 8. Left (middle): the average background (signal) images after reweighting the β_3 distribution and the ΔR (between b-tagged subjects) distribution in signal to be the same as in background. Right: the average difference between the left and middle images.

3.3.2 Global event

While much attention has been devoted to the extraction of information from jet substructure, less has been paid to the extraction of discrimination power from the full event. At the same time, probing what is learned from event properties is more complicated than for the jet image due to the reduced symmetry. As with the jet substructure, due to the color singlet nature of the Higgs, we expect that the color flows in signal and background jets should be distinct. From our study, we find that while this information does not provide as much discrimination power as the jet substructure, it nonetheless provides an additional gain in significance. While several observables for discriminating global color flow have been proposed [51, 123–125], this is in general quite a challenging task. Furthermore, we expect that it would be quite topology dependent. Nevertheless it would be interesting study in more detail, since it has not received much attention. We believe that ML is an ideal technique for extracting complicated global event information that has not yet been exploited to its full potential in LHC searches.

We also highlight the efficacy of the padding that renders the convolutional layers of the neural network symmetric under rotations in ϕ by one pixel. This is a new feature of our neural network, which we find to be helpful for training stability. figure 9 shows a typical signal image and how the neural network output changes as the event is rotated in ϕ . As desired, the network with the padding at every convolution layer is much more stable than the ones without the padding. The reason that the padded network is not completely invariant under rotations in ϕ is that the dense layers at the end of the network break the ϕ symmetry while for rotations as scales below a single pixel the discretization breaks the invariance. figure 10 shows a similar trend after averaging over many events.

3.4 Recommendations for future searches

Due to the important of the $H \rightarrow b\bar{b}$ channel for probing the Higgs sector at the LHC, we conclude this section with some concrete recommendations for improving the LHC searches, reiterating the points found in this section. In particular, although the most power is gained from a neural network, we have shown that a large fraction of this information

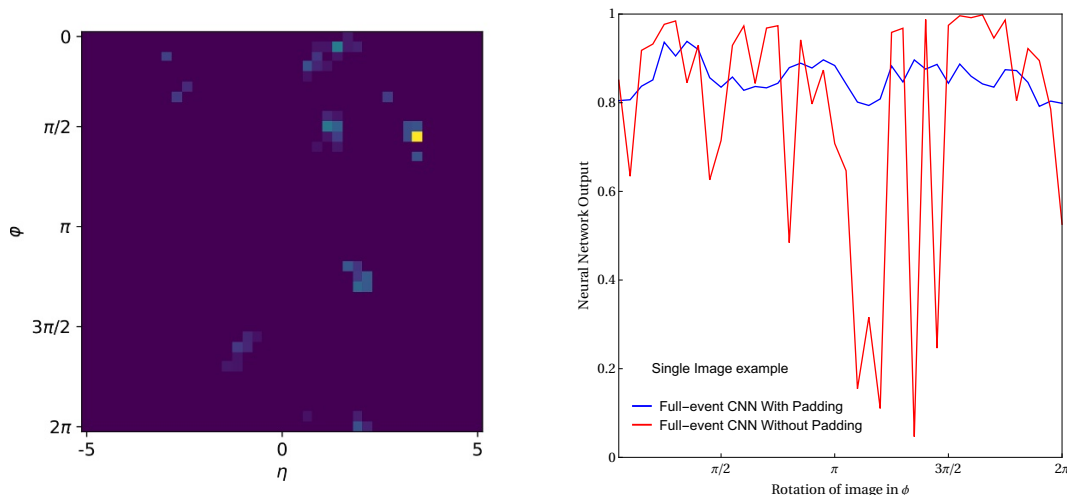


Figure 9. Left: a typical signal event image. Right: the output of the neural network on the left image, after rotation in the ϕ direction by the given number of pixels.

can be obtained through simple observables, which can immediately be implemented in current searches for boosted $H \rightarrow b\bar{b}$. A neural network using charged information could also be applied without requiring extensive calibration studies. A key component of what the network can learn since the signal and background are already in a two-prong topology is the color flow. Quantifying the additional information in the form of compact analytic observables is an interesting and important part of future work.

4 High- p_T Higgs for BSM physics

Beyond the discovery of the $H \rightarrow b\bar{b}$ decay, a major motivation for the study of boosted $H \rightarrow b\bar{b}$ final states in particular is that it allows one to study the structure of the $gg \rightarrow H$ process at high p_T . While in the Standard Model this is primarily due to the contribution of a virtual top quark loop, the total cross section $\sigma(gg \rightarrow H)$ is only sensitive to the low-energy limit of this loop, in which it is extremely well approximated by a dimension-five operator with no dependence on m_t . At $p_T \gtrsim 2m_t$, this is no longer true, as the physical momentum running through the loop is comparable to m_t , allowing potential new physics contributions to the loop to be disentangled that are not observable for the total cross section by observing the p_T dependence. This general observation has been explored in refs. [7–11]. In this section we apply our machine learning techniques and illustrate how the improved significance for $H \rightarrow b\bar{b}$ translates to improved bounds on BSM physics.

We are interested in probing new physics in the $gg \rightarrow H$ production loop that can be modeled as dimension-6 operators. Following ref. [7], the operators modifying $gg \rightarrow H$ production cross section are parameterized as

$$\begin{aligned} \mathcal{L}_{\text{eff}} = \mathcal{L}_{\text{SM}} &+ \left(c_y \frac{y_t}{v^2} |H|^2 \bar{Q}_L \tilde{H} t_R + \text{h.c.} \right) + c_H \frac{1}{2v^2} \partial_\mu |H|^2 \partial^\mu |H|^2 \\ &+ c_g \frac{\alpha_s}{12\pi v^2} |H|^2 G_{\mu\nu}^a G^{a\mu\nu} + \tilde{c}_g \frac{\alpha_s}{8\pi v^2} |H|^2 G_{\mu\nu}^a \tilde{G}^{a\mu\nu}. \end{aligned} \quad (4.1)$$

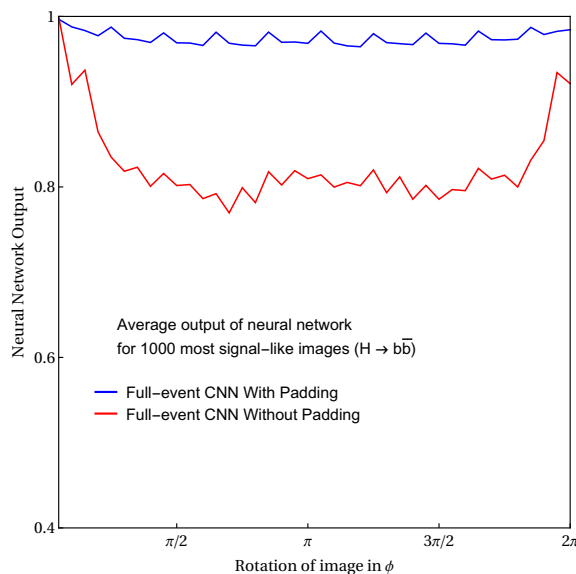


Figure 10. The average signal neural network output after rotating the images by a fixed amount.

Here $G^{a\mu\nu}$ is the QCD field strength, and $\tilde{G}^{a\mu\nu} = \frac{1}{2}\epsilon^{\mu\nu\sigma\rho}G_{\sigma\rho}^a$ its dual. After electroweak breaking, the induced operators affecting the coupling of the Higgs boson to tops and gluons take the form

$$\mathcal{L}_{\text{eff}} = \mathcal{L}_{\text{SM}} - \kappa_t \frac{m_t}{v} t\bar{t}h + i\tilde{\kappa}_t \frac{m_t}{v} \bar{t}\gamma_5 t h + \kappa_g \frac{\alpha_s}{12\pi} \frac{h}{v} G_{\mu\nu}^a G^{\mu\nu a} + \tilde{\kappa}_g \frac{\alpha_s}{8\pi} \frac{h}{v} G_{\mu\nu}^a \tilde{G}^{a\mu\nu}, \quad (4.2)$$

where

$$\kappa_t = 1 - \text{Re}(c_y) - \frac{c_H}{2}, \quad \kappa_g = c_g, \quad \tilde{\kappa}_t = \Im(c_y), \quad \tilde{\kappa}_g = \tilde{c}_g, \quad (4.3)$$

so that one degeneracy between c_H and the real part of c_y remains. In the following we will only be interested in CP -even terms, and explicitly set $\tilde{\kappa}_t = \tilde{\kappa}_g = 0$, implicitly demanding c_y be real.

Although the two couplings that we are interested in probing have distinct physical effects, namely κ_t acts as a correction to the top Yukawa, while κ_g corrects the ggH coupling, the Higgs low energy theorem [126, 127] guarantees that they contribute to the inclusive Higgs production cross section as $(\kappa_t + \kappa_g)^2$ up to corrections of $O(m_h^2/m_t^2)$. As shown in ref. [7], this degeneracy is broken by the cross-section for $H + \text{jet}$ production for a given p_T^{min} cut on the Higgs scales like

$$\frac{\sigma_{p_T^{\text{min}}}(\kappa_t, \kappa_g)}{\sigma_{p_T^{\text{min}}}^{\text{SM}}} = (\kappa_t + \kappa_g)^2 + \delta\kappa_t\kappa_g + \epsilon\kappa_g^2, \quad (4.4)$$

where ϵ and δ are terms dependent on the p_T^{min} cut placed on the Higgs. These are given in ref. [7] for tabulated values of p_T^{min} , computed at one-loop order with full m_t dependence. For $p_T^{\text{min}} \lesssim 2m_t$, behavior similar to the inclusive rate is observed as $\epsilon, \delta \ll 1$, while for $p_T^{\text{min}} \gtrsim 2m_t$, ϵ, δ become of $O(1)$, leading to a hardening in the p_T spectrum of the BSM model as compared to the SM model after taking into account the overall scaling of the

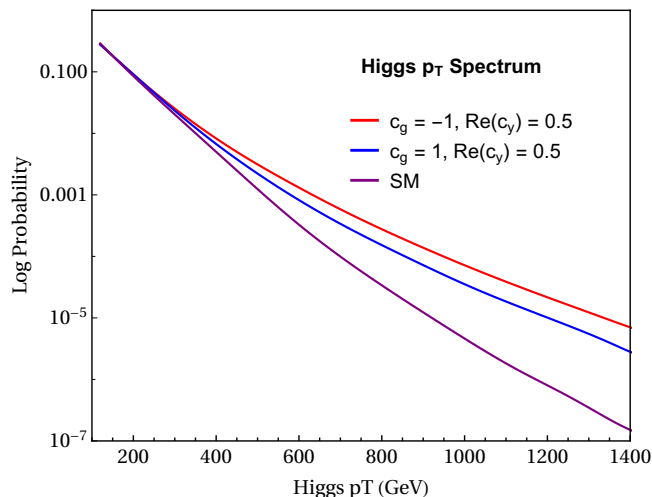


Figure 11. Typical hardening of the p_T spectrum of the Higgs for nonzero values of the Wilson coefficient c_g , defined in the text.

cross-section given by the $(\kappa_t + \kappa_g)^2$ term (provided that $c_g \neq 0$). We can see this effect in figure 11. Consequently, performing a search for boosted Higgs can provide bounds on the Wilson coefficients $\kappa_g = c_g$ and the combination $\kappa_t = 1 - \Re(c_y) - c_H/2$.

4.1 Simulation

Our signal of interest is dominated by the interference of the SM $gg \rightarrow Hj(j)$ process, with the higher dimensional operators given above. Since MADGRAPH5_AMC@NLO is currently unable to compute the interference effects of processes that start at loop-level, a modification of the typical MADGRAPH5_AMC@NLO procedure is necessary to correctly generate the processes we want to study. The effect of the operators parameterized by c_g and \tilde{c}_g are recovered by implementing a fictitious heavy top partner whose mass is set to be large enough that a contact operator approximation remains valid for all LHC processes (nominally 10 TeV) and whose coupling to the Higgs is tuned to give the correct dependence on the high-dimension operator coefficient. The other operators are implemented as actual higher-dimension operators, as is conventional.

4.2 Results and interpretation

We apply our two-stage convolutional neural network to improve the bounds on the couplings c_g and the combination $\Re(c_y) + c_H/2$ as compared with the standard search. In figure 12 we show the constraints from comparing the inclusive cross section to one with $p_T^{\min} = 650$ GeV. We see significant gains using machine learning, corresponding to the improved significance for the Higgs seen earlier.

To properly situate these results, we summarize the current bounds and future prospects for constraining these operator coefficients in the absence of a dedicated high- p_T Higgs analysis. While theoretical constraints from general principles such as causality and locality do exist [128] (in particular $\Re(c_y) + c_H/2 > 0$ seems to always be true when

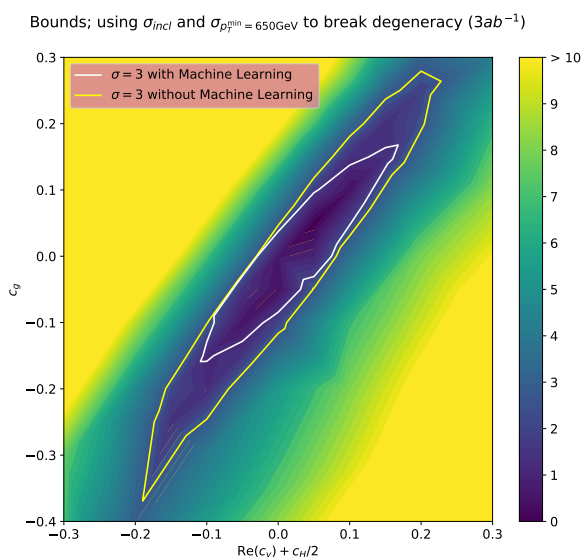


Figure 12. Expected constraints on the couplings c_g , c_y and c_H derived with and without ML. A significant improvement along the anticorrelated direction of c_g and $\Re(c_y) + c_H/2$ is observed using ML.

generated within quantum field theory), the best current bounds on the couplings c_g and $\Re(c_y) + c_H/2$ come from a combination of the most recent inclusive Higgs cross section measurements [129] and recent global fit of the Standard Model Effective Field Theory to all current Higgs and electroweak data performed in ref. [130]. The inclusive Higgs measurement constrains a combination of couplings for which the linearization $c_g - \Re(c_y) - c_H/2$ is an excellent approximation to be $0.29(46)$ at 3σ using 36.1 fb^{-1} of data. The global fit provides current world averages (again with 3σ uncertainties) of $c_g = 0.10(30)$, $\Re(c_y) = -4.7(78)$, $c_H = -1.1(18)$, all consistent with zero at 2σ . These results are primarily driven by the LHC Run II Higgs measurements, all using 36.1 fb^{-1} (35.9 fb^{-1}) of data from ATLAS (CMS), although since the effect of possible other higher-dimension operators on the backgrounds is not included these bounds should be interpreted with care. The dominant discriminating power is provided by looking for deviations in the $h \rightarrow WW^*, ZZ^*$ decays, with the most constraining bounds clearly being on the coefficient c_g .⁸

Conservatively assuming no improvement in the treatment of systematic or theoretical errors, analogous bounds to those discussed above with a full 3 ab^{-1} dataset should be able to reduce uncertainties by a factor of 2 in the inclusive Higgs cross section linear combination and by 20% to 25% for the global fits. Comparing this to the projections of figure 12, our proposed analysis has the potential to exceed these conservative extrapolations on sensitivity by a factor of a few.

⁸A recent combined fit from CMS using high- p_T $H \rightarrow b\bar{b}$ and well as differential $H \rightarrow \gamma\gamma$, and $H \rightarrow ZZ^* \rightarrow 4\ell$ decays [24] finds bounds of $c_g \simeq 0.12 \pm 0.42$, $|\Re(c_y) + c_H/2| \lesssim 0.5$ using our conventions. These are nearly competitive with the global fits on c_g and would clearly improve the fits of the other bounds if included in the global average. However, they also assume the Higgs branching ratios maintain their Standard Model values and are significantly weaker if these ratios are allowed to float.

5 Conclusions

In this paper, we have applied modern machine learning techniques to improve the search for the $H \rightarrow b\bar{b}$ decay at the LHC. This decay offers a powerful probe of BSM contributions to the $gg \rightarrow H$ loop at high p_T . Using our techniques, this process may be discoverable at the LHC (prior to the HL-LHC).

A new feature of our analysis is that we have used a two stream convolutional neural network, with one stream acting on the double b -tagged jet, and the other stream acting on the global event information. This enables us to not only exploit the maximal information in the event, combining both jet substructure information and global information, but also allows us to more easily identify the dominant physics features that the neural network is learning. In particular, we find that a significant fraction of this information is not contained in the recently proposed β_3 observable. Disentangling these differing sources of information is challenging in standard analyses, which substructure observables nominally designed to identify two-prong substructure, although in the course of optimization they may become sensitive to other features as well. Resolving an event at multiple scales and in various regions of phase space is a generic technique that should enable significant improvements in other LHC searches. By probing the neural network in detail, it may also be possible to use the neural networks as a guide to building compact, analytical, simple observables that nearly saturate the machine learning performance. With such tools in hand, increasingly extreme regions of phase space can be thoroughly explored.

Acknowledgments

We thank the LBNL theory and experimental groups, especially Heather Gray, for lively discussions on this topic at the weekly “cookie time.” We thank Kaustuv Datta, Phil Harris, Andrew Larkoski, and Caterina Vernieri for comments and suggestions on the analysis and manuscript. We also thank Wahid Bhimji and Steve Farrell at the National Energy Research Scientific Computing Center (NERSC) for helpful discussions, Phil Harris for useful discussions and technical help on the sample generation, Olivier Mattelaer for help with MADGRAPH5_AMC@NLO, and Frank Tackmann for stimulating discussions at Les Houches 2017 that started this project. Finally, we are grateful for the opportunity to use the CORI and EDISON supercomputing resources at NERSC. This work was supported by the U.S. Department of Energy, Office of Science under contract DE-AC02-05CH11231. MF is supported by the U.S. Department of Energy under contract DE-SC0011640. JL is supported by the Center for Computational Excellence, a project funded by the Computational HEP program in U.S. Department of Energy Office of Science.

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

References

- [1] ATLAS collaboration, *Observation of a new particle in the search for the standard model Higgs boson with the ATLAS detector at the LHC*, *Phys. Lett. B* **716** (2012) 1 [[arXiv:1207.7214](#)] [[INSPIRE](#)].
- [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** (2012) 30 [[arXiv:1207.7235](#)] [[INSPIRE](#)].
- [3] LHC HIGGS CROSS SECTION WORKING GROUP collaboration, J.R. Andersen et al., *Handbook of LHC Higgs cross sections: 3. Higgs properties*, [arXiv:1307.1347](#) [[INSPIRE](#)].
- [4] ATLAS collaboration, *Observation of $H \rightarrow b\bar{b}$ decays and VH production with the ATLAS detector*, *ATLAS-CONF-2018-036* (2018).
- [5] ATLAS collaboration, *Evidence for the $H \rightarrow b\bar{b}$ decay with the ATLAS detector*, *JHEP* **12** (2017) 024 [[arXiv:1708.03299](#)] [[INSPIRE](#)].
- [6] CMS collaboration, *Evidence for the Higgs boson decay to a bottom quark–antiquark pair*, *Phys. Lett. B* **780** (2018) 501 [[arXiv:1709.07497](#)] [[INSPIRE](#)].
- [7] C. Grojean, E. Salvioni, M. Schlaffer and A. Weiler, *Very boosted Higgs in gluon fusion*, *JHEP* **05** (2014) 022 [[arXiv:1312.3317](#)] [[INSPIRE](#)].
- [8] A. Azatov and A. Paul, *Probing Higgs couplings with high p_T Higgs production*, *JHEP* **01** (2014) 014 [[arXiv:1309.5273](#)] [[INSPIRE](#)].
- [9] M. Buschmann et al., *Resolving the Higgs-gluon coupling with jets*, *Phys. Rev. D* **90** (2014) 013010 [[arXiv:1405.7651](#)] [[INSPIRE](#)].
- [10] M. Schlaffer et al., *Boosted Higgs shapes*, *Eur. Phys. J. C* **74** (2014) 3120 [[arXiv:1405.4295](#)] [[INSPIRE](#)].
- [11] A. Azatov, C. Grojean, A. Paul and E. Salvioni, *Resolving gluon fusion loops at current and future hadron colliders*, *JHEP* **09** (2016) 123 [[arXiv:1608.00977](#)] [[INSPIRE](#)].
- [12] ATLAS collaboration, *Search for the standard model Higgs boson produced in association with top quarks and decaying into a $b\bar{b}$ pair in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, *Phys. Rev. D* **97** (2018) 072016 [[arXiv:1712.08895](#)] [[INSPIRE](#)].
- [13] CMS collaboration, *Search for $t\bar{t}H$ production in the $H \rightarrow b\bar{b}$ decay channel with leptonic $t\bar{t}$ decays in proton–proton collisions at $\sqrt{s} = 13$ TeV*, [arXiv:1804.03682](#) [[INSPIRE](#)].
- [14] A.J. Larkoski, I. Moult and B. Nachman, *Jet substructure at the Large Hadron Collider: a review of recent advances in theory and machine learning*, [arXiv:1709.04464](#) [[INSPIRE](#)].
- [15] L. Asquith et al., *Jet substructure at the Large Hadron Collider: experimental review*, [arXiv:1803.06991](#) [[INSPIRE](#)].
- [16] ATLAS collaboration, *Measurement of the cross section of high transverse momentum $Z \rightarrow b\bar{b}$ production in proton–proton collisions at $\sqrt{s} = 8$ TeV with the ATLAS Detector*, *Phys. Lett. B* **738** (2014) 25 [[arXiv:1404.7042](#)] [[INSPIRE](#)].
- [17] ATLAS collaboration, *Measurement of the cross-section of high transverse momentum vector bosons reconstructed as single jets and studies of jet substructure in pp collisions at $\sqrt{s} = 7$ TeV with the ATLAS detector*, *New J. Phys.* **16** (2014) 113013 [[arXiv:1407.0800](#)] [[INSPIRE](#)].
- [18] ATLAS collaboration, *Measurements of $t\bar{t}$ differential cross-sections of highly boosted top quarks decaying to all-hadronic final states in pp collisions at $\sqrt{s} = 13$ TeV using the ATLAS detector*, *Phys. Rev. D* **98** (2018) 012003 [[arXiv:1801.02052](#)] [[INSPIRE](#)].

- [19] CMS collaboration, *Inclusive search for a highly boosted Higgs boson decaying to a bottom quark-antiquark pair*, *Phys. Rev. Lett.* **120** (2018) 071802 [[arXiv:1709.05543](#)] [[INSPIRE](#)].
- [20] CMS collaboration, *Search for low mass vector resonances decaying to quark-antiquark pairs in proton-proton collisions at $\sqrt{s} = 13$ TeV*, *Phys. Rev. Lett.* **119** (2017) 111802 [[arXiv:1705.10532](#)] [[INSPIRE](#)].
- [21] CMS collaboration, *Search for low mass vector resonances decaying into quark-antiquark pairs in proton-proton collisions at $\sqrt{s} = 13$ TeV*, *JHEP* **01** (2018) 097 [[arXiv:1710.00159](#)] [[INSPIRE](#)].
- [22] ATLAS collaboration, *Search for light resonances decaying to boosted quark pairs and produced in association with a photon or a jet in proton-proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, [arXiv:1801.08769](#) [[INSPIRE](#)].
- [23] J.M. Butterworth, A.R. Davison, M. Rubin and G.P. Salam, *Jet substructure as a new Higgs search channel at the LHC*, *Phys. Rev. Lett.* **100** (2008) 242001 [[arXiv:0802.2470](#)] [[INSPIRE](#)].
- [24] CMS collaboration, *Combined measurement and interpretation of differential Higgs boson production cross sections at $\sqrt{s} = 13$ TeV*, [CMS-PAS-HIG-17-028](#) (2017).
- [25] M. Dasgupta, A. Fregoso, S. Marzani and G.P. Salam, *Towards an understanding of jet substructure*, *JHEP* **09** (2013) 029 [[arXiv:1307.0007](#)] [[INSPIRE](#)].
- [26] M. Dasgupta, A. Fregoso, S. Marzani and A. Powling, *Jet substructure with analytical methods*, *Eur. Phys. J. C* **73** (2013) 2623 [[arXiv:1307.0013](#)] [[INSPIRE](#)].
- [27] A.J. Larkoski, S. Marzani, G. Soyez and J. Thaler, *Soft drop*, *JHEP* **05** (2014) 146 [[arXiv:1402.2657](#)] [[INSPIRE](#)].
- [28] J. Thaler and K. Van Tilburg, *Identifying boosted objects with N -subjettiness*, *JHEP* **03** (2011) 015 [[arXiv:1011.2268](#)] [[INSPIRE](#)].
- [29] I. Moult, L. Necib and J. Thaler, *New angles on energy correlation functions*, *JHEP* **12** (2016) 153 [[arXiv:1609.07483](#)] [[INSPIRE](#)].
- [30] A.J. Larkoski, G.P. Salam and J. Thaler, *Energy correlation functions for jet substructure*, *JHEP* **06** (2013) 108 [[arXiv:1305.0007](#)] [[INSPIRE](#)].
- [31] A.J. Larkoski, I. Moult and D. Neill, *Power counting to better jet observables*, *JHEP* **12** (2014) 009 [[arXiv:1409.6298](#)] [[INSPIRE](#)].
- [32] ATLAS collaboration, *In-situ measurements of the ATLAS large-radius jet response in 13 TeV pp collisions*, [ATLAS-CONF-2017-063](#) (2017).
- [33] CMS Collaboration, *W and top tagging scale factors*, [CMS-DP-2017-026](#) (2017).
- [34] ATLAS collaboration, *Identification and tagging of double b -hadron jets with the ATLAS Detector*, [ATLAS-CONF-2012-100](#) (2012).
- [35] CMS collaboration, *Performance of b tagging at $\sqrt{s} = 8$ TeV in multijet, $t\bar{t}$ and boosted topology events*, [CMS-PAS-BTV-13-001](#) (2013).
- [36] ATLAS collaboration, *b -tagging in dense environments*, [ATL-PHYS-PUB-2014-014](#) (2014).
- [37] ATLAS collaboration, *Expected performance of boosted Higgs ($\rightarrow b\bar{b}$) boson identification with the ATLAS detector at $\sqrt{s} = 13$ TeV*, [ATL-PHYS-PUB-2015-035](#) (2015).
- [38] CMS collaboration, *Identification of b quark jets at the CMS experiment in the LHC Run 2*, [CMS-PAS-BTV-15-001](#) (2015).
- [39] CMS collaboration, *Identification of heavy-flavour jets with the CMS detector in pp collisions at 13 TeV*, [2018 JINST 13 P05011](#) [[arXiv:1712.07158](#)] [[INSPIRE](#)].

- [40] ATLAS collaboration, *Calibration of ATLAS b-tagging algorithms in dense jet environments*, [ATLAS-CONF-2016-001](#) (2016).
- [41] ATLAS collaboration, *Studies of b-tagging performance and jet substructure in a high p_T $g \rightarrow b\bar{b}$ rich sample of large-R jets from pp collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector*, [ATLAS-CONF-2016-002](#) (2016).
- [42] J. Thaler and K. Van Tilburg, *Maximizing boosted top identification by minimizing N-subjettiness*, *JHEP* **02** (2012) 093 [[arXiv:1108.2701](#)] [[INSPIRE](#)].
- [43] A.J. Larkoski, I. Moult and D. Neill, *Building a better boosted top tagger*, *Phys. Rev. D* **91** (2015) 034035 [[arXiv:1411.0665](#)] [[INSPIRE](#)].
- [44] A.J. Larkoski, I. Moult and D. Neill, *Analytic boosted boson discrimination*, *JHEP* **05** (2016) 117 [[arXiv:1507.03018](#)] [[INSPIRE](#)].
- [45] G.P. Salam, L. Schunk and G. Soyez, *Dichroic subjettiness ratios to distinguish colour flows in boosted boson tagging*, *JHEP* **03** (2017) 022 [[arXiv:1612.03917](#)] [[INSPIRE](#)].
- [46] P.T. Komiske, E.M. Metodiev and J. Thaler, *Energy flow polynomials: A complete linear basis for jet substructure*, *JHEP* **04** (2018) 013 [[arXiv:1712.07124](#)] [[INSPIRE](#)].
- [47] A.J. Larkoski, I. Moult and D. Neill, *Analytic boosted boson discrimination at the Large Hadron Collider*, [arXiv:1708.06760](#) [[INSPIRE](#)].
- [48] A.J. Larkoski, I. Moult and D. Neill, *Factorization and resummation for groomed multi-prong jet shapes*, *JHEP* **02** (2018) 144 [[arXiv:1710.00014](#)] [[INSPIRE](#)].
- [49] K. Datta and A.J. Larkoski, *Novel jet observables from machine learning*, *JHEP* **03** (2018) 086 [[arXiv:1710.01305](#)] [[INSPIRE](#)].
- [50] L. de Oliveira et al., *Jet-images — Deep learning edition*, *JHEP* **07** (2016) 069 [[arXiv:1511.05190](#)] [[INSPIRE](#)].
- [51] J. Gallicchio and M.D. Schwartz, *Seeing in color: jet superstructure*, *Phys. Rev. Lett.* **105** (2010) 022001 [[arXiv:1001.5027](#)] [[INSPIRE](#)].
- [52] S.H. Lim and M.M. Nojiri, *Spectral analysis of jet substructure with neural network: boosted Higgs case*, [arXiv:1807.03312](#) [[INSPIRE](#)].
- [53] CMS collaboration, *Performance of deep tagging algorithms for boosted double quark jet topology in proton-proton collisions at 13 TeV with the Phase-0 CMS detector*, [CMS-DP-2018-046](#) (2018).
- [54] SLD collaboration, T. Abe et al., *A Preliminary measurement of the gluon splitting rate into $b\bar{b}$ pairs in hadronic Z^0 decays*, in the proceedings of the American Physical Society (APS) Meeting of the Division of Particles and Fields (DPF 99), January 5–9, Los Angeles, U.S.A. (1999), [hep-ex/9905057](#) [[INSPIRE](#)].
- [55] CMS collaboration, *Identification of double-b quark jets in boosted event topologies*, [CMS-PAS-BTV-15-002](#) (2015).
- [56] J. Cogan, M. Kagan, E. Strauss and A. Schwartzman, *Jet-images: computer vision inspired techniques for jet tagging*, *JHEP* **02** (2015) 118 [[arXiv:1407.5675](#)] [[INSPIRE](#)].
- [57] L.G. Almeida et al., *Playing tag with ANN: boosted top identification with pattern recognition*, *JHEP* **07** (2015) 086 [[arXiv:1501.05968](#)] [[INSPIRE](#)].
- [58] P. Baldi et al., *Jet substructure classification in high-energy physics with deep neural networks*, *Phys. Rev. D* **93** (2016) 094034 [[arXiv:1603.09349](#)] [[INSPIRE](#)].

- [59] J. Barnard, E.N. Dawe, M.J. Dolan and N. Rajcic, *Parton shower uncertainties in jet substructure analyses with deep neural networks*, *Phys. Rev. D* **95** (2017) 014018 [[arXiv:1609.00607](#)] [[INSPIRE](#)].
- [60] P.T. Komiske, E.M. Metodiev and M.D. Schwartz, *Deep learning in color: towards automated quark/gluon jet discrimination*, *JHEP* **01** (2017) 110 [[arXiv:1612.01551](#)] [[INSPIRE](#)].
- [61] L. de Oliveira, M. Paganini and B. Nachman, *Learning particle physics by example: location-aware generative adversarial networks for physics synthesis*, *Comput. Softw. Big Sci.* **1** (2017) 4 [[arXiv:1701.05927](#)] [[INSPIRE](#)].
- [62] G. Kasieczka, T. Plehn, M. Russell and T. Schell, *Deep-learning top taggers or the end of QCD?*, *JHEP* **05** (2017) 006 [[arXiv:1701.08784](#)] [[INSPIRE](#)].
- [63] G. Louppe, K. Cho, C. Becot and K. Cranmer, *QCD-aware recursive neural networks for jet physics*, [arXiv:1702.00748](#) [[INSPIRE](#)].
- [64] J. Pearkes, W. Fedorko, A. Lister and C. Gay, *Jet constituents for deep neural network based top quark tagging*, [arXiv:1704.02124](#) [[INSPIRE](#)].
- [65] K. Datta and A. Larkoski, *How much information is in a jet?*, *JHEP* **06** (2017) 073 [[arXiv:1704.08249](#)] [[INSPIRE](#)].
- [66] P.T. Komiske, E.M. Metodiev, B. Nachman and M.D. Schwartz, *Pileup Mitigation with Machine Learning (PUMML)*, *JHEP* **12** (2017) 051 [[arXiv:1707.08600](#)] [[INSPIRE](#)].
- [67] A. Butter, G. Kasieczka, T. Plehn and M. Russell, *Deep-learned top tagging with a Lorentz layer*, *SciPost Phys.* **5** (2018) 028 [[arXiv:1707.08966](#)] [[INSPIRE](#)].
- [68] E.M. Metodiev, B. Nachman and J. Thaler, *Classification without labels: learning from mixed samples in high energy physics*, *JHEP* **10** (2017) 174 [[arXiv:1708.02949](#)] [[INSPIRE](#)].
- [69] ALICE collaboration, R. Haake, *Machine and deep learning techniques in heavy-ion collisions with ALICE*, in the proceedings of the *2017 European Physical Society Conference on High Energy Physics (EPS-HEP 2017)*, July 5–12, Venice, Italy (2017), [arXiv:1709.08497](#) [[INSPIRE](#)].
- [70] T. Cheng, *Recursive neural networks in quark/gluon tagging*, *Comput. Softw. Big Sci.* **2** (2018) 3 [[arXiv:1711.02633](#)] [[INSPIRE](#)].
- [71] S. Egan et al., *Long Short-Term Memory (LSTM) networks with jet constituents for boosted top tagging at the LHC*, [arXiv:1711.09059](#) [[INSPIRE](#)].
- [72] P.T. Komiske, E.M. Metodiev, B. Nachman and M.D. Schwartz, *Learning to classify from impure samples with high-dimensional data*, *Phys. Rev. D* **98** (2018) 011502 [[arXiv:1801.10158](#)] [[INSPIRE](#)].
- [73] S. Macaluso and D. Shih, *Pulling out all the tops with computer vision and deep learning*, [arXiv:1803.00107](#) [[INSPIRE](#)].
- [74] Y.-T. Chien and R. Kunnawalkam Elayavalli, *Probing heavy ion collisions using quark and gluon jet substructure*, [arXiv:1803.03589](#) [[INSPIRE](#)].
- [75] K. Fraser and M.D. Schwartz, *Jet charge and machine learning*, [arXiv:1803.08066](#) [[INSPIRE](#)].
- [76] A. Andreassen, I. Feige, C. Frye and M.D. Schwartz, *JUNIPR: a framework for unsupervised machine learning in particle physics*, [arXiv:1804.09720](#) [[INSPIRE](#)].
- [77] J.H. Collins, K. Howe and B. Nachman, *CWoLa hunting: extending the bump hunt with machine learning*, [arXiv:1805.02664](#) [[INSPIRE](#)].

- [78] S. Choi, S.J. Lee and M. Perelstein, *Infrared safety of a neural-net top tagging algorithm*, [arXiv:1806.01263](#) [INSPIRE].
- [79] J. Guo, J. Li, T. Li, F. Xu and W. Zhang, *Deep learning for the R-parity violating supersymmetry searches at the LHC*, [arXiv:1805.10730](#) [INSPIRE].
- [80] J.W. Monk, *Deep learning as a parton shower*, [arXiv:1807.03685](#) [INSPIRE].
- [81] D. Guest et al., *Jet flavor classification in high-energy physics with deep neural networks*, *Phys. Rev. D* **94** (2016) 112002 [[arXiv:1607.08633](#)] [INSPIRE].
- [82] W. Bhimji et al., *Deep neural networks for physics analysis on low-level whole-detector data at the LHC*, in the proceedings of the 18th *International Workshop on Advanced Computing and Analysis Techniques in Physics Research (ACAT 2017)*, August 21–25, Seattle, U.S.A. (2017), [arXiv:1711.03573](#) [INSPIRE].
- [83] ATLAS collaboration, *Identification of jets containing b-hadrons with recurrent neural networks at the ATLAS experiment*, [ATL-PHYS-PUB-2017-003](#) (2017).
- [84] ATLAS collaboration, *Optimisation and performance studies of the ATLAS b-tagging algorithms for the 2017-18 LHC run*, [ATL-PHYS-PUB-2017-013](#) (2017).
- [85] ATLAS collaboration, *Quark versus gluon jet tagging using jet images with the ATLAS detector*, [ATL-PHYS-PUB-2017-017](#) (2017).
- [86] CMS collaboration, *New developments for jet substructure reconstruction in CMS*, [CMS-DP-2017-027](#) (2017).
- [87] ATLAS collaboration, *Identification of hadronically-decaying W bosons and top quarks using high-level features as input to boosted decision trees and deep neural networks in ATLAS at $\sqrt{s} = 13$ TeV*, [ATL-PHYS-PUB-2017-004](#) (2017).
- [88] ATLAS collaboration, *Performance of top quark and W boson tagging in run 2 with ATLAS*, [ATLAS-CONF-2017-064](#) (2017).
- [89] J. Duarte et al., *Fast inference of deep neural networks in FPGAs for particle physics*, *2018 JINST* **13** P07027 [[arXiv:1804.06913](#)] [INSPIRE].
- [90] A. Lenz, M. Spannowsky and G. Tetlalmatzi-Xolocotzi, *Double-charming Higgs boson identification using machine-learning assisted jet shapes*, *Phys. Rev. D* **97** (2018) 016001 [[arXiv:1708.03517](#)] [INSPIRE].
- [91] P. Baldi, P. Sadowski and D. Whiteson, *Searching for exotic particles in high-energy physics with deep learning*, *Nature Commun.* **5** (2014) 4308 [[arXiv:1402.4735](#)] [INSPIRE].
- [92] P. Baldi, P. Sadowski and D. Whiteson, *Enhanced Higgs boson to $\tau^+\tau^-$ search with deep learning*, *Phys. Rev. Lett.* **114** (2015) 111801 [[arXiv:1410.3469](#)] [INSPIRE].
- [93] M. Abdughani, J. Ren, L. Wu and J.M. Yang, *Probing stop with graph neural network at the LHC*, [arXiv:1807.09088](#) [INSPIRE].
- [94] M.D. Zeiler, *ADADELTA: an adaptive learning rate method*, [arXiv:1212.5701](#).
- [95] J. Duchi, E. Hazan and Y. Singer, *Adaptive subgradient methods for online learning and stochastic optimization*, *J. Mach. Learn. Res.* **12** (2011) 2121.
- [96] D.P. Kingma and J. Ba, *Adam: a method for stochastic optimization*, [arXiv:1412.6980](#) [INSPIRE].
- [97] N. Srivastava et al., *Dropout: a simple way to prevent neural networks from overfitting*, *J. Machine Learn. Res.* **15** (2014) 1929.
- [98] F. Chollet, *Keras*, <https://github.com/fchollet/keras>, (2017).

- [99] M. Abadi et al., *Tensorflow: a system for large-scale machine learning*, *OSDI* **16** (2016) 265 [[arXiv:1605.08695](#)].
- [100] J. Alwall 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** (2014) 079 [[arXiv:1405.0301](#)] [[INSPIRE](#)].
- [101] T. Sjöstrand, S. Mrenna and P.Z. Skands, *A brief introduction to PYTHIA 8.1*, *Comput. Phys. Commun.* **178** (2008) 852 [[arXiv:0710.3820](#)] [[INSPIRE](#)].
- [102] M. Bahr et al., *HERWIG++ physics and manual*, *Eur. Phys. J. C* **58** (2008) 639 [[arXiv:0803.0883](#)] [[INSPIRE](#)].
- [103] T. Gleisberg et al., *Event generation with SHERPA 1.1*, *JHEP* **02** (2009) 007 [[arXiv:0811.4622](#)] [[INSPIRE](#)].
- [104] M.L. Mangano, M. Moretti, F. Piccinini and M. Treccani, *Matching matrix elements and shower evolution for top-quark production in hadronic collisions*, *JHEP* **01** (2007) 013 [[hep-ph/0611129](#)] [[INSPIRE](#)].
- [105] X. Chen et al., *NNLO QCD corrections to Higgs boson production at large transverse momentum*, *JHEP* **10** (2016) 066 [[arXiv:1607.08817](#)] [[INSPIRE](#)].
- [106] T. Neumann and C. Williams, *The Higgs boson at high p_T* , *Phys. Rev. D* **95** (2017) 014004 [[arXiv:1609.00367](#)] [[INSPIRE](#)].
- [107] R. Bonciani et al., *Two-loop planar master integrals for Higgs \rightarrow 3 partons with full heavy-quark mass dependence*, *JHEP* **12** (2016) 096 [[arXiv:1609.06685](#)] [[INSPIRE](#)].
- [108] S.P. Jones, M. Kerner and G. Luisoni, *Next-to-leading-order QCD corrections to Higgs boson plus jet production with full top-quark mass dependence*, *Phys. Rev. Lett.* **120** (2018) 162001 [[arXiv:1802.00349](#)] [[INSPIRE](#)].
- [109] J.M. Lindert, K. Kudashkin, K. Melnikov and C. Wever, *Higgs bosons with large transverse momentum at the LHC*, *Phys. Lett. B* **782** (2018) 210 [[arXiv:1801.08226](#)] [[INSPIRE](#)].
- [110] K. Kudashkin, K. Melnikov and C. Wever, *Two-loop amplitudes for processes $gg \rightarrow Hg, qg \rightarrow Hq$ and $q\bar{q} \rightarrow Hg$ at large Higgs transverse momentum*, *JHEP* **02** (2018) 135 [[arXiv:1712.06549](#)] [[INSPIRE](#)].
- [111] T. Neumann, *NLO Higgs+jet at large transverse momenta including top quark mass effects*, *J. Phys. Comm.* **2** (2018) 095017 [[arXiv:1802.02981](#)] [[INSPIRE](#)].
- [112] M. Cacciari, G.P. Salam and G. Soyez, *FastJet user manual*, *Eur. Phys. J. C* **72** (2012) 1896 [[arXiv:1111.6097](#)] [[INSPIRE](#)].
- [113] M. Cacciari, G.P. Salam and G. Soyez, *The anti- k_t jet clustering algorithm*, *JHEP* **04** (2008) 063 [[arXiv:0802.1189](#)] [[INSPIRE](#)].
- [114] ATLAS collaboration, *Measurements of b-jet tagging efficiency with the ATLAS detector using $t\bar{t}$ events at $\sqrt{s} = 13$ TeV*, *JHEP* **08** (2018) 089 [[arXiv:1805.01845](#)] [[INSPIRE](#)].
- [115] M. Cacciari, G.P. Salam and G. Soyez, *The catchment area of jets*, *JHEP* **04** (2008) 005 [[arXiv:0802.1188](#)] [[INSPIRE](#)].
- [116] J. Dolen et al., *Thinking outside the ROCs: Designing Decorrelated Taggers (DDT) for jet substructure*, *JHEP* **05** (2016) 156 [[arXiv:1603.00027](#)] [[INSPIRE](#)].
- [117] I. Moult, B. Nachman and D. Neill, *Convolved substructure: analytically decorrelating jet substructure observables*, *JHEP* **05** (2018) 002 [[arXiv:1710.06859](#)] [[INSPIRE](#)].
- [118] J.A. Aguilar-Saavedra, J.H. Collins and R.K. Mishra, *A generic anti-QCD jet tagger*, *JHEP* **11** (2017) 163 [[arXiv:1709.01087](#)] [[INSPIRE](#)].

- [119] C. Shimmin et al., *Decorrelated jet substructure tagging using adversarial neural networks*, *Phys. Rev. D* **96** (2017) 074034 [[arXiv:1703.03507](#)] [[INSPIRE](#)].
- [120] ATLAS collaboration, *Reconstruction and modelling of jet pull with the ATLAS detector*, *ATLAS-CONF-2014-048* (2014).
- [121] 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** (2015) 475 [[arXiv:1506.05629](#)] [[INSPIRE](#)].
- [122] ATLAS collaboration, *Measurement of colour flow using jet-pull observables in $t\bar{t}$ events with the ATLAS experiment at $\sqrt{s} = 13$ TeV*, [arXiv:1805.02935](#) [[INSPIRE](#)].
- [123] D. Krohn, L. Randall and L.-T. Wang, *On the feasibility and utility of ISR tagging*, [arXiv:1101.0810](#) [[INSPIRE](#)].
- [124] M.A. Ebert et al., *Exploiting jet binning to identify the initial state of high-mass resonances*, *Phys. Rev. D* **94** (2016) 051901 [[arXiv:1605.06114](#)] [[INSPIRE](#)].
- [125] L.A. Harland-Lang, V.A. Khoze, M.G. Ryskin and M. Spannowsky, *Jet activity as a probe of high-mass resonance production*, *Eur. Phys. J. C* **76** (2016) 623 [[arXiv:1606.04902](#)] [[INSPIRE](#)].
- [126] J.R. Ellis, M.K. Gaillard and D.V. Nanopoulos, *A phenomenological profile of the Higgs boson*, *Nucl. Phys. B* **106** (1976) 292 [[INSPIRE](#)].
- [127] M.A. Shifman, A.I. Vainshtein, M.B. Voloshin and V.I. Zakharov, *Low-energy theorems for Higgs boson couplings to photons*, *Sov. J. Nucl. Phys.* **30** (1979) 711 [[INSPIRE](#)].
- [128] I. Low, R. Rattazzi and A. Vichi, *Theoretical constraints on the Higgs effective couplings*, *JHEP* **04** (2010) 126 [[arXiv:0907.5413](#)] [[INSPIRE](#)].
- [129] ATLAS collaboration, *Measurement of the Higgs boson coupling properties in the $H \rightarrow ZZ^* \rightarrow 4\ell$ decay channel at $\sqrt{s} = 13$ TeV with the ATLAS detector*, *JHEP* **03** (2018) 095 [[arXiv:1712.02304](#)] [[INSPIRE](#)].
- [130] J. Ellis, C.W. Murphy, V. Sanz and T. You, *Updated global SMEFT fit to Higgs, diboson and electroweak data*, *JHEP* **06** (2018) 146 [[arXiv:1803.03252](#)] [[INSPIRE](#)].