

Higgs boson tagging with the Lund jet plane

Charanjit K. Khosa^{*} and Simone Marzani[†]

*Dipartimento di Fisica, Università di Genova and INFN, Sezione di Genova,
Via Dodecaneso 33, 16146, Italy*

 (Received 14 May 2021; accepted 31 August 2021; published 28 September 2021)

We construct a procedure to separate boosted Higgs bosons decaying into hadrons, from the background due to strong interactions. We employ the Lund jet plane to obtain a theoretically well-motivated representation of the jets of interest and we use the resulting images as the input to a convolutional neural network classifier. In particular, we consider two different decay modes of the Higgs boson, namely into a pair of bottom quarks or into light jets, against the respective backgrounds. For each case, we consider both a moderate- and high-boost scenario. The performance of the tagger is compared to what is achieved using a traditional single-variable analysis which exploits a QCD inspired color-singlet tagger, namely the jet color ring observable.

DOI: [10.1103/PhysRevD.104.055043](https://doi.org/10.1103/PhysRevD.104.055043)

I. INTRODUCTION

High-energy collision events at the CERN Large Hadron Collider (LHC) are characterized by copious hadronic activity. Not only protons are strongly-interacting, but also elementary particles produced in their collisions often carry color charges, resulting into final-states characterized by many hadrons. One powerful way to deal with the complex environment of hadronic activity at the LHC is to note that final-state hadrons tend to be produced in fairly collimated sprays, along directions that we can think of as being the ones of the originating hard (i.e., with large transverse momentum) particles. These collimated sprays of hadrons are called jets.

The vast majority of high transverse-momentum (p_T) jets are QCD jets, i.e., they originate from the fragmentation of a high- p_T parton (a quark or a gluon). However, every particle that decays hadronically can, if sufficiently boosted, give rise to a collimated spray of hadrons, which is reconstructed as a jet. Therefore, one key aspect in the context of LHC phenomenology is to correctly identify the nature of the particles originating these jets. This endeavor is often referred to as jet tagging. For example, top quarks, electroweak gauge bosons, and the Higgs boson, if produced at high- p_T , can be identified using jet tagging techniques. These algorithms can also be applied in

searches for new physics. Therefore, every tiny gain in jet tagging efficiency is important for both for measurements and new physics searches at the LHC. Tremendous progress over the past decade (see, e.g., [1]) has led to the development of jet substructure techniques, including tagging algorithms, that are very efficient in distinguishing signal jets from background ones, where the latter ones are often QCD jets. Furthermore, the application of field-theory methods to jet physics has provided us with a deeper understanding of jet substructure. This, in turn, has allowed us to develop algorithms that are not only efficient, but robust.

In the last few years a complementary approach to jet substructure has risen to prominence. The rapid development of machine-learning (ML) techniques is deeply changing the way particle physics analyses are conducted. In recent years, many groups have been exploring the potential of ML in the context of LHC phenomenology. There have been several suggestions, ranging from top tagging [2], constraining Wilson coefficients of higher dimensional operators in effective field theory (EFT) framework [3–6], quark-gluon tagging [7], model agnostic new physics searches [8], jet substructure [9], likelihood encoding [10] and many others. In addition, there is a continuous effort to adapt these techniques for the various steps of data analysis, i.e., trigger, event reconstruction, particle identification, heavy flavor tagging, jet tagging, and signal and background classification (for recent reviews see, e.g., [11,12]).

In the context of jet substructure, classification algorithms that exploit deep neural networks (NN) have been shown to often out-perform the more traditional tagging techniques, which indeed can be thought of as lower-dimensional projections of the complex sets of inputs that are fed to the NN (see, e.g., [13]).

^{*}charanjit.kaur@ge.infn.it

[†]simone.marzani@ge.infn.it

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI. Funded by SCOAP³.

One intriguing aspect of ML techniques is their ability to make data-driven decisions without using the prior knowledge of the underlying theory. This has sparked an interesting debate on whether one should take a rather agnostic approach and favor raw data, such as particles' kinematics, as inputs to NN, or whether one should make good use of the expert-knowledge developed thanks to our theoretical understanding of the underlying physical processes, and therefore exploit higher-level, theory-inspired, objects as inputs to the NN.¹ In this context, particle physics in general, and jet physics in particular, find themselves in a rather unique position to address these types of questions, because, thanks to the Standard Model, we have a deep understanding of the physical processes we are studying. For instance, we can make use of this knowledge to better understand what kind of information ML classifiers are exploiting [9,14]. Therefore a growing number of studies that combine the power of ML classification with theory-inspired variables has appeared in recent times. These include the use of N -subjettiness variables [15–17], energy flow variables [18,19], two-point energy correlations [20], and jet charge [21] for multiprong jet tagging.

A very interesting observable, both in the context of boosted object tagging as well as Standard Model measurements is the Lund jet plane [22]. Similarly to calorimeter jet images [23,24], the (primary) Lund jet plane is a two-dimensional representation of a jet. However, differently to the calorimetric approach, which constructs images in the pseudorapidity-azimuth (η, ϕ) plane, the Lund jet plane image is given in terms of kinematic variables that better describe the emissions in the jet. By construction, the Lund jet plane allows one to separate different physical effects. Indeed, the Lund plane was originally designed to describe the way phase-space is filled by Monte Carlo parton-showers. It is also often used in the context of analytic resummation. Recently, the all-order structure of the Lund jet plane density has also been computed [25]. The Lund jet plane has been used together with long short-term memory (LSTM) networks and, more recently, with graph neural networks for W boson and top tagging [22,26]. It is also a plausible choice for event generation using generative models [27]. Lund basis for the jet declustering history is also used in the context of unsupervised new physics search method [28]. From the experimental side, the Lund jet plane density has been measured by the ATLAS [29] and ALICE [30] collaborations.

In the following, we will study the performance of the Lund jet plane in the context of boosted Higgs boson identification. We shall consider both the decay of the Higgs boson in a pair bottom (b) quarks that subsequently fragment into two b -(sub)jets, henceforth $H \rightarrow b\bar{b}$, and the

decay of the Higgs boson into light (sub)jets, henceforth $H \rightarrow gg$.

The Lund jet plane offers us the opportunity to address interesting theoretical questions and simultaneously provides good classification performance. In particular, we are going to use the primary Lund plane, which is a proxy for the two-dimensional phase space of the leading emission in the jet. Resulting Lund jet plane images for signal and background are used as inputs to a convolutional neural network (CNN). CNNs are known to be very efficient for image datasets. However, they have been mostly tested in cases where images are constructed, unlike our case, from raw data or, in other words, unprocessed dataset. Within the high-energy physics community, CNNs are showing exciting potential for instance, in the context of top tagging [31,32], dark matter searches (see, e.g., [33,34]), disentangling Higgs production modes [35], and anomaly detection [36].

We believe that this study is an interesting addition to the rather extensive literature on ML-based Higgs taggers. Some of these methods exploit low-level inputs to construct jet images and use CNN [37–40] or interaction networks [41]. Other approaches combine the use of NN with N -subjettiness and related variables [42,43]. The potential of ML techniques has also been explored for a better reconstruction of the Higgs (decaying to $b\bar{b}$) for the trilinear coupling measurements [44,45]. Other ML based Higgs studies involve, for instance, disentangling Higgs production modes [35], invisible Higgs decay [46], Higgs width measurement [47], Higgs cascade decay [48], bottom Yukawa couplings [49].

The primary aim of this work is to explore the Lund jet plane tagging performance for the boosted Higgs boson decays $H \rightarrow b\bar{b}$ and $H \rightarrow gg$. The main backgrounds to these processes are given, respectively, by high transverse-momentum jets that are doubly b -tagged, mostly driven by $g \rightarrow b\bar{b}$ splittings, or by generic light jet production (jj). Therefore, we can view Higgs tagging as an example of a more general problem that aims to identify color-singlet states that are decaying hadronically from states that belong to other representations of the $SU(3)$ color group. Recently, the jet color ring [50] observable was proposed for color-singlet identification. This observable was derived using ratios of signal and background matrix elements, in the soft limit. By construction, the jet color ring is monotonic with respect to the likelihood ratio and therefore, it represents an optimal tagger, given the approximations made in its derivation. Indeed, it was found [50] that this simple observable offers good tagging efficiency for $H \rightarrow b\bar{b}$. However, it fails miserably when considering the $H \rightarrow gg$ case. In this study, we compare the tagging efficiency obtained with the Lund jet plane to the jet color ring, with the aim of identifying better ways to harvest the information on color radiation, in order to build efficient taggers.

This paper is organized as follows. In Sec. II, we describe the main building blocks of this work including the analysis

¹We note that a comparison between different approaches has been recently performed in the context of top-tagging [2].

set-up. Sections III and IV are devoted to the $H \rightarrow b\bar{b}$ and the $H \rightarrow gg$ analyses, respectively. We perform our studies in two different kinematic region: moderate-boost ($p_T > 250$ GeV) and high-boost ($p_T > 550$ GeV). We summarise our results in the last section.

II. ANALYSIS STRATEGY

Before discussing the details of our analyses, we describe the two main theoretical building blocks of the work, i.e., the Lund jet plane and the jet color ring.

A. Lund jet plane

The idea of constructing a Lund plane for an individual jet was proposed a couple of years ago in Ref. [22]. That work sparked new interest in the use of the Lund plane in jet physics, beyond its traditional application for the description of the emissions' phase space, primarily in the context of parton showers and resummation.

The Lund jet plane is a physically motivated, QCD-driven, representation of a jet. It is formed parsing backwards the Cambridge-Aachen (C/A) [51,52] 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. The most useful representation of the recorded information is given in terms of a double-logarithmic plane. Following [22], we choose as variables the azimuth-rapidity separation of the branches involved in the splitting and the transverse momentum of the *emission* with respect to the *emitter*.² Although other representations are possible, the $(\ln \frac{1}{\Delta}, \ln \frac{k_t}{\text{GeV}})$ plane has the advantage of a clear separation between different physical effects, such as collinear ($\Delta \simeq 0$) from large-angle ($\Delta \simeq 1$) emissions, and perturbative ($k_t \gg 1$ GeV) from nonperturbative ($k_t \lesssim 1$ GeV).

B. Jet color ring

The jet color ring (\mathcal{O}) [50] is a color-singlet tagger for boosted two-prong decays. It is defined as

$$\mathcal{O} = \frac{\Delta_{ka}^2 + \Delta_{kb}^2}{\Delta_{ab}^2} \quad (1)$$

where a and b are primary subjects (i.e., the leading subjects or the subjects that have been tagged according to the decay's properties, e.g., b -tagged), while k is leading remaining subject. This third subject is taken as a proxy for soft-gluon emission in the jet. As before, Δ measures

²Given the splitting $p \rightarrow p_a p_b$, with $p_{Ta} > p_{Tb}$, we consider p_b to be the emission and p the emitter, while $\Delta = \sqrt{(\phi_a - \phi_b)^2 + (y_a - y_b)^2}$.

the separation, in the rapidity-azimuth plane, between the subject pairs. Color conservation dictates that a and b are color-connected if the decaying state is a color-singlet. In such case, k will be predominantly emitted in between the legs of the ab dipole and, as a result, the distribution of the color ring will be peaked at small \mathcal{O} .

Both the Lund jet plane and the jet color ring receive their inspiration from a first-principle analysis of the physical process we are interested in. However, it is clear that the Lund jet plane provides us with more information than the jet color ring. The Lund jet plane is a two-dimensional representation of the jet of interest, in which each splitting is mapped into a point of the $(\ln \frac{1}{\Delta}, \ln \frac{k_t}{\text{GeV}})$ plane. On the other hand, the color ring is a more standard, and much simpler, jet observable, which associates a jet to a single value of \mathcal{O} . In the following, we will compare the tagging performance of the Lund jet plane, used as the input to a CNN, and of a much simpler tagging strategy based on just the jet color ring. On the one hand, we are interested in exploring possible gain brought by the Lund jet plane and the use of ML in the $H \rightarrow b\bar{b}$ case, where we know we can already achieve good performance using the jet color ring. On the other hand, we also want to study the $H \rightarrow gg$ case, where the color ring offers essentially no discrimination power.

C. Simulation setup

We now discuss the methodology of our work. First, we discuss the event generation setup, then the analysis cuts and the steps to construct our observables of interest. Finally, the generic features of the CNN architecture are described.

1. Event generation

We use MadGraph 2.7.2 [53] to generate the events for signal and background processes at $\sqrt{s} = 14$ TeV for both the $H \rightarrow b\bar{b}$ and $H \rightarrow gg$ analyses. The considered signal process is $pp \rightarrow ZH$ where $Z \rightarrow \mu^+ \mu^-$ and $H \rightarrow b\bar{b}$ or gg . Background processes are Zbb and Zjj for the $H \rightarrow b\bar{b}$ and $H \rightarrow gg$ analyses, respectively. A generation-level $p_T^{\mu\mu}$ cut of 200 and 500 GeV is imposed for the moderate- and high-boost scenarios, respectively. Further cuts on pseudo-rapidity of leptons ($|\eta_l| < 2.5$) and jets ($|\eta_j| < 5.0$) are imposed. PYTHIA 8 [54,55] is used to simulate the parton-shower and the hadronization process. We consider the particles with $|\eta| < 5$ to form jets.

We cluster jets using the anti- k_r algorithm [56] with $R = 1.0$, using its implementation in FASTJET 3.3.3 [57]. The large jet radius should ensure that the decay products of Higgs are reconstructed in a single jet, in both boosted scenarios. Hard muons from the Z boson decay are excluded from the clustering. We further check jet-lepton separation. The leading jet of $p_T > 250$ GeV ($p_T > 550$ GeV for the high-boost benchmark) is considered for further analysis.

Following standard practice, jets with an invariant mass close to the Higgs mass are considered. In particular, we keep signal and background jets with $(110 < m_J < 140 \text{ GeV})$. The Lund jet plane and the jet color ring are then measured on the leading jet, as will be detailed below.

We note that the analysis efficiency is different for the signal and background processes as well as for the different p_T values considered here. Consequently, the total number of events generated differs in all the cases and a large enough samples are generated in order to obtain at least 100 K events for each case.

2. Constructing the jet color ring

In order to construct the jet color ring, we further need to identify subjets within the leading jet.

Following Ref. [50], we consider charged particles with $p_T > 500 \text{ MeV}$ and $|\eta| < 5$, and construct $R = 0.2$ track jets using the anti- k_r algorithm. Track-jets with $p_T > 5 \text{ GeV}$ and $\Delta < 0.8$ with respect to the leading jet are considered as inputs for the jet color ring. For $H \rightarrow b\bar{b}$ analysis, these track-jets are further identified as b -jets or light-jets. Here, we employ a very crude approximation for the b -tagging procedure. We use the truth information of b -partons and calculate the Δ separation of track jets with b -partons. If $\Delta_{jb} < 0.2$ or $\Delta_{j\bar{b}} < 0.2$ then we b -tag them.

3. Mapping events to the Lund jet plane

We construct the primary Lund plane of the leading jet for the same events that pass the selection cuts described for the jet color ring. First, we recluster the leading jet using the C/A algorithm with the maximum allowed jet radius. The Lund generator module of the FASTJET CONTRIB code is used to get the Lund coordinates of the declustering of C/A jet. The primary plane images are created using $\ln(1/\Delta)$ and $\ln \frac{k_r}{\text{GeV}}$ of the declustering history of the hardest branch, as described above. In particular, in order to construct the image, we choose 25 by 25 pixels for both coordinates. Only the pixels corresponding to the declustering history are turned on. This way, image pixels have mainly two values either 1 or 0 and we do not need to normalize the data before using it for the CNN.

4. CNN architecture

We employ convolutional NN for the signal-background classification, using the Lund plane images dataset. KERAS [58] package is used for the CNN implementation. The CNN architecture is optimized for each benchmark. We use a balanced dataset of 200 K events for all the benchmarks. Each dataset is divided into 60:20:20 proportions for training, validation and test set. In this section, we mention the generic setup and details of the architecture which are common for all the benchmarks. A cartoon of the architecture is shown in Fig. 1. We find that CNN with 4 convolutional layers of filter size 3 is the best choice for all




FIG. 1. This is a schematic diagram of the CNN architecture used. Here $N_i, i = 1, \dots, 4$ are the number of filters in each layer, N_F is the filter size and N_n is the number of neurons used in the flat layer.

cases. The number of filters for the convolutional layers is different in each case. After the convolutional layers, we have a flat layer with 800 neurons in all the cases. After the second and fourth convolutional layer, a pooling layer is used. For the robust training of the CNN, the down sampling of the feature maps is achieved by the pooling layers. For the pooling layer we used the Max Pooling function. Another hyper parameter in the training is dropouts.³ We also turned on the dropout option with different strengths after the second, third and fourth convolutional layers. Further nonlinearity is introduced in the model by using activation function “relu.” We use cross-entropy loss function. CNN training model parameters are updated using Adam optimizer [59] to minimize the loss function. Since for each dataset CNN architecture

³It is a type of regularization to avoid overfitting by dropping some connections in the network.




FIG. 2. The first and second columns show primary Lund jet plane images (averaged over 100 K events) for $H \rightarrow b\bar{b}$ and $g \rightarrow b\bar{b}$ events, with jet $p_T > 250$ GeV (upper panel) and $p_T > 550$ GeV (lower panel). The image resolution is 25×25 for these images. The third column shows the normalized jet color ring distributions for the $H \rightarrow b\bar{b}$ and $g \rightarrow b\bar{b}$ events with jet $p_T > 250$ GeV and $p_T > 550$ GeV. In all cases, only jets with invariant mass $m_J \in [110, 140]$ GeV are considered.

is optimized separately, further details about batch sizes and epochs and other information regarding the CNN will be mentioned later, in the respective sections.

III. ANALYSIS $H \rightarrow b\bar{b}$

In this section we discuss the $H \rightarrow b\bar{b}$ analysis. As previously mentioned, we are going to consider two scenarios: moderate-boost and high-boost. We will refer to the lower transverse momentum case as *benchmark point 1* (BP1), while the high-transverse momentum one will be BP2.

A. Moderate-boost scenario

First, we analyze the moderate-boost scenario, where we require the transverse-momentum of the selected jet to be $p_T > 250$ GeV. As discussed in Sec. II, we construct the primary Lund jet plane and the color ring observable, event by event. Averaged primary Lund jet plane images are then obtained considering 100K events. The resulting images

are shown in the first row of Fig. 2, for signal jets (on the left) and background jets (in the center). The dominant differences between the Higgs image and the background one is clearly visible for large Δ and high k_t ($\ln \frac{k_t}{\text{GeV}} \approx 4.5$). Jet color ring distributions are instead shown in the third column of Fig. 2. Both signal and background distributions are normalized to unity. We note that signal events mainly populate the $\mathcal{O} < 1$ region, while the background distribution is flatter, as expected.⁴

We use CNN for the Lund jet images dataset to perform the binary signal-background classification. The optimized CNN architecture has 4 convolutional layers, with filter size 3 and one flat layer with 800 neurons. The number of filters used is 16, for the first two convolutional layers, and 32 for the third and fourth layer. We use a batch size of 1000 and 15 epochs, i.e., the number of times the total dataset is

⁴Up to small differences in the simulation setup, our color ring distributions are in good agreement with the ones obtained in Ref. [50].




FIG. 3. ROC curves for CNN using the Lund jet plane images and for color ring for both the benchmarks of the $H \rightarrow b\bar{b}$ analysis considering jets with invariant mass $m_J \in [110, 140]$ GeV.

shown to the network, for CNN training. We did not need to train for a larger number of epochs because CNN trains faster than other deep learning methods. See Table I for more details of architecture (BP1). The receiver operating characteristic (ROC) curves for CNN classification and for the color ring are shown in Fig. 3, on the left. For the color ring case, we vary the threshold in small steps, considering the signal (background) below (above) the threshold, and calculate the signal efficiency and false positive rate for each value of the threshold to get the ROC curve.

A standard metric used to assess the classification performance is the area under the ROC curve (AUC). With our definition of ROC curves, optimal performance corresponds to $AUC = 0$. In the following we will use the metric $\mathcal{A} = 1 - AUC$, where now $\mathcal{A} = 1$ corresponds to optimal performance. We find that the performance of the CNN using the Lund jet plane dataset is 3% better than the single-variable approach using the color ring observable. Assuming that the optimization of the CNN has been performed appropriately, the rather small improvement can be explained by the already good performance of the jet color ring, which was originally designed as an optimal color-singlet tagger.

B. High-boost scenario

We consider the second, boosted, scenario, with the leading jet transverse momentum required to be $p_T > 550$ GeV. With the exception of the generation-level p_T^{jet} and leading-jet p_T cut in the analysis, all the set-up is the same as in the previous case. The overall analysis selection efficiency for the background events is higher in this case as compared to the previous case. The averaged (over 100 K events) primary Lund jet plane images for the signal and

background, with jet mass cut $110 < m_J < 140$ GeV, are shown in the second row of Fig. 2.

The most prominent difference in the signal image, with respect to the moderate-boost scenario, is the shift of the high- k_T patch toward smaller values of Δ , and hence larger values of $-\ln \Delta$. This happens because the decay products of the Higgs boson tend to be more collimated. The background image also noticeably changes with respect to the lower- p_T case. The color ring signal and background distributions are shown in the lower right column in Fig. 2.

The CNN architecture details are mentioned in the Table I under BP2 column. This architecture has a lesser number of filters in the third and fourth convolutional layer than compared to the moderate-boost scenario. In the right panel of Fig. 3, we show the ROC curves for the CNN predictions for the Lund images and for the color ring observable. In this case, the classification accuracy of the Lund plane + CNN combination is significantly better (7%) than the single-variable approach using color ring observable. The color ring \mathcal{A} is 0.02 more than the one of BP1, but the CNN classification accuracy is 6% better than the moderate boost case.

IV. ANALYSIS $H \rightarrow gg$

We now move to analyze the other decay channel of the Higgs boson considered in this study, namely the light-jet final-state. This is a very challenging decay channel and, as mentioned earlier, the jet color ring is known to perform poorly in this context.

We use the same tools to simulate the events for $H \rightarrow gg$ benchmark as in the previous case except the model used in MadGraph. This decay mode of Higgs is mediated by the heavy quark loop and $H \rightarrow gg$ effective coupling is




FIG. 4. The first and second columns show primary Lund jet plane images (averaged over 100K events) for $H \rightarrow gg$ and Zjj events, with jet $p_T > 250$ GeV (upper panel) and $p_T > 550$ GeV (lower panel). The image resolution is 25×25 for these images. The third column shows the normalized jet color ring distributions for the $H \rightarrow gg$ and Zjj events with jet $p_T > 250$ GeV and $p_T > 550$ GeV. In all cases, only jets with invariant mass $m_J \in [110, 140]$ GeV are considered.

implemented in HEFT model [60]. For consistency, we use the same model to simulate the background events, i.e., $pp \rightarrow Z(\mu^+\mu^-) + jj$. In this section, we follow the same ordering as $H \rightarrow b\bar{b}$ analysis, i.e., first $p_T > 250$ GeV (BP3), and followed by $p_T > 550$ GeV (BP4).

A. Moderate-boost scenario

Using the analysis setup of Sec. II, we form a primary Lund jet plane and color ring for signal and background events. The averaged Lund jet plane images and color ring distributions are shown in the first row of Fig. 4 for the $p_T > 250$ GeV benchmarks. We start by considering the jet color ring. Similarly, to the $H \rightarrow b\bar{b}$ analysis, the signal distribution is falling sharply, i.e., signal events mainly populate in $\mathcal{O} < 1$ region, in accordance to our expectation for color-singlets. However, as it was found in Ref. [50], the behavior of the background distribution is very different from the corresponding $H \rightarrow b\bar{b}$ case and it is in fact almost

overlapping with the signal distribution. In this case, several possible color configurations are contributing, while the previous $g \rightarrow b\bar{b}$ case was characterized by the octet configuration.

It is then interesting to check whether the Lund jet plane can instead highlight differences between signal and background in $H \rightarrow gg$. As we can see from the first two plots in Fig. 4 this is indeed the case. We can make a more quantitative statement about the tagger performance by looking at the ROC curves for the Lund jet plane+CNN and color ring. This is done in Fig. 5, on the left. Then CNN using Lund jet plane dataset provides very good classification performance, while the color ring ROC is very poor, close to the random classifier.

B. High-boost scenario

A very similar picture holds for the high-boost scenario, $p_T > 550$ GeV case. Corresponding results are reported in




FIG. 5. ROC curves for CNN using the Lund jet plane images and for color ring for both the benchmarks of the $H \rightarrow gg$ analysis considering jets with invariant mass $m_J \in [110, 140]$ GeV.

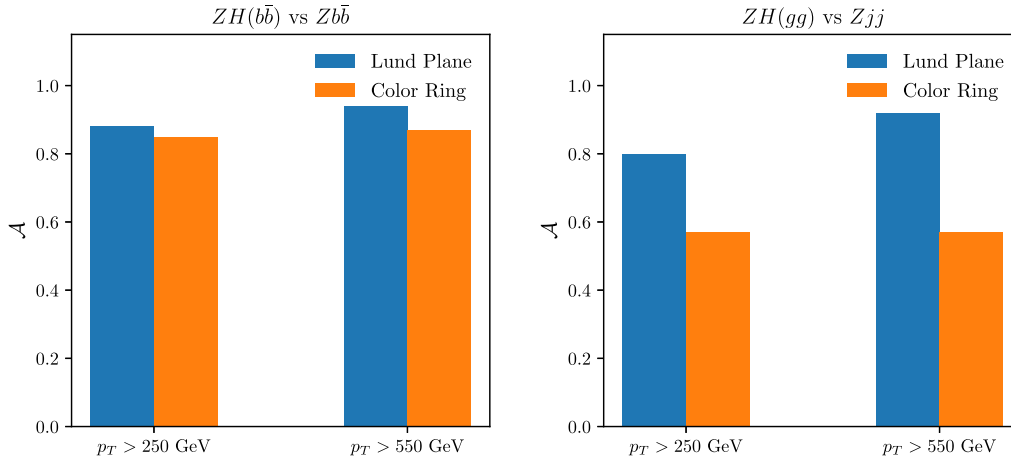


FIG. 6. Left: Higgs signal and background classification performance comparison using color ring and Lund plane + CNN combination for the $H \rightarrow b\bar{b}$ analysis. Right: same comparison for the $H \rightarrow gg$ analysis.

the second row of Fig. 4. In particular, the differences in the averaged Lund jet plane images are clearly visible by eye. For instance, as in the $H \rightarrow b\bar{b}$ case, there is a bright patch at high k_T . Lund jet plane and color ring ROC curves for this benchmark are shown in Fig. 5, on the right.

V. CONCLUSION AND OUTLOOK

In this work, we have studied the primary Lund jet plane in the context of tagging hadronically decaying Higgs bosons, in the boosted regime, where the Higgs' decay products are reconstructed into a single large-radius jet. In particular, we have considered two challenging, but crucial,

decay channels: the heavy flavor (bottom) one and the Higgs decay into light jets. We have concentrated on two transverse momentum benchmark scenarios, namely moderate and high boost of the Higgs boson.

Inspired by previous work on W and top tagging using the Lund jet plane [22,26], we have built images that are used as inputs to a convolutional neural network for classification. We have compared the performance of this tagger to a more standard approach: a single-variable analysis that exploits a theoretical-motivated observable, namely the jet color ring.

Our findings, for different transverse momentum scenarios, are summarized in Fig. 6, for the $H \rightarrow b\bar{b}$ and $H \rightarrow gg$

analysis, respectively. We have taken $\mathcal{A} = 1 - \text{AUC}$ as figure of merit to assess the taggers' performance. We can see that the Lund jet plane and CNN combination has the best separation power for all the benchmarks studied. Its performance is equally good for $H \rightarrow b\bar{b}$ and the $H \rightarrow gg$ analyses, thus providing us with some confidence about its robustness. This is in stark contrast with the jet color ring, which almost equals the CNN performance in the $H \rightarrow b\bar{b}$ case, while fails completely for the $H \rightarrow gg$ process. We note that the difference between the AUC for CNNs and color ring cases is higher for the high-boost case, in both analyses.

Despite the fact that raw information from the jet (or even particles') kinematics can serve as inputs to ML algorithms, the use of the Lund jet plane is intriguing from a theoretical perspective. It provides a physically-motivated picture of a jet that can be naturally used as input to a NN. At the same time, the Lund jet plane can be described with perturbative field theory. Indeed, the Lund plane density has been recently computed in QCD, for light jets [25]. It would be extremely interesting to perform analogous calculations for signal jets and b -jets, in order to shed light on the taggers' performance, as it was done, for instance, in the case of single-variable taggers almost a decade ago, e.g., [61–63].

Furthermore, one important feature of the Lund jet plane, especially when expressed in terms of the variable k_t , as in our case, is that it makes the separation between regions dominated by perturbative and nonperturbative physics rather clear. If we assume that the latter is characterized by energy scales corresponding to relative transverse momenta of the order of 1 GeV, then the boundary between the two regions is a straight, horizontal line at $\ln(k_t/\text{GeV}) = 0$. Even by eye, we note that, in all the cases considered in this study, the bulk of the difference between signal and background Lund jet plane images is found above this horizontal line. Thus, we expect that the CNN is mostly exploiting perturbative information. We have confirmed this intuition by performing an analysis where we input to the CNN only the upper section of the Lund jet plane, i.e., $\ln(k_t/\text{GeV}) > 0$. We have found the difference between the values of \mathcal{A} in the two cases to be below the percent level.⁵

Finally, in this work we used simulated data without incorporating detector effects. In a realistic situation, we expect some degradation in the reported results. It would be then important to study the resilience [64] of the color ring

and of the Lund jet plane against this type of contributions. On the other hand, one could improve the performance of the taggers by exploiting more information from the clustering history. This can be achieved, for instance, by going beyond the primary plane approximation. Furthermore, the use of other ML architectures, such as, e.g., graph neural networks, may lead to a further gain in the classification performance [26]. We plan to explore these directions in the future work.

ACKNOWLEDGMENTS

We thank Andrea Coccaro, Frédéric Dreyer, Andrew Larkoski, and Giovanni Stagnitto for comments on the manuscript. We also thank the members of the ATLAS groups in Genova and Pavia for many useful discussions on Higgs tagging and color flow. This work 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.

APPENDIX: CNN ARCHITECTURES

The main structure of the architecture (see Fig. 1) is the same for all cases. For each benchmark, we ran several models with different choices of hyper-parameters and presented the result for the best ones, which we call optimized models. As an example, model training curves for BP1 and BP2 are shown in Fig. 7. Since all benchmark datasets share some common features, the resulting optimized architectures are also quite similar. We optimise the architecture for each benchmark because our purpose is to see the maximum performance gain over the single variable case. Thus, even if in an actual experimental analysis, one decides to use one model for all the data (which means for all the benchmarks considered here), we do not expect significant reduction in the classification performance. In Table I, we include all the details about the architecture for each benchmark. The first two benchmarks are for the $H \rightarrow b\bar{b}$ analysis and last two are for the $H \rightarrow gg$ analysis. They are ordered by p_T of the leading jet. N_1, N_2, N_3 , and N_4 are the number of filters in the first, second, third and fourth convolutional layer, respectively. A filter size of 3×3 is used in all the cases. Further filter movement to analyze the image is controlled by the stride parameter, which we choose as one unit along both directions. In all the cases, we zero padded the images at third and fourth layer in such a way that the input and output image dimensionality remains the same.

⁵We thank Andrew Larkoski for suggesting this study.

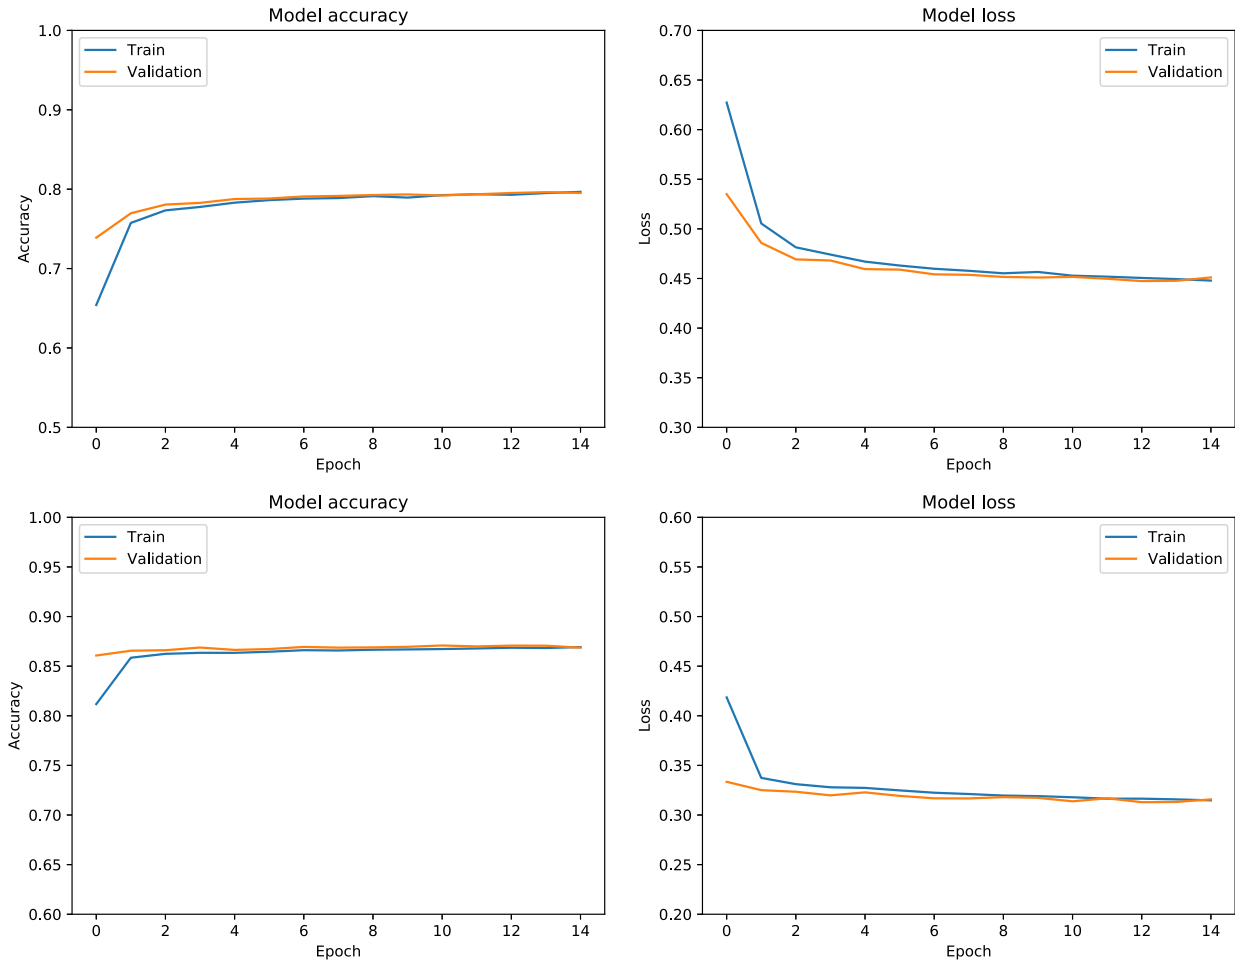


FIG. 7. CNN training curves for BP1, on the top, and for BP2, at the bottom.

TABLE I. CNN architectures used for different datasets.

	BP1	BP2	BP3	BP4
N_1 Conv2D	16	16	16	16
N_2 Conv2D	16	16	16	16
Dropouts	0.25	0.05	0.20	0.20
N_3 Conv2D	32	16	16	32
Dropouts
N_4 Conv2D	32	16	16	32
Dropouts	...	0.05	0.30	0.30
Flat layer	800	800	800	800
Epochs	15	15	20	20
Batch size	1000	1000	800	700

- [1] S. Marzani, G. Soyez, and M. Spannowsky, *Lect. Notes Phys.* **958**, 1 (2019).
- [2] G. Kasieczka, T. Plehn, A. Butter, K. Cranmer, D. Debnath, B. M. Dillon, M. Fairbairn, D. A. Faroughy, W. Fedorko, C. Gay *et al.*, *SciPost Phys.* **7**, 014 (2019).
- [3] F. F. Freitas, C. K. Khosa, and V. Sanz, *Phys. Rev. D* **100**, 035040 (2019).
- [4] J. Brehmer, K. Cranmer, G. Louppe, and J. Pavez, *Phys. Rev. Lett.* **121**, 111801 (2018).
- [5] J. Brehmer, K. Cranmer, G. Louppe, and J. Pavez, *Phys. Rev. D* **98**, 052004 (2018).
- [6] S. Chen, A. Glioti, G. Panico, and A. Wulzer, *J. High Energy Phys.* **05** (2021) 247.
- [7] G. Kasieczka, N. Kiefer, T. Plehn, and J. M. Thompson, *SciPost Phys.* **6**, 069 (2019).
- [8] G. Kasieczka, B. Nachman, D. Shih, O. Amram, A. Andreassen, K. Benkendorfer, B. Bortolato, G. Brooijmans, F. Canelli, J. H. Collins *et al.*, [arXiv:2101.08320](https://arxiv.org/abs/2101.08320).
- [9] G. Kasieczka, S. Marzani, G. Soyez, and G. Stagnitto, *J. High Energy Phys.* **09** (2020) 195.
- [10] A. Coccaro, M. Pierini, L. Silvestrini, and R. Torre, *Eur. Phys. J. C* **80**, 664 (2020).
- [11] A. Radovic, M. Williams, D. Rousseau, M. Kagan, D. Bonacorsi, A. Himmel, A. Aurisano, K. Terao, and T. Wongjirad, *Nature (London)* **560**, 41 (2018).
- [12] D. Guest, K. Cranmer, and D. Whiteson, *Annu. Rev. Nucl. Part. Sci.* **68**, 161 (2018).
- [13] A. J. Larkoski, I. Moulton, and B. Nachman, *Phys. Rep.* **841**, 1 (2020).
- [14] A. J. Larkoski and E. M. Metodiev, *J. High Energy Phys.* **10** (2019) 014.
- [15] K. Datta and A. Larkoski, *J. High Energy Phys.* **06** (2017) 073.
- [16] J. A. Aguilar-Saavedra and B. Zaldivar, *Eur. Phys. J. C* **80**, 530 (2020).
- [17] J. A. Aguilar-Saavedra, F. R. Joaquim, and J. F. Seabra, *J. High Energy Phys.* **03** (2021) 012.
- [18] P. T. Komiske, E. M. Metodiev, and J. Thaler, *J. High Energy Phys.* **04** (2018) 013.
- [19] P. T. Komiske, E. M. Metodiev, and J. Thaler, *J. High Energy Phys.* **01** (2019) 121.
- [20] A. Chakraborty, S. H. Lim, M. M. Nojiri, and M. Takeuchi, *J. High Energy Phys.* **07** (2020) 111.
- [21] K. Fraser and M. D. Schwartz, *J. High Energy Phys.* **10** (2018) 093.
- [22] F. A. Dreyer, G. P. Salam, and G. Soyez, *J. High Energy Phys.* **12** (2018) 064.
- [23] L. de Oliveira, M. Kagan, L. Mackey, B. Nachman, and A. Schwartzman, *J. High Energy Phys.* **07** (2016) 069.
- [24] P. T. Komiske, E. M. Metodiev, and M. D. Schwartz, *J. High Energy Phys.* **01** (2017) 110.
- [25] A. Lifson, G. P. Salam, and G. Soyez, *J. High Energy Phys.* **10** (2020) 170.
- [26] F. A. Dreyer and H. Qu, *J. High Energy Phys.* **03** (2021) 052.
- [27] S. Carrazza and F. A. Dreyer, *Eur. Phys. J. C* **79**, 979 (2019).
- [28] B. M. Dillon, D. A. Faroughy, J. F. Kamenik, and M. Szewc, *J. High Energy Phys.* **10** (2020) 206.
- [29] G. Aad *et al.* (ATLAS Collaboration), *Phys. Rev. Lett.* **124**, 222002 (2020).
- [30] ALICE, <https://cds.cern.ch/record/2759456>.
- [31] G. Kasieczka, T. Plehn, M. Russell, and T. Schell, *J. High Energy Phys.* **05** (2017) 006.
- [32] S. Macaluso and D. Shih, *J. High Energy Phys.* **10** (2018) 121.
- [33] C. K. Khosa, L. Mars, J. Richards, and V. Sanz, *J. Phys. G* **47**, 095201 (2020).
- [34] C. K. Khosa, V. Sanz, and M. Soughton, *SciPost Phys.* **10**, 151 (2021).
- [35] Y. L. Chung, S. C. Hsu, and B. Nachman, [arXiv:2009.05930](https://arxiv.org/abs/2009.05930).
- [36] C. K. Khosa and V. Sanz, [arXiv:2007.14462](https://arxiv.org/abs/2007.14462).
- [37] J. Guo, J. Li, and T. Li, *Phys. Rev. D* **103**, 116025 (2021).
- [38] J. Lin, M. Freytsis, I. Moulton, and B. Nachman, *J. High Energy Phys.* **10** (2018) 101.
- [39] J. Li, T. Li, and F. Z. Xu, *J. High Energy Phys.* **04** (2021) 156.
- [40] A. Alves and F. F. Freitas, *Mach. Learn. Sci. Tech.* **1**, 045025 (2020).
- [41] E. A. Moreno, T. Q. Nguyen, J. R. Vlimant, O. Cerri, H. B. Newman, A. Periwai, M. Spiropulu, J. M. Duarte, and M. Pierini, *Phys. Rev. D* **102**, 012010 (2020).
- [42] K. Datta and A. J. Larkoski, *J. High Energy Phys.* **03** (2018) 086.
- [43] K. Datta, A. Larkoski, and B. Nachman, *Phys. Rev. D* **100**, 095016 (2019).
- [44] J. Amacker, W. Balunas, L. Beresford, D. Bortoletto, J. Frost, C. Issever, J. Liu, J. McKee, A. Micheli, S. Paredes Saenz *et al.*, *J. High Energy Phys.* **12** (2020) 115.
- [45] M. Abdughani, D. Wang, L. Wu, J. M. Yang, and J. Zhao, [arXiv:2005.11086](https://arxiv.org/abs/2005.11086).
- [46] V. S. Ngairangbam, A. Bhardwaj, P. Konar, and A. K. Nayak, *Eur. Phys. J. C* **80**, 1055 (2020).
- [47] P. C. Harris, D. S. Rankin, and C. Mantilla Suarez, [arXiv:1910.02082](https://arxiv.org/abs/1910.02082).
- [48] C. Englert, M. Fairbairn, M. Spannowsky, P. Stylianou, and S. Varma, *Phys. Rev. D* **102**, 095027 (2020).
- [49] C. Grojean, A. Paul, and Z. Qian, *J. High Energy Phys.* **04** (2021) 139.
- [50] A. Buckley, G. Callea, A. J. Larkoski, and S. Marzani, *SciPost Phys.* **9**, 026 (2020).
- [51] Y. L. Dokshitzer, G. D. Leder, S. Moretti, and B. R. Webber, *J. High Energy Phys.* **08** (1997) 001.
- [52] M. Wobisch and T. Wengler, [arXiv:hep-ph/9907280](https://arxiv.org/abs/hep-ph/9907280).
- [53] J. Alwall, R. Frederix, S. Frixione, V. Hirschi, F. Maltoni, O. Mattelaer, H. S. Shao, T. Stelzer, P. Torrielli, and M. Zaro, *J. High Energy Phys.* **07** (2014) 079.
- [54] T. Sjostrand, S. Mrenna, and P. Z. Skands, *J. High Energy Phys.* **05** (2006) 026.
- [55] T. Sjostrand, S. Mrenna, and P. Z. Skands, *Comput. Phys. Commun.* **178**, 852 (2008).
- [56] M. Cacciari, G. P. Salam, and G. Soyez, *J. High Energy Phys.* **04** (2008) 063.
- [57] M. Cacciari, G. P. Salam, and G. Soyez, *Eur. Phys. J. C* **72**, 1896 (2012).
- [58] F. Chollet, Keras, <https://keras.io>, 2015.
- [59] D. P. Kingma and J. Ba, [arXiv:1412.6980](https://arxiv.org/abs/1412.6980).

- [60] <https://feynrules.irmp.ucl.ac.be/wiki/HiggsEffectiveTheory>.
- [61] M. Dasgupta, A. Fregoso, S. Marzani, and G. P. Salam, *J. High Energy Phys.* **09** (2013) 029.
- [62] M. Dasgupta, A. Fregoso, S. Marzani, and A. Powling, *Eur. Phys. J. C* **73**, 2623 (2013).
- [63] M. Dasgupta, A. Powling, and A. Siodmok, *J. High Energy Phys.* **08** (2015) 079.
- [64] J. R. Andersen, J. Bellm, J. Bendavid, N. Berger, D. Bhatia, B. Biedermann, S. Bräuer, D. Britzger, A. G. Buckley, R. Camacho *et al.*, [arXiv:1803.07977](https://arxiv.org/abs/1803.07977).