

Mapping machine-learned physics into a human-readable space

Taylor Faucett¹, Jesse Thaler^{2,3}, and Daniel Whiteson¹

¹Department of Physics and Astronomy, UC Irvine, Irvine, CA 92627

²Center for Theoretical Physics, Massachusetts Institute of Technology, Cambridge, MA 02139

³The NSF AI Institute for Artificial Intelligence and Fundamental Interactions



(Received 14 November 2020; accepted 5 February 2021; published 26 February 2021)

We present a technique for translating a black-box machine-learned classifier operating on a high-dimensional input space into a small set of human-interpretable observables that can be combined to make the same classification decisions. We iteratively select these observables from a large space of high-level discriminants by finding those with the highest decision similarity relative to the black box, quantified via a metric we introduce that evaluates the relative ordering of pairs of inputs. Successive iterations focus only on the subset of input pairs that are misordered by the current set of observables. This method enables simplification of the machine-learning strategy, interpretation of the results in terms of well-understood physical concepts, validation of the physical model, and the potential for new insights into the nature of the problem itself. As a demonstration, we apply our approach to the benchmark task of jet classification in collider physics, where a convolutional neural network acting on calorimeter jet images outperforms a set of six well-known jet substructure observables. Our method maps the convolutional neural network into a set of observables called energy flow polynomials, and it closes the performance gap by identifying a class of observables with an interesting physical interpretation that has been previously overlooked in the jet substructure literature.

DOI: 10.1103/PhysRevD.103.036020

I. INTRODUCTION

It is widely appreciated that neural networks (NNs) and related machine-learning (ML) tools can provide powerful solutions to important and difficult problems in high-energy physics [1,2]. Examples of tasks that have benefitted from NNs include triggering [3], event reconstruction [4,5], object identification [6–8], and event selection [9–11]. In all of these contexts, though, the physicist wonders: *what has the machine learned?* Satisfaction with improved performance is tempered by frustration with the “black-box” nature of NN strategies.

The advent of deep learning has made this question more urgent, as the data dimensionality of the tasks has increased dramatically and ML approaches have outperformed human-engineered strategies for problems that, until recently, were deemed too difficult for ML. Examples include event classification [12–17], jet substructure studies [18–22], jet flavor classification [23–27], detector unfolding [28–30], and uncertainty estimation [31–35]. In each case, a deep NN has successfully utilized more information from the

high-dimensional low-level input data than was captured by a smaller number of physics-motivated high-level (HL) observables. When there is a performance gap between machine-learned and human-engineered strategies, physicists want to understand how the NN is using the low-level information to improve its performance. This desire also applies to situations where performance of the ML solution matches (but does not exceed) the human-engineered approach. Has the machine learned the same strategy that humans devised, or has it found an alternative solution with equal efficacy?

In this paper, we present a technique for translating a black-box ML strategy based on low-level inputs into a human-readable space of HL observables. Instead of trying to directly interpret an NN classifier, our approach is to use the NN as a guide for the construction of a simpler classifier which makes the same decisions but relies on a small number of physics-inspired human-interpretable inputs, selected iteratively from a large space of observables. As a demonstration of our method, we present a case study in jet classification [20] using a convolutional neural network (CNN) to guide the selection of a small set of HL observables called *energy flow polynomials* (EFPs) [36]. We find that the final set of HL observables provides the same classification performance as the CNN acting on the low-level inputs, but in a more compact, interpretable, and physically meaningful format. Our study suggests that

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.

physicists should consider an overlooked set of HL observables that are relevant to jet substructure classification.

This desire to gain insight into the nature of ML strategies is important on several levels. First, it is critical that information used by the NNs be validated as real and physical, rather than an artifact of the training samples or procedure. Even in cases where the ML training does not rely on simulated samples [37–39], it is important to understand what information is being used. Translating a black-box ML strategy into a simpler network that relies on a smaller number of physically meaningful observables allows for effective validation. Second, having an interpretable strategy based on HL inputs allows for more reliable estimates of systematic uncertainties. If the HL space is small enough, the HL inputs themselves can be individually studied and calibrated, which is a more straightforward task than handling low-level inputs directly. Third, if one can replace a sophisticated ML strategy with a simpler network based on preprocessed inputs, this can enable faster inference and lower memory requirements at run-time [40,41]. Finally, if previously overlooked information in the low-level data is physical, identifying it can provide new insights into the nature of the problem. Indeed, in our jet classification case study, we show how the six HL observables studied in Ref. [20] can be augmented by a seventh HL observable that had not been previously considered in the literature. By scrutinizing the structure of the EFPs, we can provide a physical interpretation for this new observable.

There have been previous proposed strategies to draw connections between a learned NN strategy and existing HL observables. This can be done by comparing the performance with and without the HL observables [42] or projecting the decision surfaces along those observables [12]. An alternative strategy is to expand the NN function in a basis of the input features [43–45]. These strategies are valuable, but are primarily limited to studying the structure of the NN in terms of already-identified HL observables. The goal of the present work is to identify new HL observables relevant for ML tasks, starting from a large space of HL observables that is as broad and comprehensive as possible (yet still interpretable) and systematically mapping a black-box ML strategy into that space.

Our mapping strategy suggests a new approach to the application of deep learning to high-energy physics data. In this approach, training a powerful deep neural network (DNN) on low-level inputs is just the first step, which helps gauge the effective upper limit on possible ML performance and determine asymptotically optimal decision boundaries. The new second step is translating as much of the ML strategy as possible to a well-understood set of HL observables. This allows for physical interpretation of the information being used, validation of the modeling, definition of reasonable systematic uncertainties, as well as computational benefits due to dimensionality reduction.

An extended outline of this paper is as follows. In Sec. II, we present a general approach to map an ML model's learned solution into a human-readable space. This mapping requires that we construct a large space of candidate HL observables and use a reliable similarity metric for comparing these observables to the learned solution. As our similarity metric, we introduce the *average decision ordering* (ADO), which is related to Kendall's rank correlation coefficient [46] and quantifies to what extent two classifiers make the same (even if incorrect) decisions. We then present a mapping strategy that leverages the ADO. Our *black-box guided strategy* iteratively finds the maximum ADO between the HL observables and a fixed black-box ML algorithm.

In Sec. III, we review the collider task of discriminating between jets originating from boosted W bosons and those originating from light quarks and gluons. This is a well-studied problem in the field of jet substructure [21,47–58], where both HL [59–64] and NN strategies [18–20] have proven effective. Our starting point will be the analysis of Ref. [20], which found a small but persistent improvement in classification performance with a deeply-connected CNN when compared to a boosted decision tree (BDT) of HL observables. To augment the set of HL observables for jet tagging, we search the space of EFPs [36], which form an (over)complete basis of collider observables that are infrared and collinear (IRC)-safe. We also introduce EFP variants inspired by IRC-unsafe observables that have been successful in other jet tagging studies [65–68].

The results of our case study are presented in Secs. IV and V. We start with the six HL observables from Ref. [20], and use the black-box guided strategy to identify a seventh HL observable that closes the performance gap with the CNN. We then apply the black-box guided strategy starting from just the mass and transverse momentum of the jet, comparing the results to a brute force strategy of directly searching the space of EFPs and a guided search based on ground truth labels. The black-box guided strategy significantly outperforms the label guided search, reaching comparable performance to the brute force strategy with considerably reduced computational costs. At the end of each of these sections, we provide a physical interpretation of the translated ML strategy, and we draw broader conclusions in Sec. VI.

II. TRANSLATING FROM MACHINE TO HUMAN

In our mapping approach, we seek to identify a small set of physically-motivated HL observables that, when combined into a joint classifier, make the same classification decisions as a deep network operating on the low-level features. Crucially, our set of HL features is designed to, when combined, maximize the classification performance by following the learned strategy of the black-box NN, which we argue below is more efficient than training

directly on ground truth information. If this translation is successful, then we will have expressed the ML strategy more simply and transparently in terms of a few HL observables.

The first step in our approach is to identify a comprehensive set of HL observables that are potentially relevant for solving the ML task at hand. This in turn requires (human) knowledge about the physical system being studied and about the kinds of HL observables that have interpretable meaning. We know of no way to automate this step, though automation may not be desirable, as the choice and structure of the HL observable space defines the physical interpretation. For our jet substructure classification case study, the EFPs have already been identified as a suitable basis of HL observables [36], as discussed further in Sec. III B.¹ Other ML tasks in high-energy physics might require the development of alternative bases of HL observables.

In the rest of this section, we describe the aspects of our approach that are generic to any ML task, focusing on the case of binary classification. To evaluate the relative performance of our simple HL network and a black-box NN, we need a metric to evaluate whether two classifiers make the same classification decisions. There are many such metrics one could use, but we introduce the ADO, in part because it shares the conceptual simplicity of the area under the curve (AUC) metric often used to benchmark classifiers against ground truth. Armed with an explorable set of HL observables and a metric for assessing learning similarity, we then present a guiding strategy to map black-box NNs into a physically meaningful space.

A. Average decision ordering

Our guided strategy requires a similarity metric that compares the output of two decision functions $f(x)$ and $g(x)$. Here, x represents the full high-dimensional low-level data, which are inputs to both the black-box NN and the physically-motivated HL observables. Of course, the definition of similarity must reflect the task for which these decision functions are applied, which in this case is binary classification. Because classification performance is invariant under any nonlinear monotonic transformation of f or g , our similarity metric cannot be affected by such a transformation. This rules out naive metrics like functional overlap or linear correlation.

Furthermore, it is not sufficient to simply compare the overall performance of two classifiers over a given dataset, since that does not provide information about how the low-level inputs are being used. As discussed in Ref. [67], two decision functions might use information from different regions of the low-level input space and make conflicting

¹Beyond jet substructure, the EFPs are also formally a complete basis for event-wide IRC-safe observables, but their utility in that context has not yet been established.

classification decisions case by case, yet still achieve similar overall performance. The key to our guided strategies is that we aim to match not just the classification performance of the black-box NN, but also the specific classification decisions.

We assume that the decision functions $f(x)$ and $g(x)$ are real valued and that the final binary classification is determined by a threshold on the decision function output. Objects on one side of the threshold are labeled “signal,” while objects on the other side are labeled “background.” Depending on the particular application, this threshold can be tuned to different points on the receiver operating characteristic curve to optimize the signal acceptance versus background rejection. To quantify overall classification performance, we use the AUC, which is equivalent to the probability that a randomly selected signal/background pair is correctly ordered by a decision function $f(x)$:

$$\text{AUC}[f] = \int dx dx' p_{\text{sig}}(x) p_{\text{bkg}}(x') \Theta(f(x) - f(x')). \quad (1)$$

Here, Θ is the Heaviside theta function [i.e., $\Theta(x < 0) = 0$ and $\Theta(x \geq 0) = 1$] and p_{sig} and p_{bkg} are the ground truth signal and background probability distributions. A perfect decision function has $\text{AUC} = 1$ and random guessing yields $\text{AUC} = \frac{1}{2}$.

To compare the classification behavior of two decision functions, we consider the *decision surface* defined by a threshold on the function output. For each function, the set of such thresholds defines a set of surfaces in x space. If two decision functions have the same decision surfaces, then they are effectively using the same low-level information for classification. Note that the absolute output values of the decision functions are not relevant for determining whether the decision surfaces are similar. The relative locations of the decision surfaces are determined by the relative ordering of the two decision functions when evaluated at pairs of points in the input space. We can encapsulate this information via the *decision ordering* (DO) for a pair of points x and x' :

$$\text{DO}[f, g](x, x') = \Theta((f(x) - f(x'))(g(x) - g(x'))), \quad (2)$$

where 1 corresponds to f and g having the same ordering and 0 corresponds to inverted ordering.

If two decision functions have $\text{DO} = 1$ for all pairs of x and x' , then they are monotonically related to each other, have identical decision surfaces, and are therefore identical for the purposes of classification. To build a summary statistic, we average over all possible values of x and x' , weighted by the signal and background distributions, yielding the ADO:

$$\text{ADO}[f, g] = \int dx dx' p_{\text{sig}}(x) p_{\text{bkg}}(x') \text{DO}[f, g](x, x'). \quad (3)$$

This evaluates to 1 if the decision functions make the same relative classification decision for every pair, to 0 if the functions make the opposite classification for every pair, and to $\frac{1}{2}$ if there is no consistency in their orderings. Since a decision function can be trivially inverted, the cases of $\text{ADO} = 0$ and $\text{ADO} = 1$ have the same meaning, so we map $\text{ADO} \rightarrow 1 - \text{ADO}$ whenever it is less than $\frac{1}{2}$. The ADO has a similar philosophy to Kendall's rank correlation coefficient [46], with the key difference that we are comparing inputs drawn from separate signal and background distributions.

To gain intuition for the ADO, note that it has a very similar structure to the AUC in Eq. (1). The AUC is the probability that a single decision function orders objects correctly relative to ground truth. The ADO is the probability that two decision functions order objects in the same way, even if incorrectly. In the case that $f(x) = p_{\text{sig}}(x)/p_{\text{bkg}}(x)$ is the likelihood ratio, then $f(x)$ is an optimal classifier by the Neyman–Pearson lemma, so an $\text{ADO} = 1$ implies that $g(x)$ defines the same optimal decision boundaries as $f(x)$. In most ML applications, one is trying to maximize the AUC or other similar metric of absolute classification performance. Our guided strategies, by contrast, aim to maximize the ADO relative to an already trained ML tool.

There are other similarity metrics that one could use, but they are not as easy to interpret in terms of classification

decisions. One way to capture information similarity is to use mutual information, or more appropriate to binary classification, mutual information with the truth [67]. For our guided strategies, though, we are less interested in whether two decision functions have the same quantity of information available to a classification task, and more interested in quantifying the degree to which two decision functions make classification decisions in the same way. Even if $f(x)$ and $g(x)$ contain a high level of mutual information, they do not necessarily define the same decision boundaries. This is especially important to keep in mind given the flexibility of deep networks, which allow the same discrimination power to be encoded in many informationally equivalent ways. Beyond the ADO, there are other summary statistics one could use based on the raw DO information, and we leave a study of those alternatives to future work.

B. Black-box guided search strategy

The idea behind the black-box guided strategy is shown in Fig. 1, where the goal is to find HL observables that maximize the ADO relative to an already trained ML tool. We denote the reference black-box network as “BBN” (with apologies to cosmologists), and it will typically be some kind of deep network acting on the low-level inputs. Starting from a large set \mathcal{S} of human-defined HL observables motivated by physics considerations, our aim is to build a high-level network (HLN) with the same decision surfaces as the BBN. The initial step ($n = 0$) in the black-box guided approach is to identify the observable HL_1 that has the largest ADO with the BBN

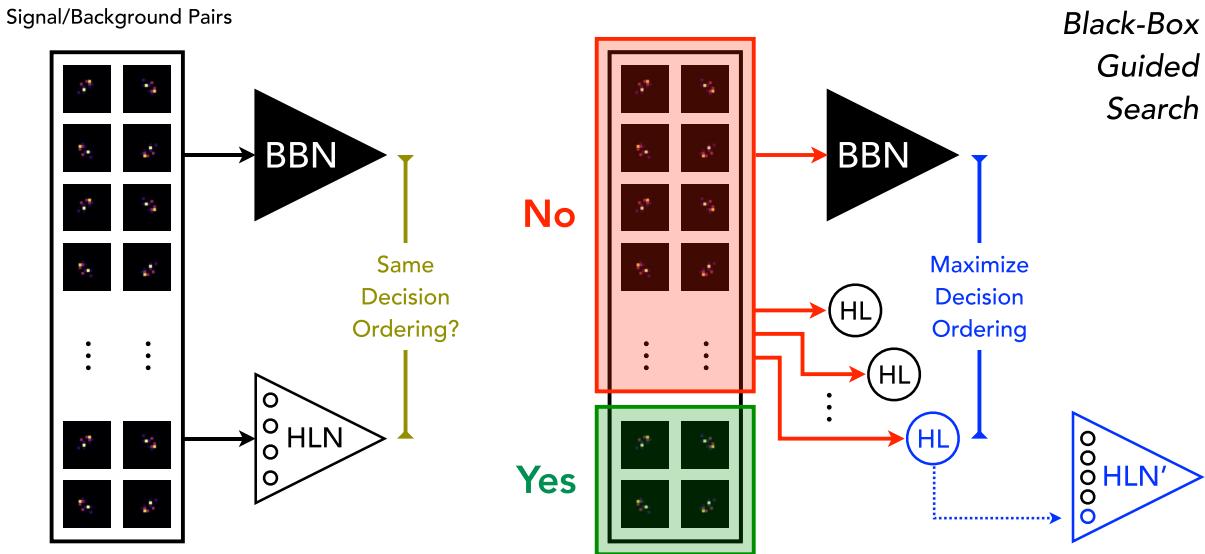


FIG. 1. Schematic of the black-box guided search in Sec. II B. In each iteration of this strategy, the relative decision ordering of signal/background pairs between the fixed black-box network (BBN, black triangle) and a trainable network of HL observables (white triangle) is used to identify the subset (red box) in which pairs are differently ordered. From a large space of HL observables (circles), the one with the largest ADO in the misordered space (blue circle) is selected for the next iteration. The schematic above corresponds to the $n = 4$ iteration. Note that the BBN is not retrained in each iteration, but the network of HL observables is.

$$\text{HL}_1 = \text{argmax}_{\text{HL} \in \mathcal{S}} \text{ADO}[\text{BBN}, \text{HL}]_{X_{\text{all}}}. \quad (4)$$

Here, the X_{all} subscript indicates that we are using the full set of signal/background training pairs (x, x') when computing the ADO. The observable HL_1 is therefore the physics-motivated observable in the set \mathcal{S} that best approximates the decision surfaces of the black box.

In the next step ($n = 1$), we focus our search for HL observables in regions of the feature space where the black-box network disagrees with our current set of HL observables, by isolating the subset of signal/background pairs X_1 that are ordered differently by the BBN and HL_1 :

$$X_1 = \{(x, x') | \text{DO}[\text{BBN}, \text{HL}_1](x, x') = 0\}. \quad (5)$$

We then identify the observable HL_2 that has the largest ADO with the BBN when restricted to the X_1 subset:

$$\text{HL}_2 = \text{argmax}_{\text{HL} \in \mathcal{S}} \text{ADO}[\text{BBN}, \text{HL}]_{X_1}. \quad (6)$$

For each subsequent step $n > 1$, we combine the HL observables already identified in the previous steps into a joint network,

$$\text{HLN}_n = \text{NN}[\text{HL}_1, \dots, \text{HL}_n], \quad (7)$$

where NN indicates a neural network trained on the full signal/background training set with just the n HL observables as inputs. From this joint HLN, the misordered subset X_n is defined via

$$X_n = \{(x, x') | \text{DO}[\text{BBN}, \text{HLN}_n](x, x') = 0\}. \quad (8)$$

Because a new HLN is trained in each iteration, X_n may not be a strict subset of X_{n-1} . The next observable HL_{n+1} is determined via

$$\text{HL}_{n+1} = \text{argmax}_{\text{HL} \in \mathcal{S}} \text{ADO}[\text{BBN}, \text{HL}]_{X_n}. \quad (9)$$

Note that the same black-box network is used in each iteration, but the changing subset X_n means that different decision surface information is tested at each step. These steps are repeated until the $\text{ADO}[\text{BBN}, \text{HLN}_{n+1}]$ gets as close to 1 as desired.

Isolating the differently-classified pairs in Eq. (8) is similar in spirit to the boosting step of the BDTs [69,70]. This approach focuses attention only on the subspace of pairs where the BBN disagrees with the current set of HL observables, allowing us to identify new HL observables that make signal-background ordering decisions most similar to the BBN in that subspace. It is worth emphasizing that the ADO, or some other metric for network decision similarity, is essential for this approach to work.

Later in Sec. V C, we will compare this black-box guided approach to a label guided approach. Instead of using the

ADO, the label guided approach uses the AUC with respect to ground truth information. It is straightforward to understand why the ADO is superior to the AUC for guiding purposes. To the extent that the BBN is well trained, it represents a good approximation to the Neyman–Pearson optimal classifier. Achieving the correct DO relative to the optimal classifier for every signal/background training pair is the best one could ever hope to do. Therefore, if the black-box guiding strategy is working correctly, then the subsets X_n will get smaller and smaller until almost all signal/background pairs have been correctly ordered relative to the BBN.

By contrast, the AUC captures DO relative to truth labels. Unless the BBN is able to achieve $\text{AUC} = 1$, there will inevitably be signal/background pairs that are incorrectly ordered even by the theoretically optimal classifier. Instead of getting smaller and smaller, the subsets X_n will stall at the set of signal/background pairs that can never be ordered correctly. This in turn means that the classification performance of HLN_n will stall well below the theoretical maximum in the label guided approach. That is why we advocate for the selection of HL observables to be guided by the ADO, since then the classification performance of the HLN_n will eventually match that of the BBN, as desired.

As with any “greedy algorithm,” our black-box guided strategy cannot identify situations where two HL observables could be combined simultaneously to match the BBN decision surfaces. This means that we might miss sets of observables that are individually poor classifiers but perform well jointly. If the goal was just to maximize performance, this would be an undesirable feature. In the context of mapping a black-box ML strategy to a physically-interpretable space, though, we are indeed looking for individual observables with high information content relevant for classification, so this greedy strategy is the one most likely to yield physical insight.

III. A CASE STUDY IN JET SUBSTRUCTURE

We now apply the technique introduced in Sec. II to a specific case study involving jet classification at the LHC. In this section, we review boosted W boson classification using jet substructure and highlight the elements of Ref. [20] that will be used in our case study. We then introduce the EFPs [36] as our set of HL physics-motivated observables. Details about our NN architectures and training parameters are provided in the Appendix A.

A. Boosted boson classification

Massive objects produced at the LHC often have enough transverse momentum that their decay products become collimated. For an object with a hadronic decay mode, such as a W boson decaying to a quark-antiquark pair ($W \rightarrow q\bar{q}'$), the resulting jet in the detector consists of two

clusters of energy, one from each of the fragmenting quarks. The substructure of these jets is distinct from those that arise from the fragmentation of a single hard quark or gluon. Identification of jets with nontrivial substructure has become an essential tool for probing the nature of collisions at the LHC [21,47–58].

There are many different ways to represent the information in a jet. At the most fundamental level, a jet is a variable-length collection of four-vectors with associated particle properties, motivating set-based ML tools [71–76]. Another popular approach is to describe a jet as a grid of calorimeter cells with energy depositions, giving rise to a “jet image” [19,77]. In any of these low-level representations, the jet data is high-dimensional. This motivates the development of HL observables that intelligently summarize the low-level information to reduce the effective dimensionality of the task. Physicists have engineered numerous HL observables tasks that incorporate domain knowledge about jet formation (see Refs. [51,59–64,78–85] for an incomplete list). Typical usage is to apply cuts on one or more of these HL observables, or to combine several of them using a shallow ML classifier.

In the context of jet classification, ML tools based on low-level inputs have outperformed traditional strategies based on HL observables [86]. Of course, the HL observables themselves are just functions of the low-level inputs, so it should be possible to find a large enough set of physics-motivated HL observables that can match the performance of these ML classifiers [87–89]. This is indeed the intuition behind the guided strategy in Sec. II, where the goal is to leverage a black-box ML method to identify the most effective HL observables.

Our case study is based on the same datasets as Ref. [20]. These datasets correspond to $\sqrt{s} = 14$ TeV proton-proton collision, where hard scattering and resonance decay were generated using MADGRAPH 5 v2.2.3 [90], showering and hadronization were generated with PYTHIA v6.426 [91], and the response of the detectors was simulated with DELPHES v3.2.0 [92]. The boosted W boson signal process is diboson production ($pp \rightarrow W^+W^-$), which yields two fat jets each with 2-prong substructure. The background process is QCD dijet production ($pp \rightarrow qq, qg, gg$), which typically yields 1-prong jets. These samples do not include contamination from pileup (multiple proton-proton collisions per beam crossing). Jets are clustered using the anti- k_t algorithm [93] with a radius parameter of $R = 1.2$, using FASTJET 3.1.2 [94]. The dataset contains 5×10^6 events, split equally between signal and background. Following the approach in Ref. [20], each jet is pixelated into a 32×32 grid in the rapidity-azimuth plane, and a jet image is formed from the transverse momentum (p_T) deposits in each cell. The jet image is then trimmed [95], where subjets of radius $R_{\text{sub}} = 0.2$ are discarded if their p_T is less than 3% of the original jet. The final jet selection takes jets with trimmed momentum $p_T^{\text{trim}} \in [300, 400]$ GeV within the rapidity

range $|\eta| < 5.0$. While important jet information is lost by pixelation and trimming, we include these steps in our analysis in order to perform an apples-to-apples comparison to Ref. [20].

The trimmed jet’s constituents are used to compute six HL jet substructure observables: the trimmed jet mass (M_{jet}), four ratios of energy correlation functions ($C_2^{\beta=1}, C_2^{\beta=2}, D_2^{\beta=1}, D_2^{\beta=2}$) [62,64], and the N -subjettiness ratio ($\tau_2^{\beta=1}$) [60,61]. These observables are well-established in the context of boosted W boson classification, including studies at ATLAS [96,97] and CMS [98]. The W boson classification performance of these six HL observables is summarized in Table I. The trimmed jet mass is the most powerful single observable, since the 80.4 GeV mass peak is a characteristic feature of boosted W bosons.

We can use the ADO from Eq. (3) to gain additional insight into these six HL observables. In Fig. 2, we assess the pairwise ADO between each of the HL observables considered. The observable pairs that make the most similar decisions (i.e., $\text{ADO} \rightarrow 1$) are $C_2^{\beta=1}$ with $C_2^{\beta=2}$ and $D_2^{\beta=1}$ with $D_2^{\beta=2}$. This is expected since these observables have relatively similar structures except for the choice of β coefficient, which controls the weighting of angular information within the jets. These pairs also have similar AUC values, as seen in Table I, since pairs that make common classification decisions should exhibit similar classification power. Comparing the AUC and ADO values provides a more detailed picture about the degree of correlation in classification.

TABLE I. Classification performance of the six HL observables studied in Ref. [20], as well as a 6HL joint classifier. The six HL observables face a small but significant performance gap compared to the benchmark CNN. As discussed later in Sec. IV A, this performance gap is bridged by a seventh feature discovered using our black-box guided strategy. The 488HL case involving a large set of EFPs is discussed at the end of Sec. V B. The uncertainty on the AUC is computed from one standard deviation of 10-fold cross-validation. The decision similarity (ADO) to the benchmark CNN is also shown. Details of the NN architectures are provided in Appendix A.

Observable	AUC	ADO[CNN, Obs.]
M_{jet}	0.898 ± 0.004	0.807
$C_2^{\beta=1}$	0.660 ± 0.006	0.584
$C_2^{\beta=2}$	0.604 ± 0.007	0.548
$D_2^{\beta=1}$	0.790 ± 0.005	0.743
$D_2^{\beta=2}$	0.807 ± 0.005	0.762
$\tau_2^{\beta=1}$	0.662 ± 0.006	0.600
6HL	0.9504 ± 0.0002	0.971
CNN	0.9531 ± 0.0002	1.000
488HL	0.9535 ± 0.0002	0.978
7HL _{black-box}	0.9528 ± 0.0003	0.971

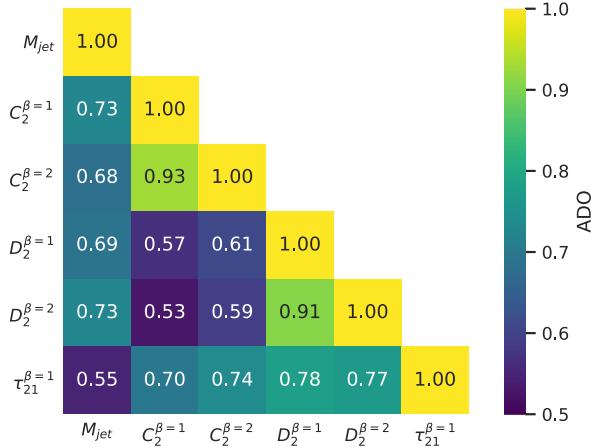


FIG. 2. Similarity of the classification decisions between the six HL observables, as quantified by the ADO. A value of $ADO = 1$ indicates identical decision ordering for all signal/background pairs, while $ADO = \frac{1}{2}$ corresponds to no similarity. In this way, the ADO has a similar interpretation to the AUC, but with respect to classification decisions instead of ground truth.

The observable pairs that make the least similar decisions (i.e., $ADO \rightarrow \frac{1}{2}$) are M_{jet} with $\tau_{21}^{\beta=1}$ and C_2^{β} with D_2^{β} . For M_{jet} versus $\tau_{21}^{\beta=1}$, this is expected since the N -subjettiness probes the degree of pronglike collimation, whereas mass is sensitive to the energies of the prongs and their relative angles. For C_2^{β} versus D_2^{β} , this is expected since they have different scalings under boosts along the jet direction [64]. Pairs that make dissimilar decisions can often be combined into more powerful joint classifiers. This is shown in Fig. 3, where we consider the pairwise classifiers $NN[HL_i, HL_j]$, where the details of the NN parameters are presented in Appendix A 2. More comprehensive studies of these six jet substructure observables can be found in Ref. [56].

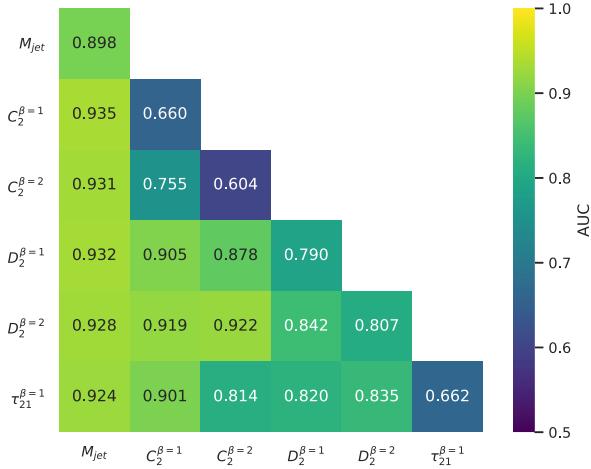


FIG. 3. Classification performance of the six HL observables in Table I on the diagonal entries, along with AUC values for pairwise classification between coupled HL inputs on the off-diagonal entries.

While these six engineered HL features are powerful jet substructure discriminants, they do not capture the full information relevant for W boson tagging. Viewing the calorimeter cells as pixels of a two-dimensional image, we can try to enhance the discrimination power using computer vision techniques [18–20,22,27,32,77,99,100]. Indeed, Ref. [20] demonstrated that a deeply-connected CNN using the low-level jet image inputs yielded better classification performance than the six HL observables combined with a BDT. The performance gain was modest though persistent, making it an excellent benchmark problem for studying the interpretation of ML strategies. We repeat this exercise in Table I, now using the NN parameters in Appendix A for the CNN and for the 6HL combination,

$$6\text{HL} \equiv NN[M_{jet}, C_2^{\beta=1}, C_2^{\beta=2}, D_2^{\beta=1}, D_2^{\beta=2}, \tau_{21}^{\beta=1}], \quad (10)$$

where we find a 0.0027 ± 0.0003 gap in the AUC, as seen in Table I. Using our guided strategy, we seek to understand the nature of this performance gap, and whether the CNN has found a strategy similar to the existing HL observables or something distinct. Does the gap indicate a mild optimization of the same basic HL ideas, or does it hint at the existence of a new HL observable not appearing previously in the jet substructure literature?

B. Energy flow polynomials

In order to map the CNN from Ref. [20] to a human-readable space, we first define a suitable set of physics-motivated HL observables for use in the guided strategies. This requires domain knowledge about the underlying physics as well as intuition about the kinds of information that might be missing from existing HL observables. This also requires identifying HL observables that are likely to work well as classifiers individually, since the black-box guided strategy in Sec. II B is based on finding a single observable that maximizes the ADO in each step.

Our set of HL observables is based on the EFPs [36]. The EFPs are a large (formally infinite) set of parametrized engineered functions, inspired by previous work on energy correlation functions [62,64,101–104]. In the jet image representation, they are defined in terms of the momentum fraction z_a of calorimeter cell a , as well as the pairwise angular distance θ_{ab} between cells a and b . The EFPs are built in increasing levels of complexity, from simple sums over single cells to many higher-order combinations of momentum and pair-wise angles. An EFP can be represented as a multigraph, where

$$\text{each node} \Rightarrow \sum_{a=1}^N z_a, \quad (11)$$

$$\text{each } k\text{-fold edge} \Rightarrow (\theta_{ab})^k. \quad (12)$$

As one example, we have



$$= \sum_{a=1}^N \sum_{b=1}^N \sum_{c=1}^N \sum_{d=1}^N z_a z_b z_c z_d \theta_{ab}^3 \theta_{bc} \theta_{ac} \theta_{ad}^2. \quad (13)$$

From these graph representations, we can express both connected and disconnected graphs. For the purposes of our study, we only consider connected graphs.

The observable corresponding to each graph can be modified with parameters (κ, β) . These parameters determine the specific meaning of z_a and θ_{ab} , where

$$z_a^{(\kappa)} = \left(\frac{p_{Ta}}{\sum_b p_{Tb}} \right)^\kappa, \quad (14)$$

$$\theta_{ab}^{(\beta)} = (\Delta \eta_{ab}^2 + \Delta \phi_{ab}^2)^{\beta/2}. \quad (15)$$

Here, p_{Ta} is the transverse momentum of cell a , and η_{ab} (ϕ_{ab}) is the pseudorapidity (azimuth) difference between cells a and b . The original IRC-safe EFPs require $\kappa = 1$, however we include examples with $\kappa \neq 1$ to explore a broader space of HL observables, motivated by Refs. [65–68].² Note that $\kappa > 0$ generically corresponds to IR-safe but C-unsafe observables. We intentionally include zero and negative values of κ to explore both IR-unsafe and C-unsafe observables as well.³ For our study, we use the ENERGYFLOW PYTHON package [105] to translate jet-image p_T , η , and ϕ information to the EFPs with varying graphs and choices of κ and β .

For our guided search, we consider all combinations of (κ, β) where $\kappa = [-1, 0, \frac{1}{2}, 1, 2]$ and $\beta = [\frac{1}{2}, 1, 2]$. Each of the 15 combinations of (κ, β) are applied to the complete set of connected graphs with degree (i.e., number of edges) $d \leq 7$ along with all connected graphs with degree $d \leq 8$ and chromatic number $c = 4$ (to be defined in Sec. IV B below), which comes to 509 in total. This yields a space of 7635 HL observables to search from. Note that β in Eq. (15) can sometimes be traded for κ in Eq. (12); we remove degenerate graphs from our space, reducing our pool of unique observables to 7545.

It is important to emphasize that, although the EFP space constitutes a formally complete basis for (IRC-safe) jet classification, we are primarily concerned with the pragmatics of isolating individual observables that can map out the CNN behavior. The ideal case is that the CNN strategy maps to a single EFP, indicating that it can be expressed compactly in terms that can be easily interpreted by humans. Failing that, though, it is still of considerable value if a similar mapping can be made using a small collection of observables [87–89]. This would still provide a significantly more physically meaningful interpretation of the data and a marked reduction in data complexity and dimensionality.

²We thank Patrick Komiske and Eric Metodiev for discussions related to the normalization of the $\kappa \neq 1$ EFPs.

³For $\kappa < 0$, empty cells are omitted from the sum in Eq. (11). This is not as discontinuous as one might naively think due to the pixelation and trimming steps.

Beyond this specific benchmark example, if one is unable to map the CNN strategy into a small number of EFPs, this could mean one of two things. First, it could mean that the CNN strategy simply does not operate in this HL space, requiring us to revisit the assumption that the HL space was sufficiently complete to capture the essential information for jet classification. Second, it could mean that the CNN strategy is still encodable in terms of these HL observables, but in a more complex combination. As an example of this second possibility, consider the C_2 [62] and D_2 [64] observables discussed in Sec. III A. These can be written as EFPs with $\kappa = 1$:



$$C_2^{(\beta)} = \text{triangle graph} / \left(\text{line graph} \right)^2, \quad (16)$$



$$D_2^{(\beta)} = \text{line graph} / \left(\text{triangle graph} \right)^3, \quad (17)$$

where the graphs corresponds to



$$= \sum_{a=1}^N \sum_{b=1}^N \sum_{c=1}^N z_a z_b z_c \theta_{ab}^{(\beta)} \theta_{bc}^{(\beta)} \theta_{ca}^{(\beta)}, \quad (18)$$



$$= \sum_{a=1}^N \sum_{b=1}^N z_a z_b \theta_{ab}^{(\beta)}. \quad (19)$$

The guided strategies, however, would not necessarily be able to identify these ratio combinations unless they were defined ahead of time.⁴ Therefore, whether or not the guided mapping is effective, one learns something about the nature of the physics problem either way.

IV. SUPPLEMENTING EXISTING OBSERVABLES

In this section, we demonstrate the success of the mapping strategy from Sec. II in searching for an additional HL observable in the context of boosted W classification.

From Table I, we saw that the difference in classification power between a CNN acting on jet images and an NN combination of six HL observables is relatively small, but genuine and statistically significant. As such, it is interesting to ask whether this existing set of six HL observables could be supplemented by a new single HL observable which has not yet been considered by human physicists. We employ our black-box guiding strategy to find such a seventh HL observable, which closes the performance gap previously identified in Ref. [20].

⁴It might be interesting to combine the guided strategy with some kind of symbolic regression to find interesting combinations [106]. If this symbolic regression allowed for index contraction, then one could use the energy flow moments [107] to more efficiently search the space of $\beta = 2$ EFPs.

TABLE II. A selection of EFPs, sorted by their similarity with the CNN, evaluated using the ADO in the differently-ordered subspace X_6 . This corresponds to step one in the black-box guiding technique depicted in Fig. 1. After the top ten, EFPs are shown if they correspond to a dot graph, appear in the C_2/D_2 observables from Eqs. (16) and (17), or have the highest ADO among graphs with a given value of κ , β , or chromatic number.

Rank	EFP	κ	β	Chrom #	$\text{ADO}[\text{EFP, CNN}]_{X_6}$	$\text{AUC}[\text{EFP}]$	$\text{ADO}[6\text{HL} + \text{EFP, CNN}]_{X_{\text{all}}}$	$\text{AUC}[6\text{HL} + \text{EFP}]$
1		2	$\frac{1}{2}$	3	0.6207	0.8031	0.9714	0.9528 ± 0.0003
2		2	$\frac{1}{2}$	3	0.6205	0.8203	0.9714	0.9524
3		0	...	1	0.6205	0.6737	0.9715	0.9525
4		2	$\frac{1}{2}$	3	0.6199	0.8301	0.9715	0.9527
5		2	$\frac{1}{2}$	3	0.6197	0.8290	0.9714	0.9527
6		2	$\frac{1}{2}$	3	0.6196	0.8251	0.9715	0.9522
7		0	$\frac{1}{2}$	2	0.6187	0.7511	0.9715	0.9526
8		2	$\frac{1}{2}$	3	0.6184	0.8257	0.9712	0.9527
9		2	$\frac{1}{2}$	3	0.6182	0.8090	0.9714	0.9527
10		2	$\frac{1}{2}$	3	0.6180	0.8314	0.9714	0.9526
60		0	1	2	0.6163	0.7194	0.9715	0.9525
341		-1	$\frac{1}{2}$	4	0.6142	0.6286	0.9714	0.9509
589		0	2	2	0.6109	0.7579	0.9714	0.9523
3106		-1	...	1	0.5891	0.5882	0.9714	0.9510
3519		$\frac{1}{2}$	$\frac{1}{2}$	2	0.5664	0.7698	0.9715	0.9524
3521		$\frac{1}{2}$...	1	0.5663	0.7093	0.9714	0.9522
5531		1	2	1	0.5290	0.7454	0.9714	0.9507
5554		1	$\frac{1}{2}$	2	0.5279	0.8210	0.9713	0.9505
5610		2	...	1	0.5245	0.7117	0.9714	0.9507
5657		1	1	3	0.5224	0.8257	0.9712	0.9506
5793		1	1	2	0.5191	0.8640	0.9714	0.9505
6052		1	2	3	0.5153	0.8500	0.9716	0.9504
7438		1	2	2	0.5011	0.8835	0.9716	0.9506

A. Black-box guiding

The first step of the black-box guided strategy from Sec. II B is to identify the subset of signal/background pairs that are differently ordered by the CNN and the 6HL combination:

$$X_6 = \{(x, x') | \text{DO}[\text{CNN}, 6\text{HL}](x, x') = 0\}. \quad (20)$$

Though we have 6.25×10^{12} signal/background pairs, we construct X_6 from a randomly selected subset of 5×10^7 pairs. We then search for the EFP that has the greatest decision similarity to the CNN in the X_6 subspace, which ideally captures all of the remaining discrepancies between the CNN and 6HL approaches:

$$\text{HL}_7^{\text{black-box}} = \text{argmax}_{\text{HL} \in \text{EFPs}} \text{ADO}[\text{CNN, HL}]_{X_6}. \quad (21)$$

The results are shown in the first row of Table II. The EFP with the largest ADO with the CNN in the X_6 subspace is

$$\text{HL}_7^{\text{black-box}}(\kappa=2, \beta=\frac{1}{2}) = \sum_{a=1}^N \sum_{b=1}^N \sum_{c=1}^N \sum_{d=1}^N z_a^2 z_b^2 z_c^2 z_d^2 \sqrt{\theta_{ab} \theta_{bc} \theta_{ac} \theta_{ad}}. \quad (22)$$

By itself, Eq. (22) only has an AUC of 0.8031, but when used as the seventh feature of an NN that also uses the previously identified 6HL observables,

$$7\text{HL}_{\text{black-box}} \equiv \text{NN} \left[M_{\text{jet}}, \dots, \tau_2^{\beta=1}, \text{graph} \left(\kappa=2, \beta=\frac{1}{2} \right) \right], \quad (23)$$

it closes the performance gap with the CNN by achieving $\text{AUC} = 0.9528 \pm 0.0003$, as previewed in Table I. Interestingly, this happens even though the ADO between the CNN and $7\text{HL}_{\text{black-box}}$ is only 0.971, implying that they still make inconsistent decisions around 3% of the time. So while the black box has guided the selection of a new HL observable that closes the AUC performance gap, the remaining ADO gap implies that there is additional information not being captured.

The remaining rows of Table II show portions of the ranked list of 7545 EFPs, ordered by their ADO values. Note that the statistical uncertainties on the ADO are large enough that the precise ranking is not so meaningful, though the overall trends are. One striking feature is that many observables have a similar ADO to Eq. (22), and they often feature $\kappa = 2$ and $\beta = \frac{1}{2}$. Recall that $\kappa = 2$ corresponds to IRC-unsafe EFPs, which suggests that IRC-unsafe information is valuable (though perhaps not uniquely so) for mapping the CNN strategy. Similarly, the appearance of $\beta = \frac{1}{2}$ suggests the importance of probing small-angle behavior. Other IRC-unsafe EFPs with $\kappa = 0$ and $\kappa = -1$ also perform well, especially the constituent multiplicity appearing third on this list. The best performing IRC-safe $\kappa = 1$ observable appears 5531th on this list and is not able to close the performance gap with the CNN. Specifically, the EFPs in Eqs. (18) and (19) with $\kappa = 1$ have a relatively small ADO in the X_6 subspace, never getting above 0.5279. This is as one might expect, since these observables already effectively appear in the C_2 and D_2 combinations. Further discussions of the physics implications are provided below.

For completeness, we show distributions for the top three EFPs from Table II in Fig. 4, both in the full space as well as in X_6 . The first two observables show good separation between signal and background in the full space, as expected given that their AUC is around 0.8. The third observable, constituent multiplicity, is a relatively poor discriminant by itself. When restricted to the X_6 subspace, there is only modest residual separation power shown by these three observables, but enough to close the performance gap with the CNN.

B. Physics interpretation

Our first physics conclusion is that the κ -augmented EFP space is sufficiently comprehensive to close the performance gap between the 6HL and the CNN. Had we restricted our attention to just the IRC-safe EFPs, this would not have been the case, since the top ranked $\kappa = 1$ EFP in both strategies can only achieve $\text{AUC} = 0.9507$ when combined with the six previous HL observables.

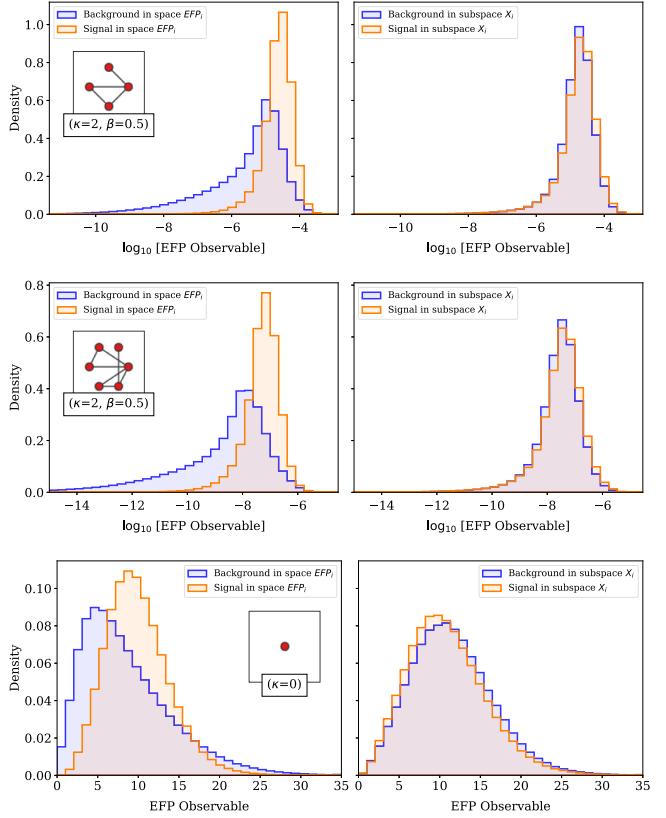


FIG. 4. The top three EFPs from the black-box guided search for a seventh HL observable; see Table II. Shown are the EFP distributions for signal and background events, both in the full set of events X_{all} (left column) and in X_6 (right column), i.e., the differently ordered space between the 6HL and the CNN. The top two observables, while not identical, have very similar functional forms, up to an overall rescaling. The third observable is the jet constituent multiplicity.

Thus, IRC-unsafe information seems essential for closing the performance gap.⁵

Fascinatingly, $\kappa = 2$ appears prominently in the top ten EFPs, though in a different form than previously considered in the literature. The key feature of the $\kappa = 2$ EFPs is that they weight higher energy particles more than lower energy particles. Looking at the top $\kappa = 2$ observables in Table II, they all have the common feature of corresponding to chromatic number $c = 3$ graphs. Chromatic number is the minimum number of colors needed to decorate the nodes of a graph such that no edge connects same-color nodes. If an EFP has chromatic number c , then it is only nonzero if the jet has at least c distinct particle directions, making it an effective probe for deviations from a $(c - 1)$ -prong substructure. The $\kappa = 2$ and $c = 3$ EFPs found by our guided

⁵As discussed in Ref. [108], a CNN is formally IRC-safe. To map this IRC-safe behavior to the EFPs, however, would require very high-point correlators, with in principle as many nodes as pixels in the original jet image. With IRC-unsafe EFPs, we can instead match the CNN decision boundaries with low-point correlators.

strategy therefore probe IRC-unsafe deviations from 2-prong substructure (as one might expect for boosted W boson tagging), with a particular emphasis on the higher energy particles inside the jet.

By contrast, the only $\kappa = 2$ observable that has received any significant attention in the jet substructure literature is p_T^D [65,66]. In the EFP language, p_T^D is a $c = 1$ graph with no edges:

$$\bullet^{(\kappa=2)} = \sum_{a=1}^N z_a^2. \quad (24)$$

Here, we see that p_T^D is only ranked 5610th by the ADO. Apparently, generic IRC-unsafe information is not, by itself, useful for boosted W boson classification, but must be paired with the correct angular dependence to highlight the physics of interest. It is interesting that $\beta = \frac{1}{2}$ is preferred as the angular exponent, since this choice appeared previously in the context of the Les Houches angularity for quark/gluon discrimination [68].

There are also $\kappa = 0$ observables in the top ten EFPs, including the well-known constituent multiplicity:

$$\bullet^{(\kappa=0)} = \sum_{a=1}^N 1. \quad (25)$$

The fact that a $\kappa = 0$, $c = 1$ observable yields nearly the same performance as a class of $\kappa = 2$, $c = 3$ observables is a surprising result of our study. Our tentative interpretation is that this represents two complementary approaches to solving this jet classification task. On the one hand, boosted W bosons are 2-prong objects, so one expects $c = 3$ observables to be most relevant. Indeed, the numerators of Eqs. (16) and (17) are $c = 3$ graphs that probe 2-prong substructure, which was part of the original motivation for the C_2 and D_2 observables. On the other hand, the background quark and gluon jets are 1-prong objects, and constituent multiplicity is well-known to be a powerful quark/gluon discriminant [68] (though sensitive to detector effects [109]). The next $\kappa = 0$ observable on the list has $c = 2$ and $\beta = \frac{1}{2}$, which is an IRC-unsafe probe of 1-prong substructure with an emphasis on collinear physics, which should also be an effective quark/gluon discriminant. This suggests that one can improve the classification performance either with a refined probe of the W boson signal or a refined probe of the quark/gluon background, which happens to have a similar effect on the decision boundaries.

In summary, by translating an ML strategy into a human-readable space, we have identified an important class of jet substructure observables that have been missing from previous boosted W boson studies. This motivates further studies of IRC-unsafe observables, especially high degree EFPs with $\kappa = 2$. In Sec. VI, we discuss the implications of this result on future work with jet substructure observables.

V. ITERATIVELY MAPPING FROM MINIMAL FEATURES

In the previous section, our aim was to supplement an existing set of HL features with one new feature to bridge the gap with the CNN performance. This jet substructure case study is unusual, however, in that it benefits from a highly mature literature of theoretically motivated features. Other applications of our black-box guided strategy may have to begin from a more minimal starting point and build an HL classification strategy essentially from scratch.

In this section, we start from just the most basic jet properties—transverse momentum p_T and jet mass M_{jet} —and iteratively identify a small set of EFPs relevant for boosted W boson classification. Using the black-box guided strategy, we are able to match the CNN performance using around seven EFPs. This is a similar dimensionality to the 7HL combination (which did not include p_T), though the physics features being probed will turn out to be interestingly different. We then show that this black-box strategy is more computationally efficient than a brute force search and more effective than a label-guided search.

A. Black-box guiding

Here, we apply the same black-box approach as Sec. IV A, starting just from the smaller set of observables, p_T and M_{jet} . The motivation for this starting point is as follows. The W boson mass at 80.4 GeV is one of the most important (and obvious) features of boosted W bosons. Because of the choice of z_a variable in Eq. (14), though, the EFPs are dimensionless. Therefore, we need at least one dimensionful HL observable to capture the W boson mass peak, and either p_T or M_{jet} would suffice for this purpose.

We begin from both p_T and M_{jet} for two reasons. The first is that they are ubiquitous jet observables appearing in myriad collider studies. The second is to streamline the selection of the EFPs. Naively, M_{jet} could be derived from p_T using the EFP in Eq. (15) with $\kappa = 1$ and $\beta = 2$:

$$\bullet^{(\kappa=1, \beta=2)} \approx \frac{M_{\text{jet}}^2}{p_T^2}. \quad (26)$$

With the choice of the θ_a variable in Eq. (15), though, Eq. (26) is only approximately true, so multiple EFPs are needed to map out the mass information if p_T is the only dimensionful scale. We checked that the black-box strategy is still effective starting from just p_T or from just M_{jet} , but the chosen EFPs tend to be more masslike in their structure. By contrast, starting from both p_T and M_{jet} yields more variation in the types of EFPs selected.

We start by training an NN on just the p_T and M_{jet} information:

$$\text{HLN}_0 \equiv \text{NN}[p_T, M_{\text{jet}}]. \quad (27)$$

This yields an AUC of 0.9119, which is substantially below the CNN performance for boosted W boson tagging. We then restrict our attention to the subset of events that are differently ordered by these minimal features relative to the CNN:

$$X_0 = \{(x, x') | \text{DO}[\text{CNN}, \text{HLN}_0](x, x') = 0\}. \quad (28)$$

The ADO between the CNN and HLN_0 is 0.9150, so X_0 contains 8.5% of the original X_{all} sample, though we only consider a random subset of 5×10^7 pairs in X_0 to reduce the computational burden. Our aim is to find a set of EFPs that can order these signal and background events the same as the CNN decision boundaries.

To identify the n -th EFP, we use the black-box guided strategy of Sec. II B, adapted to the current notation:

$$\text{EFP}_n = \text{argmax}_{\text{EFP} \in \mathcal{S}} \text{ADO}[\text{CNN}, \text{EFP}]_{X_{n-1}}. \quad (29)$$

We construct a new joint classifier that includes this EFP:

$$\text{HLN}_n \equiv \text{NN}[p_{\text{T}}, M_{\text{jet}}, \text{EFP}_1, \dots, \text{EFP}_n]. \quad (30)$$

This allows us to identify the remaining differently-ordered subset of events as

$$X_n = \{(x, x') | \text{DO}[\text{CNN}, \text{HLN}_n](x, x') = 0\}, \quad (31)$$

where in each iteration we only keep a random subset of 5×10^7 pairs. The main computational cost in this procedure is in training the joint classifier in Eq. (30).

The AUC and ADO values for this black-box guided procedure are shown in Fig. 5 versus the EFP scan iteration, and in Fig. 6 versus the cumulative computing time. More details about the selected EFPs are given in Table III. By the 5th iteration, the AUC performance matches that of the 6HL combination. By the seventh iteration, the AUC performance matches that of the CNN with an ADO of 0.974, indicating closer agreement with the CNN decisions than we found with the 7HL_{black-box} strategy. Since we started from a minimal set of jet features, it is not surprising that the EFPs identified here are qualitatively different from the ones in Sec. IV. The physics interpretation of these various EFPs will be presented in Sec. V D.

B. Comparison to brute force search

An alternative approach to maximizing the ADO is to perform a brute force search through the space of EFPs to find the set that maximally matches the decisions of the CNN. This is much more computationally expensive than the black-box guided strategy, but it has the potential to converge to a smaller number of EFPs if there are important correlations between the observables. In an absolute brute force search, one would construct all possible sets of EFPs, and evaluate the ADO of each relative to the CNN; given

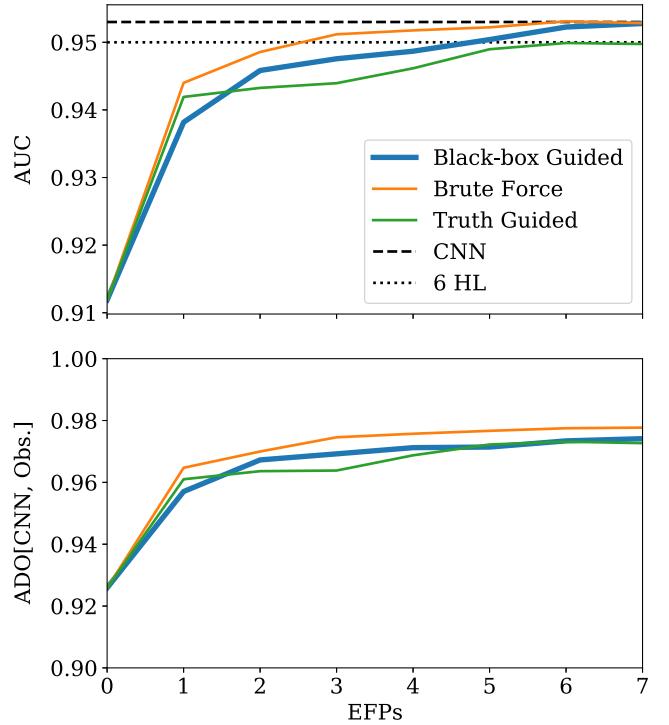


FIG. 5. Performance of the black-box guided search strategy to map the CNN solution into human-interpretable observables. Here, we start from just the basic jet features p_{T} and M_{jet} and iteratively add one EFP at a time. The performance is shown in terms of the AUC (top) and the ADO (bottom) as a function of the scan number. The performance of a brute force scan of the EFP space (Sec. V B) and a truth-label guided search (Sec. V C) are also shown. For reference, the performance of the CNN and of the existing 6HL features are indicated by horizontal lines.

the number of graphs and combinations, this would be completely intractable. Instead, we attempt an iterative greedy algorithm, which incrementally builds the EFP set. This is still computationally expensive, but (borderline) tractable.

We again start from the jet p_{T} and M_{jet} information, but immediately train a joint classifier using each of the EFPs as an input:

$$\text{NN}[p_{\text{T}}, M_{\text{jet}}, \text{EFP}]. \quad (32)$$

We then select the EFP that yields the largest ADO with the CNN, evaluated on the full training set. In the first iteration, we select the EFP via

$$\text{EFP}_1 = \text{argmax}_{\text{EFP} \in \mathcal{S}} \text{ADO}[\text{CNN}, \text{NN}[p_{\text{T}}, M_{\text{jet}}, \text{EFP}]]_{X_{\text{all}}}. \quad (33)$$

We repeat this procedure in each subsequent iteration, choosing the EFP that yields the largest improvement in the ADO when combined with the previous observables:

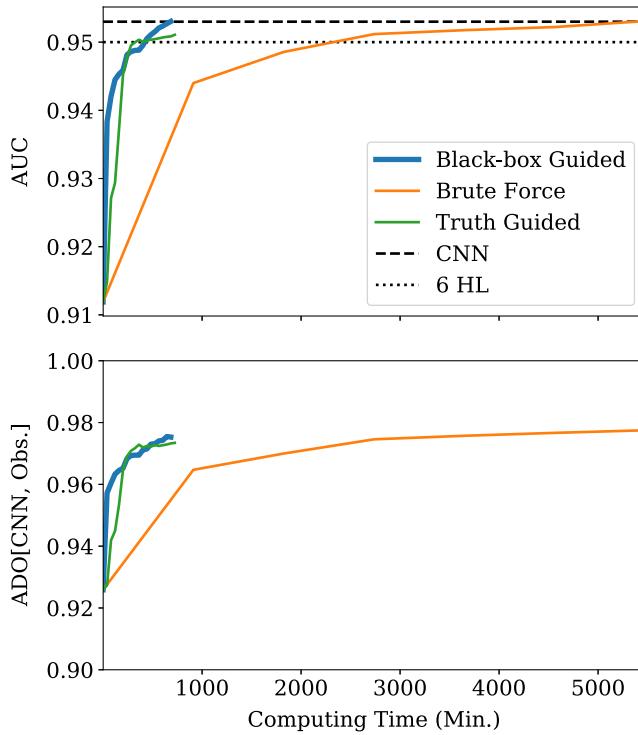


FIG. 6. The same as Fig. 5, but now plotted in terms of the cumulative computing time.

$$\text{EFP}_n = \operatorname{argmax}_{\text{EFP} \in S} \text{ADO}[\text{CNN}, \text{NN}[\dots, \text{EFP}_{n-1}, \text{EFP}]]_{X_{\text{all}}}. \quad (34)$$

The key difference from the black-box guided strategy is that the joint classifier is trained before evaluating the ADO, and the ADO is evaluated on the full training set, instead of just the differently-ordered subset.

The primary computational cost of the brute force approach comes from training the joint classifier appearing in Eq. (34), which combines each EFP with the current set of observables. This has to be done for each EFP under consideration, and it is too computationally expensive to

examine all 7545 EFP graphs over multiple iterations. Therefore, we only consider a subset of graphs at each iteration, which means there is no guarantee that the brute force method will perform better than the black-box guided strategy. For our purposes, our subset consists of the 54 connected graphs of degree $d \leq 5$ and (κ, β) choices of $\kappa = [\frac{1}{2}, 1, 2]$ and $\beta = [\frac{1}{2}, 1, 2]$. This reduces our original search space down to a more manageable 486 choices.

The results from this brute force procedure are shown in Fig. 5 in terms of the ADO and AUC values after each iteration. In the first few iterations, the AUC and ADO values are higher than for black-box guiding, achieving a comparable performance to the original 6HL result after the inclusion of a third EFP. The brute force process continues until it matches the CNN performance with six EFPs (eight HL inputs total). As one would expect, the brute force approach performs well as it is effectively trying every possible combination of inputs and selecting the best. This computational cost, however, must be weighed against the marginal decrease in the number of EFPs required to match the CNN as well as the need to restrict the input space prior to exploring the performance. As shown in Fig. 6, the brute force approach does not complete even a single iteration before the guided approaches have converged to a complete solution.

Finally, for completeness, we consider the alternative brute force approach of a network trained with all available EFPs. Using the baseline DNN architecture in Appendix A1, a network was trained with M_{jet} , p_{T} , and the reduced set of 486 EFPs as input features. The performance of “488HL” is shown in Table I, with a marginally better performance than the CNN, indicating that the EFPs are effectively a complete basis for this task.

C. Comparison to truth-label guiding

In the black-box guided strategy, the CNN and the ADO similarity metrics are auxiliary tools to help identify a set of EFPs that maximizes the classification performance. Assuming the EFP space is sufficiently complete and

TABLE III. The EFPs selected during each iteration of the black-box guiding strategy beginning from HLN_0 , which uses just p_{T} and M_{jet} . For each iteration, the selected EFP is the one with the largest ADO with the CNN in the differently-ordered subspace X_{n-1} .

Iteration (n)	EFP	κ	β	Chrom #	$\text{ADO}[\text{EFP}, \text{CNN}]_{X_{n-1}}$	$\text{AUC}[\text{EFP}]$	$\text{ADO}[\text{HLN}_n, \text{CNN}]_{X_{\text{all}}}$	$\text{AUC}[\text{HLN}_n]$
0	$M_{\text{jet}} + p_{\text{T}}$	—	—	—	—	—	0.9259	0.9119
1		2	$\frac{1}{2}$	2	0.8144	0.8190	0.9570	0.9382
2		0	2	2	0.6377	0.8106	0.9673	0.9458
3		0	—	1	0.5460	0.6737	0.9692	0.9476
4		1	$\frac{1}{2}$	2	0.5274	0.8464	0.9712	0.9487
5		-1	—	1	0.5450	0.5882	0.9714	0.9504
6		1	$\frac{1}{2}$	4	0.5382	0.7678	0.9734	0.9523
7		-1	$\frac{1}{2}$	2	0.5561	0.5957	0.9741	0.9528

labeled samples exist, one could dispense with the CNN entirely and simply search the space of EFPs for the most powerful set that maximizes the AUC, in an approach similar to that of Ref. [110]. As a computationally tractable alternative to the brute force search, we test what happens when the selection of the EFP is guided by the truth labels, instead of by the ADO.

Analogously to decision ordering in Eq. (2), we define *truth ordering* (TO) for a pair of signal/background points x and x' and a decision function f as

$$\text{TO}[f](x, x') = \Theta(f(x) - f(x')), \quad (35)$$

where 1 corresponds to f , correctly ordering the points, and 0 corresponds to inverted ordering. Starting again from the jet p_T and M_{jet} information, we identify the subset of event pairs that are incorrectly ordered:

$$Y_0 = \{(x, x') | \text{TO}[\text{HLN}_0](x, x') = 0\}. \quad (36)$$

In each iteration, we find the EFP that has the highest AUC in the incorrectly-ordered subspace,

$$\text{EFP}_n = \underset{\text{EFP} \in \mathcal{S}}{\text{argmax}} \text{AUC}[\text{EFP}]_{Y_{n-1}}, \quad (37)$$

construct a new joint classifier $\text{HLN}_n \equiv \text{HLN}_0 + n\text{EFP}$, and identify the next incorrectly-ordered subset of events,

$$Y_n = \{(x, x') | \text{TO}[\text{HLN}_n](x, x') = 0\}. \quad (38)$$

Note that this procedure is completely independent of the CNN.

The results from this truth-label guided procedure are shown in Fig. 5 in terms of the AUC and ADO. In the first iteration, the classification performance is better than in the black-box guided search, which makes sense since the label guided method is trying to optimize AUC directly. After seven iterations, though, the classification performance never rises above $\text{AUC} = 0.951$. As mentioned in Sec. II B, isolating the incorrectly-ordered pairs turns out to be counter productive, since some of these pairs could never be ordered correctly even by the optimal classifier. This emphasizes the value of using the ADO relative to an already-trained network, to make sure attention is focused on event pairs that have a chance to be correctly ordered.

D. Physics interpretation

By translating the CNN into a space of physically-motivated observables, we can gain physical insight into the observables used in the classification decision. In particular, the first few observables in Table III give us a glimpse at a possible alternative history for the field of jet substructure, if combinations like C_2 and D_2 had not been previously identified. Distributions of the EFPs found in the first four iterations are shown in Fig. 7.

After p_T and M_{jet} , the first EFP selected by the black-box guided strategy is

$$\text{EFP} \quad (\kappa=2, \beta=\frac{1}{2}). \quad (39)$$

The fact that a $\kappa = 2$ observable shows up early in the iterative procedure bolsters the evidence from Sec. IV A that these kinds of observables are important for mapping the CNN strategy. This is a chromatic number $c = 2$ graph, so just like jet mass, it probes deviations from 1-prong substructure. However, it uses a 5-point correlator (unlike mass which is a 2-point correlator) and it uses the $\beta = \frac{1}{2}$ angular exponent (unlike mass which uses $\beta = 2$). Putting these together, Eq. (39) is an IRC-unsafe probe of hard, small-angle radiation.

The second EFP is also IRC-unsafe and also corresponds to a $c = 2$ graph:

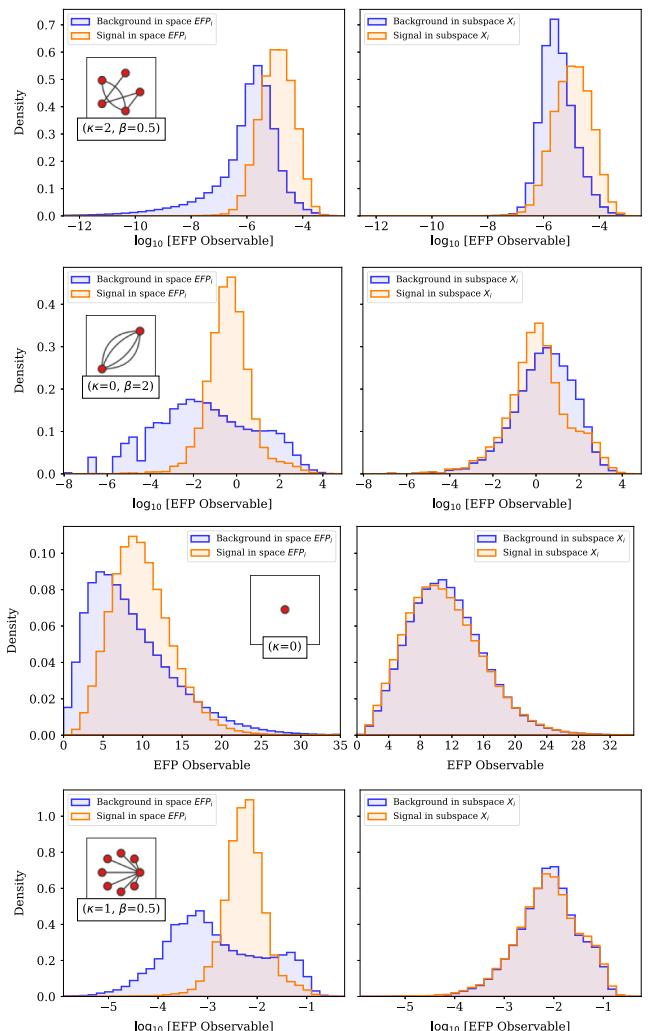


FIG. 7. The first four EFP graphs selected by the black-box guided strategy beginning from the minimal set of HL observables, p_T and M_{jet} ; see Table III. Shown are the EFP distributions for signal and background events, both in the full set of events X_{all} (left column) and in X_n (right column), i.e., the differently ordered space between the HLN_n and the CNN after n iterations.

$$(40)$$

Here, though, we have $\kappa = 0$ and $\beta = 2$, which is a probe of soft wide-angle radiation. It is interesting that the black-box guided strategy selects these two complementary $c = 2$ observables in the first two iterations, indicating the importance of 1-prong substructure probes even if the goal is to identify 2-prong boosted W bosons.

The third EFP is constituent multiplicity, as seen before in Eq. (25), which reinforces the idea that controlling the composition of the quark/gluon background is important for W boson tagging. These three observables, together with p_T and M_{jet} , yield an AUC of 0.9476. This is not as good as the 6HL combination, but still quite encouraging given that we did not give the black-box guided strategy any information about the ratio structures used to construct C_2 and D_2 .

The main surprise from this study is that IRC-safe information was not selected by the black-box guided search until the fourth iteration,

$$(41)$$

Moreover, it is a $c = 2$ graph, so still a probe of 1-prong substructure. Only in interaction six do we finally see a higher chromatic number graph, but the guided search skips over the C_2/D_2 -like graphs with $c = 3$ and goes straight to $c = 4$. The black-box guided strategy has identified a very different strategy for boosted W boson tagging that nevertheless matches the 6HL combination with a comparable number of observables.

One interpretation of this result is that it simply reflects the “entropy” of our HL space. There are four times as many IRC-unsafe observables in our HL collection than IRC-safe ones, so just by random chance one expects to see more unsafe observables in the scan. Indeed, there are IRC-safe observables that are highly ranked in the first three iterations, just not at the top of the list. Another interpretation is that the black-box guided strategy is teaching us that IRC-unsafe information is more relevant for boosted W boson tagging than one might naively think. A related observation was made in Ref. [85], which introduced a color ring observable to identify color-singlet configurations. Intriguingly, when restricted to three particles, the angular structure of Eq. (39) defines similar decision boundaries to the color ring.⁶ Either way, by searching through a large space of HL observables in a systematic way, the black-box guided strategy has given us a new perspective on an old problem in a human-readable format.

VI. DISCUSSION

The ever increasing complexity of new ML strategies has produced better classification performance for various physics problems. At the same time, the increasing opaqueness of these methods has widened the gap between our understanding of a problem and our appreciation of the ML solution. In this paper, we have proposed a new technique for mapping an ML solution into a space of human-interpretable observables. Our guided strategies mitigate some of the well-founded concerns about black-box approaches, while still allowing us to capitalize on the black-box performance to efficiently guide the selections of HL observables. The end result is a set of HL observables that have a more direct physical interpretation and allow for a more transparent treatment of systematic uncertainties.

In our jet substructure case study, we have shown that the black-box guided strategy could be used to isolate information that is not captured by previous HL representations. Remarkably, only a single observable was needed to close the performance gap identified in Ref. [20], nearly duplicating the CNN strategy with a low-dimensional input representation. Beginning from a minimal set of basic jet observables (p_T and M_{jet}), we successfully condensed the CNN behavior to a small set of EFP observables which reproduce its performance and very nearly match its decisions. It would be interesting to study the utility of the EFPs in more complicated contexts, such as event-wide classification tasks.

Interestingly, the structure of the selected EFPs differ in qualitative ways from the C_2 and D_2 jet substructure observables custom designed for boosted W boson classification. While these previous observables are based on fully connected graphs, the guided strategy picked out multinode EFP graphs with a relatively low chromatic number. While these previous observables use the IRC-safe choice of $\kappa = 1$, the guided strategy emphasized the importance of unsafe $\kappa = 2$ observables, particularly ones with nontrivial angular dependence. This motivates further physics studies of these exotic EFPs. It is worth emphasizing that we were only able to identify these new observables because we considered a sufficiently large space of HL observables. There may be other hidden organizing principles to exploit for jet substructure studies, which motivates the construction of alternative sets of observables based on different physical principles than the EFPs. In particular, we did not capitalize on the power counting and scaling properties of ratio/product observables [64,104,110–112], which may reveal more efficient HL observables for jet classification. It may also be beneficial to leverage first-principles knowledge about signal/background likelihood ratios [85,113–118] to identify promising HL observables.

These results have important implications for what we should regard as the “best practices” of the application of ML methods to high-energy physics problems. Primarily,

⁶We thank Andrew Larkoski for discussions of this point.

we should be more wary of utilizing opaque ML strategies which obscure how a problem is solved in exchange for relatively small classification improvements. In general, one should evaluate whether “additional” information captured by DNNs represents genuine patterns or is the byproduct of something unintentionally pressed into the data during simulation and then rediscovered by the network.

The informational gap in our benchmark problem could be closed using a single HL observable, suggesting that the CNN strategy was not relying on subtle correlations among the low-level features, but rather exploiting information encodable into a $\kappa = 2$ EFP. Thus, instead of a purely performance-oriented approach, we suggest a strategy of using deep networks to establish performance benchmarks, but always seek to translate ML strategies into a more tractable space of well-motivated physical observables. If this proves to be impossible or impractical, it might be that the ML approach really is identifying genuinely new information, or more likely, that the space of physical observables needs to be augmented or optimized.

More broadly, our view is that the ultimate goal of ML research in high-energy physics should not be to develop artificial-intelligence physicists which (or should we say *who?*) can blindly process raw data and make statements about the Universe without being able to communicate the intermediate steps. The power of modern ML can certainly be used to identify gaps in our knowledge where existing human-engineered approaches are insufficient. At the end of the day, though, we should insist that data analysis strategies used to make statements about physics should be understandable to human physicists.

ACKNOWLEDGMENTS

We thank Pierre Baldi, Timothy Cohen, Michael Fenton, Patrick Komiske, Andrew Larkoski, Eric Metodiev, Benjamin Nachman, Tilman Plehn, Peter Sadowski, and Edison Weik for helpful conversations and feedback. J. T. was supported by the U.S. Department of Energy (DOE), Office of High Energy Physics under Grant No. DE-SC0012567. D. W. and T. F. are supported by the U.S. Department of Energy (DOE), Office of Science under Grant No. DE-SC0009920. T. F. was supported by the National Science Foundation under Award No. 1633631.

APPENDIX: NETWORK ARCHITECTURES AND HYPERPARAMETERS

For consistency, all neural networks, regardless of input data type, use a set of common settings and training procedures:

- (1) $N = 5 \times 10^6$ training samples, broken down as:
 - (a) 70% training,
 - (b) 15% validation,
 - (c) 15% testing.

- (2) Training samples are pre-processed via SCI-KIT’s `StandardScaler`.
- (3) Adam optimizer used with default TENSORFLOW settings:
 - (a) `learning_rate` = 0.001,
 - (b) `beta1` = 0.9,
 - (c) `beta2` = 0.999,
 - (d) `epsilon` = $1e-07$,
 - (e) `amsgrad` = `False`.
- (4) Output layer uses a sigmoid activation.
- (5) Uncertainties are calculated via a 10-fold cross-validation (see Appendix A 3).
- (6) Early stopping on `validation_loss` with a patience of 30 epochs.
- (7) Model checkpoints saved (for best results only) on minimum validation loss.

1. Baseline convolutional neural network

The following training details are specific to all convolutional neural networks trained on jet-images:

- (1) Input data undergoes a log transformation prior to `StandardScaler` preprocessing.
- (2) Hidden layers consist of five hidden convolutional layers with the parameters:
 - (a) 500 nodes,
 - (b) `kernel_size` = (4, 4),
 - (c) `strides` = (1, 1),
 - (d) `padding` = ‘valid’,
 - (e) `kernel_initializer` = ‘glorot_normal’,
 - (f) `activation` = ‘relu’,
 - (g) `kernel_constraint` = `max_norm(3)`.
- (3) Three dropout layers (i.e., one between each convolutional layer) with a rate of 0.2.

2. Baseline dense neural network

The following training details are specific to all dense neural networks acting on jet substructure observables, including EFPs:

- (1) Hidden layers consist of three hidden dense layers with the parameters:
 - (a) 300 nodes,
 - (b) `activation` = ‘relu’,
 - (c) `kernel_constraint` = `max_norm(3)`.
- (2) Two dropout layers (i.e., one between each dense layer) with a rate of 0.5.

3. K-fold validation

To derive uncertainties on the trained model prediction accuracy, we use the bootstrap cross-validation package in SCI-KIT to equally divide the test set ten times and measure the performance across ten bootstrapping iterations. Averages and standard deviations are then taken from these ten iterations to define the central value and uncertainties of the AUC.

[1] D. Guest, K. Cranmer, and D. Whiteson, Deep learning and its application to LHC physics, *Annu. Rev. Nucl. Part. Sci.* **68**, 161 (2018).

[2] A. Radovic, M. Williams, D. Rousseau, M. Kagan, D. Bonacorsi, A. Himmel, A. Aurisano, K. Terao, and T. Wongjirad, Machine learning at the energy and intensity frontiers of particle physics, *Nature (London)* **560**, 41 (2018).

[3] J. K. Kohne *et al.*, Realization of a second level neural network trigger for the H1 experiment at HERA, *Nucl. Instrum. Methods Phys. Res., Sect. A* **389**, 128 (1997).

[4] G. Aad *et al.* (ATLAS Collaboration), A neural network clustering algorithm for the ATLAS silicon pixel detector, *J. Instrum.* **9**, P09009 (2014).

[5] C. Peterson, Track finding with neural networks, *Nucl. Instrum. Methods Phys. Res., Sect. A* **279**, 537 (1989).

[6] H. Abramowicz, A. Caldwell, and R. Sinkus, Neural network based electron identification in the ZEUS calorimeter, *Nucl. Instrum. Methods Phys. Res., Sect. A* **365**, 508 (1995).

[7] V. Khachatryan *et al.* (CMS Collaboration), Observation of the diphoton decay of the Higgs boson and measurement of its properties, *Eur. Phys. J. C* **74**, 3076 (2014).

[8] V. M. Abazov *et al.* (D0 Collaboration), First measurement of $\sigma(p\bar{p} \rightarrow Z) \cdot \text{Br}(Z \rightarrow \tau\tau)$ at $\sqrt{s} = 1.96$ TeV, *Phys. Rev. D* **71**, 072004 (2005); Erratum, *Phys. Rev. D* **77**, 039901 (2008).

[9] DELPHI Collaboration, Classification of the hadronic decays of the Z^0 into b and c quark pairs using a neural network, *Phys. Lett. B* **295**, 383 (1992).

[10] DZero Collaboration, Search for single top quark production at $d\bar{q}$ using neural networks, *Phys. Lett. B* **517**, 282 (2001).

[11] CDF Collaboration, Search for the Higgs Boson Using Neural Networks in Events with Missing Energy and b -Quark Jets in $p\bar{p}$ Collisions at $\sqrt{s} = 1.96$ TeV, *Phys. Rev. Lett.* **104**, 141801 (2010).

[12] P. Baldi, P. Sadowski, and D. Whiteson, Searching for exotic particles in high-energy physics with deep learning, *Nat. Commun.* **5**, 4308 (2014).

[13] P. Baldi, P. Sadowski, and D. Whiteson, Enhanced Higgs Boson to $\tau^+\tau^-$ Search with Deep Learning, *Phys. Rev. Lett.* **114**, 111801 (2015).

[14] R. Santos, M. Nguyen, J. Webster, S. Ryu, J. Adelman, S. Chekanov, and J. Zhou, Machine learning techniques in searches for $t\bar{t}h$ in the $h \rightarrow b\bar{b}$ decay channel, *J. Instrum.* **12**, P04014 (2017).

[15] A. Aurisano, A. Radovic, D. Rocco, A. Himmel, M. D. Messier, E. Niner, G. Pawloski, F. Psihas, A. Sousa, and P. Vahle, A convolutional neural network neutrino event classifier, *J. Instrum.* **11**, P09001 (2016).

[16] T. Cohen, M. Freytsis, and B. Ostdiek, (Machine) learning to do more with less, *J. High Energy Phys.* **02** (2018) 034.

[17] M. Andrews, M. Paulini, S. Gleyzer, and B. Poczos, End-to-end event classification of high-energy physics data, *J. Phys. Conf. Ser.* **1085**, 042022 (2018).

[18] L. G. Almeida, M. Backović, M. Cliche, S. J. Lee, and M. Perelstein, Playing tag with ANN: Boosted top identification with pattern recognition, *J. High Energy Phys.* **07** (2015) 086.

[19] L. de Oliveira, M. Kagan, L. Mackey, B. Nachman, and A. Schwartzman, Jet-images—Deep learning edition, *J. High Energy Phys.* **07** (2016) 069.

[20] P. Baldi, K. Bauer, C. Eng, P. Sadowski, and D. Whiteson, Jet substructure classification in high-energy physics with deep neural networks, *Phys. Rev. D* **93**, 094034 (2016).

[21] 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, *Phys. Rep.* **841**, 1 (2020).

[22] G. Kasieczka, T. Plehn, M. Russell, and T. Schell, Deep-learning top taggers or the end of QCD?, *J. High Energy Phys.* **05** (2017) 006.

[23] D. Guest, J. Collado, P. Baldi, S.-C. Hsu, G. Urban, and D. Whiteson, Jet flavor classification in high-energy physics with deep neural networks, *Phys. Rev. D* **94**, 112002 (2016).

[24] Identification of jets containing b -hadrons with recurrent neural networks at the ATLAS experiment, CERN Tech. Report No. ATL-PHYS-PUB-2017-003, 2017.

[25] CMS Collaboration, Heavy flavor identification at CMS with deep neural networks, CERN Tech. Report No. CMS-DP-2017-005, 2017.

[26] A. M. Sirunyan *et al.* (CMS Collaboration), Identification of heavy-flavour jets with the CMS detector in pp collisions at 13 TeV, *J. Instrum.* **13**, P05011 (2018).

[27] P. T. Komiske, E. M. Metodiev, and M. D. Schwartz, Deep learning in color: Towards automated quark/gluon jet discrimination, *J. High Energy Phys.* **01** (2017) 110.

[28] A. Andreassen, P. T. Komiske, E. M. Metodiev, B. Nachman, and J. Thaler, OmniFold: A Method to Simultaneously Unfold All Observables, *Phys. Rev. Lett.* **124**, 182001 (2020).

[29] K. Datta, D. Kar, and D. Roy, Unfolding with generative adversarial networks, *arXiv:1806.00433* [J. High Energy Phys. (to be published)].

[30] M. Bellagente, A. Butter, G. Kasieczka, T. Plehn, A. Rousselot, R. Winterhalder, L. Ardizzone, and U. Kthe, Invertible networks or partons to detector and back again, *SciPost Phys.* **9**, 074 (2020).

[31] C. Englert, P. Galler, P. Harris, and M. Spannowsky, Machine learning uncertainties with adversarial neural networks, *Eur. Phys. J. C* **79**, 4 (2019).

[32] 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**, 014018 (2017).

[33] S. Bollweg, M. Haußmann, G. Kasieczka, M. Luchmann, T. Plehn, and J. Thompson, Deep-learning jets with uncertainties and more, *SciPost Phys.* **8**, 006 (2020).

[34] B. Nachman, A guide for deploying deep learning in LHC searches: How to achieve optimality and account for uncertainty, *SciPost Phys.* **8**, 090 (2020).

[35] G. Kasieczka, M. Luchmann, F. Otterpohl, and T. Plehn, Per-object systematics using deep-learned calibration, *SciPost Phys.* **9**, 089 (2020).

[36] P. T. Komiske, E. M. Metodiev, and J. Thaler, Energy flow polynomials: A complete linear basis for jet substructure, *J. High Energy Phys.* **04** (2018) 013.

[37] E. M. Metodiev, B. Nachman, and J. Thaler, Classification without labels: Learning from mixed samples in high energy physics, *J. High Energy Phys.* **10** (2017) 174.

[38] A. Andreassen, I. Feige, C. Frye, and M. D. Schwartz, JUNIPR: A framework for unsupervised machine learning in particle physics, *Eur. Phys. J. C* **79**, 102 (2019).

[39] J. H. Collins, K. Howe, and B. Nachman, Anomaly Detection for Resonant New Physics with Machine Learning, *Phys. Rev. Lett.* **121**, 241803 (2018).

[40] C. Buciluundefined, R. Caruana, and A. Niculescu-Mizil, Model compression, in *Proceedings of the 12th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD 06* (Association for Computing Machinery, New York, 2006), p. 535–541.

[41] J. Duarte *et al.*, Fast inference of deep neural networks in FPGAs for particle physics, *J. Instrum.* **13**, P07027 (2018).

[42] S. Chang, T. Cohen, and B. Ostdiek, What is the machine learning?, *Phys. Rev. D* **97**, 056009 (2018).

[43] A. A. Alemi, I. Fischer, J. V. Dillon, and K. Murphy, Deep variational information bottleneck, [arXiv:1612.00410](https://arxiv.org/abs/1612.00410).

[44] S. Wunsch, R. Friese, R. Wolf, and G. Quast, Identifying the relevant dependencies of the neural network response on characteristics of the input space, *Comput. Softw. Big Sci.* **2**, 5 (2018).

[45] T. Roxlo and M. Reece, Opening the black box of neural nets: case studies in stop/top discrimination, [arXiv:1804.09278](https://arxiv.org/abs/1804.09278).

[46] M. G. Kendall, A new measure of rank correlation, *Biometrika* **30**, 81 (1938).

[47] M. H. Seymour, Tagging a heavy Higgs boson, in *ECFA Large Hadron Collider (LHC) Workshop: Physics and Instrumentation* (1991), pp. 557–569, <https://inspirehep.net/literature/29197>.

[48] M. H. Seymour, Searches for new particles using cone and cluster jet algorithms: A comparative study, *Z. Phys. C* **62**, 127 (1994).

[49] J. M. Butterworth, B. E. Cox, and J. R. Forshaw, WW scattering at the CERN LHC, *Phys. Rev. D* **65**, 096014 (2002).

[50] J. M. Butterworth, J. R. Ellis, and A. R. Raklev, Reconstructing sparticle mass spectra using hadronic decays, *J. High Energy Phys.* **05** (2007) 033.

[51] 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**, 242001 (2008).

[52] A. Abdesselam *et al.*, Boosted objects: A probe of beyond the standard model physics, *Eur. Phys. J. C* **71**, 1661 (2011).

[53] A. Altheimer *et al.*, Jet substructure at the tevatron and LHC: New results, new tools, new benchmarks, *J. Phys. G* **39**, 063001 (2012).

[54] J. Shelton, Jet substructure, in *Theoretical Advanced Study Institute in Elementary Particle Physics: Searching for New Physics at Small and Large Scales* (2013), pp. 303–340, https://doi.org/10.1142/9789814525220_0007.

[55] A. Altheimer *et al.*, Boosted objects and jet substructure at the LHC. Report of BOOST2012, held at IFIC Valencia, 23rd-27th of July 2012, *Eur. Phys. J. C* **74**, 2792 (2014).

[56] D. Adams *et al.*, Towards an understanding of the correlations in jet substructure, *Eur. Phys. J. C* **75**, 409 (2015).

[57] R. Kogler *et al.*, Jet substructure at the large hadron collider: Experimental review, *Rev. Mod. Phys.* **91**, 045003 (2019).

[58] S. Marzani, G. Soyez, and M. Spannowsky, *Looking Inside Jets: An Introduction to Jet Substructure and Boosted-Object Phenomenology* (Springer, New York, 2019), Vol. 958.

[59] Y. Cui, Z. Han, and M. D. Schwartz, W-jet tagging: Optimizing the identification of boosted hadronically-decaying W bosons, *Phys. Rev. D* **83**, 074023 (2011).

[60] J. Thaler and K. Van Tilburg, Identifying boosted objects with N-subjettiness, *J. High Energy Phys.* **03** (2011) 015.

[61] J. Thaler and K. Van Tilburg, Maximizing boosted top identification by minimizing N-subjettiness, *J. High Energy Phys.* **02** (2012) 093.

[62] A. J. Larkoski, G. P. Salam, and Jesse Thaler, Energy correlation functions for jet substructure, *J. High Energy Phys.* **06** (2013) 108.

[63] G. Aad *et al.* (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**, 113013 (2014).

[64] A. J. Larkoski, I. Moult, and D. Neill, Power counting to better jet observables, *J. High Energy Phys.* **12** (2014) 009.

[65] F. Pandolfi, Search for the standard model Higgs boson in the $H \rightarrow ZZ \rightarrow l^+l^-q\bar{q}$ decay channel at CMS, Ph.D. thesis, Zurich, ETH, New York, 2012.

[66] S. Chatrchyan *et al.* (CMS Collaboration), Search for a Higgs boson in the decay channel $H \rightarrow ZZ^{(*)} \rightarrow q\bar{q}\ell^-\ell^+$ in pp collisions at $\sqrt{s} = 7$ TeV, *J. High Energy Phys.* **04** (2012) 036.

[67] A. J. Larkoski, J. Thaler, and W. J. Waalewijn, Gaining (Mutual) information about quark/gluon discrimination, *J. High Energy Phys.* **11** (2014) 129.

[68] P. Gras, S. Hche, D. Kar, A. Larkoski, L. Lnnblad, S. Pltzer, A. Sidmok, P. Skands, G. Soyez, and J. Thaler, Systematics of quark/gluon tagging, *J. High Energy Phys.* **07** (2017) 091.

[69] R. E. Schapire, The strength of weak learnability, *Mach. Learn.* **5**, 197 (1990).

[70] Y. Freund, Boosting a weak learning algorithm by majority, *Inf. Comput.* **121**, 256 (1995).

[71] P. T. Komiske, E. M. Metodiev, and J. Thaler, Energy flow networks: Deep sets for particle jets, *J. High Energy Phys.* **01** (2019) 121.

[72] H. Qu and L. Gouskos, ParticleNet: Jet tagging via particle clouds, *Phys. Rev. D* **101**, 056019 (2020).

[73] E. A. Moreno, O. Cerri, J. M. Duarte, H. B. Newman, T. Q. Nguyen, A. Periwal, M. Pierini, A. Serikova, M. Spiropulu, and J.-R. Vlimant, JEDI-net: A jet identification algorithm based on interaction networks, *Eur. Phys. J. C* **80**, 58 (2020).

[74] V. Mikuni and F. Canelli, ABCNet: An attention-based method for particle tagging, *Eur. Phys. J. Plus* **135**, 463 (2020).

[75] A. Bogatskiy, B. Anderson, J. T. Offermann, M. Roussi, D. W. Miller, and R. Kondor, Lorentz Group Equivariant neural network for particle physics, [arXiv:2006.04780](https://arxiv.org/abs/2006.04780).

[76] J. Shlomi, S. Ganguly, E. Gross, K. Cranmer, Y. Lipman, H. Serviansky, H. Maron, and N. Segol, Secondary vertex finding in jets with neural networks, [arXiv:2008.02831](https://arxiv.org/abs/2008.02831).

[77] J. Cogan, M. Kagan, E. Strauss, and A. Schwartzman, Jet-images: Computer vision inspired techniques for jet tagging, *J. High Energy Phys.* **02** (2015) 118.

[78] D. E. Kaplan, K. Rehermann, M. D. Schwartz, and B. Tweedie, Top Tagging: A Method for Identifying Boosted Hadronically Decaying Top Quarks, *Phys. Rev. Lett.* **101**, 142001 (2008).

[79] L. G. Almeida, S. J. Lee, G. Perez, G. F. Sterman, I. Sung, and J. Virzi, Substructure of high- p_T jets at the LHC, *Phys. Rev. D* **79**, 074017 (2009).

[80] S. D. Ellis, C. K. Vermilion, and J. R. Walsh, Recombination algorithms and jet substructure: Pruning as a tool for heavy particle searches, *Phys. Rev. D* **81**, 094023 (2010).

[81] T. Plehn, M. Spannowsky, M. Takeuchi, and D. Zerwas, Stop reconstruction with tagged tops, *J. High Energy Phys.* **10** (2010) 078.

[82] S. D. Ellis, A. Hornig, T. S. Roy, D. Krohn, and M. D. Schwartz, Qjets: A Non-Deterministic Approach to Tree-Based Jet Substructure, *Phys. Rev. Lett.* **108**, 182003 (2012).

[83] M. Dasgupta, A. Fregoso, S. Marzani, and G. P. Salam, Towards an understanding of jet substructure, *J. High Energy Phys.* **09** (2013) 029.

[84] A. J. Larkoski, S. Marzani, G. Soyez, and J. Thaler, Soft drop, *J. High Energy Phys.* **05** (2014) 146.

[85] A. Buckley, G. Callea, A. J. Larkoski, and S. Marzani, An optimal observable for color singlet identification, *SciPost Phys.* **9**, 026 (2020).

[86] G. Kasieczka, T. Plehn, A. Butter *et al.*, The machine learning landscape of top taggers, *SciPost Phys.* **7**, 014 (2019).

[87] K. Datta and A. Larkoski, How much information is in a jet?, *J. High Energy Phys.* **06** (2017) 073.

[88] L. Moore, K. Nordström, S. Varma, and M. Fairbairn, Reports of my demise are greatly exaggerated: N -subjettiness taggers take on jet images, *SciPost Phys.* **7**, 036 (2019).

[89] J. A. Aguilar-Saavedra and B. Zaldvar, Jet tagging made easy, *Eur. Phys. J. C* **80**, 530 (2020).

[90] J. Alwall, M. Herquet, F. Maltoni, O. Mattelaer, and T. Stelzer, MADGRAPH 5: Going beyond, *J. High Energy Phys.* **06** (2011) 128.

[91] T. Sjostrand, S. Mrenna, and P. Z. Skands, PYTHIA 6.4 physics and manual, *J. High Energy Phys.* **05** (2006) 026.

[92] J. de Favereau, C. Delaere, P. Demin, A. Giannmanco, V. Lemaître, A. Mertens, and M. Selvaggi (DELPHES 3 Collaboration), DELPHES 3, A modular framework for fast simulation of a generic collider experiment, *J. High Energy Phys.* **02** (2014) 057.

[93] M. Cacciari, G. P. Salam, and G. Soyez, The Anti- $k(t)$ jet clustering algorithm, *J. High Energy Phys.* **04** (2008) 063.

[94] M. Cacciari, G. P. Salam, and G. Soyez, FASTJET user manual, *Eur. Phys. J. C* **72**, 1896 (2012).

[95] D. Krohn, J. Thaler, and L.-T. Wang, Jet trimming, *J. High Energy Phys.* **02** (2010) 084.

[96] G. Aad *et al.* (ATLAS Collaboration), Identification of boosted, hadronically decaying W bosons and comparisons with ATLAS data taken at $\sqrt{s} = 8$ TeV, *Eur. Phys. J. C* **76**, 154 (2016).

[97] M. Aaboud *et al.* (ATLAS Collaboration), Performance of top-quark and W -boson tagging with ATLAS in Run 2 of the LHC, *Eur. Phys. J. C* **79**, 375 (2019).

[98] V. Khachatryan *et al.* (CMS Collaboration), Identification techniques for highly boosted W bosons that decay into hadrons, *J. High Energy Phys.* **12** (2014) 017.

[99] Quark versus Gluon Jet Tagging Using jet images with the ATLAS detector, CERN Tech. Rep. No. ATL-PHYS-PUB-2017-017 (2017).

[100] S. Macaluso and D. Shih, Pulling out all the tops with computer vision and deep learning, *J. High Energy Phys.* **10** (2018) 121.

[101] A. Banfi, G. P. Salam, and G. Zanderighi, Principles of general final-state resummation and automated implementation, *J. High Energy Phys.* **03** (2005) 073.

[102] G. Gur-Ari, M. Papucci, and G. Perez, Classification of energy flow observables in narrow jets, [arXiv:1101.2905](https://arxiv.org/abs/1101.2905).

[103] M. Jankowiak and A. J. Larkoski, Jet substructure without trees, *J. High Energy Phys.* **06** (2011) 057.

[104] I. Moult, L. Necib, and J. Thaler, New angles on energy correlation functions, *J. High Energy Phys.* **12** (2016) 153.

[105] P. T. Komiske, <https://pypi.org/project/EnergyFlow/>.

[106] S.-M. Udrescu and M. Tegmark, AI Feynman: A physics-inspired method for symbolic regression, *Sci. Adv.* **6**, eaay2631 (2020).

[107] P. T. Komiske, E. M. Metodiev, and J. Thaler, Cutting multiparticle correlators down to size, *Phys. Rev. D* **101**, 036019 (2020).

[108] S. Choi, S. J. Lee, and M. Perelstein, Infrared safety of a neural-net top tagging algorithm, *J. High Energy Phys.* **02** (2019) 132.

[109] G. Kasieczka, N. Kiefer, T. Plehn, and J. M. Thompson, Quark-gluon tagging: Machine learning vs detector, *SciPost Phys.* **6**, 069 (2019).

[110] K. Datta and A. J. Larkoski, Novel jet observables from machine learning, *J. High Energy Phys.* **03** (2018) 086.

[111] K. Datta, A. Larkoski, and B. Nachman, Automating the construction of jet observables with machine learning, *Phys. Rev. D* **100**, 095016 (2019).

[112] A. M. Sirunyan *et al.* (CMS Collaboration), Identification of heavy, energetic, hadronically decaying particles using machine-learning techniques, *J. Instrum.* **15**, P06005 (2020).

[113] D. E. Soper and M. Spannowsky, Finding physics signals with shower deconstruction, *Phys. Rev. D* **84**, 074002 (2011).

[114] D. E. Soper and M. Spannowsky, Finding top quarks with shower deconstruction, *Phys. Rev. D* **87**, 054012 (2013).

[115] D. E. Soper and M. Spannowsky, Finding physics signals with event deconstruction, *Phys. Rev. D* **89**, 094005 (2014).

[116] D. F. de Lima, P. Petrov, D. Soper, and M. Spannowsky, Quark-gluon tagging with shower deconstruction: Un-

earthing dark matter and Higgs couplings, [Phys. Rev. D **95**, 034001 \(2017\)](#).

[117] A. J. Larkoski and E. M. Metodiev, A theory of quark vs. gluon discrimination, [J. High Energy Phys. 10 \(2019\) 014](#).

[118] G. Kasieczka, S. Marzani, G. Soyez, and G. Stagnitto, Towards machine learning analytics for jet substructure, [J. High Energy Phys. 09 \(2020\) 195](#).