Boosted W and Z tagging with jet charge and deep learning

Yu-Chen Janice Chen¹, Cheng-Wei Chiang^{1,2,3}, Giovanna Cottin^{1,4,5} and David Shih⁶

¹*Department of Physics, National Taiwan University, Taipei 10617, Taiwan*

²*Institute of Physics, Academia Sinica, Taipei 11529, Taiwan*

³*Department of Physics and Center of High Energy and High Field Physics, National Central University, Chungli 32001, Taiwan*

⁴*Instituto de Física, Pontificia Universidad Católica de Chile, Avenida Vicuña Mackenna 4860, Santiago, Chile*

⁵*Departamento de Ciencias, Facultad de Artes Liberales, Universidad Adolfo Ibáñez, Diagonal Las Torres 2640 Santiago, Chile*

⁶*NHETC, Dept. of Physics and Astronomy, Rutgers, The State University of New Jersey, Piscataway, New Jersey 08854, USA*

(Received 30 August 2019; accepted 20 February 2020; published 17 March 2020)

We demonstrate that the classification of boosted, hadronically decaying, weak gauge bosons can be significantly improved over traditional cut-based and boosted decision tree-based methods using deep learning and the jet charge variable. We construct binary taggers for W^+ vs W^- and Z vs W discrimination, as well as an overall ternary classifier for $W^+/W^-/Z$ discrimination. Besides a simple convolutional neural network, we also explore a composite of two simple convolutional neural networks, with different numbers of layers in the jet p_T and jet charge channels. We find that this novel structure boosts the performance particularly when considering the Z boson as a signal. The methods presented here can enhance the physics potential in Standard Model measurements and searches for new physics that are sensitive to the electric charge of weak gauge bosons.

DOI: [10.1103/PhysRevD.101.053001](https://doi.org/10.1103/PhysRevD.101.053001)

I. INTRODUCTION

Boosted heavy resonances play a central role in the study of physics at the Large Hadron Collider (LHC). These include both Standard Model particles such as W 's, Z 's, tops and Higgses, as well as hypothetical new physics (NP) particles such as Z' 's. The decay products of the boosted heavy resonance are typically collimated into a single “fat jet” with a nontrivial internal substructure. A vast amount of effort has been devoted to the important problem of “tagging” (i.e., identifying and classifying) boosted resonances through the understanding of jet substructure. (For recent reviews and original references, see, e.g., [1–3].)

Recently, there has been enormous interest in the application of modern deep learning techniques to boosted resonance tagging [4–24]. By enabling the use of high-dimensional, low-level inputs (such as jet constituents), deep learning automates the process of feature engineering. Many works have demonstrated the enormous potential of

deep learning to construct extremely powerful taggers, vastly improving on previous methods.

So far, most of the attention has focused on distinguishing various boosted resonances from QCD background in a binary classification task. Less attention has been paid to multiclass classification, i.e., a tagger that would categorize jets in a multitude of possibilities. (Notable exceptions include Refs. [23,24].) In this work, we will examine an important multiclass classification task: distinguishing W^+ , W^- , and Z bosons. Having a W/Z classifier that can also recognize charge could have many interesting applications. For instance, one could use such a classifier to measure charge asymmetries and same-sign diboson production at the LHC. Or, there are many potential applications to NP scenarios, such as the reconstruction of doubly charged Higgs bosons from its like-sign diboson decay in models with an extended scalar sector.

Since we are interested in distinguishing W^+ and W^- bosons from each other, a key element in our work will be the jet charge observable Q_{jet} . It was first introduced in Ref. [25] and its theoretical potential was discussed further in Ref. [26]. Such an observable has also been measured at the LHC [27,28]. When used in conjunction with other quantities, such as the invariant mass $\mathcal{M}$, it can help distinguish between hadronically decaying W from Z bosons [29]. Moreover, its performance as a charge tagger

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International](https://creativecommons.org/licenses/by/4.0/) 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.

TABLE I. Selections imposed on the jet sample used in our analyses.

Jet sample	$p_T \in (350, 450)$ GeV, $ \eta \leq 1$
	Jets with anti- k_T and $R = 0.7$
	$V - V$ merging: $\Delta R(V_1, V_2) < 0.6$
	V -jet matching: $\Delta R(V, j) < 0.1$

was assessed in Ref. [30] for jets produced in semileptonic $t\bar{t}$, W + jets, and dijet processes. Most recently, the authors of Ref. [31] incorporated jet charge into various machine learning quark/gluon taggers, including boosted decision trees (BDTs), convolutional neural networks (CNNs), and recurrent neural networks. They showed that including jet charge in the input channels improved quark/gluon discrimination and up vs down quark discrimination.

In our study of $W^+/W^-/Z$ tagging, we will compare a number of techniques, from simple cut-based methods, to BDTs, to deep learning methods based on CNNs and jet images. As in Ref. [31], we will include jet charge as one of the input channels and examine the gain in performance from including this additional input. We will go beyond Refs. [15,29,31] and construct a ternary classifier that can discriminate among W^+ , W^- , and Z , depending on the physics process of interest. We will study the overall performance of our ternary tagger as well as its specialization to binary classification. For the latter we will compare its performance to specifically trained binary classifiers and show that the ternary classifier reproduces their performance, and in this sense it is optimal. Overall, we will demonstrate that deep learning with jet charge offers a significant boost in performance, around $\sim 30\%$ – 40% improvement in background rejection rate at fixed signal efficiency.

In addition to a simple CNN, we will also develop a novel composite algorithm consisting of two CNNs, one for each of the Q_κ and p_T channels, combined in a merge layer, which we refer to as CNN². This allows us to separately optimize the hyperparameters of the CNNs for the two input channels. We show that this new CNN² architecture further boosts the performance for most combinations.

The rest of the paper is organized as follows. In Sec. II we describe the jet samples and jet images used in this study, and review the definition of the jet charge variable.

TABLE II. Summary of the jet sample sizes used for training and testing, after the selections in Table I. The entries in the sum correspond to the (W^+ , W^- , Z) samples, respectively. In the training stage, the training set is further divided into two subsets: 9/10 is the actual training set and 1/10 serves as the validation set. All the receiver operating characteristic (ROC) and significance improvement characteristic (SIC) curves shown in this paper are evaluated using the testing sets described here.

	Training set	Testing set
Jet sample size	$188k + 198k + 175k$	$38k + 40k + 35k$

FIG. 1. Reconstructed jet mass of W and Z samples.

In Sec. III, we describe the different taggers studied in this paper. These include cut-based and BDT taggers used as baselines for comparison, as well as two different taggers based on convolutional neural networks. We show results for a binary W^-/W^+ classification problem in Sec. IV and compare our performance with the recent work in Ref. [31]. In Sec. V, we discuss the Z/W^+ discrimination problem, focusing on the benefit from including jet charge, and compare our performance with the ATLAS boson tagger in Ref. [29]. Finally in Sec. VI, we extend our results to the full ternary $Z/W^+/W^-$ classification problem, and comment on the reduction from a three-class tagger to a two-class one. In Sec. VII, we attempt to shed some light on what the deep neural networks learned. We summarize our findings and conclude in Sec. VIII.

II. JET SAMPLES AND INPUTS

For this study, we use MADGRAPH5v2.6.1 [32] at leading order to simulate events at the 13 TeV LHC for vector boson fusion production of doubly charged Higgses $H_5^{\pm\pm}$ and heavy neutral Higgses H_5 , with decays $H_5^{\pm\pm} \rightarrow W^\pm W^\pm \rightarrow jjjj$ and $H_5 \rightarrow ZZ \rightarrow jjjj$, respectively. We take the spectrum of the exotic Higgs bosons in the Georgi-Machacek model generated by GMCalc [33] as the input. For simplicity, we fix $m_{H_5} = 800$ GeV, so that the p_T of each vector boson is typically ~ 400 GeV.¹ All events are further processed in PYTHIA 8.2.19 [34] for showering and hadronization and passed onto DELPHES 3.4.1 [35] with the default CMS card for detector simulation.² Jets are reconstructed with FASTJET 3.1.3 [36] using the anti- k_T clustering algorithm [37] with the jet radius parameter $R = 0.7$. This jet

¹We have also studied the scenario where $m_{H_5} = 2$ TeV, which leads to the p_T of each boson around 1 TeV, and found all results qualitatively the same.

²A dataset based on HERWIG showering and hadronization is also generated for the purpose of checking the reliability of the deep learning jet-tagging results. Details are given in the Appendix.

FIG. 2. Q_κ distributions for the three samples under study. Representative κ values are shown.

radius is appropriate to the p_T range indicated above, as it will capture most of the hadronic W and Z decay products.

The jets selected in this work are required to satisfy $|\eta| \leq 1$. The jets must also be truth matched to a W/Z boson by requiring their distance between the jet and the vector boson, known from the Monte Carlo truth, in the (η, ϕ) plane be less than 0.1. A summary detailing the full set of selections for our jet sample is given in Table I. The jet sample sizes of the training and testing sets for our taggers are summarized in Table II.

In what follows, we consider a number of inputs in the construction of our BDT and CNN taggers.

A. Jet mass and jet charge ($\mathcal{M}$, Q_κ) observables

For the problem at hand, there are two obvious high-level observables: jet invariant mass and jet charge. The jet invariant mass $\mathcal{M}$ is defined as

$$\mathcal{M}^2 = \left(\sum_{i \in J} E_i \right)^2 - \left(\sum_{i \in J} \mathbf{p}_i \right)^2. \quad (1)$$

The sum runs over all constituents i inside the jet J (i.e., all tracks and calorimeter towers in the jet) with four-momentum $(E_i, \mathbf{p}_i)$ and $p_T > 500$ MeV.

Figure 1 shows the reconstructed jet mass after all selections in Table I. The broader widths in the mass

distribution originate from a combination of showering, hadronization, jet clustering, and detector effects.

It has been known for some time [25] that a useful observable for distinguishing jets initiated by particles of different charges is the jet charge:

$$Q_\kappa = \frac{1}{(p_{T,J})^\kappa} \sum_{i \in J} q_i \times (p_T^i)^\kappa, \quad (2)$$

where q_i corresponds to the integer charge of the jet constituent in units of the proton charge, and κ is a free parameter. The Q_κ observable is computed in this p_T -weighted scheme to minimize mismeasurements from low- p_T particles.

Figure 2 shows the Q_κ distributions for jets coming from the W^+ , W^- , and Z samples. Distributions are shown for different κ values. The choice of κ together with the p_T range of the vector bosons will affect the tagging performance.

B. Jet images

As described in the Introduction, the deep learning based taggers studied in this work are based on jet images and convolutional neural networks. In this work, our jet images are made from jets reconstructed in a $\Delta\eta = \Delta\phi = 1.6$ box, with 75×75 pixels. This choice yields a resolution

TABLE III. Summary of the configurations of our CNN taggers.

	CNN	CNN ²	
Image	(75 × 75) pixels within ($ \eta \leq 0.8$, $ \phi \leq 0.8$)		
Channels	p_T, Q_κ	p_T	Q_κ
Architecture	BN-32C6-MP2-128C4-MP2-256C6-MP2-512N-512N	BN-32C3-32C3-MP2-64C3-MP2-64C3-MP2-64C3-64C3-128C5-256C5-256N-256N	BN-32C3-32C3-MP2-64C4-64C4-MP2-256C6-MP2-256N
Settings	Relu activation, Padding = same, Dropout = 0.5, 12 Regularizer = 0.01		
Preprocessing	<i>Centralization, rotation, flipping</i>		
Training	Adam optimizer, Minibatchsize = 512, <i>cross entropy</i> loss

FIG. 3. The left plot shows the average of W^+ jet images in the p_T channel using the preprocessed testing set sample. The right plot shows the difference between Z and W^+ average jet images in the p_T channel.

consistent with that of the CMS ECal (see Table III). The input variables or channels are now Q_κ and p_T per pixel. Therefore, the sum in the Q_κ definition in Eq. (2) in this case goes over all jet constituents in each pixel.

To improve the performance of our taggers, we preprocess each image, following a similar procedure as in Ref. [16]: *centralization*, *rotation*, and *flipping*. In Figs. 3 and 4, we use ϕ' and η' to denote the new coordinate system for the images after preprocessing. Note that we do

not perform *normalization* in the preprocessing, though it is another common preprocessing operation. Instead, we introduce a BatchNormalization layer [38] and observe a comparable (or even better) performance.

Figure 3 shows the average of jet images in the p_T channel after preprocessing. For comparison, we show the W^+ average jet images in left plot and the difference between Z and W^+ average jet images in the right plot. It is observed that the average Z jet image is wider in the η' - ϕ'

FIG. 4. Average of jet images in the Q_κ channel, with $\kappa = 0.15$, for the three jet samples in our testing sets, after preprocessing.

FIG. 5. The left plot shows the true distributions for $W^+/W^-/Z$ in the $(Q_\kappa, \mathcal{M})$ plane (for $\kappa = 0.3$). The T-shaped lines in the left plot mark the decision boundaries of the cut-based tagger. The right plot shows the output prediction of the ternary BDT classifier in the $(Q_\kappa, \mathcal{M})$ plane, whose Y-shaped color boundaries match our intuition for the optimal border.

plane than the average W jet image, as expected from their difference in the invariant mass.

Figure 4 shows the average of jet images in the Q_κ channel for a fixed κ value of 0.15 and after preprocessing. On average, the jet charge images are consistent with what we expect from the W^+ , W^- and Z -boson charges. In particular, the average Z -boson jet charge image is very close to zero, as the charges of the constituents tend to cancel out on the average.

III. METHODS

In the following, we will investigate possible classification tasks, including an overall ternary problem and binary ones (W^- vs W^+ and Z vs W^+). Throughout this paper,

unless otherwise specified, we will treat W^+ as the background in all binary classification tasks.

Here we describe the various taggers used in our work: baseline cut-based and BDT taggers that take high-level inputs $(\mathcal{M}, Q_\kappa)$, as well as taggers based on CNNs, which are trained on jet images formed by lower level inputs.

A. Cut-based and BDT taggers

We first construct a cut-based tagger. We will identify the optimal values of κ for the different binary discrimination tasks, W^- vs W^+ and Z vs W^+ . For the latter, we will explore all possible simple 2D rectangular cuts in the jet charge and jet mass plane for the optimal cut-based discriminator.

FIG. 6. CNN model architecture.

We will also study a “single- κ BDT.” It is built out of Q_κ with κ held fixed together with the jet mass $\mathcal{M}$. Both observables are fed into a gradient BDT implemented with the SKLEARN [39] package and assuming the default parameters. We can also combine models of different κ values to form another BDT tagger, dubbed “multi- κ BDT,” similar to the “multi- κ ” jet tagger constructed in Ref. [31]. For this multi- κ BDT, $\mathcal{M}$, Q_κ , and $\kappa = 0.2, 0.3, 0.4$ are specified as inputs. We find that the single- κ BDT, when taking the optimal κ value, generally has a comparable performance as the multi- κ BDT. In the following sections, the single- and multi- κ BDTs are shown as benchmark models to compare with our deep neural networks.

We use these BDT taggers in both binary and ternary classifications. The prediction of the ternary single- κ BDT classifier on the testing set can be visualized in Fig. 5. The three blobs (red, green, and blue) correspond to the three classes of jets (Z , W^- , W^+). The single- κ BDT nicely separates and distributes the jets as intuitively expected. That is, to mark out the border among the three classes, the best choice would be the Y-shaped cut in the two-dimensional plane.

B. CNN-based taggers

We now describe the architectures of two deep CNN models based on jet images developed in this study. We feed the two-color jet images (p_T , Q_κ), as described in Sec. II B, into the CNNs. We use the Keras library with the TensorFlow backend for the implementation of the networks.

Throughout the paper, the “CNN” label describes a network composed of three *convolutional layers* followed by two *fully connected layers*. The padding option is activated to enable the network to go deeper. To prevent overfitting, regularizers are used as well as *dropout layers*. For training, we use Adam [40] as the optimizer algorithm.³ In Fig. 6, we show the architecture of the CNN, with detailed model configuration parameters summarized in Table III.

The second deep neural network we have considered is a composite design consisting of two CNNs with asymmetrical depths, which we call “CNN².” The machine learning in the p_T and Q_κ channels is done in parallel; i.e., we feed the p_T images and the Q_κ images into two separate CNNs that differ in structure. These are combined at the end in a merge layer.

The CNN² architecture is motivated by our finding that the depth of the CNN network is limited by the classification task between W^+ and W^- . If the network is made too deep, the W^+/W^- tagger tends to overfit. This makes sense, since the W^+ and W^- jet images are identical except in the sign of the Q_κ channel. Therefore, it requires fewer

³We also tried the SGD optimizer and the performance is roughly the same as Adam.

FIG. 7. Detailed model architecture of CNN².

convolutional layers to capture the difference between the two. On the other hand, a deeper network structure does help a lot in successfully identifying the Z boson as the signal. The Z samples differ from the other two in the spatial distribution (substructure) of the constituents. If we enhance the resolution/ability of the CNN’s pattern recognition, it is natural that the CNN could do better in the Z discrimination. Therefore, there seems to be a trade-off in

FIG. 8. Summary of the performance as a function of κ in the range $[0.1, 0.6]$ for a binary classification task to discriminate W^- from W^+ for all taggers. The three metrics are the AUC (left), accuracy (middle), and background rejection (right).

the ternary classification problem. The CNN² architecture is an attempt to have our cake and eat it too.

After investigations on the model structure and seeing performance trends in the different classification problems, we have chosen the CNN² architecture detailed in Fig. 7 and Table III. The one dealing with the Q_κ images is shallower. Based on the observed performance, we keep on stacking up to eight convolutional layers for the other CNN that processes the p_T images.

IV. W^-/W^+ BINARY CLASSIFICATION

We begin with a study of the W^-/W^+ classification. Since the signal and background differ only in their charges, the only useful quantity here is the jet charge Q_κ . An important aspect of the jet charge variable defined in Eq. (2) is that it depends on a parameter κ which specifies how the contributing charges are p_T weighted. Nothing *a priori* tells us which value of κ to use, and different tasks may prefer different values of κ . In the following subsection we examine this issue for our various taggers.

A. Determining κ

In Fig. 8, we show the trend in performance when varying κ for the W^- vs W^+ classification. We show three typical metrics for evaluating the algorithm's performance: area under the curve (AUC), best accuracy (ACC), and background rejection at a 50% signal efficiency working point ($1/\epsilon_b|_{\epsilon_s=50\%}$, denoted by R50). Deep learning taggers as well as the cut-based reference tagger are plotted together for comparison. Note, in particular, that the performance of single- κ BDT, not shown in these plots, would trivially reduce to that of the cut-based tagger in the W^-/W^+ binary classification task, in which case the $\mathcal{M}$ information does not provide additional discriminative power.

We find in Fig. 8 that the performance does depend on the choice of κ . It is also interesting to see that the deep learning taggers have a qualitatively different κ dependence from the traditional (cut-based) taggers, while little difference exists between the two CNN models. The former are always better than the latter and have a smaller optimal κ .

Based on these results, we determine the optimal κ in each of the tagger definitions in the following sections. A value of $\kappa = 0.3$ is fixed for the single- κ BDT reference tagger, and $\kappa = 0.15$ for the CNN taggers.

B. Comparison of taggers

Having fixed the value of κ in each tagger, we are now ready to compare the various tagging methods. Table IV shows the performance metrics described above (AUC, ACC, and R50). In Fig. 9, the left plot shows the receiver operating characteristic (ROC) curves and the right plot shows the significance improvement characteristic (SIC) curves. While all the ROC curves seem to be close to one another, one can readily see a clear benefit from employing deep learning in the SIC curves. The improvement in background rejection rate is around 30%–40% across a wide range of signal efficiencies. It is also worth noting that for the W^-/W^+ task, there is no particular gain in using the CNN² structure; in fact, the performance is slightly worse.

Although differing in the details, it is still instructive to compare with the previous work on jet charge and deep learning [31], which focuses on up/down quark jet discrimination. Our performance gain from BDT to deep learning turns out to be quite comparable. For the 1000-GeV benchmark scenario, their CNN's background rejection rate at 50% signal efficiency grows by about 40% relative to their “ κ and λ BDT” reference tagger, as given in Table 1 of Ref. [31]. Our result shows the same amount of enhancement, as seen in Table IV.

TABLE IV. Performance metrics for all taggers, except for the single- κ BDT, in a W^-/W^+ binary classification task.

	R50	AUC	ACC
Cut-based	16.1372	0.8600	0.7811
Multi- κ BDT	16.0960	0.8615	0.7820
CNN	21.9559	0.8855	0.8042
CNN ²	20.5057	0.8800	0.8000

FIG. 9. ROC (left) and SIC (right) curves for the binary classification to discriminate W^- from W^+ for all taggers, except for the single- κ BDT.

V. Z/W^+ BINARY CLASSIFICATION

Next we turn to the Z vs W^+ binary classification task. (Z vs W^- would obviously have the same result because of charge symmetry.) Here we are primarily

interested in how much the jet charge observable adds to the discriminative power, compared to just the information computable from the four vectors (e.g., the jet mass).

FIG. 10. Same as Fig. 8, but for a binary classification task to discriminate Z from W^+ for all taggers.

FIG. 11. ROC (left) and SIC (right) curves for a binary classification discriminating Z from W^+ for all taggers.

TABLE V. Performance metrics for all taggers in a binary Z/W^+ classification task.

	R50	AUC	ACC
Cut-based	9.9590	0.8118	0.7705
Single- κ BDT	14.1638	0.8608	0.7875
Multi- κ BDT	14.2383	0.8611	0.7880
CNN	40.4205	0.9091	0.8345
CNN ²	52.6028	0.9206	0.8452

A. Determining κ

The performance of various taggers as a function of κ is shown in Fig. 10. Unlike the W^-/W^+ classification, the dependence in κ is very mild here. We will keep using the same values of optimal κ as before for simplicity.

B. Comparison of taggers

Figure 11 shows the ROC curves (left plot) and the SIC curves (right plot) of the different taggers in the binary task of distinguishing Z from W^+ . Table V gives the three performance metrics.

Compared to the W^-/W^+ classification task, the benefit from deep learning in the current case is much greater, with an improved background rejection rate at 50% signal efficiency of as much as ~ 2.85 . This is perhaps unsurprising, since our cut-based and BDT methods do not include any jet substructure variables. As noted before in Fig. 3, in the W^-/W^+ classification task, the samples have identical average distribution in the p_T images and differ only in pixel intensity in the Q_κ channel. However, the Z and W^+ events have distinct spatial distributions in the jet's

constituents, in addition to the charge difference. Since the CNN naturally learns spatial differences, it will naturally show a big improvement over methods that do not include any substructure information.

We also note that for the Z/W^+ task, the CNN² tagger with the architecture tailored for better Z identification outperforms the CNN tagger by a sizable amount.

We now scrutinize the role of jet charge in Z vs W^+ discrimination. Figure 12 shows the performance of taggers with either one or two input channels. We split our predictions from $\mathcal{M}$ to $(\mathcal{M}, Q_\kappa)$ in the BDT case, or from p_T to (p_T, Q_κ) for the CNNs. Comparing the solid curves to the dotted ones, we see that all the three taggers get significant improvements after adding the Q_κ information.

The role of jet charge in boosted Z vs W^+ discrimination was previously studied by ATLAS [29]. However, the ATLAS study is different from ours in the following details: it focuses on a different signal process, $W' \rightarrow WZ$; its jet samples are defined differently; and what they construct is a likelihood tagger with the $\mathcal{M}$ and Q_κ as the inputs. Nevertheless, it is instructive to compare our results to theirs. Throughout a wide range of working points, both the ATLAS tagger and our CNN tagger attain an additional 30% enhancement in the background rejection rate after further incorporating Q_κ information. In the high signal efficiency region, our taggers seem to enjoy a larger gain by introducing the Q_κ channel.

VI. TERNARY $W^+/W^-/Z$ CLASSIFICATION

Finally, we turn to the ultimate task of a full classification of boosted weak gauge bosons ($W^+/W^-/Z$). We will

FIG. 12. ROC (left) and SIC (right) curves for Z/W^+ binary classification using taggers with different numbers of input channels. The dotted lines are for one channel only, either $\mathcal{M}$ (for the reference cut-based and BDT taggers) or p_T (for the CNN taggers), and the solid lines are for those with two channels ($\mathcal{M} + Q_\kappa$ or $p_T + Q_\kappa$) and correspond to the results in the previous section. The single- κ BDT tagger is compared with the cut-based tagger because the BDT with only one input channel reduces to the cut-based tagger.

FIG. 13. ROC (left) and SIC (right) curves for a ternary classification discriminating W^- from (W^+, Z) for all the taggers. FIG. 14. ROC (left) and SIC (right) curves for a ternary classification discriminating Z from W s for all the taggers.

quantify how well the full ternary classification performs. Then we will see how to reduce to the binary classifications described in the previous sections. For the choice of κ for the CNN taggers, we will continue to use the same values of κ as before because they are also optimal for the ternary classifier.

We summarize and compare the performance of the ternary taggers according to two metrics: their overall accuracy, defined as (number of correct predictions)/(total number of instances),⁴ and a “one-against-all” metric which binarizes the task, i.e., singling out one class as the “signal” and treating all the others as the “background.”

⁴For BDT- and CNN-based taggers, the largest class probability is taken to be the prediction.

Figures 13 and 14 plot the ROC and SIC curves resulting from treating W^- and Z as the signals, respectively.⁵ The corresponding metrics are given in Table VI. Our results show that the CNN² generally performs better than CNN. Thanks to the parallel structure, the network has a sufficient depth for a sizable improvement in the Z -signal performance, while having comparable or better performance than the CNN in the W^- (or W^+) discrimination.

A. Comparison with binary taggers

We expect that a multiclass classification task should be able to fully recover the binary classification performance

⁵The W^+ -against-all metric is omitted here, because the results and tendencies are very comparable to the W^- -against-all metric by symmetry.

TABLE VI. Performance metrics for all taggers in the ternary classification task.

	Overall	Signal: W^-			Signal: Z		
	ACC	R50	AUC	ACC	R50	AUC	ACC
Cut-based	0.6581	8.0262	0.7893	0.7643	10.0882	0.8233	0.7839
Single- κ BDT	0.6667	12.5230	0.8339	0.7576	11.0726	0.8363	0.7725
Multi- κ BDT	0.6675	12.7115	0.8348	0.7579	11.0678	0.8366	0.7726
CNN	0.7197	17.3403	0.8715	0.7890	32.8981	0.8936	0.8170
CNN ²	0.7318	19.0907	0.8764	0.7950	42.1927	0.9088	0.8334

after an appropriate projection. If the multiclass NN output is supposed to approximate the class probability $P_i(x)$, where x is a data point and $i = 1, \dots, N$ is the class label, then the projection to binary classification between class i and class j is simply

$$P_i^{i \text{ or } j}(x) = \frac{P_i(x)}{P_i(x) + P_j(x)}. \quad (3)$$

In Fig. 15, we plot the results of this projection in solid curves. The dotted curves are the reproduced binary BDT and CNN results from Figs. 9 and 11 for comparison. In the left plot, the solid and dotted curves are almost on top of each other. For the Z/W^+ projection, however, we observe that the ternary CNN outperforms the binary CNN after the projection in the low signal efficiency region.

VII. WHAT DID THE MACHINE LEARN?

In this section we will attempt to shed some light on how our CNN and CNN² taggers learn to classify $W^+/W^-/Z$

bosons, and the differences between them. Although a complete understanding is not possible—they are still very much “black boxes”—we will find that with the help of some visualization techniques we can understand better what the machine has learned.

A. Saliency maps

Here we will use “saliency maps” [41] to compare the CNN and CNN² networks in an attempt to understand why the latter outperforms the former. We will use the toolkit of Keras-vis [42] to compute the saliency maps. Here the class saliency is extracted by computing the pixel-wise derivative of the class probability $P_i(x)$, as denoted in Eq. (3):

$$w_i = \frac{\partial P_i(x)}{\partial x}, \quad (4)$$

where the gradient is obtained by back propagation. By making a map of the gradients (4) across an image, we can identify the regions of the image where the decision of the CNN (to be class i or not) depends most sensitively.

FIG. 15. SIC curves of the ternary classification for a W^-/W^+ discrimination (left) and for a Z/W^+ discrimination (right), when projected to binary according to Eq. (3). The dashed curves are for binary classifications, and the solid curves are for the projected ternary results.

FIG. 16. Saliency maps for the p_T channel of the CNN (upper plots) and CNN² (lower plots) networks on nine W^- jet images from the test sample (for which both networks give correct output predictions).

The saliency maps for nine W^- jet images from the test sample are shown in Figs. 16 (the p_T channel) and 17 (the Q_κ channel) for the CNN and the CNN² networks. The color in each pixel indicates the magnitude of the gradient value.

The difference in the saliency maps between CNN and CNN² is very striking. We can see that the attention of the CNN² is generally concentrated on much smaller regions than the CNN. Evidently, the resolving power of the CNN² is much better than that of the CNN. Given that the CNN²

FIG. 17. Same as Fig. 16 but for the Q_κ channel.

network goes much deeper than the CNN network, this is perhaps expected: a deeper network structure with more convolutional layers is supposed to be more capable of capturing more subtle features in the training data. Altogether, this could explain why the CNN² mostly outperforms the CNN.

B. Phase transition in CNN's learning

Another interesting difference between our networks is in their learning curves. We observe that the learning curve of the CNN always has a sudden jump in the performance. Such a *phase transition* in learning comes from the fact that the CNN tends to first learn

FIG. 18. Accuracy of a one-against-all metric, ACC, at different callback points for CNN (left) and CNN² (right). A phase transition in ACC during the CNN training stage occurs around the 25th epoch.

characteristics of the Z sample, and then those of the W^+ (or W^-) sample.

In order to investigate this phase transition behavior, we monitor the network performance during its intermediate stages. During the training, we set a check point every three epochs and record the network's weights, then analyze to see how the discrimination ability evolves. At each check point, we evaluate the network performance using one-against-all metrics on the testing jet samples. The results of the ACC metric are shown in Fig. 18. We can see that the CNN develops the ability of Z discrimination at an earlier stage. However, the network has not learned how to discriminate between W^- and W^+ until after the 24th epoch, around which a phase transition is seen in their learning curves.

Even though CNN² shows a more “steady” learning process, the possibility of phase transition phenomenon cannot be ruled out. It is possible that the CNN² learns so fast that the performance in all classes saturates within one epoch and, therefore, the phase transition is not manifest.

VIII. CONCLUSIONS

In this work, we apply modern deep learning techniques to build better taggers of boosted, hadronically decaying weak gauge bosons. We demonstrate and provide the results for the boosted weak bosons with $p_T \sim 400$ GeV throughout this paper. (We have also studied the scenario with even more boosted W and Z bosons, with $p_T \sim 1$ TeV, and found results similar to what are presented here.) Going beyond previous works, we incorporate jet charge information in order to discriminate between positively and negatively charged W bosons, and between W and Z bosons. We study all

possible binary classification tasks as well as the full ternary classification problem. Taking BDT and cut-based taggers as our baselines for comparison, we construct a simple CNN tagger that takes jet images as the input, and show that it leads to significant gains in classification accuracy and background rejection.

In addition to the simple CNN tagger, we also construct a novel CNN structure consisting of two parallel CNNs, which we call CNN². The key feature of this structure is to assign different network depths to each of the p_T and Q_κ channels. This further improves the performance of nearly all the classification tasks.

We see various ways in which our work could be extended and improved. First, traditional CNNs may have some drawbacks due to the fact that the receptive field of every neuron, which is the field of view that one unit can perceive, is fixed by the assigned kernel sizes and depth of the network. But the complexity in detecting the patterns or features in realistic problems generally differ. As the $W^-/W^+/Z$ discrimination problem in our analysis shows, the complexity/depth for dealing with W^-/W^+ and W/Z is different. To optimize the performance as well as the computational costs, a “ResNet” network architecture with “skip connections” [19,43,44] may be a desirable solution. It would be interesting to study this further, along with other architectures and jet representations such as point clouds and sequences.

Another direction where it may be possible to improve on this work is to come up with an architecture that takes p_T and charge information as totally separate channels, such that the network could learn the ideal combination of them for measuring the charge of the boosted heavy resonance. The fact that our tagger performance depends on the value of κ is a symptom

that it is not truly learning the optimal combination of p_T and charge information.

Comparing the performance gains due to the deep learning seen in this work with recent related works in the literature [29,31], we have seen similar improvements over more conventional methods. We caution that these are merely rough comparisons, as the nature and details of the classification problems are different from ours. Nevertheless, this gives further evidence for the enormous potential of deep learning for the study of jet substructure and boosted resonance tagging.

ACKNOWLEDGMENTS

We gratefully acknowledge the support of the NVIDIA Corporation with the donation of the Titan Xp GPU used for this research. We thank Gregor Kasieczka and Ben Nachman for helpful discussions. Y.-C.C. thanks Yi-Ting Lee for the advice on the technical details in deep learning. C.-W.C. was supported in part by the Ministry of Science and Technology (MOST) of Taiwan under Grant No. MOST-104-2628-M-002-014-MY4. G.C. acknowledges support by Grant No. MOST-107-2811-M-002-3120 and CONICYT-Chile FONDECYT Grant No. 3190051. The work of D.S. was supported by the U.S. Department of Energy under Grant No. DE-SC0010008.

APPENDIX: CHECK OF SOFT SENSITIVITY USING POWER SHOWERING FROM HERWIG

In this appendix, we show the performance of our taggers on a separate set of jet samples using a different

TABLE VII. Summary of the Herwig-jet sample sizes used for training and testing, after the selections in Table I. The entries in the sums correspond to the (W^+, W^-, Z) samples, respectively. In the training stage, the training set is also divided into two subsets as described in Table II.

	Training set	Testing set
Jet sample size	191k + 201k + 176k	38k + 40k + 35k

parton shower model. This is to study whether our deep learning jet-tagging method is vulnerable to the modeling of parton showering and hadronization. For this dataset, the parton-level hard process events generated by MADGRAPH5, the detector simulation carried out by DELPHES and the clustering procedure by FastJet are controlled to be identical to those introduced in Sec. II. The only difference is that here we use HERWIG 7.1.5 [45,46] for showering and hadronization. Detailed tuning is referred to the CMS software collection.

For the quality cuts, the same selection criteria (Table I) are performed on the raw jet data. The number of jets for later training and testing purpose are listed in Table VII. And overall, the jet variables and jet images of this Herwig-jet sample have been checked to be consistent with those obtained from the Pythia-jet sample.

In the following, the results of CNN and CNN², whose structures are described in Table III, working on the Herwig-jet sample for binary W^-/W^+ , Z/W^+ , and ternary classifications are demonstrated. For Figs. 19 and 20, the SIC curves for the Herwig dataset (darker colors) are overlapped on top of the Pythia results (lighter colors), which are already shown in Figs. 9, 11 and 13, 14. And the corresponding performance metrics are provided in

FIG. 19. SIC curves for the binary classifications to discriminate W^- from W^+ (left) and Z from W^+ (right) using the CNN and CNN² taggers.

FIG. 20. SIC curves for the ternary classifications to discriminate W^- from the rest (left) and Z from the rest (right) using the CNN and CNN^2 taggers.

TABLE VIII. Performance metrics for CNN and CNN^2 taggers on the Herwig-jet sample in the W^-/W^+ and Z/W^+ binary classification tasks.

	W^-/W^+			Z/W^+		
	R50	AUC	ACC	R50	AUC	ACC
CNN	19.5819	0.8788	0.7970	35.7788	0.9013	0.8253
CNN^2	18.3411	0.8728	0.7920	43.6485	0.9109	0.8336

TABLE IX. Performance metrics for CNN and CNN^2 taggers on the Herwig-jet sample in the ternary classification tasks.

	Overall	Signal: W^-			Signal: Z		
	ACC	R50	AUC	ACC	R50	AUC	ACC
CNN	0.6943	14.2221	0.8556	0.7742	24.8287	0.8734	0.7980
CNN^2	0.7197	16.6117	0.8677	0.7865	33.6355	0.8961	0.8212

Tables VIII and IX. The results are slightly worse than those from the Pythia-jet dataset, based upon which all the taggers in this work are optimized. Nevertheless, the tagger performance indicated by the SIC curves roughly agrees

with what is observed in Pythia-jet dataset. This shows that the tagging abilities of our CNN and CNN^2 taggers are independent of showering and hadronization models employed in the analysis.

- [1] 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).
 [2] L. Asquith *et al.*, Jet substructure at the large hadron collider: Experimental review, *Rev. Mod. Phys.* **91**, 045003 (2019).

- [3] S. Marzani, G. Soyez, and M. Spannowsky, Looking inside jets: An introduction to jet substructure and boosted-object phenomenology, *Lect. Notes Phys.* **958** (2019).
 [4] 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.
- [5] L. de Oliveira, M. Kagan, L. Mackey, B. Nachman, and A. Schwartzman, Jet-images—deep learning edition, *J. High Energy Phys.* **07** (2016) 069.
- [6] 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).
- [7] 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).
- [8] 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.
- [9] 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.
- [10] G. Louppe, K. Cho, C. Becot, and K. Cranmer, QCD-aware recursive neural networks for jet physics, *J. High Energy Phys.* **01** (2019) 057.
- [11] J. Pearkes, W. Fedorko, A. Lister, and C. Gay, Jet constituents for deep neural network based top quark tagging, [arXiv:1704.02124](https://arxiv.org/abs/1704.02124).
- [12] A. Butter, G. Kasieczka, T. Plehn, and M. Russell, Deep-learned top tagging with a lorentz layer, *SciPost Phys.* **5**, 028 (2018).
- [13] T. Cheng, Recursive neural networks in Quark/Gluon tagging, *Comput. Softw. Big Sci.* **2**, 3 (2018).
- [14] S. Egan, W. Fedorko, A. Lister, J. Pearkes, and C. Gay, Long Short-Term Memory (LSTM) networks with jet constituents for boosted top tagging at the LHC, [arXiv:1711.09059](https://arxiv.org/abs/1711.09059).
- [15] ATLAS Collaboration, Identification of hadronically-decaying w bosons and top quarks using high-level features as input to boosted decision trees and deep neural networks in ATLAS at $\sqrt{s} = 13$ TeV, Tech. Rep. ATL-PHYS-PUB-2017-004, CERN, Geneva, 2017.
- [16] S. Macaluso and D. Shih, Pulling out all the tops with computer vision and deep learning, *J. High Energy Phys.* **10** (2018) 121.
- [17] K. Datta, A. Larkoski, and B. Nachman, Automating the construction of jet observables with machine learning, *Phys. Rev. D* **100**, 095016 (2019).
- [18] H. Qu and L. Gouskos, ParticleNet: Jet tagging via particle clouds, [arXiv:1902.08570](https://arxiv.org/abs/1902.08570).
- [19] G. Kasieczka *et al.*, The machine learning landscape of top taggers, *SciPost Phys.* **7**, 014 (2019).
- [20] B. M. Dillon, D. A. Faroughy, and J. F. Kamenik, Uncovering latent jet substructure, *Phys. Rev. D* **100**, 056002 (2019).
- [21] B. Bhattacharjee, S. Mukherjee, and R. Sengupta, Discrimination between prompt and long-lived particles using convolutional neural network, *J. High Energy Phys.* **11** (2019) 156.
- [22] S. Diefenbacher, H. Frost, G. Kasieczka, T. Plehn, and J. M. Thompson, CapsNets continuing the convolutional quest, *SciPost Phys.* **8**, 023 (2020).
- [23] E. A. Moreno, O. Cerri, J. M. Duarte, H. B. Newman, T. Q. Nguyen, A. Periwé, 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).
- [24] CMS Collaboration, Machine learning-based identification of highly Lorentz-boosted hadronically decaying particles at the CMS experiment, Tech. Rep. CMS-PAS-JME-18-002, CERN, Geneva, 2019.
- [25] R. D. Field and R. P. Feynman, A parametrization of the properties of quark jets, *Nucl. Phys.* **B136**, 1 (1978).
- [26] D. Krohn, M. D. Schwartz, T. Lin, and W. J. Waalewijn, Jet Charge at the LHC, *Phys. Rev. Lett.* **110**, 212001 (2013).
- [27] G. Aad *et al.* (ATLAS Collaboration), Measurement of jet charge in dijet events from $\sqrt{s} = 8$ TeV pp collisions with the ATLAS detector, *Phys. Rev. D* **93**, 052003 (2016).
- [28] A. M. Sirunyan *et al.* (CMS Collaboration), Measurements of jet charge with dijet events in pp collisions at $\sqrt{s} = 8$ TeV, *J. High Energy Phys.* **10** (2017) 131.
- [29] G. Aad *et al.* (ATLAS Collaboration), A new method to distinguish hadronically decaying boosted Z bosons from W bosons using the ATLAS detector, *Eur. Phys. J. C* **76**, 238 (2016).
- [30] ATLAS Collaboration, Jet Charge Studies with the ATLAS Detector Using $\sqrt{s} = 8$ TeV Proton-Proton Collision Data, Tech. Rep. ATLAS-CONF-2013-086, Geneva, 2013.
- [31] K. Fraser and M. D. Schwartz, Jet charge and machine learning, *J. High Energy Phys.* **10** (2018) 093.
- [32] J. Alwall, R. Frederix, S. Frixione, V. Hirschi, F. Maltoni, O. Mattelaer, H.-S. Shao, T. Stelzer, P. Torrielli, and M. Zaro, The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations, *J. High Energy Phys.* **07** (2014) 079.
- [33] K. Hartling, K. Kumar, and H. E. Logan, GMCALC: A calculator for the Georgi-Machacek model, [arXiv:1412.7387](https://arxiv.org/abs/1412.7387).
- [34] T. Sjöstrand, S. Ask, J. R. Christiansen, R. Corke, N. Desai, P. Ilten, S. Mrenna, S. Prestel, C. O. Rasmussen, and P. Z. Skands, An introduction to PYTHIA 8.2, *Comput. Phys. Commun.* **191**, 159 (2015).
- [35] J. de Favereau, C. Delaere, P. Demin, A. Giammanco, 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.
- [36] M. Cacciari, G. P. Salam, and G. Soyez, FastJet user manual, *Eur. Phys. J. C* **72**, 1896 (2012).
- [37] M. Cacciari, G. P. Salam, and G. Soyez, The anti- k_t jet clustering algorithm, *J. High Energy Phys.* **04** (2008) 063.
- [38] S. Ioffe and C. Szegedy, Batch normalization: Accelerating deep network training by reducing internal covariate shift, [arXiv:1502.03167](https://arxiv.org/abs/1502.03167).
- [39] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel *et al.*, Scikit-learn: Machine learning in Python, *J. Mach. Learn. Res.* **12**, 2825 (2011).
- [40] D. P. Kingma and J. Ba, Adam: A method for stochastic optimization, [arXiv:1412.6980](https://arxiv.org/abs/1412.6980).

- [41] K. Simonyan, A. Vedaldi, and A. Zisserman, Deep inside convolutional networks: Visualising image classification models and saliency maps, [arXiv:1312.6034](#).
- [42] R. Kotikalapudi and contributors, keras-vis, <https://github.com/raghakot/keras-vis> (2017).
- [43] K. He, X. Zhang, S. Ren, and J. Sun, Identity mappings in deep residual networks, [arXiv:1603.05027](#).
- [44] W. Luo, Y. Li, R. Urtasun, and R. S. Zemel, Understanding the effective receptive field in deep convolutional neural networks, [arXiv:1701.04128](#).
- [45] M. Bahr *et al.*, Herwig++ physics and manual, *Eur. Phys. J. C* **58**, 639 (2008).
- [46] J. Bellm *et al.*, Herwig 7.0/Herwig++ 3.0 release note, *Eur. Phys. J. C* **76**, 196 (2016).