Equivariant energy flow networks for jet tagging

Matthew J. Dolan^{*} and Ayodele Ore[†]

*ARC Centre of Excellence for Dark Matter Particle Physics,
School of Physics, The University of Melbourne, Victoria 3010, Australia*

(Received 22 December 2020; accepted 18 March 2021; published 27 April 2021)

Jet tagging techniques that make use of deep learning show great potential for improving physics analyses at colliders. One such method is the energy flow network (EFN)—a recently introduced neural network architecture that represents jets as permutation-invariant sets of particle momenta while maintaining infrared and collinear safety. We develop a variant of the EFN architecture based on the deep sets formalism, incorporating permutation-equivariant layers. We derive conditions under which infrared and collinear safety can be maintained, and study the performance of these networks on the canonical example of W -boson tagging. We find that equivariant EFNs have similar performance to particle flow networks, which are superior to standard EFNs. Due to convergence and generalization issues, the same improvement was not observed when extending particle flow networks with equivariant layers. Finally, we study how equivariant networks sculpt the jet mass and provide some initial results on decorrelation using planing.

DOI: [10.1103/PhysRevD.103.074022](https://doi.org/10.1103/PhysRevD.103.074022)

I. INTRODUCTION

In the past half-decade, there has been a burst of activity surrounding the application of deep learning methods (particularly neural networks) to a variety of problems in collider physics. Much of this research has been focused on tagging jets [1–25]. Jets are collimated sprays of particles produced from energetic partons and contribute significantly to the complexity of a collider event. Optimal methods of jet tagging are therefore highly desirable in order to maximally exploit the physics capabilities of the LHC and future colliders. Examples include quark/gluon discrimination [26–32] and the discrimination of QCD jets from overlapping jets produced by the hadronic decay of a boosted heavy resonance [33–38], a topic of much interest at the LHC. Indeed, the “standard candle” for jet tagging studies is the identification of boosted W bosons from their decay products, a process which we study in this work.

The new machine learning based methods leverage the fine resolution of modern detectors through their ability to automatically construct powerful discrimination variables from low-level inputs. By choosing a suitable representation for the inputs, existing deep learning models can be

directly applied to the task of classifying jets. For example, multilayer perceptrons (MLPs) can be trained on lists of jets’ constituent particle momenta [12], and convolutional neural networks (CNNs) can be trained on jet images [7–11]. These methods have led to substantial improvements on the performances of traditional taggers. However, the jet representations they use have drawbacks. For example, jet images tend to be sparsely populated by particles, making a large number of the input pixels redundant. On the other hand, some methods based on the treatment of jets as lists (such as recurrent neural networks) fail to account for the fact that the particles in a jet have no intrinsic ordering. Accordingly, it is important to continue to thoroughly explore the landscape of machine learning-based taggers, and better understand their performance through analytic calculations [39].

A specific recent example is the energy flow network (EFN), introduced in Ref. [4]. EFNs are based on the deep sets framework [40], which treats data as point clouds—sets of points in space. Point clouds do not have a fixed length or notion of ordering, similar to particle jets at colliders. The deep sets framework demonstrates how to incorporate permutation invariance and equivariance of the data—jet constituents in our case. The permutation-invariant case was extensively studied in Ref. [4], which treats jets as sets of particle four momenta, and was also able to show that the resulting observables can be made infrared and collinear (IRC) safe. More general non-IRC safe networks were also studied, and dubbed particle flow networks (PFNs).

Along with permutation-invariant networks the deep sets framework can incorporate permutation equivariance.

^{*}matthew.dolan@unimelb.edu.au

[†]ayodele@student.unimelb.edu.au

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.

A permutation-equivariant function is one which commutes with the permutation operation. That is, for input data $\vec{x}$, a permutation π and a function f equivariance requires $f(\pi(\vec{x})) = \pi(f(\vec{x}))$. Clearly this requires the dimensionality of the output of f to be the same as the input. For binary classifiers the output is usually one dimensional. Accordingly, it is straightforward to have an entire network which exhibits invariance of its output under permutation of the input parameters, but equivariance is implemented on individual internal layers of the network. While equivariance under transformations of inputs is less obviously motivated than invariance, it has proven to benefit performance on some tasks. Most notably, the convolutions performed within a CNN are translation equivariant [41] and are responsible for such models' suitability to images, which are naturally symmetric to translations. Equivariance to more complex symmetries has also been investigated such as in Refs. [42–44]. Reference [44] constructs networks that are equivariant under the action of the Lorentz group. Since the Lorentz group is a Lie group and hence continuous, this necessitates a different approach from the finite permutation groups relevant to this work.

In this paper, we continue the study of energy flow networks by implementing permutation-equivariant network layers into the EFN architecture. We are able to find a variant which maintains infrared and collinear safety of the observables learned by the resulting network. We study the performance of equivariant energy flow and particle flow networks on boosted W tagging, including some networks that use information about the jet constituent ID. While the equivariant EFNs slightly outperform the invariant EFNs, we find that this can change under decorrelation of the network output using planing. We also consider some possible future directions.

II. ENERGY FLOW NETWORKS

In principle, the output of any neural network that takes jets or jet constituents as input can be interpreted as a jet observable. This output could be sensitive to the ordering of the constituents, contrary to the notion that collider events can be viewed as unordered list of particles. This has motivated the development and application of neural network architectures which are insensitive to the ordering of the jet constituents, where jets are viewed as point clouds: a set of data points in space. The representation of functions on sets from which neural network architectures may be modeled has been much studied recently in the machine learning community [40,45–48]. Collider physics studies in this vein include the use of architectures based on graph neural networks such as Ref. [6] and related message passing neural networks [16,49,50].

We will explore the EFN, an example which is based closely on the deep sets model introduced in Ref. [40]. Among deep learning models for jet tagging EFNs have the important feature that they are engineered to enforce IRC

safety on the network output. IRC safety is the requirement that a jet observable be invariant under soft and collinear emissions. We will demonstrate IRC safety of equivariant EFNs in Sec. III.

A. Deep sets framework

Deep sets [40] provides a framework for constructing neural networks that incorporate permutation invariance of the inputs. An energy flow network is an implementation of this model which acts on jets (considering jets as sets of particles with associated momenta), with the additional requirement of IRC safety.

The deep sets framework represents permutation-invariant functions f in the following way. For a function acting on a set X , permutation invariance is ensured by first mapping the individual elements of X to an intermediate *latent space* using a function Φ . The results $\Phi(x_i)$ are then summed in the latent space. Finally, a different map F takes the result of the summation in the latent space to the desired range of $f(X)$. Given that the sum is invariant under permutations of elements in X , the same property is inherited by the total function f . The sum also allows such networks to accept inputs of varying length.

The authors of Ref. [40] prove the following theorem.

1. Deep sets theorem

A function f operating on a set X having elements from a countable universe $\mathfrak{X}$ is invariant to the permutation of instances in X if and only if it can be decomposed in the form

$$f(X) = F\left(\sum_{x \in X} \Phi(x)\right), \quad (1)$$

for suitable transformations Φ and F .

Reference [4] applies this decomposition to jet observables, which we may consider as functions on the set of constituent particle features (four-momentum, charge, flavor etc.). Such observables can therefore be written in the form

$$\mathcal{O}(\{p_1, \dots, p_M\}) = F\left(\sum_{i=1}^M \Phi(p_i)\right), \quad (2)$$

where each $p_i \in \mathbb{R}^d$ is a vector containing d features of particles in the jet, $\Phi: \mathbb{R}^d \rightarrow \mathbb{R}^\ell$ is a per-particle map to the latent space, and $F: \mathbb{R}^\ell \rightarrow \mathbb{R}$ is a continuous mapping. The index i runs over the M particles in the jet. In fact, even before applying the function F , each of the ℓ components of the summed per-particle maps is a jet observable. The final observable $\mathcal{O}$ is some combination of these, dictated by F . A number of familiar examples can be written in this way. For example, the jet mass corresponds to choices $\Phi(p^\mu) = p^\mu$ and $F(x) = \sqrt{x^\mu x_\mu}$.

FIG. 1. Examples of the two multilayer perceptrons comprising an EFN. The per-particle mapping (left) produces ℓ filters and the function F (right) takes the observables $\mathcal{O}_a = \sum_i z_i \Phi_a(y_i, \phi_i)$ to a probability that the input was a signal jet (F_S), or a background jet (F_B).

Reference [4] makes use of the Stone-Weierstrass approximation theorem [51] to derive the specialization of Eq. (2) to the case where $\mathcal{O}$ is IRC safe. They arrive at

$$\mathcal{O}(\{p_1, \dots, p_M\}) = F\left(\sum_{i=1}^M z_i \Phi(\hat{p}_i)\right), \quad (3)$$

where the per-particle mapping $\Phi: \mathbb{R}^2 \rightarrow \mathbb{R}^\ell$ now depends only on the angular coordinates of the jet constituent momenta through $\hat{p} \equiv (y, \phi)$, and the sum is weighted by the transverse momentum fractions $z_i = p_{T,i} / \sum_j p_{T,j}$.

Finally, we comment on the dimensionality of the latent space. While any IRC safe observable can be represented this way, this may require a latent space of high dimensionality. In Ref. [40], the deep sets theorem was only proven for the case where instances of the set X are elements of a countable domain. Reference [52] argues that extending consideration to uncountable sets is necessary since continuity on countable set (such as $\mathbb{Q}$) is a vastly different condition to continuity on the (uncountable) reals. They show that the deep sets theorem can only be extended to *finite subsets* of uncountable domains, with the condition that the latent dimension should be at least the number of elements in the set. That is, in order to ensure universal function approximation of a network in the deep sets framework, one should have $\ell \geq M$. This is consistent with the observation in Ref. [4] that an EFN's performance saturates near $\ell = 64$, since the multiplicity distributions of the jet samples peak around 40–50. We will present results below for $\ell = 128$ and 256.

B. Architecture

To realize Eq. (3) as a neural network, Ref. [4] uses two MLPs to approximate the functions Φ and F . A schematic of this construction for an EFN with latent dimension ℓ is presented in Fig. 1. The network on the left represents the per-particle mapping and the number of nodes in its output

is equal to the dimension of the latent space. Each of these output nodes is called a “filter” and is a function of the jet constituents’ angular coordinates y and ϕ . Each filter is transformed into a jet observable by performing the p_T -weighted sum $\mathcal{O}_a = \sum_i z_i \Phi_a(y_i, \phi_i)$, where the sum runs over particles in the input jet. This set of observables is then passed as input to the second network which represents the map F . As is customary for the application of neural networks to classification tasks, the output layer of F has two nodes and uses the softmax activation, which correlates these nodes such that their outputs take values in the interval $[0, 1]$ and sum to unity. The output of the network may then be interpreted as probabilities, F_S and F_B , that the input is a W (signal) or QCD (background) jet, respectively. However, since the outputs are correlated, the network really only represents a single IRC safe jet observable for which W jets have a value close to 1, and QCD jets, 0 (or vice versa).

III. AUGMENTED ENERGY FLOW NETWORK ARCHITECTURE

In addition to presenting a permutation-invariant neural network architecture, on which the EFN is based, Ref. [40] also defines permutation-equivariant neural network layers. Equivariance of a function means that it commutes with permutations, rather than being invariant to them. Here we detail the way in which an equivariant layer can be incorporated into the existing structure of an EFN and seek an architecture that is able to maintain IRC safety of the network output.

A. Equivariant network layers

Permutation equivariance constrains the form of neural network layers. Consider a network layer $E_\Theta(\mathbf{x}) = \sigma(\Theta\mathbf{x})$, having activation σ , weight matrix $\Theta \in \mathbb{R}^{M \times M}$, and acting on a set $\mathbf{x}$ with elements $x_i \in \mathbb{R}$. A result from Ref. [40] states that such a layer is permutation equivariant if and only if Θ can be decomposed as

FIG. 2. An example implementation of two equivariant layers (blue) into the EFN architecture. The first equivariant layer has ℓ input channels and the second has ℓ' output channels. A pooling operation is required to merge the vectors E_i into a single vector $\vec{O}$ containing jet observables.

$$\Theta = \lambda \mathbf{I} + \gamma (\mathbf{1}\mathbf{1}^T), \quad (4)$$

with $\lambda, \gamma \in \mathbb{R}$, $\mathbf{I} = id \in \mathbb{R}^{M \times M}$ and $\mathbf{1} = (1, \dots, 1)^T \in \mathbb{R}^M$. That is, the matrix Θ has all diagonal terms equal and all off-diagonal terms equal. The function E therefore depends only on a linear combination of the input $\mathbf{x}$ itself and $\mathbf{1}\mathbf{1}^T \mathbf{x}$ —a vector with each component being the sum of components of $\mathbf{x}$. The equivariance of the layer arises from the fact that this summation respects permutation symmetry. One can also consider functions E whose arguments are not of the linear form $\Theta \mathbf{x}$. For example, the function

$$E(\mathbf{x}) = \sigma(\lambda \mathbf{I} \mathbf{x} + \gamma \max(\mathbf{x}) \mathbf{1}), \quad (5)$$

also defines an equivariant network layer. Equations (4) and (5) are valid when the elements of the set $x_i \in \mathbb{R}$. However, the x_i could themselves be vectorial $x_i \in \mathbb{R}^D$, so that $\mathbf{x} \in \mathbb{R}^{M \times D}$. In that case the parameters λ and γ can be promoted to matrices $\Lambda, \Gamma \in \mathbb{R}^{D \times D'}$ with D' being the dimension of each output node. This leads to

$$E(\mathbf{x}) = \sigma(\mathbf{x} \Lambda + \mathbf{1}\mathbf{1}^T \mathbf{x} \Gamma), \quad (6)$$

$$E(\mathbf{x}) = \sigma(\mathbf{x} \Lambda + \mathbf{1} \max\text{pool}(\mathbf{x}) \Gamma), \quad (7)$$

where $\max\text{pool}$ is an operation that merges feature vectors in the set $\mathbf{x} \in \mathbb{R}^{M \times D}$ to a single summary vector in $\mathbb{R}^D$ with i th component equal to $\max \{x_j^i\}_{j=1}^M$.

We will see that EFNs using these layers are IRC unsafe in general. In our studies below we will use layers of the form Eq. (7) as an example of such a network, and further discuss IRC safety below.

B. Layer implementation

As well as introducing the equivariant layer itself, Ref. [40] also discusses how such layers can be

implemented in a specific model. Two properties of the equivariant layer will be exploited. The first is that the composition of equivariant layers is also equivariant. This can be seen by considering equivariant functions f, g acting on a set $\mathbf{x}$ with n elements and $\pi \in S_n$ a permutation. Then

$$\begin{aligned} (f \circ g)(\pi \mathbf{x}) &= f(g(\pi \mathbf{x})) = f(\pi g(\mathbf{x})) \\ &= \pi(f(g(\mathbf{x}))) = \pi(f \circ g)(\mathbf{x}), \end{aligned} \quad (8)$$

so the composition of f and g is equivariant. This implies that equivariant layers can be stacked to build deep models.

The second property is that following an equivariant layer with any commutative pooling operation (such as sum, mean, or max) across set elements yields a permutation-invariant function. This allows us to replace the sum over per-particle maps (the permutation-invariant function in the standard EFN architecture) with a sequence of equivariant layers followed by a pooling operation. The network is then able to learn its own particular permutation-invariant function by which to combine the per-particle information.

A schematic of the architecture is presented in Fig. 2. Just as with the standard permutation-invariant EFN, on the left we have a dense network representing a per-particle map $\Phi: \mathbb{R}^2 \rightarrow \mathbb{R}^\ell$ that acts on the y - ϕ plane. In the center of the diagram the vector representation of each particle under this map is then passed as input to a number of equivariant layers shown with blue shaded nodes, the first of which has ℓ input channels. The number of input (output) channels in an equivariant layer is defined as the dimensionality of each input (output) node.

It is important to note that the size of the equivariant layers must be at least the number of jet constituents in the input, which may vary during training. As such, our new architecture is not naturally able to accept variable-length sets. This is addressed by fixing the layer size, M , to be

suitably large and zero padding the inputs. Such treatment is not necessary for an EFN, since its architecture is completely independent of the input jet multiplicity. However, the number of parameters contained in an equivariant layer does not vary with M , so zero padding inputs to this new architecture is not as troublesome as it is in some other examples.

The final equivariant layer, having ℓ' output channels, is followed by a pooling operation that projects the $(M \times \ell')$ -dimensional output of the layer to a vector $\vec{\mathcal{O}}$ of length ℓ' . There is freedom to choose this operation, for example some of the models described in Sec. IV B use a maxpooling to construct the components of the vector according to $\mathcal{O}^a = \max \{E_i^a\}_{i=1}^M$, where E_i^a is a component of an output node in the final equivariant layer, with i labeling jet constituents and a indexing the latent dimension. These $\mathcal{O}^a$ are permutation-invariant jet observables, in analogy with the p_T -weighted sum over filters in an EFN, and are passed as input to the MLP representing the function $F: \mathbb{R}^{\ell'} \rightarrow \mathbb{R}$.

C. IRC safety

The networks generated from the method above are distinct from the IRC safe EFN studied in Ref. [4], which corresponds to taking the pooling operation as a z -weighted sum and the function E as the identity. As such, IRC safety of the jet observables from the new networks is not guaranteed. In this section, we aim to find an architecture that maintains IRC safety. This allows for a closer comparison with the EFNs of Ref. [4].

There are a number of potential architectures of the type described in the previous section. We will consider an equivariant energy flow network denoted by EV-EFN, which for simplicity we assume has a just a single equivariant layer $E: \mathbb{R}^{M \times \ell} \rightarrow \mathbb{R}^{M \times \ell'}$. We will consider the implications of multiple equivariant layers later in the section. We may express the action of the network as

$$\text{EV-EFN}(\{p_i\}_{i=1}^M) = (F \circ P \circ E)(\{\Phi(\hat{p}_i)\}_{i=1}^M), \quad (9)$$

where $P: \mathbb{R}^{M \times \ell'} \rightarrow \mathbb{R}^{\ell'}$ is the pooling operation, and $F: \mathbb{R}^{\ell'} \rightarrow \mathbb{R}$ is the final function. To specify a network, we must make choices for the pooling operation P and the form of the equivariant layer [using Eqs. (6), (7) above for instance etc.). The maps F and Φ are given by fully connected MLPs. We can write the vector of jet observables $\vec{\mathcal{O}}$ by dropping F from Eq. (9). This gives

$$\vec{\mathcal{O}} = (P \circ E)(\{\Phi(\hat{p}_i)\}_{i=1}^M). \quad (10)$$

Since we know that taking P as a z -weighted sum $P(\mathbf{x}) = \sum_i z_i x_i$ is an IRC safe choice in an EFN, we will only consider this operation for this section. This may not be the unique choice for ensuring IRC safety, but we were able to rule out some other possibilities. For example, a max

function is not necessarily invariant to the collinear splitting of a particle.

In an EFN with equivariant layers, this choice of P alone is not enough to ensure that the observables vector $\vec{\mathcal{O}}$ is IRC safe. The reason for this is that in the new architecture, the output nodes of the equivariant layer are not strictly per-particle maps. That is, the vector E_i has dependence on the momentum of particles other than i . However, the fact that this mixing comes through an equivariant layer imposes helpful constraints.

Specifically, permutation equivariance demands that the i th output of the layer depends on inputs other than i only through permutation-invariant functions of the inputs (in linear combination). For example, in Eq. (4) for $(\Theta \mathbf{x})_i$, the first term depends only on x_i and the second term is proportional to the permutation-invariant sum $\sum_i x_i$. Thus, despite the outputs E_i of the equivariant layer not being per-particle maps, they do have a unique dependence on the i th particle.

In addition, the permutation-invariant functions that give the dependence on all particles can be identified as jet observables. We show below that the IRC safety of the z -weighted sum of such an output, which depends on the linear combination of single particle information and a jet observable, is tied to the IRC safety of the jet observables in the equivariant layer. We therefore parametrize an equivariant layer as

$$E(\{\Phi(\hat{p}_i)\}_{i=1}^M) = \sigma(\Phi(\hat{p}_i) \cdot \Lambda + \mathcal{Q}(\{\Phi(\hat{p}_i)\}_{i=1}^M) \cdot \Gamma), \quad (11)$$

where $\mathcal{Q} \in \mathbb{R}^{\ell}$ holds jet observables that depend on the set of per-particle latent vectors, the matrices Λ and Γ are defined as in Eq. (6), and σ is the activation function. However, as it stands, the equivariant layer only takes directional information $\hat{p}$ as input through the filters Φ . If $\mathcal{Q}$ is to be IRC safe, it must also receive the momentum fractions z as input. Such a dependence can be added to the layer E through $\mathcal{Q}$ while maintaining equivariance if $\mathcal{Q}$ is understood to be permutation invariant with regard to pairs $(\Phi(\hat{p}_i), z_i)$. This yields

$$\begin{aligned} E(\{\Phi(\hat{p}_i)\}_{i=1}^M, z) \\ = \sigma(\Phi(\hat{p}_i) \cdot \Lambda + \mathcal{Q}(\{(\Phi(\hat{p}_i), z_i)\}_{i=1}^M) \cdot \Gamma), \end{aligned} \quad (12)$$

Taking P as a z -weighted sum and the discussion above into account, Eq. (10) for a network with a single equivariant layer can be written out in more detail as

$$\vec{\mathcal{O}} = \sum_{i=1}^M z_i \sigma(\Phi(\hat{p}_i) \cdot \Lambda + \mathcal{Q}(\{(\Phi(\hat{p}_i), z_i)\}_{i=1}^M) \cdot \Gamma). \quad (13)$$

We wish to show that the observable $\vec{\mathcal{O}}$ is IRC safe if every component of $\mathcal{Q}$ is IRC safe. To study infrared safety, we add an arbitrarily soft particle labeled by index s that can have momentum in any direction:

$$\begin{aligned}
\vec{\mathcal{O}}_{\text{IR}} &= \lim_{z_s \rightarrow 0} z_s \sigma(\Phi(\hat{p}_s) \cdot \Lambda + \mathcal{Q}_{\text{IR}} \cdot \Gamma) \\
&\quad + \sum_{i=1}^M z_i \sigma(\Phi(\hat{p}_i) \cdot \Lambda + \mathcal{Q}_{\text{IR}} \cdot \Gamma), \\
&= \sum_{i=1}^M z_i \sigma(\Phi(\hat{p}_i) \cdot \Lambda + \mathcal{Q}_{\text{IR}} \cdot \Gamma), \tag{14}
\end{aligned}$$

where $\mathcal{Q}_{\text{IR}} = \mathcal{Q}(\{(\Phi(\hat{p}_s), z_s)\} \cup \{(\Phi(\hat{p}_i), z_i)\}_{i=1}^M)$.

For collinear safety a particle splits with momentum fractions λz and $(1-\lambda)z$ with $\lambda \in [0, 1]$. Without loss of generality we take this to be the first particle. Under a collinear splitting the observable becomes

$$\begin{aligned}
\vec{\mathcal{O}}_C &= \lambda z_1 \sigma(\Phi(\hat{p}_1) \cdot \Lambda + \mathcal{Q}_C \cdot \Gamma) \\
&\quad + (1-\lambda) z_1 \sigma(\Phi(\hat{p}_1) \cdot \Lambda + \mathcal{Q}_C \cdot \Gamma) \\
&\quad + \sum_{i=2}^M z_i \sigma(\Phi(\hat{p}_i) \cdot \Lambda + \mathcal{Q}_C \cdot \Gamma) \\
&= \sum_{i=1}^M z_i \sigma(\Phi(\hat{p}_i) \cdot \Lambda + \mathcal{Q}_C \cdot \Gamma), \tag{15}
\end{aligned}$$

and $\mathcal{Q}_C = \mathcal{Q}(\{(\Phi(\hat{p}_1), \lambda z_1), (\Phi(\hat{p}_1), (1-\lambda)z_1)\} \cup \{(\Phi(\hat{p}_i), z_i)\}_{i=2}^M)$.

If $\mathcal{Q}_{\text{IR}} = \mathcal{Q}_C = \mathcal{Q}$, then Eq. (13) is recovered from both Eqs. (14) and (15). Thus, $\vec{\mathcal{O}}$ is IRC safe if $\mathcal{Q}$ is IRC safe. One obvious choice for the components $\mathcal{Q}^a$ is to again use the z -weighted sum over per-particle maps, $\mathcal{Q}^a = \sum_i z_i \Phi^a(\hat{p}_i)$ leading to

$$E(\{\Phi(\hat{p}_i)\}_{i=1}^M, z) = \sigma\left(\{\Phi(\hat{p}_i)\}_{i=1}^M \cdot \Lambda + \sum_{j=1}^M z_j \Phi(\hat{p}_j) \cdot \Gamma\right). \tag{16}$$

The IRC safe jet observables are then obtained by composing with the z -weighted sum. In full, they read

$$\begin{aligned}
\vec{\mathcal{O}} &= \sum_{i=1}^M z_i E_i(\{\Phi(\hat{p}_i)\}_{i=1}^M, z), \\
&= \sum_{i=1}^M z_i \sigma\left(\Phi(\hat{p}_i) \cdot \Lambda + \sum_{j=1}^M z_j \Phi(\hat{p}_j) \cdot \Gamma\right). \tag{17}
\end{aligned}$$

Interestingly, one finds that in certain cases the network is equivalent to a standard EFN. For example, replacing activations by the identity gives

$$\begin{aligned}
\vec{\mathcal{O}} &= \sum_{i=1}^M z_i \left(\Phi(\hat{p}_i) \cdot \Lambda + \sum_{j=1}^M z_j \Phi(\hat{p}_j) \cdot \Gamma \right), \\
&= \sum_{i=1}^M z_i \Phi(\hat{p}_i) \cdot \Lambda + \sum_{j=1}^M z_j \Phi(\hat{p}_j) \cdot \Gamma, \\
&= \sum_{i=1}^M z_i \Phi(\hat{p}_i) \cdot (\Lambda + \Gamma), \\
&\equiv \sum_{i=1}^M z_i \Psi(\hat{p}_i), \tag{18}
\end{aligned}$$

where the first step follows from the fact that the p_T fractions sum to unity, and we have redefined the per-particle map $\Psi^a(\hat{p}_i) \equiv \sum_{b=1}^{\ell} \Phi^b(\hat{p}_i) (\Lambda^{ba} + \Gamma^{ba})$. This is a consequence of the fact that IRC safety imposes a linear structure on the transformations prior to σ . Removing the activation allows the action of the equivariant layer to be absorbed into the definition of the per-particle map. A similar result is reached even when including activations if one restricts the equivariant layer(s) with $\Gamma = 0$ or $\Lambda = 0$. See the Appendix for further comment on this.

Next, we must check that IRC safety is maintained in a network consisting of more than one such equivariant layer. Composing two layers leads to observables

$$\begin{aligned}
\vec{\mathcal{O}} &= (P \circ E \circ E')(\{\Phi(\hat{p}_i)\}_{i=1}^M), \\
&= \sum_{i=1}^M z_i E_i(E'(\{\Phi(\hat{p}_i)\}_{i=1}^M, z), z), \\
&= \sum_{i=1}^M z_i \sigma\left(E'_i(\{\Phi(\hat{p}_i)\}_{i=1}^M, z) \cdot \Lambda \right. \\
&\quad \left. + \sum_{j=1}^M z_j E'_j(\{\Phi(\hat{p}_i)\}_{i=1}^M, z) \cdot \Gamma\right), \\
&= \sum_{i=1}^M z_i \sigma\left[\sigma\left(\Phi(\hat{p}_i) \cdot \Lambda' + \sum_{k=1}^M z_k \Phi(\hat{p}_k) \cdot \Gamma'\right) \cdot \Lambda \right. \\
&\quad \left. + \sum_{j=1}^M z_j \sigma\left(\Phi(\hat{p}_j) \cdot \Lambda' + \sum_{k=1}^M z_k \Phi(\hat{p}_k) \cdot \Gamma'\right) \cdot \Gamma\right], \tag{19}
\end{aligned}$$

where E' denotes the first equivariant layer, which has matrix parameters Λ' and Γ' . Thus the vector $\vec{\mathcal{O}}$ is IRC safe. To see this, notice that the summand for each sum over particles is weighted by z and that these weights constitute all of the z dependence in the expression. The same treatment can be applied for additional stacked layers. Although it would be impractical to present the equations here, one finds that the sums over particles are always z weighted, leading to IRC safety. This is because adding an additional equivariant layer corresponds to replacing each $\Phi(\hat{p}_{\{i,j,k\}})$ by $E_{\{i,j,k\}}$. Since $E_{\{i,j,k\}}$ contains only IRC safe

observables and $\Phi(\hat{p}_{\{i,j,k\}})$, this replacement leaves the same z -weighted structure of the latent vectors.

IV. TRAINING AND PERFORMANCE

Here we evaluate the performance of equivariant EFNs in discriminating jets from boosted W bosons against the QCD background. We will consider both IRC safe and unsafe equivariant networks. The IRC unsafe networks will be contrasted with the performance of PFNs [4]. PFNs are similar to EFNs, but are based around the general observable decomposition of Eq. (2) and are thus IRC unsafe. We begin by detailing the data generation and preparation steps before outlining the construction of each model.

A. Data generation and preprocessing

The W and QCD jets used for this work are obtained respectively from $pp \rightarrow Z + W$ and $pp \rightarrow Z + \text{jet}$ events generated and hadronized with PYTHIA 8.226 [53] at the center of mass energy $\sqrt{s} = 13$ TeV. For the W events, we use the `WeakDoubleBoson:ffbar2ZW` channel with the W decaying to quarks. For the QCD events we use `WeakBosonAndParton:qg2gmZq` and `qqbar2gmZg` with the photon contribution switched off. The Feynman diagrams for the tree-level contributions to these processes are shown in Fig. 3. In both cases, the Z is required to decay invisibly to neutrinos. Multiple parton interactions as well as initial and final state radiation are left on.

To study network performance on jets of different transverse momenta we produce two datasets: the first containing jets with $p_T \in [250, 300]$ GeV and the second with $p_T \in [500, 550]$ GeV. The datasets are produced as follows: after each event is generated, particles with p_T less than 500 MeV are removed with the intent of mitigating the influence of contamination from the underlying event. The event is then clustered by the anti- k_t algorithm [54] using FastJet 3.3.3 [55] with a jet radius of $R = 1.0$ for the $p_T \in [250, 300]$ GeV dataset and $R = 0.6$ for the $p_T \in [500, 550]$ GeV set.

If the hardest jet in the event has mass $m \in [65, 105]$ GeV, rapidity $|y| < 2.5$, and transverse momentum within the relevant range, the set of its constituent particle feature vectors is added to the dataset, with each vector in the format: $(p_T, y, \phi, m, \text{pdgid})$. The field `pdgid` follows the particle data group Monte Carlo particle

numbering scheme [56]. We center the selection window slightly beyond the mass of the W since the final jet will contain some additional radiation from the underlying event, translating the distribution. The final datasets contain 1 million W and QCD jets in equal proportion and are split into three subsets for training (75%), validation (15%), and testing (10%).

We employ four preprocessing steps intended to isolate the differences between the W and QCD jets. Each jet in the dataset is first reclustered using the k_t algorithm [57] with a jet radius of $R = 0.2$ or $R = 0.3$ for the high- or low- p_T datasets, respectively. The resulting subsets are ordered by p_T and the following steps applied:

1. Normalize the p_T of all particles in the jet such that their sum is unity.
2. Translate all particles in the y - ϕ plane such that the leading subjet lies at the origin.
3. Rotate all particles in the jet such that the secondary subjet lies directly below the origin.
4. Reflect all particles about the axis $\phi = 0$ such that the third subjet has $\phi > 0$. If there is no third subjet, reflect such that the p_T -weighted sum of ϕ over all particles in the jet is positive.

The normalization step serves to partially reduce the dependence on the p_T of the jet. Although the radial scale of the event still varies with the transverse momentum, we do not remove this dependence (through enlarging [11] or a similar procedure) since the angular scale of the radiation pattern is correlated with the jet type whereas the jet p_T itself is not. The remaining steps remove redundancies arising from spatial symmetries, namely the location and orientation of the event in the detector.

B. Model details and performance

To evaluate the performance of the equivariant architecture, we employ six neural network models which we call EFN, PFN(-ID), EV-EFN, and EV-PFN(-ID). The details of each model are given below. All nodes in the networks (except for those in an output layer) use the rectified linear unit (ReLU) activation [58] and all weights and biases are initialized by He-uniform parameter initialization [59]. The Adam optimization algorithm [60] is used to minimize the binary cross-entropy loss function with data split into batches of 250. Equivariant layers are fixed to a size of $M = 140$ (larger than the number of particles in

FIG. 3. Tree-level Feynman diagrams for the W (blue) and QCD (red) contributions to the processes of interest. Rearrangement of the final diagram gives an additional s -channel contribution, $qg \rightarrow qZ$.

any of jets in the datasets). The EFN, PFN, and PFN-ID models are implemented using the ENERGYFLOW package [61] for PYTHON [4]. The remaining networks, which involve equivariant layers, were trained through KERAS 2.2.4 [62] with a TENSORFLOW 1.15 [63] backend the code used to construct these models can be found at Ref. [64].

We decide on the number of epochs over which to train a model by monitoring the accuracy and loss on the validation set during training and stopping when the results begin to diverge from the test set. The choices of number and size of layers in the EV models are such that the networks contain a comparable number of parameters as their unaugmented partners. We confirm that equalizing the parameter counts (either by increasing or decreasing ℓ in the EFN or EV-EFN, respectively) yields no significant changes to the results presented in this paper.

The networks are defined as follows.

EFN: An $\ell = 256$ energy flow network where the network representing the per-particle map has two hidden layers of 100 nodes, and the network representing F has three, also with 100 nodes each. The model contains 82,358 parameters and is trained for 30 epochs.

PFN and PFN-ID: These two networks have the same general construction as the EFN network above, however the constituent particles' p_T fractions are given as input to the per-particle map instead of weighting the sum over filters. These networks represent the observable decomposition in Eq. (2), which is not IRC safe in general. In the ID variant, the input feature vector contains both momentum and particle identification. The identities are represented by small floats beginning at 0 for an absent particle, then increasing by 0.05 for each category of constituent: γ , p , $\bar{p}$, n , $\bar{n}$, π^+ , π^- , K^+ , K^- , K_L^0 , e^+ , e^- , μ^+ , μ^- . The models contain 82,458 parameters and are trained for 25 epochs.

EV-EFN: This is an $\ell = 128$ EFN containing two IRC safe equivariant layers as defined in Eq. (16). We use a smaller latent dimension of $\ell = 128$ and reduce the depth of the MLP representing the function F such that the total number of parameters in the model is comparable to that of the $\ell = 256$ EFN. Specifically, the networks for Φ and F both contain two hidden layers of 100 nodes. The equivariant layers have size $M = 140$ (corresponding to the maximum number of jet constituents) and 100 output channels to match the size of the layers in F . The model contains 89,330 parameters and is trained for 30 epochs.

EV-PFN and EV-PFN-ID: These are similar to the $\ell = 128$ PFN(-ID) networks, but extended with two equivariant layers in the same fashion as the EV-EFN. The equivariant layers are of the type in Eq. (7) (which is not IRC safe) and max-pooling is used as the projection operation. Additionally, we use

TABLE I. Performance metrics obtained by the models listed in text on both p_T datasets. The background rejection metric R is defined as $1/\epsilon_{\text{QCD}}$ at 50% signal efficiency. For both the area under the curve (AUC) score and R , a greater value corresponds to better classification. The reported uncertainties are interquartile ranges of values from ten separately trained models.

Model	Jet p_T [GeV]	AUC	$R_{\epsilon_W=0.5}$
EFN	500–550	0.9339 ± 0.0007	87.4 ± 5.7
	250–300	0.9092 ± 0.0014	49.1 ± 1.6
EV-EFN	500–550	0.9367 ± 0.0009	93.4 ± 2.5
	250–300	0.9111 ± 0.0015	51.5 ± 1.9
PFN	500–550	0.9366 ± 0.0012	93.2 ± 4.6
	250–300	0.9117 ± 0.0011	51.4 ± 1.6
EV-PFN	500–550	0.9333 ± 0.0026	80.1 ± 3.8
	250–300	0.9083 ± 0.0008	46.2 ± 1.6
PFN-ID	500–550	0.9439 ± 0.0014	119.0 ± 9.0
	250–300	0.9180 ± 0.0008	60.0 ± 1.8
EV-PFN-ID	500–550	0.9402 ± 0.0010	96.0 ± 4.0
	250–300	0.9152 ± 0.0012	53.5 ± 2.6

14 dimensional basis vectors to represent each category of constituent particle in the ID variant. Although these choices for PID encoding and projection are different to those in the PFN(-ID) models, we find that they lead to better stability of training. Models using other choices [such as sum pooling or the use of equivariant layers of the type in Eq. (6)] often failed to converge, having AUC scores of 0.5 after training. In the instances where training did converge, results similar to those achieved by the architecture outlined here were observed. The EV-PFN(-ID) models contain 89,430 (90,830) parameters and, due to poor generalization, are trained for 20 epochs.

The AUC scores and background rejections, $R \equiv 1/\epsilon_{\text{QCD}}$, for the models trained on the low- and high- p_T datasets are presented in Table I. Comparing the IRC safe models, we see that the EV-EFN achieves slightly better classification than the EFN across both measures in both p_T ranges. We also show the receiver operating characteristic (ROC) curves averaged over 10 networks for these models (as well as the PFN) in Fig. 4. The curves here again reflect the classification advantage that the EV-EFN achieves over the EFN. This suggests that the addition of equivariant layers aids the network's ability to learn from IRC safe information. In fact, the improvement obtained by the EV-EFN is greater than the improvement seen from the PFN model, which is not IRC safe and uses a latent space of twice the dimensionality. However, the results in Table I show that the EV-PFN does not discriminate better than the PFN model. In fact, on both datasets it performs worse than even the IRC safe EFN across both metrics. The result for the EV-PFN-ID model is not as poor—although it does not match the classification power of the PFN-ID networks, it does at least outperform the PFN.

FIG. 4. ROC curves averaged over 10 networks for each of the EV-EFN (green), EFN (blue), and PFN (orange) models trained on the $p_T \in [500, 550]$ GeV dataset. The bands show one standard deviation. The lower plot shows the ratios relative to the EFN ROC curve.

A major issue that the EV-PFN(-ID) networks face is poor generalization. To illustrate this, in Fig. 5 we plot the classification accuracy of the PFN and EV-PFN models on the training and validation datasets over 50 epochs. The test accuracy for the EV-PFN diverges from the validation

accuracy after just 15 epochs of training, whereas this behavior is not seen to such a strong degree for the PFN despite both models containing approximately the same number of parameters.

The plots reveal that the EV-PFN does indeed classify jets from the training set more accurately than the PFN, however its poor generalization means that this performance is not transferred to test results. This could possibly be addressed by regularization techniques. One such option is to apply dropout—a method wherein nodes in a layer are randomly selected to be switched off during training [65]. We investigate application of dropout on layers in the network for F , but find that for a variety of dropout rates, only the performance on the training set is affected. The same result is observed for the ID variants. In the interest of avoiding potentially unfruitful trial and error, we leave fine-tuning of the EV-PFN(-ID) models to future work.

C. Dependence on jet mass

In this section, we compare the effect of the EV-EFN and EFN on the QCD jet mass distribution. We follow closely the approach taken in Ref. [66]. An unrestricted neural network will make use of all available information, jet mass included. Consequently, the output of such networks will be correlated with the jet mass, leading to a sculpting effect of the jet mass distribution. This is not a desirable feature, since background modeling is often done by extrapolating from a region containing no signal, relying on an approximately flat background. A background distribution that is smooth and featureless is associated with less systematic uncertainty. There exists a trade-off between the raw discrimination power a neural network offers and the

FIG. 5. Classification accuracies of the PFN (left) and EV-PFN (right) models on the training and validation subsets of the $p_T \in [500, 550]$ GeV data over 50 epochs.

FIG. 6. Jet mass histograms (normalized) for the $p_T \in [500, 550]$ GeV QCD jets that are tagged as W s by successively harder cuts on the EFN (left) or EV-EFN (right) output. The distribution for W jets is shown in gray.

degree to which it reshapes the background jet mass distribution. For this reason, jet observables that are strongly correlated with the jet mass are less valuable for experimental use compared to those that are uncorrelated. This problem has been studied extensively from both analytic [67,68] and machine learning perspectives [25,66,69–75].

Here we compare results first for the standard training procedure described above and then including the application of a method for decorrelating network output from the jet mass. To compare the two models, we train 10 EFNs and 10 EV-EFNs on the $p_T \in [500, 550]$ GeV dataset, split into training (75%) and testing (25%) subsets. By increasing the size of the testing set, we reduce the impact of statistical fluctuations on the results. This is important since, at low signal efficiencies, a very small number of QCD jets are tagged. In Fig. 6 we visualize the sculpting effect of the EFN and EV-EFN models on the training set by showing the jet mass distributions of QCD jets that are tagged as signal by successively tighter cuts the network output. Both models strongly sculpt the QCD distribution to the W mass peak. In order to quantify the degree of sculpting we make use of the Hellinger distance, defined as

$$H(p, q) = \sqrt{1 - \sum_i \sqrt{p(i)q(i)}}, \quad (20)$$

where p and q are discrete probability distributions. Hellinger distances are normalized between 0 and 1, with a distance of 0 corresponding to p and q being identical histograms. For each trained model, the Hellinger distances between the jet mass distributions of tagged QCD jets and the original are calculated (using histograms of 25 bins) for a range of cuts on network output. We then average the distances over network instances.

A plot of these average distances as a function of the network background rejection for the two models is shown on the left of Fig. 7. In these terms a better classifier corresponds to a curve that is closer to the bottom-right corner, reaching greater background rejection at smaller Hellinger distance. The general behavior of the two models is very similar, however the EV-EFNs achieve smaller Hellinger distances across all cuts on the network output. The greatest difference is seen at a background rejection of $1/\epsilon_{\text{QCD}} \sim 90$ (corresponding to a signal efficiency of $\epsilon_W = 0.5$), where they improve upon the EFNs by approximately 5%.

While the EV-EFNs have a smaller effect on the background distribution compared to the EFNs, both models sculpt to such a great degree that this improvement can be considered somewhat inconsequential. Having said this, a more meaningful result can be obtained by investigating the networks' sculpting after decorrelating the network output from the jet mass. Reference [66] found that for neural network taggers, techniques that achieve this through adjustments to the input data provide decorrelation almost on par with those that alter the training procedure, while having significantly lower computational cost. For this reason, we compare the response of EFNs and EV-EFNs to mass planing—a particular case of data planing [8,69]. A dataset can be planed in some observable, ρ , by assigning each instance of the training data x_i with a weight w according to

$$[w(x_i)]^{-1} = \left. \frac{d\sigma_x}{d\rho} \right|_{\rho=\rho(x_i)}, \quad (21)$$

where $\frac{d\sigma_x}{d\rho}$ is the differential cross-section for jets in x . In practice, this simply corresponds to inverting the histogram

FIG. 7. Hellinger distances between the tagged QCD jet mass distribution and the full distribution as a function of the background rejection achieved by cutting on the output of an EFN or EV-EFN. Results are shown for the standard training procedure (left) and with the inclusion of mass-planing (right). We indicate with markers the positions along the curves corresponding to fixed signal efficiencies of $\epsilon_W = 0.3$ or $\epsilon_W = 0.5$. The ratios to the EFN line are shown beneath each plot.

for ρ at the bin in which x_i lies. After such a weighting, the data has a uniform distribution in the planed observable.

To implement mass planing in the networks' training procedure, we calculate weights from the W and QCD mass distributions histogrammed into 100 bins. The weights are used to give the relative importance of a jet's contribution to the loss function during training. The Hellinger distances are then calculated and averaged in the same fashion as for the un-planed case. On the right side of Fig. 7, we again present a plot of the calculated distances against background rejection for the newly trained networks. Interestingly, the advantage that the EV-EFN displayed without planing is no longer apparent. In fact, the standard EFN exhibits less background sculpting across most cuts on the network output. One should therefore be cautious about interpreting the absolute performance of some given taggers without considering decorrelation, as this may change the order of the rankings. However, the EV-EFN still achieves greater background rejection at fixed signal efficiency. This implies that in a bump hunt scenario, neither model has a clear advantage in terms of achieved significance.

V. CONCLUSIONS

The development of new and increasingly accurate jet taggers remains a high priority for the collider physics community. The use of techniques based on

machine-learning methods has provided a number of advances in this area in the past few years. Recent research has explored incorporating permutation invariance of the constituent particle momenta into the network architecture via the deep sets formalism, leading to EFNs. However, deep sets also allows for the possibility of permutation equivariance, which is what we have explored in this paper.

We have successfully implemented permutation-equivariant neural network layers into the energy flow network architecture. Furthermore, we identified a specific construction, the EV-EFN, that maintains the infrared and collinear safety of the learned jet observable. When trained to discriminate W from QCD jets in simulated samples in two boosted p_T windows these EV-EFNs exhibited a marginal improvement in performance over permutation-invariant EFNs. The performance of the EV-EFN was equal to that achieved by a particle flow network, which is not restricted by IRC safety. Some issues were observed when extending PFNs by equivariant layers. Specifically, certain choices of equivariant operation or projection prevented the model from converging during training. Further, models that did converge generalized poorly to the testing sample. These issues could possibly be remedied by alternative constructions or regularization techniques, however we leave that task for future work.

Finally, we compared the action of the networks on the QCD jet mass distribution before and after applying mass

planing. The planing decorrelates the tagger performance from the jet mass, at some cost in performance. However, it also changes the relative performances of the EFN and EV-EFN networks. While the EV-EFN initially had slightly less impact on the background distribution compared to the EFN, this advantage was lost after mass planing the input jets. Consequently it would be interesting to apply recent ideas such as distance correlation [73], adversarial networks [71], mass unspecific [25] (or agnostic [66]) supervised tagging to see if networks with equivariant layers continue to outperform invariant ones.

A number of different future directions present themselves. Future work could investigate the use of other types of equivariant layers and their impact on IRC safety including those which are equivariant to continuous symmetries of the Lorentz group as in Ref. [44] and similar to Ref. [24]. It may also be interesting to attempt to translate the EV-EFN into a low-dimensional human-interpretable space using the method outlined in Ref. [76]. This could potentially illuminate the effect that equivariant layers have on the representation of IRC safe jet observables compared to the original EFN.

ACKNOWLEDGMENTS

We thank John Gargalionis for useful comments and discussions. This work was supported in part by the Australian Research Council and the Australian Government Research Training Program Scholarship initiative.

TABLE II. Classification metrics on the $p_T \in [500, 550]$ GeV dataset achieved by EV-EFNs with restricted equivariant layers in comparison to the full EV-EFN and the base EFN.

Model	AUC	$R_{\epsilon_W=0.5}$
EV-EFN $_{\sigma=id}$	0.9320 ± 0.0009	83.5 ± 5.2
EV-EFN $_{\Gamma=0}$	0.9327 ± 0.0013	85.0 ± 1.9
EV-EFN $_{\Lambda=0}$	0.9337 ± 0.0008	87.6 ± 2.4
EFN	0.9339 ± 0.0007	87.4 ± 5.7
EV-EFN	0.9367 ± 0.0009	93.4 ± 2.5

APPENDIX: IMPACT OF EQUIVARIANCE

As mentioned in Sec. III C, under certain restrictions to the equivariant layers, the architecture of an EV-EFN reduces to that of a standard EFN. These restrictions include taking one of $\sigma \equiv id$, $\Gamma = 0$, or $\Lambda = 0$. The first case corresponds to rescaling the output of Φ as shown in Eq. (18). The other two cases allow the equivariant layers to be absorbed as additional layers in Φ or F , respectively.

Here we evaluate the performance of EV-EFNs (with architecture as described in Sec. IV B) trained under each of these restrictions. The results are presented in Table II and show that the performance of the EV-EFN degrades in all cases. This demonstrates that the improvement of the EV-EFN over the EFN is not simply due to the model's increased depth.

-
- [1] L. Lonnblad, C. Peterson, and T. Rognvaldsson, Using neural networks to identify jets, *Nucl. Phys.* **B349**, 675 (1991).
 - [2] T. Cheng, Recursive neural networks in quark/gluon tagging, *Comput. Softw. Big Sci.* **2**, 3 (2018).
 - [3] H. Luo, M.-x. Luo, K. Wang, T. Xu, and G. Zhu, Quark jet versus gluon jet: Fully-connected neural networks with high-level features, *Sci. China Phys. Mech. Astron.* **62**, 991011 (2019).
 - [4] P. T. Komiske, E. M. Metodiev, and J. Thaler, Energy flow networks: Deep sets for particle jets, *J. High Energy Phys.* **01** (2019) 121.
 - [5] G. Kasieczka, N. Kiefer, T. Plehn, and J. M. Thompson, Quark-gluon tagging: Machine learning vs detector, *SciPost Phys.* **6**, 069 (2019).
 - [6] H. Qu and L. Gouskos, ParticleNet: Jet tagging via particle clouds, *Phys. Rev. D* **101**, 056019 (2020).
 - [7] 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.
 - [8] L. de Oliveira, M. Kagan, L. Mackey, B. Nachman, and A. Schwartzman, Jet-images—deep learning edition, *J. High Energy Phys.* **07** (2016) 069.
 - [9] 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.
 - [10] 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).
 - [11] 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).
 - [12] J. Pearkes, W. Fedorko, A. Lister, and C. Gay, Jet Constituents for Deep Neural Network Based Top Quark Tagging, *arXiv:1704.02124*.
 - [13] L. G. Almeida, M. Backovi, 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.
 - [14] 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.
 - [15] A. Chakraborty, S. H. Lim, and M. M. Nojiri, Interpretable deep learning for two-prong jet classification with jet spectra, *J. High Energy Phys.* **07** (2019) 135.

- [16] I. Henrion, K. Cranmer, J. Bruna, K. Cho, J. Brehmer, G. Louppe, and G. Rochette, Neural message passing for jet physics, in *DLPS* (2017).
- [17] E. A. Moreno, O. Cerri, J. M. Duarte, H. B. Newman, T. Q. Nguyen, A. Perival, 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).
- [18] G. Louppe, K. Cho, C. Becot, and K. Cranmer, QCD-aware recursive neural networks for jet physics, *J. High Energy Phys.* **01** (2019) 057.
- [19] F. A. Dreyer, G. P. Salam, and G. Soyez, The lund jet plane, *J. High Energy Phys.* **12** (2018) 064.
- [20] T. Heimel, G. Kasieczka, T. Plehn, and J. M. Thompson, QCD or What?, *SciPost Phys.* **6**, 030 (2019).
- [21] M. Farina, Y. Nakai, and D. Shih, Searching for new physics with deep autoencoders, *Phys. Rev. D* **101**, 075021 (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] S. Macaluso and D. Shih, Pulling out all the tops with computer vision and deep learning, *J. High Energy Phys.* **10** (2018) 121.
- [24] A. Butter, G. Kasieczka, T. Plehn, and M. Russell, Deep-learned Top tagging with a lorentz layer, *SciPost Phys.* **5**, 028 (2018).
- [25] J. Aguilar-Saavedra, F. Joaquim, and J. Seabra, Mass unspecific supervised tagging (MUST) for boosted jets, *J. High Energy Phys.* **03** (2021) 012.
- [26] J. Gallicchio and M. D. Schwartz, Quark and Gluon tagging at the LHC, *Phys. Rev. Lett.* **107**, 172001 (2011).
- [27] J. Gallicchio and M. D. Schwartz, Quark and Gluon jet substructure, *J. High Energy Phys.* **04** (2013) 090.
- [28] D. Ferreira de Lima, P. Petrov, D. Soper, and M. Spannowsky, Quark-Gluon tagging with Shower Deconstruction: Unearthing dark matter and Higgs couplings, *Phys. Rev. D* **95**, 034001 (2017).
- [29] C. Frye, A. J. Larkoski, J. Thaler, and K. Zhou, Casimir meets poisson: Improved Quark/Gluon discrimination with counting observables, *J. High Energy Phys.* **09** (2017) 083.
- [30] J. Davighi and P. Harris, Fractal based observables to probe jet substructure of quarks and gluons, *Eur. Phys. J. C* **78**, 334 (2018).
- [31] P. T. Komiske, E. M. Metodiev, and J. Thaler, An operational definition of quark and gluon jets, *J. High Energy Phys.* **11** (2018) 059.
- [32] A. J. Larkoski and E. M. Metodiev, A theory of quark vs. gluon discrimination, *J. High Energy Phys.* **10** (2019) 014.
- [33] 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).
- [34] G. P. Salam, Towards jetography, *Eur. Phys. J. C* **67**, 637 (2010).
- [35] T. Plehn and M. Spannowsky, Top tagging, *J. Phys. G* **39**, 083001 (2012).
- [36] M. Dasgupta, A. Powling, and A. Siodmok, On jet substructure methods for signal jets, *J. High Energy Phys.* **08** (2015) 079.
- [37] J. Thaler and K. Van Tilburg, Identifying Boosted Objects with N-subjettiness, *J. High Energy Phys.* **03** (2011) 015.
- [38] X. Ju and B. Nachman, Supervised jet clustering with graph neural networks for Lorentz boosted bosons, *Phys. Rev. D* **102**, 075014 (2020).
- [39] G. Kasieczka, S. Marzani, G. Soyez, and G. Stagnitto, Towards machine learning analytics for jet substructure, *J. High Energy Phys.* **09** (2020) 195.
- [40] M. Zaheer, S. Kottur, S. Ravanbakhsh, B. Póczos, R. Salakhutdinov, and A. J. Smola, Deep sets, *arXiv:1703.06114*.
- [41] Y. LeCun, B. Boser, J. S. Denker, D. Henderson, R. E. Howard, W. Hubbard, and L. D. Jackel, Backpropagation applied to handwritten zip code recognition, *Neural Comput.* **1**, 541 (1989).
- [42] T. S. Cohen and M. Welling, Group equivariant convolutional networks, *arXiv:1602.07576*.
- [43] N. Thomas, T. Smidt, S. M. Kearnes, L. Yang, L. Li, K. Kohlhoff, and P. Riley, Tensor field networks: Rotation- and translation-equivariant neural networks for 3d point clouds, *arXiv:1802.08219*.
- [44] 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*.
- [45] C. R. Qi, H. Su, K. Mo, and L. J. Guibas, Pointnet: Deep learning on point sets for 3d classification and segmentation, *arXiv:1612.00593*.
- [46] I. Korshunova, J. Degraeve, F. Huszár, Y. Gal, A. Gretton, and J. Dambre, in *Bruno: A deep recurrent model for exchangeable data*, *NeurIPS* (2018) [*arXiv:1802.07535*].
- [47] R. L. Murphy, B. Srinivasan, V. A. Rao, and B. Ribeiro, Janossy pooling: Learning deep permutation-invariant functions for variable-size inputs, *arXiv:1811.01900*.
- [48] C. Pabbaraju and P. Jain Learning functions over sets via permutation adversarial networks, *arXiv:1907.05638*.
- [49] J. Ren, L. Wu, and J. M. Yang, Unveiling CP property of top-Higgs coupling with graph neural networks at the LHC, *Phys. Lett. B* **802**, 135198 (2020).
- [50] M. Abdughani, D. Wang, L. Wu, J. M. Yang, and J. Zhao, Probing triple Higgs coupling with machine learning at the LHC, *arXiv:2005.11086*.
- [51] M. H. Stone, The generalized weierstrass approximation theorem, *Math. Mag.* **21**, 167 (1948).
- [52] E. Wagstaff, F. B. Fuchs, M. Engelcke, I. Posner, and M. A. Osborne, On the limitations of representing functions on sets, *arXiv:1901.09006*.
- [53] T. Sjstrand, 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).
- [54] M. Cacciari, G. P. Salam, and G. Soyez, The anti- k_t jet clustering algorithm, *J. High Energy Phys.* **04** (2008) 063.
- [55] M. Cacciari, G. P. Salam, and G. Soyez, FastJet user manual, *Eur. Phys. J. C* **72**, 1896 (2012).
- [56] M. Tanabashi *et al.* (Particle Data Group), Review of particle physics, *Phys. Rev. D* **98**, 030001 (2018).
- [57] S. Catani, Y. L. Dokshitzer, M. H. Seymour, and B. R. Webber, Longitudinally invariant k_t clustering algorithms for hadron hadron collisions, *Nucl. Phys. B* **406**, 187 (1993).
- [58] V. Nair and G. E. Hinton, in Rectified linear units improve restricted boltzmann machines, *ICML 10* (2010), pp. 807–814.

- [59] K. He, X. Zhang, S. Ren, and J. Sun, Delving deep into rectifiers: Surpassing human-level performance on imagenet classification, [arXiv:1502.01852](#).
- [60] D. P. Kingma and J. Ba, in Adam: A method for stochastic optimization, *ICLR* (2015) [arXiv:1412.6980](#).
- [61] <https://energyflow.network>
- [62] F. Chollet *et al.*, Keras, <https://keras.io> (2015).
- [63] M. Abadi, P. Barham, J. Chen, Z. Chen, A. Davis, J. Dean, M. Devin, S. Ghemawat, G. Irving, M. Isard, M. Kudlur, J. Levenberg, R. Monga, S. Moore, D. G. Murray, B. Steiner, P. Tucker, V. Vasudevan, P. Warden, M. Wicke, Y. Yu, and X. Zheng, in Tensorflow: A system for large-scale machine learning, *OSDI 16* (2016), pp. 265–283 [[arXiv:1605.08695](#)].
- [64] <https://github.com/ayo-ore/equivariant-efns>
- [65] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, Dropout: A simple way to prevent neural networks from overfitting, *J. Mach. Learn. Res.* **15**, 1929 (2014).
- [66] L. Bradshaw, R. K. Mishra, A. Mitridate, and B. Ostdiek, Mass agnostic jet taggers, *SciPost Phys.* **8**, 011 (2020).
- [67] J. Dolen, P. Harris, S. Marzani, S. Rappoccio, and N. Tran, Thinking outside the ROCs: Designing Decorrelated Taggers (DDT) for jet substructure, *J. High Energy Phys.* **05** (2016) 156.
- [68] I. Moul, B. Nachman, and D. Neill, Convoluted substructure: Analytically decorrelating jet substructure observables, *J. High Energy Phys.* **05** (2018) 002.
- [69] S. Chang, T. Cohen, and B. Ostdiek, What is the machine learning?, *Phys. Rev. D* **97**, 056009 (2018).
- [70] J. Aguilar-Saavedra, J. H. Collins, and R. K. Mishra, A generic anti-QCD jet tagger, *J. High Energy Phys.* **11** (2017) 163.
- [71] C. Shimmin, P. Sadowski, P. Baldi, E. Weik, D. Whiteson, E. Goul, and A. Sgaard, Decorrelated jet substructure tagging using adversarial neural networks, *Phys. Rev. D* **96**, 074034 (2017).
- [72] J. Stevens and M. Williams, uBoost: A boosting method for producing uniform selection efficiencies from multivariate classifiers, *J. Instrum.* **8**, P12013 (2013).
- [73] G. Kasieczka and D. Shih, Robust Jet Classifiers through Distance Correlation, *Phys. Rev. Lett.* **125**, 122001 (2020).
- [74] Performance of mass-decorrelated jet substructure observables for hadronic two-body decay tagging in ATLAS, Technical Report No. ATL-PHYS-PUB-2018-014, CERN, Geneva, 2018.
- [75] T. Dorigo and P. De Castro Manzano, Dealing with Nuisance Parameters using Machine Learning in High Energy Physics: a Review, [arXiv:2007.09121](#).
- [76] T. Faucett, J. Thaler, and D. Whiteson, , *Phys. Rev. D* **103**, 036020 (2021).