

# An efficient Lorentz equivariant graph neural network for jet tagging

Shiqi Gong<sup>a,e,1</sup> Qi Meng,<sup>b</sup> Jue Zhang,<sup>b</sup> Huilin Qu,<sup>c</sup> Congqiao Li,<sup>d</sup> Sitian Qian,<sup>d</sup>  
Weitao Du,<sup>a</sup> Zhi-Ming Ma<sup>a</sup> and Tie-Yan Liu<sup>b</sup>

<sup>a</sup>Academy of Mathematics and Systems Science, Chinese Academy of Sciences,  
Zhongguancun East Road, Beijing 100190, China

<sup>b</sup>Microsoft Research Asia,  
Danling Street, Beijing 100080, China

<sup>c</sup>CERN, EP Department,  
CH-1211 Geneva 23, Switzerland

<sup>d</sup>School of Physics, Peking University,  
Chengfu Road, Beijing 100871, China

<sup>e</sup>School of Mathematical Sciences, University of Chinese Academy of Sciences,  
No. 19 Yuquan Road, Beijing 100049, China

E-mail: [gongshiqi15@mails.ucas.ac.cn](mailto:gongshiqi15@mails.ucas.ac.cn), [meq@microsoft.com](mailto:meq@microsoft.com),  
[jue.zhang@microsoft.com](mailto:jue.zhang@microsoft.com), [huilin.qu@cern.ch](mailto:huilin.qu@cern.ch), [licongqiao@pku.edu.cn](mailto:licongqiao@pku.edu.cn),  
[stqian@pku.edu.cn](mailto:stqian@pku.edu.cn), [duweitao@amss.ac.cn](mailto:duweitao@amss.ac.cn), [mazm@amt.ac.cn](mailto:mazm@amt.ac.cn),  
[tyliu@microsoft.com](mailto:tyliu@microsoft.com)

**ABSTRACT:** Deep learning methods have been increasingly adopted to study jets in particle physics. Since symmetry-preserving behavior has been shown to be an important factor for improving the performance of deep learning in many applications, Lorentz group equivariance — a fundamental spacetime symmetry for elementary particles — has recently been incorporated into a deep learning model for jet tagging. However, the design is computationally costly due to the analytic construction of high-order tensors. In this article, we introduce LorentzNet, a new symmetry-preserving deep learning model for jet tagging. The message passing of LorentzNet relies on an efficient Minkowski dot product attention. Experiments on two representative jet tagging benchmarks show that LorentzNet achieves the best tagging performance and improves significantly over existing state-of-the-art algorithms. The preservation of Lorentz symmetry also greatly improves the efficiency and generalization power of the model, allowing LorentzNet to reach highly competitive performance when trained on only a few thousand jets.

**KEYWORDS:** Jets and Jet Substructure, Top Quark

**ARXIV EPRINT:** [2201.08187](https://arxiv.org/abs/2201.08187)

<sup>1</sup>This work was done when the first author was visiting Microsoft Research Asia.

---

## Contents

<b>1</b>	<b>Introduction</b>	<b>1</b>
<b>2</b>	<b>Preliminary</b>	<b>3</b>
2.1	Notations	3
2.2	Theory of the Lorentz group	3
2.3	Graph neural network for particle cloud	4
<b>3</b>	<b>Network architecture</b>	<b>5</b>
3.1	LorentzNet	5
3.2	Theoretical analysis	7
3.3	Implementation details	8
<b>4</b>	<b>Experiment</b>	<b>8</b>
4.1	Datasets	8
4.2	Baselines and tagging performance	9
4.3	Sample efficiency	11
4.4	Equivariance test	11
4.5	Ablation study	12
4.6	Computational cost	13
<b>5</b>	<b>Conclusion</b>	<b>14</b>
<b>A</b>	<b>Details about LGN</b>	<b>15</b>
<b>B</b>	<b>Relation with EGNN</b>	<b>16</b>
<b>C</b>	<b>Computational cost in comparison to LGN</b>	<b>16</b>

---

## 1 Introduction

Machine learning techniques, especially deep learning, have been widely adopted to solve a variety of problems in collider physics. Jet tagging, i.e., identifying which type of elementary particle initiates the jet, is one of the most notable applications. It provides a powerful handle for analyzing collision events and serves as a crucial tool for new physics searches and standard model measurements. Based on the theory and phenomenology of jet physics, many expert-designed high-level jet substructure observables are constructed for jet tagging (See reference [1–7] for reviews of the development and use of jet substructure technique). Using different jet representations, various deep learning approaches have been investigated in recent years: Linear classifier [8] and multilayer perceptrons (MLPs) [9–13] can be trained

on a collection of jet-level observables; Fisher discriminant analysis [14], MLPs [15], locally connected networks [16], and 2D convolutional neural networks (CNNs) [17–29] are applied to jet images; MLPs [30–33], 1D CNNs [29], and recurrent neural networks [28, 30, 34, 35] are used to process a jet as a sequence of its constituent particles; Graph neural networks (GNNs) [36–46] are developed for the “particle cloud”, i.e., an unordered set of particles; Recursive neural networks and GNNs are adopted to study jet clustering trees [47–50]. Compared with classic jet substructure methods, deep learning approaches can better leverage the fine resolution of modern detectors and automatically discover complex patterns from low-level inputs, leading to substantial improvements in jet tagging performance.

However, a key aspect neglected by most of these deep learning models is the fundamental spacetime symmetry. The jet tagging result should not depend on the spatial orientation or the Lorentz boost of a jet. Starting from the work [51], symmetry-preserving deep learning models have been developed and shown their power in improving performance [52–58] and enhancing the model’s interpretability [59] in many applications, e.g., the Deep Sets framework is proposed for preserving the permutation invariance and equivariance [52, 54, 55] design gauge equivariant flow for sampling for lattice gauge theory, and [56–58] propose group equivariant models for physical or molecular dynamics prediction. For jet tagging, symmetry-preserving deep learning models have not been fully explored. Although a preprocessing step is applied in many existing models to reduce such dependence, e.g., by performing a Lorentz boost along the collision axis and a rotation in the transverse plane such that the jet momentum always points to the same direction, this only functions to a limited extent and cannot achieve full invariance to arbitrary Lorentz transformations. There exist models such as LoLa [32] and LBN [60] that attempt to exploit the underlying Lorentz symmetry with specially designed neural network layers, but they cannot guarantee the invariance of the entire model. A better way to achieve independence from any Lorentz transformation is via equivariant neural networks. A neural network layer is Lorentz-equivariant if its output transforms accordingly when the input undergoes a Lorentz transformation. Therefore, equivariant layers can be stacked to build a symmetry-preserving deep neural network, of which the classification result is unchanged under any Lorentz transformation. In [61], LGN, a Lorentz equivariant neural network architecture built with tensor products of Lorentz group representations is developed and demonstrated for the first time. However, the tagging performance of LGN reported in [61] is unsatisfactory. One bottleneck is the slow computational speed due to the explicit calculation of the high-order tensors through consecutive tensor product operations over hidden layers.

In this article, we propose a new design of symmetry-preserving deep learning models for jet tagging. Using the particle cloud representation of a jet, we propose LorentzNet, which directly scalarizes the input 4-vectors to realize Lorentz symmetry. Specifically, we design Minkowski dot product attention, which aggregates the 4-vectors with weights learned from Lorentz-invariant geometric quantities under the Minkowski metric. The construction of LorentzNet is guided by the universal approximation theory on Lorentz equivariant mapping, which ensures the equivariance and the universality of the LorentzNet. Compared to LGN which requires computationally expensive tensor products of the geometric quantities to achieve expressiveness, LorentzNet only requires the Minkowski inner product of two vectors.

Thus, it is more efficient in terms of both training and inference. Benchmarked on two representative jet tagging datasets, LorentzNet achieves the best tagging performance and improves significantly over existing state-of-the-art algorithms such as ParticleNet. Moreover, we test the tagging performance with 5%, 1% and 0.5% of training samples, the degradation in performance for LorentzNet is smaller than that for ParticleNet, which clearly shows the benefit of the inductive bias brought by the Lorentz symmetry.

The rest of this paper is organized as follows. Section 2 reviews the theory of the Lorentz group. Section 3 introduces the architecture of LorentzNet. Section 4 details the experiments carried out on two jet tagging benchmarks to show the effectiveness of LorentzNet. Section 5 serves as the conclusion.

## 2 Preliminary

In this section, we introduce the foundation of the Lorentz group and describe the graph neural network for modelling particle clouds.

### 2.1 Notations

Following the jet representation in [37], we regard the constituent particles as a point cloud, which is an unordered, permutation invariant set of particles. Let  $V = (v_1, \dots, v_N) \in \mathbb{R}^{N \times 4}$  denote the point clouds living in  $\mathbb{R}^4$  and  $N$  denotes the number of particles in a jet. Note that the number of particles for jets may be different.  $\langle \cdot, \cdot \rangle$  and  $\| \cdot \|$  denote the Minkowski inner product and Minkowski norm.  $\oplus$  denotes the direct sum and  $\otimes$  denotes the tensor product.  $[N]$  denotes the set  $\{1, 2, \dots, N\}$ .

### 2.2 Theory of the Lorentz group

**Minkowski metric.** Consider the 4-dimensional space-time  $\mathbb{R}^4$  with basis  $\{e_i\}_{i=0}^3$ . We define a bilinear form  $\eta : \mathbb{R}^4 \times \mathbb{R}^4 \rightarrow \mathbb{R}$  as follows. For  $u, v \in \mathbb{R}^4$ , we set  $\eta(u, v) = u^T J v$  where  $J = \text{diag}(1, -1, -1, -1)$  is the Minkowski metric. The Minkowski inner product of two vectors  $u = (t, x, y, z)$  and  $v = (t', x', y', z')$  is defined as  $\langle u, v \rangle = \eta(u, v) = tt' - xx' - yy' - zz'$ . The Minkowski norm of a vector  $u = (t, x, y, z)$  is defined to be  $\|u\| = \sqrt{\eta(u, u)} = \sqrt{t^2 - x^2 - y^2 - z^2}$ .

**Lorentz transformation.** The Lorentz group is defined to be the set of all matrices that preserve the bilinear form  $\eta$ . Restricting the inertial frames to be positively oriented and positively time-oriented, we obtain proper orthochronous Lorentz group, denoted as  $\text{SO}(1, 3)^+$ . The 3d spatial rotation group  $\text{SO}(3)$  is a subgroup of  $\text{SO}(1, 3)^+$ . Additionally, the Lorentz boost is also included. Given two inertial frames  $\{e_i\}_{i=0}^3$  and  $\{e'_i\}_{i=0}^3$ , the relative velocity vector  $\beta$  and the boost factor  $\gamma$  are defined by  $e'_0 = \gamma e_0 + \sum_{i=1}^3 \gamma \beta_i e_i$  where  $\gamma = (1 - \beta^2)^{-1/2}$ .

If we perform Lorentz boost along the x-spatial axis, then the Lorentz transformation between these two frames is the matrix

$$Q = \begin{pmatrix} \gamma & -\gamma\beta & 0 & 0 \\ -\gamma\beta & \gamma & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}.$$

**Lorentz group equivariance.** Let  $T_g : V \rightarrow V$  and  $S_g : U \rightarrow U$  be group actions of  $g \in G$  on sets  $V$  and  $U$ , respectively. We say a function  $\phi : V \rightarrow U$  is equivariant to group  $G$  if

$$\phi(T_g(v)) = S_g(\phi(v)) \tag{2.1}$$

holds for all  $v \in V$  and  $g \in G$ . In this work, we only consider the case that the type of the output is a scalar or vector. Therefore, we explore the following equivariance on a set of particles  $V \in \mathbb{R}^{N \times 4}$ . Let  $Q$  be the Lorentz transformation, the Lorentz equivariance of  $\phi(\cdot)$  means:

$$Q\phi(v) = \phi(Qv), \quad \text{for } \phi(v) \in \mathbb{R}^4; \tag{2.2}$$

$$\phi(v) = \phi(Qv), \quad \text{for } \phi(v) \in \mathbb{R}. \tag{2.3}$$

Note that when the output is a scalar, the group equivariance equals the group invariance.

### 2.3 Graph neural network for particle cloud

A jet can be denoted as a graph when we regard the constituent particles as nodes. For the particle with index  $i$ , we use its 4-momentum vector  $v_i = (E^i, p_x^i, p_y^i, p_z^i)$  as the coordinate of node  $i$  in Minkowski space. We use  $s_i = (s_1^i, s_2^i, \dots, s_\alpha^i)$  to denote the scalars, such as mass, charge and particle identity information, etc, which compose the node attributes. Now  $f_i = v_i \oplus s_i$  contains essential features for tagging. The graph can be denoted as  $G = (V, E)$  where  $V$  is the set of nodes and  $E$  is the set of edges. The edges characterize the message passing between two particles, hence the interaction of two individual sets of particle-wise features. If there is no such interaction, there will be no edge between the two corresponding nodes. Here, we regard the graph as a fully connected graph as we do not assume that we have any prior on the interactions among these particles.

Graph neural networks are natural to learn representations for graph-structured data [62]. Given a graph  $G = (V, E)$ , assuming  $L$  steps in total, the  $l$ -th message passing step on the graph can be described as [63]:

$$m_i^{l+1} = \sum_{j \in \mathcal{N}(i)} M_l(h_i^l, h_j^l, e_{ij}); \tag{2.4}$$

$$h_i^{l+1} = U_l(h_i^l, m_i^{l+1}); \tag{2.5}$$

where  $h_i^0 = f_i$  is the input feature,  $e_{ij}$  is the edge feature,  $\mathcal{N}(i)$  is the set of neighbors of the node  $i$ , and  $M_l, U_l$  are neural networks. For a classification problem, the output  $\hat{y}$  can be obtained by applying the softmax function after decoding  $\{h_i^L; i \in [N]\}$ .

Directly applying the message passing to the jet can not ensure the Lorentz group equivariance because the non-linear neural networks are applied directly to the whole input

features and ignore its intrinsic structure. There are some variants of message passing designed to satisfy continuous group symmetry. A common way is to project the inputs to the basis of irreducible representations of the corresponding group, e.g., the LGN [61] which satisfies the Lorentz group equivariance. To ensure the universality, high-order geometric tensors should be realized by LGN, which brings high computational cost. (See detailed discussions in appendix A).

In the next section, we will introduce the architecture of LorentzNet which is built upon a universal approximation theorem of the Lorentz equivariant function and does not require explicitly calculating the higher-order representations.

### 3 Network architecture

In this section, we illustrate the architecture of LorentzNet. The construction of the LorentzNet is based on the following universal approximation theorem for the Lorentz group equivariant continuous function.

**Proposition 3.1.** [64] *A continuous function  $\phi : (\mathbb{R}^{N \times 4}) \rightarrow \mathbb{R}^4$  is Lorentz-equivariant if and only if*

$$\phi(v_1, v_2, \dots, v_N) = \sum_{i=1}^N g_i(\langle v_i, v_j \rangle_{i,j=1}^N) v_i, \tag{3.1}$$

where  $g_i$  are continuous Lorentz-invariant scalar functions, and  $\langle \cdot, \cdot \rangle$  is the Minkowski inner product.

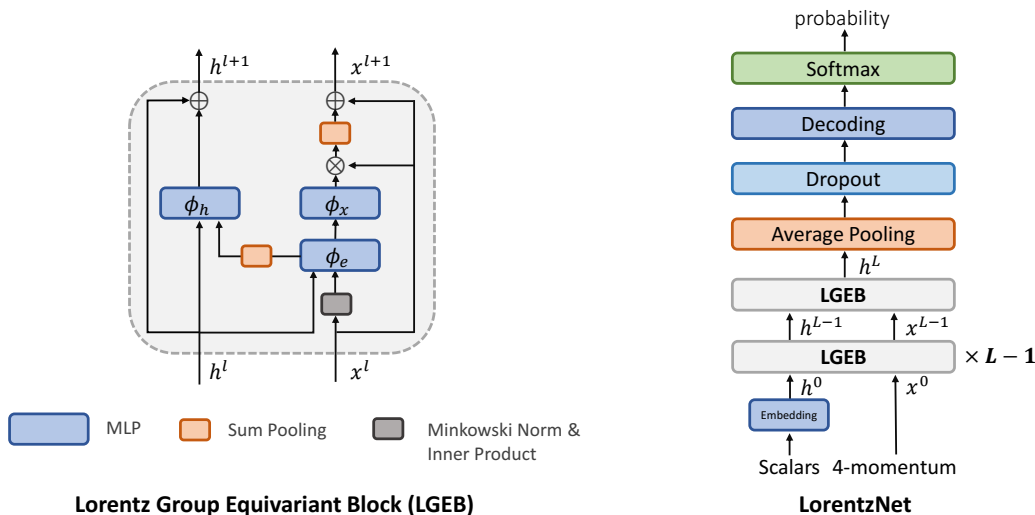
Proposition 3.1 provides a way to construct Lorentz group equivariant mapping with no need to calculate the high-order tensors. Instead, a Lorentz group equivariant continuous mapping can be constructed by the attention on  $v_i$  with encoding the Minkowski dot products of  $v_i$  with its neighbours. This motivates us to design the Minkowski dot product attention in LorentzNet, which will be introduced in the next section.

#### 3.1 LorentzNet

We introduce the blocks in LorentzNet. As described in figure 1, LorentzNet is mainly constructed by the stack of Lorentz Group Equivariant Block (LGEB) along with encoder and decoder layers.

**Input layer.** The inputs into the network are 4-momenta of particles from a collision event, and may include scalars associated with them (such as label, charge, etc.). Using the notations defined in section 2.3, the input is the set of vectors  $f_i = v_i \oplus s_i$ . In the experiments in this paper, the scalars include the mass of the particle (i.e.,  $(E^i)^2 - (p_x^i)^2 - (p_y^i)^2 - (p_z^i)^2$ ) or particle identification (PID) information directly if it is available.

**Lorentz Group Equivariant Block.** We use  $h^l = (h_1^l, h_2^l, \dots, h_N^l)$  to denote the node embedding scalars, and  $x^l = (x_1^l, x_2^l, \dots, x_N^l)$  to denote the coordinate embedding vectors in the  $l$ -th LGEB layer. When  $l = 0$ ,  $x_i^0$  equals the input of the 4-momenta  $v_i$  and  $h_i^0$  equals



**Figure 1.** (left): the structure of the Lorentz Group Equivariant Block (LGEb). (right): the network architecture of the LorentzNet.

the embedded input of the scalar variables  $s_i$ . LGEb aims to learn deeper embeddings  $h^{l+1}, x^{l+1}$  via current  $h^l, x^l$ . Motivated by equation (3.1), the message passing of LorentzNet is written as follows. We use  $m_{ij}$  to denote the edge message between particle  $i$  and  $j$ , and it encodes the scalar information of the particle  $i$  and  $j$ , i.e.,

$$m_{ij}^l = \phi_e \left( h_i^l, h_j^l, \psi(\|x_i^l - x_j^l\|^2), \psi(\langle x_i^l, x_j^l \rangle) \right), \quad (3.2)$$

where  $\phi_e(\cdot)$  is a neural network and  $\psi(\cdot) = \text{sgn}(\cdot) \log(|\cdot| + 1)$  in equation (3.2) is to normalize large numbers from broad distributions for ease of optimization. Except for the embedding of the scalar features  $h_i^l$  and  $h_j^l$ , according to proposition 3.1, the input of the neural network contains the Minkowski dot product  $\langle x_i, x_j \rangle$ . The  $\|x_i^l - x_j^l\|^2$  is also included because the interaction between particles relies on this term and we include it as a prior feature for ease of learning.

According to equation (3.1), we design *Minkowski dot product attention* as

$$x_i^{l+1} = x_i^l + c \sum_{j \in [N]} \phi_x(m_{ij}^l) \cdot x_j^l \quad (3.3)$$

where  $\phi_x(\cdot) \in \mathbb{R}$  is a scalar function modeled by neural networks. To ensure the equivariance, we can not arbitrarily apply the normalization trick to control the scale of  $x_i^{l+1}$ . Therefore, the hyperparameter  $c$  is introduced to control the scale of  $x_i^{l+1}$  to avoid the scale exploding. This step captures the interactions of the  $i$ -th particle with other particles via the ensemble of the 4-momenta of all particles. Unlike most of the symmetry-preserving neural networks such as LGN and EGNN [56] (for  $E(n)$  equivariance)<sup>1</sup> which only apply non-linear transformation (e.g., neural network) to encode the radial distance  $\|x_i - x_j\|^2$ , we include the dot product  $\langle x_i, x_j \rangle$  in  $m_{ij}$  according to equation (3.1).

<sup>1</sup>The relation with EGNN is discussed in the appendix B.

The scalar features for particle  $i$  is forward as

$$h_i^{l+1} = h_i^l + \phi_h \left( h_i^l, \sum_{j \in [N]} w_{ij} m_{ij}^l \right), \tag{3.4}$$

where  $\phi_h(\cdot)$  is also modeled by neural networks whose output dimension equals the dimension of  $h_i^{l+1}$ . For efficient computation, we operate summation  $\sum_{j \in [N]} w_{ij} m_{ij}^l$  to aggregate  $m_{ij}^l$ . We introduce an neural network  $\phi_m(\cdot)$  to learn the edge significance from node  $j$  to node  $i$ , i.e.,  $w_{ij} = \phi_m(m_{ij}^l) \in [0, 1]$ . This can both ensure the permutation invariance and also ease the implementation for jets with different numbers of particles. This operation is also widely adopted in other types of graph neural networks [56, 63].

**Decoding layer.** After stacks of LGEB for  $L$  layers, we decode the node embedding  $h^L = (h_1^L, \dots, h_N^L)$ . Note that the information of  $x^{L-1}$  has been included in  $h^L$  through the  $m_{ij}^{L-1}$ . Therefore, to avoid redundant information, we only decode  $h^L$ . First we use average pooling to get

$$h^{\text{av}} = \frac{1}{N} \sum_{i \in [N]} h_i^L. \tag{3.5}$$

A subsequent dropout layer is applied to  $h^{\text{av}}$  to prevent overfitting. A decoding block with two fully connected layers, followed by a softmax function, is used to generate the output for the binary classification task.

### 3.2 Theoretical analysis

In this section, we analyze the Lorentz group equivariance of LorentzNet.

**Proposition 3.2.** *The coordinate embedding  $x^l = (x_1^l, x_2^l, \dots, x_N^l)$  are Lorentz group equivariant and the node embedding  $h^l = (h_1^l, \dots, h_N^l)$  are Lorentz group invariant.*

*Proof.* We denote  $Q$  as the Lorentz transformation. If  $m_{ij}^l$  are invariant under  $Q$  for all  $i, j, l$ ,  $x_i^{l+1}$  will be Lorentz group equivariant because

$$\begin{aligned} Qx_i^{l+1} &= Q \left( x_i^l + c \sum_{j \in [N]} x_j^l \cdot \phi_x(m_{ij}) \right) \\ &= Qx_i^l + c \sum_{j \in [N]} Qx_j^l \cdot \phi_x(m_{ij}). \end{aligned}$$

Then we illustrate the invariance of  $m_{ij}^l$ . Let's start from the input. Since Minkowski norm and Minkowski inner product are invariant to Lorentz group, we have  $\|x_i^0 - x_j^0\|^2 = \|Qx_i^0 - Qx_j^0\|^2$  and  $\langle x_i^0, x_j^0 \rangle = \langle Qx_i^0, Qx_j^0 \rangle$ . Therefore, the input of  $\phi_e$  are invariant variables under transformation  $Q$  and then  $m_{ij}^0$  are invariant. Recursively using the invariance of  $m_{ij}^l$  and the equivariance of  $x_i^l$ , we can get the conclusion.  $\square$

As for the expressiveness of the LGEB structure, we have the following discussions. Because  $h_i$  is a function of the aggregation of  $m_{ij}$ , it contains the information of all the



Minkowski dot products of the particle pairs. Therefore, the weights in Minkowski dot product attention  $\phi_x(\cdot)$  in equation (3.3) at  $l$ -th layer ( $l > 1$ ) contain all the Minkowski dot products of the particle pairs. According to proposition 3.1, the input information of the attention factor is complete for expressiveness.

### 3.3 Implementation details

The LorentzNet architecture used in this paper is shown in figure 1. It consists of 6 Lorentz group equivariant blocks ( $L = 6$ ). The scalar embedding is implemented as one fully connected layer which maps the scalars to latent space of width 72, i.e.,  $\text{LINEAR}(\text{scalar\_num}, 72)$ .  $\phi_x, \phi_e, \phi_h$  are all implemented as  $\text{LINEAR}(72, 72) \rightarrow \text{RELU} \rightarrow \text{BATCHNORM1D}(72) \rightarrow \text{LINEAR}(72, 72)$ .  $\phi_m$  is implemented as  $\text{LINEAR}(72, 1) \rightarrow \text{SIGMOID}$ . The decoding layers are  $\text{LINEAR}(72, 72) \rightarrow \text{RELU} \rightarrow \text{LINEAR}(72, 2)$ . The dropout rate here is 0.2. The hyperparameter  $c$  in equation (3.3) is chosen to be  $5 \times 10^{-3}$  and  $1 \times 10^{-3}$  for the top tagging and the quark-gluon tagging task respectively to achieve the best performances.

The networks are implemented with PyTorch and the training is performed on a cluster with four Nvidia Tesla V100 GPUs. A batch size of 32 is used on every single GPU for LorentzNet architecture. The ADAMW [65] optimizer, with a weight decay of 0.01, is used to minimize the cross-entropy loss. The LorentzNet is trained for 35 epochs in total. Firstly a linear warm-up period of 4 epochs is applied to reach the initial learning rate  $1 \times 10^{-3}$ . Then a COSINEANNEALINGWARMRESTARTS [66] learning rate schedule with  $T_0 = 4$ ,  $T_{\text{mult}} = 2$  is adopted for next 28 epochs. Finally, an exponential learning rate decay with  $\gamma = 0.5$  is used for the last 3 epochs. We test the model on the validation dataset at the end of each training epoch, and the model with the highest validation accuracy is saved as our best model for the final test.

## 4 Experiment

The performance of the LorentzNet architecture is evaluated on two representative jet tagging tasks: top tagging and quark-gluon tagging. In this section, we show the benchmark results.

### 4.1 Datasets

**Top tagging.** We first evaluate LorentzNet on top tagging classification experiments with the publicly available reference dataset [67]. This dataset contains 1.2M training entries, 400k validation entries, and 400k testing entries. Each of these entries represents a single jet whose origin is either an energetic top quark, a light quark, or a gluon. The events are produced using the PYTHIA8 Monte Carlo event generator. The ATLAS detector response is modelled with the DELPHES software package.

The jets in the reference dataset are clustered using the anti- $k_T$  algorithm, with a radius of  $R = 0.8$ . For each jet, the 4-momenta are saved in Cartesian coordinates  $(E, p_x, p_y, p_z)$  for up to 200 constituent particles selected by the highest transverse momentum. Following the settings in [61], the colliding particle beams aligned along the  $z$ -axis are added. Each jet contains an average of 50 particles, and events with less than 200 are zero-padded.

**Quark-gluon tagging.** The second dataset is for quark-gluon tagging, i.e., discriminating jets initiated by quarks and by gluons. The signal (quark) and background (gluon) jets are generated with PYTHIA8 using the  $Z(\rightarrow \nu\nu) + (u, d, s)$  and  $Z(\rightarrow \nu\nu) + g$  processes, respectively. No detector simulation is performed. The final state non-neutrino particles are clustered into jets using the anti- $k_T$  algorithm with  $R = 0.4$ . This dataset consists of 2 million jets in total, with half signal and half background. We follow the setting in [37] to split 1.6M/200K/200K for training, validation, and testing.

One difference of the quark-gluon tagging dataset is that it includes not only the 4-momentum vector but also the type of each particle (i.e., electron, photon, pion, etc.). We use the one-hot encoding of the type of each particle as input scalars into the LorentzNet. We include this task to test the performance of LorentzNet under this different type of input feature.

## 4.2 Baselines and tagging performance

We compare the performance of LorentzNet with six baseline models: ResNeXt-50 [68], P-CNN [69], PFN [36], ParticleNet [37], EGNN [56] and LGN [61]. The tagging accuracy of these six models except EGNN has been reported in [37] and [61]. In this paper, we also investigate their robustness under Lorentz transformation and the computational cost. For self-contained, we briefly introduce them here. The ResNeXt-50 model is a 50-layer convolutional neural network with skip connections for image classification. Representing jets as images, we can apply ResNeXt-50 to jet tagging. The P-CNN is a 14-layer 1D CNN using particle sequences as inputs. The P-CNN architecture is proposed in the CMS particle-based DNN boosted jet tagger and shows significant improvement in performance compared to a traditional tagger using boosted decision trees and jet-level observables. The Particle Flow Network (PFN) and the ParticleNet also treat a jet as an unordered set of particles. The PFN is based on the Deep Sets framework. The ParticleNet is based on Dynamic Graph Convolutional Neural Network with carefully designed EdgeConv operation. We also include EGNN as a representative symmetry-preserving model as our baseline which is  $E(4)$  equivariant. It is different from LorentzNet in that EGNN uses metrics in Euclidean space instead of Minkowski space. We compare LorentzNet with EGNN to show the necessity of Lorentz group symmetry. For a fair comparison, we set the number of parameters of EGNN in the same order as LorentzNet.

For ResNeXt, P-CNN, PFN, and ParticleNet, we follow the implementation in [37]. For LGN, we follow the implementation in [61] on top tagging. For the quark-gluon dataset, the implementation details are reported in appendix A.

**Tagging performance.** The results for the top tagging dataset and quark-gluon dataset are summarized in table 1 and table 2, respectively. We adopt the widely used measures to evaluate the performance of the LorentzNet model including accuracy, the Area Under the ROC Curve (AUC),<sup>2</sup> the background rejection  $1/\varepsilon_B$  at the signal efficiency of  $\varepsilon_S = 0.3$  and

<sup>2</sup>ROC curve is a plot of the true positive rate (Sensitivity) in the function of the false positive rate (100-Specificity) for different cut-off points of a parameter. Each point on the ROC curve represents a sensitivity/specificity pair corresponding to a particular decision threshold. The Area Under the ROC curve (AUC) is a measure of how well a parameter can distinguish between two diagnostic groups.

Model	Accuracy	AUC	$1/\varepsilon_B$ ( $\varepsilon_S = 0.5$ )	$1/\varepsilon_B$ ( $\varepsilon_S = 0.3$ )
ResNeXt	0.936	0.9837	$302 \pm 5$	$1147 \pm 58$
P-CNN	0.930	0.9803	$201 \pm 4$	$759 \pm 24$
PFN	0.932	0.9819	$247 \pm 3$	$888 \pm 17$
ParticleNet	0.940	0.9858	$397 \pm 7$	$1615 \pm 93$
EGNN	0.922	0.9760	$148 \pm 8$	$540 \pm 49$
LGN	0.929	0.9640	$124 \pm 20$	$435 \pm 95$
LorentzNet	<b>0.942</b>	<b>0.9868</b>	<b><math>498 \pm 18</math></b>	<b><math>2195 \pm 173</math></b>


**Table 1.** Performance comparison between LorentzNet and other representative algorithms on top tagging dataset. The results for LorentzNet and EGNN are averaged on 6 runs. The results for other baselines are referred to [36, 37, 61].

Model	Accuracy	AUC	$1/\varepsilon_B$ ( $\varepsilon_S = 0.5$ )	$1/\varepsilon_B$ ( $\varepsilon_S = 0.3$ )
ResNeXt	0.821	0.8960	30.9	80.8
P-CNN	0.827	0.9002	34.7	91.0
PFN	–	0.9005	$34.7 \pm 0.4$	–
ParticleNet	0.840	0.9116	$39.8 \pm 0.2$	$98.6 \pm 1.3$
EGNN	0.803	0.8806	$26.3 \pm 0.3$	$76.6 \pm 0.5$
LGN	0.803	0.8324	16.0	44.3
LorentzNet	<b>0.844</b>	<b>0.9156</b>	<b><math>42.4 \pm 0.4</math></b>	<b><math>110.2 \pm 1.3</math></b>

**Table 2.** Performance comparison between LorentzNet and other representative algorithms on quark-gluon tagging dataset. The results for LorentzNet, EGNN and LGN are averaged on 6 runs. The results for other baselines are referred to [36, 37].

0.5 ( $\varepsilon_B$ ,  $\varepsilon_S$  are also known as the false positive and the true positive rates, respectively), and the number of trainable parameters. Especially, the background rejection metric is widely adopted to select the best jet tagging algorithm as it is directly related to the expected contribution of the background [36, 37, 61].

From table 1 and table 2, we conclude that LorentzNet achieves the state-of-the-art performance on both the top tagging dataset and the quark-gluon in terms of accuracy, AUC, and background rejection at  $\varepsilon_S = 0.3, 0.5$ . The results verify the effectiveness of LorentzNet compared with the baselines. Figure 2 shows the background rejection at a fine-grained signal efficiency. The ROC curves of LorentzNet achieve the highest score at all the selected signal efficiency compared to the baselines. Especially, LorentzNet shows superiority compared to LGN, since it achieves 4 or 5 times improvement on the background rejection. The results verify our discussions in section 3.2.



**Figure 2.** A comparison of ROC curves between LorentzNet and other algorithms on top tagging dataset (left) and quark-gluon dataset (right).

### 4.3 Sample efficiency

The benefit of the preservation of Lorentz group symmetry in jet tagging has not been studied in the literature. In theory, the Lorentz group symmetry injects inductive bias into the deep learning model which restricts the function class of the hypothesis space. The inductive bias can help to boost the generalization and improve the sample efficiency. As the improvement in the generalization performance (i.e., the tagging accuracy) has been shown in the previous section, we show the robustness of LorentzNet trained on smaller training data to verify the sample efficiency of LorentzNet in this part.

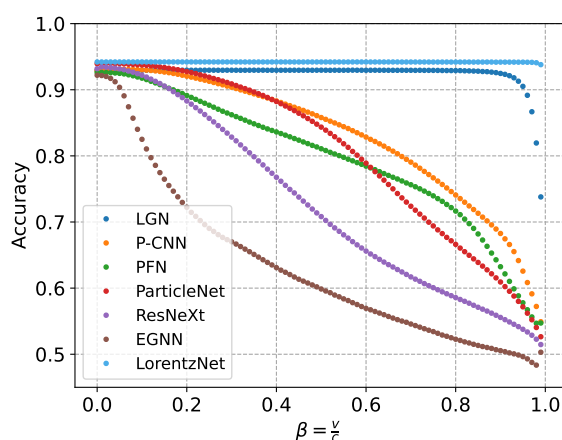
We choose the best-performed architecture among the models with and without fully Lorentz group symmetry (i.e., the LorentzNet and the ParticleNet) to compare. The inductive bias in ParticleNet is a subgroup symmetry of the Lorentz group, which only considers the Lorentz boosts on the  $z$ -axis and the rotation on the  $x - y$  plane, while LorentzNet is symmetric to the Lorentz group. We randomly select 5%, 1%, and 0.5% fraction of training data to train the LorentzNet and ParticleNet on the top tagging dataset, and we test their performances on the same test data with size 400k. The training strategy keeps the same as the experiments on the full training data. The results are reported in table 3. The gap between the tagging accuracy and AUC between LorentzNet and ParticleNet becomes larger as the number of training data becomes smaller. The results clearly show the benefit of the preservation of Lorentz group symmetry in jet tagging.

### 4.4 Equivariance test

Another advantage of symmetry-preserving deep learning models is their robustness under Lorentz transformation. To verify it, we rotate the test data by Lorentz transformation with different scales of  $\beta$  along the  $x$ -axis, i.e., the value of  $(E, p_x)$  in the 4-momentum vector will be rotated. As  $\beta$  becomes larger, the difference between the distributions of training and test data will become larger. We test the model trained on the original training data,

Training Fraction	Model	Accuracy	AUC	$1/\varepsilon_B$ ( $\varepsilon_S = 0.5$ )	$1/\varepsilon_B$ ( $\varepsilon_S = 0.3$ )
0.5%	ParticleNet	0.913	0.9687	$77 \pm 4$	$199 \pm 14$
	LorentzNet	<b>0.929</b>	<b>0.9793</b>	<b><math>176 \pm 14</math></b>	<b><math>562 \pm 72</math></b>
1%	ParticleNet	0.919	0.9734	$103 \pm 5$	$287 \pm 19$
	LorentzNet	<b>0.932</b>	<b>0.9812</b>	<b><math>209 \pm 5</math></b>	<b><math>697 \pm 58</math></b>
5%	ParticleNet	0.931	0.9807	$195 \pm 4$	$609 \pm 35$
	LorentzNet	<b>0.937</b>	<b>0.9839</b>	<b><math>293 \pm 12</math></b>	<b><math>1108 \pm 84</math></b>

**Table 3.** Performance comparison between LorentzNet and ParticleNet on top tagging dataset by a fraction of training data. The results are all averaged on 6 runs.



**Figure 3.** Equivariant test under Lorentz boosts on top tagging dataset.

and the tagging accuracy on the rotated test data is reported in figure 3. The horizontal axis of figure 3 shows the value of  $\beta$  and the vertical axis shows the tagging accuracy on the top tagging dataset under Lorentz transformation with corresponding  $\beta$ . The results show that the accuracy of LorentzNet and LGN on the test data after Lorentz transformation is robust in a large range of  $\beta$ , while the test accuracy of other non-equivariant models will drop as  $\beta$  becomes larger. According to special relativity, the fundamental quantities to clarify the particles will not be changed. Even compared with LGN, LorentzNet is more stable when  $\beta$  approaches 1, and the instability of LGN is caused by the rounding errors in float arithmetic as described in its original paper [61].

#### 4.5 Ablation study

In this section, we report the results of the ablation study to further demonstrate the effectiveness of the components in LorentzNet. To show the effectiveness of keeping the Lorentz group equivariance, we directly use  $x_i^l, x_j^l$  as inputs of the  $\phi_e$  to break the Lorentz group equivariance because  $x_i^l, x_j^l$  are not Lorentz group invariant variables, i.e.,

Model	Equivariance	Accuracy	AUC	$1/\varepsilon_B$ ( $\varepsilon_S = 0.5$ )	$1/\varepsilon_B$ ( $\varepsilon_S = 0.3$ )
LorentzNet (w/o)	<b>✗</b>	0.934	0.9832	290 ± 30	1105 ± 59
LorentzNet	<b>✓</b>	<b>0.942</b>	<b>0.9868</b>	<b>498 ± 18</b>	<b>2195 ± 173</b>

**Table 4.** Performance comparison between LorentzNet and corresponding non-equivariant version on top tagging dataset. Both of the results are averaged on 6 runs.

Model	Equivariance	Time on CPU (ms/batch)	Time on GPU (ms/batch)	#Params
ResNeXt	<b>✗</b>	5.5	0.34	1.46M
P-CNN	<b>✗</b>	<b>0.6</b>	<b>0.11</b>	348k
PFN	<b>✗</b>	<b>0.6</b>	0.12	82k
ParticleNet	<b>✗</b>	11.0	0.19	366k
EGNN	E(4)	30.0	0.30	222k
LGN	SO <sup>+</sup> (1,3)	51.4	1.66	4.5k
LorentzNet	SO <sup>+</sup> (1,3)	32.9	0.34	224k

**Table 5.** Inference time of each model on both CPU and GPU along with their parameter numbers. Models are executed on a cluster with an Intel Xeon CPU E5-2698 v4 and an Nvidia Tesla V100 32GB. All the inference times are collected with a batch size of 100.

the equation (3.2) is replaced by

$$m_{ij} = \phi_e(x_i^l, x_j^l, h_i^l, h_j^l). \tag{4.1}$$

We name this variant as LorentzNet without equivariance (abbreviated as LorentzNet (w/o)). We compare the performance of LorentzNet (w/o) with LorentzNet on the top tagging dataset. The hyperparameters of training LorentzNet (w/o) keep the same as LorentzNet and we train LorentzNet (w/o) till it converges. As shown in table 4, LorentzNet (w/o) performs worse than LorentzNet, which shows the necessity of the Lorentz group equivariance on the tagging performance.

#### 4.6 Computational cost

We report the inference time of LorentzNet and other baseline models. The inference time of each model on both CPU and GPU along with the number of their parameters are reported in table 5. The number of trainable parameters of LorentzNet is in the same order as P-CNN and ParticleNet. Models are executed on a cluster with an Intel Xeon CPU E5-2698 v4 and an Nvidia Tesla V100 32GB GPU. All the inference times are collected with a batch size of 100. As shown in table 5, the inference time on GPU of LorentzNet is slightly larger than P-CNN, PFN, and particleNet, and is comparable with ResNeXt. Especially, compared with LGN, the LorentzNet is almost 5 times faster on GPU, although the number of trainable parameters of LGN is much smaller. For the inference time on CPU,

LorentzNet is also faster than LGN. Both results demonstrate the efficiency of LorentzNet. Since there is no need to compute the high-order tensors in LorentzNet, it is more efficient. See appendix C for more computational comparisons of LorentzNet and LGN. The main computational cost of LorentzNet comes from its message passing on the fully connected graph. Its cost quadratically depends on the number of particles in a jet. This cost can be reduced by clustering the nodes and we will explore it in future studies.

## 5 Conclusion

We have presented LorentzNet, a Lorentz group equivariant graph neural network for jet tagging. Experiments on two representative jet tagging datasets show that LorentzNet achieves substantial performance improvement over all existing methods. The efficient design of the message passing to preserve the Lorentz symmetry significantly reduces both training and inference time compared to LGN. Moreover, the efficient design enhances the generalization power, allowing LorentzNet to reach highly competitive performance with only a few thousand jets for training. Although the effectiveness of LorentzNet is only demonstrated on two jet tagging tasks in this article, the symmetry-preserving architecture can be applied to a broad range of tasks in particle physics, such as regression of the true jet mass or even the full 4-momentum, generation of jets for fast simulation, classification of the entire collision event, mitigation of pileup interactions, and more. As ever-larger datasets become available in particle physics, we hope that LorentzNet will help enable more accurate prediction and accelerate the discovery of interesting new phenomena.

## Acknowledgments

We thank Prof. Qing-Hong Cao (Peking University) for his helpful discussions on AI methods in physical system modeling. We also thank Wei Chen (MSRA) for her helpful discussions on symmetry-preserving deep learning models and support. This project was funded by Microsoft Research.



## A Details about LGN

In this section, we first introduce the architecture of LGN and analyze its stability and computational efficiency.

The input of LGN is the same as LorentzNet. The node features  $h^{l+1}$  for LGN is obtained by

$$h_i^l \oplus CG(h_i^l \otimes h_i^l) \oplus CG\left(\sum_j \phi_x(\|x_i - x_j\|^2)(x_i - x_j) \otimes h_j^l\right), \quad (\text{A.1})$$

where  $CG(\cdot \otimes \cdot)$  denotes the Clebsch-Gordan tensor product and  $\phi_x(\cdot)$  denotes a neural network. The term  $h_i^l$  is the shortcut connection. The middle term reflects the self-interaction and the last term reflects the interaction of the  $i$ -th particle with other particles. A multi-layer perceptron layer will be added to operate on the scalar representations of the Lorentz group after the operating equation (A.1). The Clebsch-Gordan tensor product projects tensors to different orders of an irreducible representation of the Lorentz group, i.e.,

$$e_{l,m}^{(k,n)} \rightarrow \sum H_{(k,n),l,m}^{(k_1,n_1),l_1,m_1;(k_2,n_2),l_2,m_2} e_{l_1,m_1}^{(k_1,n_1)} \otimes e_{l_2,m_2}^{(k_2,n_2)}, \quad (\text{A.2})$$

where  $H_{(k,n),l,m}^{(k_1,n_1),l_1,m_1;(k_2,n_2),l_2,m_2}$  denotes the Clebsch-Gordan coefficient,  $e_{(l,m)}^{(k,n)}$  denotes the canonical basis of the  $(k, n)$ -order irreducible representation of the Lorentz group where  $l = |k - n|/2, \dots, (k + n)/2; m = -l, \dots, l$ .

Ideally, the universal approximation of the LGN can be achieved if the order of tensors that the LGN realized is sufficiently high. However, in practice, the order of tensors should be cut off for efficient computation and there is no direct prior to determining this cutoff in many applications.

In terms of computational efficiency, one problem is that the tensor product brings the computational cost to (eq. (A.2)) compared with the simple dot product in LorentzNet. Another problem is that the tensor product operation is a weighted sum of products of the input features. Without normalization, the output of the network will be sensitive to the scale of the input, which leads to instability (e.g., value explosion or vanishing) in both the forward and backward processes. To ensure stability, the input 4-momenta were scaled by a factor of 0.005 on the top tagging dataset in the original paper [61]. The scaling factor is a hyper-parameter to tune for different datasets. However, there is no principle way designed for the scaling or normalization which can stabilize the training process and keep the Lorentz symmetry meanwhile. Therefore, compared with LGN, LorentzNet can achieve a better tradeoff between approximation ability and computational efficiency.

In the experiments of LGN on the quark-gluon dataset, we follow its public code base in <https://github.com/fizisist/LorentzGroupNetwork>. The batch size is set to be 48 and we run 53 epochs in total. Other parameters are all aligned with its original paper [61]. The wall block time of running one epoch of LGN is more than 10 times larger than other baseline models, which is also reported in its original paper [61] (e.g., 7.5 hours per epoch with a batch size of 8 on top tagging). The computational efficiency is the main bottleneck for further evaluation of LGN.



## B Relation with EGNN

The message passing of LorentzNet is also inspired by EGNN, which is proposed to ensure  $E(n)$  group equivariance. If the output is a scalar, the message passing of EGNN is written as:

$$m_{ij}^{l+1} = \phi_e(h_i^l, h_j^l, \|x_i - x_j\|_E); \tag{B.1}$$

$$h_i^{l+1} = \phi_h \left( h_i^l, \sum_{j \neq i} m_{ij}^l \right), \tag{B.2}$$

where  $\|\cdot\|_E$  denotes the  $L^2$  norm under Euclidean metric.

We discuss the difference between LorentzNet and EGNN. First, EGNN uses the Euclidean norm of the relative distance as the only scalar information of the vector. Since relative distance can not recover the information of the angles (i.e., the inner products), its expressiveness is arguable [70]. In the context of  $E(n)$  equivariant, the inner product  $\langle x_i, x_j \rangle$  can not be added because of the translation invariance. As we target Lorentz group symmetry, we can include the Minkowski inner product for information completeness. Second, EGNN does not include the coordinate update component for scalar prediction, while LorentzNet includes the Minkowski dot product attention for expressiveness. For fair comparison in the experiments, we include the coordinate update step  $x_i^{l+1} = x_i^l + \sum_{j \neq i} (x_j - x_i) \phi_x(m_{ij}^l)$  in EGNN to ensure the difference only caused by different groups.

In the experiments of EGNN on top and quark-gluon dataset, we follow its public code base in <https://github.com/vgsatorras/egnn>. We use a batch size of 128 in the training process for 35 epochs in total. The ADAMW [65] optimizer, with a weight decay of 0.01, is used to minimize the cross-entropy loss. Firstly a linear warm-up period of 4 epochs is applied to reach the initial learning rate  $1 \times 10^{-3}$ . Then a COSINEANNEALINGWARMRESTARTS [66] learning rate schedule with  $T_0 = 4$ ,  $T_{\text{mult}} = 2$  is adopted for next 28 epochs. Finally, an exponential learning rate decay with  $\gamma = 0.5$  is used for the last 3 epochs. We test the model on the validation dataset at the end of each training epoch, and the model with the highest validation accuracy is saved as our best model for the final test.

## C Computational cost in comparison to LGN

We report the inference time of LorentzNet and LGN with two different scales on both CPU and GPU along with the parameter numbers in table 6. The inference time of LorentzNet (small) is 31 times faster than LGN (large) on CPU and 12 times faster on GPU, while the number of their trainable parameters is approximately equal. Besides, LorentzNet (small) still reaches 21 times faster than LGN on CPU and 10 times faster on GPU, even when the network of LGN is smaller than LorentzNet (small). The comparison under the approximately equal number of parameters shows that LorentzNet is much more efficient.

We also report the performance of LorentzNet and LGN in table 7. The LorentzNet (small) still shows a highly competitive performance even when the number of parameters is reduced to only 5% of the original model. We do not evaluate the performance of LGN (large) due to its slow training process.

Model	Time on CPU (ms/batch)	Time on GPU (ms/batch)	#Params
LGN	51.4	1.66	4.5k
LGN (large)	77.0	1.88	12k
LorentzNet	32.9	0.34	224k
LorentzNet (small)	2.48	0.16	12k

**Table 6.** Inference time of each model on both CPU and GPU along with their parameter numbers. Models are executed on a cluster with an Intel Xeon CPU E5-2698 v4 and an Nvidia Tesla V100 32GB. Both of the inference times are collected with a batch size of 100.

Model	Accuracy	AUC	$1/\varepsilon_B$ ( $\varepsilon_S = 0.5$ )	$1/\varepsilon_B$ ( $\varepsilon_S = 0.3$ )
LGN	0.929	0.9640	$124 \pm 20$	$435 \pm 95$
LGN (large) <sup>†</sup>	–	–	–	–
LorentzNet	0.942	0.9868	$498 \pm 18$	$2195 \pm 173$
LorentzNet (small)	0.941	0.9861	$428 \pm 29$	$1813 \pm 195$

**Table 7.** Performance comparison between LorentzNet and LGN on top tagging dataset. The results for LorentzNet are averaged on 6 runs. The results for LGN are referred to [61]. <sup>†</sup>We do not evaluate this model since it takes too long to train.

**Implementation details.** LorentzNet (small) follows the settings in section 3.3 except all the width of latent space is set to be 16. LGN (large) follows the settings in [61] except the numbers of channels are chosen as  $N_{\text{ch}}^{(0)} = 6$ ,  $N_{\text{ch}}^{(1)} = 6$ ,  $N_{\text{ch}}^{(2)} = 6$ ,  $N_{\text{ch}}^{(3)} = 6$ .

**Open Access.** This article is distributed under the terms of the Creative Commons Attribution License ([CC-BY 4.0](https://creativecommons.org/licenses/by/4.0/)), which permits any use, distribution and reproduction in any medium, provided the original author(s) and source are credited. SCOAP<sup>3</sup> supports the goals of the International Year of Basic Sciences for Sustainable Development.

## References

- [1] R. Kogler et al., *Jet Substructure at the Large Hadron Collider: Experimental Review*, *Rev. Mod. Phys.* **91** (2019) 045003 [[arXiv:1803.06991](https://arxiv.org/abs/1803.06991)] [[INSPIRE](#)].
- [2] A. Abdesselam et al., *Boosted Objects: A Probe of Beyond the Standard Model Physics*, *Eur. Phys. J. C* **71** (2011) 1661 [[arXiv:1012.5412](https://arxiv.org/abs/1012.5412)] [[INSPIRE](#)].
- [3] A. Altheimer et al., *Jet Substructure at the Tevatron and LHC: New results, new tools, new benchmarks*, *J. Phys. G* **39** (2012) 063001 [[arXiv:1201.0008](https://arxiv.org/abs/1201.0008)] [[INSPIRE](#)].
- [4] A. Altheimer et al., *Boosted Objects and Jet Substructure at the LHC. Report of BOOST2012, held at IFIC Valencia, 23rd-27th of July 2012*, *Eur. Phys. J. C* **74** (2014) 2792 [[arXiv:1311.2708](https://arxiv.org/abs/1311.2708)] [[INSPIRE](#)].

- [5] D. Adams et al., *Towards an Understanding of the Correlations in Jet Substructure*, *Eur. Phys. J. C* **75** (2015) 409 [[arXiv:1504.00679](#)] [[INSPIRE](#)].
- [6] S. Marzani, G. Soyez and M. Spannowsky, *Looking inside jets: an introduction to jet substructure and boosted-object phenomenology*, vol. 958, Springer, Cham, Germany (2019) [[DOI](#)] [[arXiv:1901.10342](#)] [[INSPIRE](#)].
- [7] 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. Rept.* **841** (2020) 1 [[arXiv:1709.04464](#)] [[INSPIRE](#)].
- [8] P.T. Komiske, E.M. Metodiev and J. Thaler, *Energy flow polynomials: A complete linear basis for jet substructure*, *JHEP* **04** (2018) 013 [[arXiv:1712.07124](#)] [[INSPIRE](#)].
- [9] K. Datta and A. Larkoski, *How Much Information is in a Jet?*, *JHEP* **06** (2017) 073 [[arXiv:1704.08249](#)] [[INSPIRE](#)].
- [10] K. Datta and A.J. Larkoski, *Novel Jet Observables from Machine Learning*, *JHEP* **03** (2018) 086 [[arXiv:1710.01305](#)] [[INSPIRE](#)].
- [11] K. Datta, A. Larkoski and B. Nachman, *Automating the Construction of Jet Observables with Machine Learning*, *Phys. Rev. D* **100** (2019) 095016 [[arXiv:1902.07180](#)] [[INSPIRE](#)].
- [12] A. Chakraborty, S.H. Lim and M.M. Nojiri, *Interpretable deep learning for two-prong jet classification with jet spectra*, *JHEP* **07** (2019) 135 [[arXiv:1904.02092](#)] [[INSPIRE](#)].
- [13] A. Chakraborty, S.H. Lim, M.M. Nojiri and M. Takeuchi, *Neural Network-based Top Tagger with Two-Point Energy Correlations and Geometry of Soft Emissions*, *JHEP* **07** (2020) 111 [[arXiv:2003.11787](#)] [[INSPIRE](#)].
- [14] J. Cogan, M. Kagan, E. Strauss and A. Schwartzman, *Jet-Images: Computer Vision Inspired Techniques for Jet Tagging*, *JHEP* **02** (2015) 118 [[arXiv:1407.5675](#)] [[INSPIRE](#)].
- [15] L.G. Almeida, M. Backović, M. Cliche, S.J. Lee and M. Perelstein, *Playing Tag with ANN: Boosted Top Identification with Pattern Recognition*, *JHEP* **07** (2015) 086 [[arXiv:1501.05968](#)] [[INSPIRE](#)].
- [16] 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** (2016) 094034 [[arXiv:1603.09349](#)] [[INSPIRE](#)].
- [17] L. de Oliveira, M. Kagan, L. Mackey, B. Nachman and A. Schwartzman, *Jet-images — deep learning edition*, *JHEP* **07** (2016) 069 [[arXiv:1511.05190](#)] [[INSPIRE](#)].
- [18] J. Barnard, E.N. Dawe, M.J. Dolan and N. Rajcic, *Parton Shower Uncertainties in Jet Substructure Analyses with Deep Neural Networks*, *Phys. Rev. D* **95** (2017) 014018 [[arXiv:1609.00607](#)] [[INSPIRE](#)].
- [19] P.T. Komiske, E.M. Metodiev and M.D. Schwartz, *Deep learning in color: towards automated quark/gluon jet discrimination*, *JHEP* **01** (2017) 110 [[arXiv:1612.01551](#)] [[INSPIRE](#)].
- [20] G. Kasieczka, T. Plehn, M. Russell and T. Schell, *Deep-learning Top Taggers or The End of QCD?*, *JHEP* **05** (2017) 006 [[arXiv:1701.08784](#)] [[INSPIRE](#)].
- [21] S. Macaluso and D. Shih, *Pulling Out All the Tops with Computer Vision and Deep Learning*, *JHEP* **10** (2018) 121 [[arXiv:1803.00107](#)] [[INSPIRE](#)].
- [22] S. Choi, S.J. Lee and M. Perelstein, *Infrared Safety of a Neural-Net Top Tagging Algorithm*, *JHEP* **02** (2019) 132 [[arXiv:1806.01263](#)] [[INSPIRE](#)].

- [23] F.A. Dreyer, G.P. Salam and G. Soyez, *The Lund Jet Plane*, *JHEP* **12** (2018) 064 [[arXiv:1807.04758](#)] [[INSPIRE](#)].
- [24] J. Lin, M. Freytsis, I. Moutl and B. Nachman, *Boosting  $H \rightarrow b\bar{b}$  with Machine Learning*, *JHEP* **10** (2018) 101 [[arXiv:1807.10768](#)] [[INSPIRE](#)].
- [25] Y.-L. Du et al., *Identifying the nature of the QCD transition in relativistic collision of heavy nuclei with deep learning*, *Eur. Phys. J. C* **80** (2020) 516 [[arXiv:1910.11530](#)] [[INSPIRE](#)].
- [26] J. Li, T. Li and F.-Z. Xu, *Reconstructing boosted Higgs jets from event image segmentation*, *JHEP* **04** (2021) 156 [[arXiv:2008.13529](#)] [[INSPIRE](#)].
- [27] J. Filipek, S.-C. Hsu, J. Kruper, K. Mohan and B. Nachman, *Identifying the Quantum Properties of Hadronic Resonances using Machine Learning*, [arXiv:2105.04582](#) [[INSPIRE](#)].
- [28] K. Fraser and M.D. Schwartz, *Jet Charge and Machine Learning*, *JHEP* **10** (2018) 093 [[arXiv:1803.08066](#)] [[INSPIRE](#)].
- [29] CMS collaboration, *Identification of heavy, energetic, hadronically decaying particles using machine-learning techniques*, *2020 JINST* **15** P06005 [[arXiv:2004.08262](#)] [[INSPIRE](#)].
- [30] 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** (2016) 112002 [[arXiv:1607.08633](#)] [[INSPIRE](#)].
- [31] J. Pearkes, W. Fedorko, A. Lister and C. Gay, *Jet Constituents for Deep Neural Network Based Top Quark Tagging*, [arXiv:1704.02124](#) [[INSPIRE](#)].
- [32] A. Butter, G. Kasieczka, T. Plehn and M. Russell, *Deep-learned Top Tagging with a Lorentz Layer*, *SciPost Phys.* **5** (2018) 028 [[arXiv:1707.08966](#)] [[INSPIRE](#)].
- [33] G. Kasieczka, N. Kiefer, T. Plehn and J.M. Thompson, *Quark-Gluon Tagging: Machine Learning vs Detector*, *SciPost Phys.* **6** (2019) 069 [[arXiv:1812.09223](#)] [[INSPIRE](#)].
- [34] 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](#) [[INSPIRE](#)].
- [35] E. Bols, J. Kieseler, M. Verzetti, M. Stoye and A. Stakia, *Jet Flavour Classification Using DeepJet*, *2020 JINST* **15** P12012 [[arXiv:2008.10519](#)] [[INSPIRE](#)].
- [36] P.T. Komiske, E.M. Metodiev and J. Thaler, *Energy Flow Networks: Deep Sets for Particle Jets*, *JHEP* **01** (2019) 121 [[arXiv:1810.05165](#)] [[INSPIRE](#)].
- [37] H. Qu and L. Gouskos, *ParticleNet: Jet Tagging via Particle Clouds*, *Phys. Rev. D* **101** (2020) 056019 [[arXiv:1902.08570](#)] [[INSPIRE](#)].
- [38] E.A. Moreno et al., *JEDI-net: a jet identification algorithm based on interaction networks*, *Eur. Phys. J. C* **80** (2020) 58 [[arXiv:1908.05318](#)] [[INSPIRE](#)].
- [39] E.A. Moreno et al., *Interaction networks for the identification of boosted  $H \rightarrow b\bar{b}$  decays*, *Phys. Rev. D* **102** (2020) 012010 [[arXiv:1909.12285](#)] [[INSPIRE](#)].
- [40] V. Mikuni and F. Canelli, *ABCNet: An attention-based method for particle tagging*, *Eur. Phys. J. Plus* **135** (2020) 463 [[arXiv:2001.05311](#)] [[INSPIRE](#)].
- [41] E. Bernreuther, T. Finke, F. Kahlhoefer, M. Krämer and A. Mück, *Casting a graph net to catch dark showers*, *SciPost Phys.* **10** (2021) 046 [[arXiv:2006.08639](#)] [[INSPIRE](#)].

- [42] J. Guo, J. Li, T. Li and R. Zhang, *Boosted Higgs boson jet reconstruction via a graph neural network*, *Phys. Rev. D* **103** (2021) 116025 [[arXiv:2010.05464](#)] [[INSPIRE](#)].
- [43] M.J. Dolan and A. Ore, *Equivariant Energy Flow Networks for Jet Tagging*, *Phys. Rev. D* **103** (2021) 074022 [[arXiv:2012.00964](#)] [[INSPIRE](#)].
- [44] V. Mikuni and F. Canelli, *Point cloud transformers applied to collider physics*, *Mach. Learn. Sci. Tech.* **2** (2021) 035027 [[arXiv:2102.05073](#)] [[INSPIRE](#)].
- [45] P. Konar, V.S. Ngairangbam and M. Spannowsky, *Energy-weighted message passing: an infra-red and collinear safe graph neural network algorithm*, *JHEP* **02** (2022) 060 [[arXiv:2109.14636](#)] [[INSPIRE](#)].
- [46] C. Shimmin, *Particle Convolution for High Energy Physics*, [arXiv:2107.02908](#) [[INSPIRE](#)].
- [47] G. Louppe, K. Cho, C. Becot and K. Cranmer, *QCD-Aware Recursive Neural Networks for Jet Physics*, *JHEP* **01** (2019) 057 [[arXiv:1702.00748](#)] [[INSPIRE](#)].
- [48] T. Cheng, *Recursive Neural Networks in Quark/Gluon Tagging*, *Comput. Softw. Big Sci.* **2** (2018) 3 [[arXiv:1711.02633](#)] [[INSPIRE](#)].
- [49] F.A. Dreyer and H. Qu, *Jet tagging in the Lund plane with graph networks*, *JHEP* **03** (2021) 052 [[arXiv:2012.08526](#)] [[INSPIRE](#)].
- [50] F. Dreyer, G. Soyez and A. Takacs, *Quarks and gluons in the Lund plane*, [arXiv:2112.09140](#) [[INSPIRE](#)].
- [51] T. Cohen and M. Welling, *Group equivariant convolutional networks*, in *International conference on machine learning*, pp. 2990–2999, PMLR (2016) [[arXiv:1602.07576](#)].
- [52] M. Zaheer, S. Kottur, S. Ravanbakhsh, B. Póczos, R.R. Salakhutdinov and A.J. Smola, *Deep sets*, *Adv. Neural Inf. Process. Syst.* **30** (2017) [[arXiv:1703.06114](#)].
- [53] M. Finzi, S. Stanton, P. Izmailov and A.G. Wilson, *Generalizing convolutional neural networks for equivariance to lie groups on arbitrary continuous data*, in *International Conference on Machine Learning*, pp. 3165–3176, PMLR (2020) [[arXiv:2002.12880](#)].
- [54] G. Kanwar et al., *Equivariant flow-based sampling for lattice gauge theory*, *Phys. Rev. Lett.* **125** (2020) 121601 [[arXiv:2003.06413](#)] [[INSPIRE](#)].
- [55] D. Boyda et al., *Sampling using  $SU(N)$  gauge equivariant flows*, *Phys. Rev. D* **103** (2021) 074504 [[arXiv:2008.05456](#)] [[INSPIRE](#)].
- [56] V.G. Satorras, E. Hoogeboom and M. Welling,  *$E(n)$  equivariant graph neural networks*, [arXiv:2102.09844](#).
- [57] M.J. Hutchinson, C. Le Lan, S. Zaidi, E. Dupont, Y.W. Teh and H. Kim, *Lietransformer: Equivariant self-attention for lie groups*, in *International Conference on Machine Learning*, pp. 4533–4543, PMLR (2021) [[arXiv:2012.10885](#)].
- [58] S. Batzner et al.,  *$E(3)$ -equivariant graph neural networks for data-efficient and accurate interatomic potentials*, [arXiv:2101.03164](#).
- [59] K. Lenc and A. Vedaldi, *Understanding image representations by measuring their equivariance and equivalence*, in *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 991–999 (2015) [[DOI](#)].
- [60] M. Erdmann, E. Geiser, Y. Rath and M. Rieger, *Lorentz Boost Networks: Autonomous Physics-Inspired Feature Engineering*, *2019 JINST* **14** P06006 [[arXiv:1812.09722](#)] [[INSPIRE](#)].

- [61] A. Bogatskiy, B. Anderson, J. Offermann, M. Roussi, D. Miller and R. Kondor, *Lorentz group equivariant neural network for particle physics*, in *International Conference on Machine Learning*, pp. 992–1002, PMLR (2020) [[arXiv:2006.04780](#)].
- [62] F. Scarselli, M. Gori, A.C. Tsoi, M. Hagenbuchner and G. Monfardini, *The graph neural network model*, *IEEE Trans. Neural Networks* **20** (2008) 61.
- [63] J. Gilmer, S.S. Schoenholz, P.F. Riley, O. Vinyals and G.E. Dahl, *Neural message passing for quantum chemistry*, in *International conference on machine learning*, pp. 1263–1272, PMLR (2017) [[arXiv:1704.01212](#)].
- [64] S. Villar, D. Hogg, K. Storey-Fisher, W. Yao and B. Blum-Smith, *Scalars are universal: Equivariant machine learning, structured like classical physics*, *Adv. Neural Inf. Process. Syst.* **34** (2021) [[arXiv:2106.06610](#)].
- [65] I. Loshchilov and F. Hutter, *Decoupled weight decay regularization*, in *International Conference on Learning Representations*, (2018) [[arXiv:1711.05101](#)].
- [66] I. Loshchilov and F. Hutter, *SGDR: Stochastic gradient descent with warm restarts*, [arXiv:1608.03983](#).
- [67] G. Kasieczka, T. Plehn and M. Russel, *Top quark tagging reference dataset*, [10.5281/zenodo.2603256](#) (2019).
- [68] S. Xie, R. Girshick, P. Dollár, Z. Tu and K. He, *Aggregated residual transformations for deep neural networks*, in *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 1492–1500 (2017) [[DOI](#)].
- [69] CMS collaboration, *Boosted jet identification using particle candidates and deep neural networks*, [CMS-DP-2017-049](#) (2017).
- [70] W. Du et al., *Equivariant vector field network for many-body system modeling*, [arXiv:2110.14811](#).