

Jet tagging with more-interaction particle transformer*

Yifan Wu (吴佚凡)¹ Kun Wang (王坤)^{1†} Congqiao Li (李聪乔)²
Huilin Qu (曲慧麟)³ Jingya Zhu (朱经亚)⁴

¹College of Science, University of Shanghai for Science and Technology, Shanghai 200093, China

²School of Physics and State Key Laboratory of Nuclear Physics and Technology, Peking University, Beijing 100871, China

³CERN, EP Department, CH-1211 Geneva 23, Switzerland

⁴School of Physics and Electronics, Henan University, Kaifeng 475004, China

Abstract: In this paper, we introduce the More-Interaction Particle Transformer (MIParT), a novel deep-learning neural network designed for jet tagging. This framework incorporates our own design, the More-Interaction Attention (MIA) mechanism, which increases the dimensionality of particle interaction embeddings. We tested MIParT using the top tagging and quark-gluon datasets. Our results show that MIParT not only matches the accuracy and AUC of LorentzNet and a series of Lorentz-equivariant methods, but also significantly outperforms the ParT model in background rejection. Specifically, it improves background rejection by approximately 25% with a signal efficiency of 30% on the top tagging dataset and by 3% on the quark-gluon dataset. Additionally, MIParT requires only 30% of the parameters and 53% of the computational complexity needed by ParT, proving that high performance can be achieved with reduced model complexity. For very large datasets, we double the dimension of particle embeddings, referring to this variant as MIParT-Large (MIParT-L). We found that MIParT-L can further capitalize on the knowledge from large datasets. From a model pre-trained on the 100M JetClass dataset, the background rejection performance of fine-tuned MIParT-L improves by 39% on the top tagging dataset and by 6% on the quark-gluon dataset, surpassing that of fine-tuned ParT. Specifically, the background rejection of fine-tuned MIParT-L improves by an additional 2% compared to that of fine-tuned ParT. These results suggest that MIParT has the potential to increase the efficiency of benchmarks for jet tagging and event identification in particle physics.

Keywords: jet tagging, collider physics, machine learning

DOI: 10.1088/1674-1137/ad7f3d **CSTR:** 32044.14.ChinesePhysicsC.49013110

I. INTRODUCTION

Jet identification has become a key area where machine learning is applied in high-energy physics, making significant progress in the past few years [1, 2]. Jets are collimated sprays of particles produced in high-energy collisions, typically from quarks, gluons, or hadronic decay of heavy particles. The process known as jet tagging, which involves identifying the particle that initiated the jet, is complex and challenging. This complexity arises because the initial particle evolves into a jet through multiple stages, increasing the number of particles within the jet and obscuring the characteristics of the initiating particle.

By analyzing the constituents of a jet, it is possible to determine the type of particle that initiated the jet. This

identification is critical for revealing fundamental physical processes and discovering new particles. Initially, jet tagging relied heavily on the quantum chromodynamics (QCD) theory, which provided methodologies for distinguishing between quark and gluon jets [3–9]. With the advent of machine learning, a variety of new jet tagging methods have been introduced that exploit different machine learning models to increase the breadth and accuracy of previous techniques [10–15]. Recent advances in deep learning have further refined jet tagging methods, allowing modern algorithms to effectively process large and complex datasets. These algorithms are suitable for identifying subtle patterns that differentiate various types of jets, significantly improving the accuracy and efficiency of jet tagging [16–20]. The exceptional ability of deep learning to manage large datasets has been instru-

Received 14 August 2024; Accepted 23 September 2024; Published online 24 September 2024

* K. Wang and J. Zhu is Supported by the National Natural Science Foundation of China (12275066, 11605123)

† E-mail: kwang@usst.edu.cn



Content from this work may be used under the terms of the Creative Commons Attribution 3.0 licence. Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI. Article funded by SCOAP³ and published under licence by Chinese Physical Society and the Institute of High Energy Physics of the Chinese Academy of Sciences and the Institute of Modern Physics of the Chinese Academy of Sciences and IOP Publishing Ltd

mental in these advances, leading to the discovery of new physical phenomena and deepening our understanding of particle interactions.

Jet tagging has undergone many changes over the years. Initially, traditional methods relied heavily on features designed by experts based on physical principles. The introduction of machine learning brought more advanced approaches, starting with the concept of jet images. These images, representing pixelated depictions of the energy deposited by particles in a detector, constituted a pivotal development in the field. The earliest application of jet images dates back to 1991, when Pumplin introduced the idea of representing jets as images [21]. Subsequent studies, starting around 2014, were inspired by computer vision. These studies used techniques such as Fisher's Linear Discriminant, originally used in face recognition technology, to improve jet tagging [10]. By 2015, deep neural networks (DNNs) were being applied to top tagging [11]; later, convolutional neural networks (CNNs) were widely adopted in jet tagging [12–15, 22], demonstrating significant improvements in jet tagging performance.

In 2016, sequence-based representations began to gain traction in the field of jet tagging, using recurrent neural networks (RNNs) to process ordered data. This period gave rise to significant advancements with the pioneering use of Long Short-Term Memory (LSTM) networks for classification purposes [23]. Subsequently, Gated Recurrent Units (GRUs) were also used for event topology classification, further extending the applications of RNNs in this domain [18]. Simultaneously, an innovative approach combining CNNs and LSTMs, known as DeepJet, was developed. This hybrid model significantly improved the performance of quark-gluon tagging [24]. Additionally, several studies using RNNs introduced new methods and insights [25, 26]. These methods have successfully overcome the limitations associated with input size in jet tagging, providing a more flexible approach for analyzing and using jet data.

In 2017, the introduction of graph-based representations using graph neural networks (GNNs) was a significant leap forward in jet tagging [27]. Subsequently, GNNs began to be widely used in particle identification, greatly expanding the capabilities in the field [28–31]. This broad application of GNNs has opened new avenues for accurately classifying and understanding complex particle interactions.

In 2018, the exploration of point cloud representations, which treat jets as unordered sets of particles, gave rise to notable advancements. Komiske *et al.* introduced the concept of Energy Flow Networks (EFNs), which can deal with variable-length unordered particle sets effectively [32]. This method exploits the "Deep Sets" concept, developed by Zaheer *et al.* [33] in 2017, which treats jets specifically as sets of particles and represents a significant

advance in jet tagging. Crucially, it made the algorithms permutation-invariant, thereby enhancing their capability to represent complex particle interactions.

In 2019, Qu *et al.* introduced ParticleNet [34], building on the Dynamic Graph Convolutional Neural Network (DGCNN) framework developed by Wang *et al.* [35] in 2018. ParticleNet, which also treats jets as unordered sets of particles, was a significant advancement in the field. In 2022, Qu *et al.* further extended their contributions by developing the Particle Transformer (ParT) [36], which is based on the Transformer architecture [37]. By incorporating pairwise particle interaction inputs, it significantly improved jet tagging performance. Furthermore, the introduction of a new large-scale dataset, JetClass, enabled pre-training of the ParT model, thereby reaching even higher performance.

However, currently, the most efficient jet tagging models, namely pre-trained ParT models, not only require pre-training, but also have a significant number of parameters. In addition, other transformer-based jet taggers fail to outperform the DGCNN-based ParticleNet owing to an insufficient number of jets in the training samples. This indicates that transformer-based models are effective at exploiting larger training datasets by using the attention mechanism. Moreover, we observed that pairwise particle interaction inputs play a crucial role in ParT. Therefore, we aimed to construct a transformer-based jet tagging model with an increased focus on particle interaction inputs, aiming for optimal results without pre-training.

In this paper, we propose a new jet tagging method based on the Transformer architecture, called More-Interaction Particle Transformer (MIParT). We enhanced the algorithm of ParT by modifying the attention mechanism and increasing the embedding dimensions of the pairwise particle interaction inputs while reducing the total number of parameters and computational complexity. We tested MIParT on two widely used jet tagging benchmarks and found that it exhibits higher performance than existing methods. Additionally, to address the challenges posed by very large datasets, we doubled the particle embedding dimensions to construct a larger model. We pre-trained this enhanced model on the 100M JetClass dataset before fine-tuning it on smaller datasets. This approach showed measurable performance gains over fine-tuned ParT, indicating the efficacy of our modifications.

The remainder of this manuscript is organized as follows. In Sec. II, we provide an overview of various deep-learning models and specifically focus on the architecture of MIParT. In Sec. III, we detail the conducted experimental process and extensively discuss the results obtained from our analysis. In Sec. IV, we end the paper by summarizing the main conclusions and discussing their implications for future research in this area.

II. MIParT MODEL ARCHITECTURE

Traditional deep-learning models such as CNNs and RNNs face significant challenges in representing jets effectively. Image representations often struggle with incorporating particle identity, which affects performance improvement [10]. Similarly, sequence [23] and tree [25] representations impose artificial ordering on jet particles, which inherently possess no sequential structure. Considering a jet as an unordered collection of its constituent particles leads to a more natural representation. This format not only facilitates the inclusion of particle-specific features, but also guarantees permutation invariance. Among models that adopt this perspective, ParticleNet describes jets as "particle clouds", similar to the point cloud technique for 3D shape analysis in computer vision. ParticleNet uses the DGCNN architecture, with its EdgeConv operations effectively using the local spatial structures of particle clouds to achieve notably higher performance.

ParT, a transformative variant based on the Class-Attention in Image Transformers (CaiT) framework [38], integrates interaction variables as a secondary input. The self-attention mechanism of this architecture uniquely addresses all positions within the input sequence, capturing extensive range dependencies efficiently and maintaining invariance to particle order. By refining the Multi-Head Attention (MHA) mechanism to include jet-particle inter-

action variables, ParT not only outperforms traditional transformer models, but also sets a new benchmark in jet tagging. These modifications position ParT as the leading model in jet tagging.

Based on the ParT framework, we developed MIParT to enhance the input of interaction data, as depicted in Fig. 1. MIParT adopts the input formats of ParT and processes jet data with two distinct inputs:

- Particle Input x_1 : It comprises a list of C features per particle, arranged into an array of shape (N, C) , where N represents the number of particles within a jet.
- Interaction Input U_1 : It includes a matrix of C' features for each particle pair, formatted as an array of shape (N, N, C') .

The particle input is first transformed by a Multilayer Perceptron (MLP) to project feature dimensions to D_1 , resulting in an array x_1 with dimensions (N, D_1) . Similarly, the interaction input undergoes pointwise 1D convolution processing, yielding U_1 with dimensions (N, N, D_1) ; x_1 then passes through K MI-Particle Attention Blocks to generate x_2 of the same shape. In each block, U_1 serves as an additional input and is dimensionally reduced by a pointwise 1D convolution to U_2 , having dimensions (N, N, D_2) .




Fig. 1. (color online) Schematic of the More-Interaction Particle Transformer (MIParT) architecture. The particle features x_1 are processed sequentially through K MI particle attention blocks and L particle attention blocks. The interaction features U_1 are first fed to K MI particle attention blocks, then dimensionally reduced by a 1D pointwise convolution to U_2 , and then fed to L particle attention blocks. The MIParT architecture ends with the application of the Class-Attention in Image Transformers (CaiT) methodology, which uses a class token x_{class} to systematically extract and summarize information from x_3 in the class attention blocks.

Following the structural framework of ParT, x_2 progresses through L Particle Attention Blocks, enhancing with U_2 at each layer, to produce x_3 . Subsequently, using the CaiT methodology, a class token x_{class} is used to systematically extract and summarize information from x_3 in the class attention blocks. Finally, this summarized information forms a single vector that constitutes the input of a linear classifier through an MLP and a softmax function to derive the classification scores.

A. Particle Attention Block

The Particle Attention Block, a crucial element of the ParT framework, was seamlessly integrated into our MI-ParT model. The architecture of this block is based on the NormFormer design [39], specifically using Layer Normalization instead of Batch Normalization. Layer Normalization optimizes normalization by adjusting each layer individually for every single sample, enhancing model stability and overall performance across diverse datasets. The architecture of the Particle Attention Block is shown in Fig. 2. Furthermore, in this configuration, the traditional Multi-Head Attention (MHA) is substituted by Particle-Multi-Head Attention (P-MHA). This modification allows for the incorporation of particle interaction features directly into the attention mechanism, enriching the capability of the model to capture complex particle




Fig. 2. (color online) Schematic of the More-Interaction Attention (MIA) architecture. The shape of U is (N, N, C) , while both the input x and output x' have a shape (N, C) . MIA maintains a one-to-one correspondence between the feature dimensions of U , x , and the heads of MHA C .

dynamics. The P-MHA mechanism, which is key to the Particle Attention Block, is mathematically expressed as

$$\text{P-MHA}(Q, K, V) = \text{SoftMax}\left(\frac{QK^T}{\sqrt{d_k}} + U\right)V, \quad (1)$$

where Q , K , and V are the linear projections of the particle embedding x , and U represents the interaction embedding. The dimensions of U are precisely aligned with the attention heads in the MHA mechanism, thereby facilitating the integration of particle interaction features. The specific implementation of P-MHA can be found in Ref. [36]. This integration significantly enhances the ability of the model to capture complex particle interactions, which is crucial in particle physics applications.

B. MI-Particle Attention Block

In the original P-MHA mechanism, the feature dimensions of U align one-to-one with the heads of MHA, both denoted as C . Increasing the feature dimensions of U necessitates a proportional increase in the number of attention heads, which significantly increases the complexity of the model. To mitigate this issue, we introduce More-Interaction Attention (MIA) and the MI-Particle Attention Block. These components replace P-MHA with MIA, as shown in Fig. 2 (MIA architecture) and Fig. 3 (MI-Particle Attention Block/Particle Attention Block architecture). The MI-Particle Attention Block incorporates Layer Normalization and the Gaussian Error Linear Unit (GELU) activation function. When the red block in Fig. 3 uses MIA, it forms the MI-Particle Attention Block. Conversely, when it uses P-MHA, it forms the Particle Attention Block. This approach allows the model to effectively use interaction inputs without significantly increasing complexity. The MIA is calculated using the following formula:

$$\text{MIA}(U, V) = \text{SoftMax}(U)V, \quad (2)$$

where V is a linear projection of the particle embedding x . In MIA, each feature dimension of U and x , as well as each head, are denoted by C , ensuring a one-to-one correspondence.

By increasing the feature dimensions of U , MIA effectively exploits the interaction inputs without significantly increasing the complexity of the model. Moreover, the MI-Particle Attention Block, which incorporates self-attention on x , acts as a supplement in front of the Particle Attention Block rather than as a replacement.

C. Class Attention Block

We incorporated the Class Attention Block from the ParT framework, inspired by the CaiT architecture. This block uses a class token x_{class} to efficiently extract in-




Fig. 3. (color online) Schematic of the MI-Particle Attention Block / Particle Attention Block architecture. Here, LN represents Layer Normalization, and GELU represents the Gaussian Error Linear Unit activation function. The block forms the MI-Particle Attention Block when using MIA and the Particle Attention Block when using P-MHA.

formation through attention mechanisms, as depicted in Fig. 4. The Multi-Head Attention inputs are defined as follows:

$$Q = W_q x_{\text{class}} + b_q, \quad (3)$$

$$K = W_k z + b_k, \quad (4)$$

$$V = W_v z + b_v, \quad (5)$$

where $z = [x_{\text{class}}, x]$, and W and b represent learnable parameters. This design ensures a lower computational overhead for the Class Attention mechanism by using the concatenated vector z .

The Class Attention Block significantly enhances feature extraction from the input x by capitalizing on the class token, thereby improving the focus of the model on essential aspects of the data. This enhancement also improves jet classification performance significantly, making the Class Attention Block a crucial component within the ParT framework.




Fig. 4. (color online) Schematic of the Class Attention Block architecture. Here, LN represents Layer Normalization, GELU represents the Gaussian Error Linear Unit activation function, and MHA stands for Multi-Head Attention block.

D. Implementation

The architecture of our MIParT model includes $K = 5$ MI-particle attention blocks, $L = 5$ particle attention blocks, and 2 class attention blocks. The choice of these hyperparameters balances complexity and accuracy; we observed an increase in accuracy with additional layers at the cost of increased complexity. Therefore, we limited the total number of attention blocks to ten. The rationale for choosing two class attention blocks is based on the CaiT framework [38], which recommends such a configuration for efficient classification. For particle embeddings x_1 , a three-layer Multi-Layer Perceptron (MLP) is used, with each layer containing 128, 512, and 64 neurons, respectively. This configuration results in embeddings with a dimensionality of $D_1 = 64$. The decision to reduce the embedding dimension compared to the ParT model was motivated by the addition of the MIA module.

This adjustment rationalizes the complexity of the model while maintaining its efficiency, thus optimizing the trade-off between performance and computational load. Each layer incorporates GELU as the activation function and Layer Normalization. Additionally, a three-layer, 64-channel pointwise 1D convolution is used for the interaction embeddings U_1 , performing convolutions only along the feature dimension. The U_1 embeddings are further processed through a single-layer, 8-channel pointwise 1D convolution to generate U_2 , achieving a dimensionality of $D_2 = 8$. This design choice maintains consistency with the ParT model, ensuring alignment with established architectural standards and facilitating comparative analysis. The MI-particle attention blocks implement MIA with 64 heads, while the P-MHA and Class Multi-Head Attention in the particle and class attention blocks use 8 heads each. A dropout rate of 0.1 is maintained in all MI-particle and particle attention blocks, with the class attention blocks being exempt from dropout.

For very large datasets, increasing the embedding dimension significantly enhances model performance. Therefore, for such datasets, we double the dimension of the particle embeddings to $D_1 = 128$. This adjustment is straightforward, requiring a change in the neuron configuration of the three-layer MLP to 128, 512, and 128. Consequently, the dimensions of \mathbf{x} and \mathbf{U} in MIA will no longer be identical; however, this discrepancy is acceptable as long as the dimension of \mathbf{x} is an integer multiple of the dimension of \mathbf{U} . We refer to this modified model as MIParT-Large (MIParT-L).

III. RESULTS AND DISCUSSION

We developed the MIParT model using the PyTorch framework [40]. Its implementation is based on Weaver¹⁾, and is referred to the implementation of ParT²⁾.

We initially evaluated the MIParT model on two widely used jet tagging benchmark datasets, namely the top tagging [16] and quark-gluon datasets [32]. The model was trained on an NVIDIA RTX 4090 GPU, using a learning rate of 0.001 and a batch size of 256. Training was limited to 15 epochs to prevent overfitting. Both datasets incorporate kinematic variables as particle input features, with particle identification information included only in the quark-gluon dataset. All these input features for the two datasets are shown in Table 1.

We then pre-trained our larger model variant, MIParT-L, on the JetClass dataset containing 100M samples [36]. This model was pre-trained on dual NVIDIA RTX 3090 GPUs using a learning rate of 0.0008 and a batch size of 384, with pre-training limited to 50 epochs to

Table 1. Summary of kinematic and particle identification variables included in the top tagging (TOP), quark-gluon (QG), and JetClass (JC) datasets. Variables present in each dataset are indicated by a star symbol (*). The table includes seven kinematic variables describing the physical characteristics of particles relative to the jet axis, six particle identification variables categorizing particles by type and charge, and four trajectory displacement features, which provide detailed information on particle trajectories.

Category	Variable	TOP	QG	JC
Kinematics	$\Delta\eta$	*	*	*
	$\Delta\phi$	*	*	*
	$\log p_T$	*	*	*
	$\log E$	*	*	*
	$\log p_T/p_T(\text{jet})$	*	*	*
	$\log E/E(\text{jet})$	*	*	*
	ΔR	*	*	*
Particle identification	charge		*	*
	Electron		*	*
	Muon		*	*
	Photon		*	*
	Charged Hadron		*	*
	Neutral Hadron		*	*
Trajectory displacement	$\tanh d_0$			*
	$\tanh d_z$			*
	σ_{d_0}			*
	σ_{d_z}			*

avoid overfitting. After pre-training, MIParT-L was fine-tuned on the top tagging and quark-gluon datasets. Interestingly, the pre-training of MIParT-L on the JetClass dataset for the top tagging dataset included only kinematic features, while for the quark-gluon dataset, both kinematic and particle identification features were included.

For fine-tuning, we replaced the last MLP for classification with a newly initialized MLP having two output nodes. All weights were then fine-tuned across the datasets for 20 epochs. We used a learning rate of 0.00016 for the pre-trained weights and 0.008 for the new MLP.

The seven kinematic input features are as follows:

- $\Delta\eta$: difference in pseudorapidity η between the particle and jet axis;
- $\Delta\phi$: difference in azimuthal angle ϕ between the particle and jet axis;

1) Weaver provides a streamlined yet flexible machine learning R&D framework for high energy physics, <https://github.com/hqucms/weaver-core>.

2) The official implementation of Particle Transformer for Jet Tagging, which includes the code and pre-trained models, https://github.com/jet-universe/particle_transformer.

- $\log p_T$: logarithm of the particle transverse momentum p_T ;
- $\log E$: logarithm of the particle energy;
- $\log p_T/p_T(\text{jet})$: logarithm of the particle p_T relative to the jet p_T ;
- $\log E/E(\text{jet})$: logarithm of the particle energy relative to the jet energy;
- ΔR : angular separation between the particle and jet axis.

The six particle identification features are as follows:

- "Charge": electric charge of the particle;
- "Electron": whether the particle is an electron;
- "Muon": whether the particle is a muon;
- "Photon": whether the particle is a photon;
- "Charged Hadron": whether the particle is a charged hadron;
- "Neutral Hadron": whether the particle is a neutral hadron.

The four trajectory displacement features in the Jet-Class are as follows:

- $\tanh d_0$: hyperbolic tangent of the transverse impact parameter value;
- $\tanh d_z$: hyperbolic tangent of the longitudinal impact parameter value;
- σ_{d_0} : error of the measured transverse impact parameter;
- σ_{d_z} : error of the measured longitudinal impact parameter.

For particle interaction features, we consider four logarithmic characteristics ($\ln \Delta, \ln k_T, \ln z, \ln m^2$) derived from the energy-momentum four-vector $p = (E, p_x, p_y, p_z)$ [41]. These features are defined as follows:

$$\Delta = \sqrt{(y_a - y_b)^2 + (\phi_a - \phi_b)^2}, \quad (6)$$

$$k_T = \min(p_{T,a}, p_{T,b})\Delta, \quad (7)$$

$$z = \min(p_{T,a}, p_{T,b}) / (p_{T,a} + p_{T,b}), \quad (8)$$

$$m^2 = (E_a + E_b)^2 - |\mathbf{p}_a + \mathbf{p}_b|^2, \quad (9)$$

where y_i is the rapidity, ϕ_i is the azimuthal angle, $p_{T,i}$ is the transverse momentum, and \mathbf{p}_i is the momentum three-vector of the particle $i = a, b$. The motivation for selecting these variables comes from their widespread adoption in several advanced neural networks [34, 36].

To evaluate the performance of the MIParT model, we conducted comparative evaluations with several popular models using the top tagging and quark-gluon datasets. Our evaluation focused on several key metrics:

- **Accuracy**: This metric quantifies the proportion of correct predictions made by the model, including both true positive and true negative identifications. Mathematically, accuracy is defined as follows:

$$\text{Accuracy} = \frac{TP + TN}{TP + TN + FN + FP}, \quad (10)$$

where TP stands for true positives, TN stands for true negatives, FN stands for false negatives, and FP stands for false positives.

- **AUC (Area Under the Curve)**: AUC is a comprehensive metric to assess model performance across all classification thresholds. This metric is derived from the Receiver Operating Characteristic (ROC) curve, which represents the true positive rate (sensitivity) against the false positive rate (1 - specificity) for various thresholds. This curve illustrates the trade-off between sensitivity and specificity. AUC values range from 0.5, which indicates no discriminatory ability (similar to random guessing), to 1.0, which represents perfect discrimination and indicates excellent ability of the model to discriminate between classes.

- **Background Rejection at a Certain Signal Efficiency, $\text{Rej}_{X\%}$** : This metric calculates the inverse of the false positive rate (FPR) when the true positive rate (TPR) is fixed at a certain percentage, commonly referred to as $\text{Rej}_{X\%}$. It is mathematically expressed as follows:

$$\text{Rej}_{X\%} = \frac{1}{\text{FPR}} \Big|_{\text{TPR}=X\%}. \quad (11)$$

For example, a $\text{Rej}_{30\%}$ value of 2500 indicates that at a TPR of 30%, the inverse of the FPR is 2500. This implies a single false positive per 2500 negative instances, highlighting the exceptional specificity and minimal er-

ror rate of the model at this level.

Top tagging is a critical task in jet tagging, which is often used in the search for new physics at the LHC. For this study, we used a top tagging dataset [16] consisting of 2M jets, with $t \rightarrow bqq'$ as the signal and q/g as the background. This dataset only provides the energy-momentum four-vectors (kinematic features) for each particle.

Figure 5 shows the performance of our MIParT model compared to other popular models on the top tagging dataset. The MIParT model achieves accuracy and AUC metrics nearly identical to those of LorentzNet [42], and

its $\text{Rej}_{50\%}$ and $\text{Rej}_{30\%}$ metrics are, within the error range, comparable to those of LorentzNet. Note that a series of Lorentz-equivariant methods demonstrate similar performance to that of LorentzNet, such as Clifford Group Equivariant Neural Networks (CGENN) [43], permutation equivariant and Lorentz invariant or covariant aggregator network (PELICAN) [44], and Lorentz-Equivariant Geometric Algebra Transformers (L-GATr) [45]. Moreover, MIParT, LorentzNet, and several Lorentz-equivariant based models significantly outperform other models, including Particle Flow Network (PFN) [32], Particle-level Convolutional Neural Network

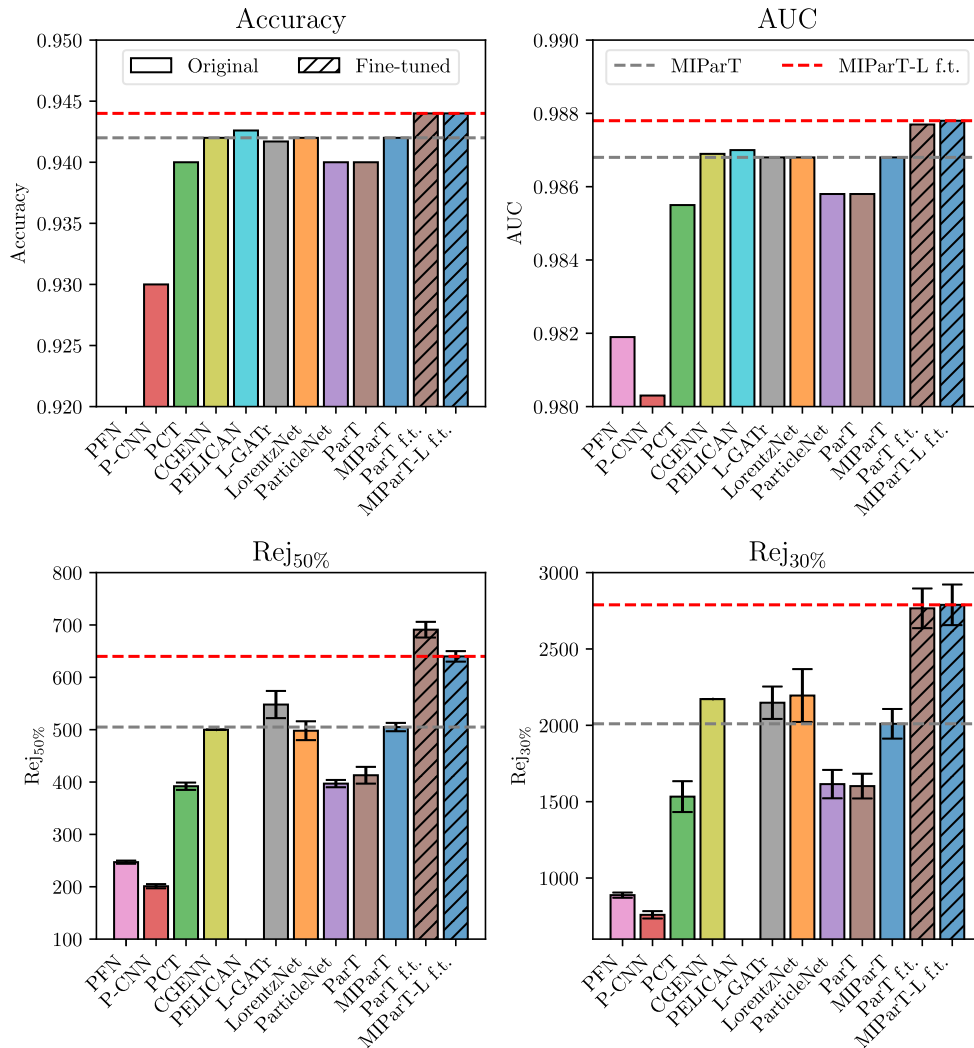


Fig. 5. (color online) Comparison of MIParT performance metrics with those of other models on the top tagging dataset. This figure shows Accuracy, AUC, $\text{Rej}_{50\%}$, and $\text{Rej}_{30\%}$ metrics for the MIParT model alongside those of Particle Flow Network (PFN) [32], Particle-level Convolutional Neural Network (P-CNN), Point Cloud Transformer (PCT) [46], Clifford Group Equivariant Neural Networks (CGENN) [43], permutation equivariant and Lorentz invariant or covariant aggregator network (PELICAN) [44], Lorentz-Equivariant Geometric Algebra Transformers (L-GATr) [45], LorentzNet [42], and ParticleNet [34], ParT [36]. Metrics of other models are extracted from published results. Detailed outcomes are provided in Table 2. Bars without slashes indicate original models without fine-tuning, while bars with slashes indicate models with fine-tuning. The gray dashed line represents the results for MIParT, whereas the red dashed line represents the results for fine-tuned MIParT-L (MIParT-L f.t.).

(P-CNN) [34], ParticleNet [34], Point Cloud Transformer (PCT) [46], and ParT [36], with metrics in the comparison extracted from published results. For the fine-tuned MIParT-L model pre-trained on the 100M Jet-Class dataset, a 39% enhancement in background rejection performance was achieved, comparable to that of the fine-tuned ParT. Detailed comparison results are presented in Table 2. The MIParT model significantly outperforms ParT in the top tagging benchmark, with approximately 25% better background rejection at a 30% signal efficiency. Among evaluated the models, MIParT, along with LorentzNet and other Lorentz-equivariant based models, ranks at the top tier, exhibiting robustness and high performance.

Quark-gluon tagging is another crucial jet tagging task. Unlike the top tagging dataset, the quark-gluon dataset [32] includes not only the kinematic features of each particle, but also particle identification information. This dataset allows for a more detailed categorization of particles, including specific distinctions among electrically charged and neutral hadrons, such as pions, kaons, and protons. Additionally, similar to the top tagging dataset, the quark-gluon dataset contains 2M jets, with quarks and gluons designated as the signal and background, respectively.

Figure 6 shows the performance of our MIParT model compared to other popular models on the quark-gluon dataset. Within this dataset, the MIParT model significantly outperforms LorentzNet across all metrics, including accuracy, AUC, $\text{Rej}_{50\%}$, and $\text{Rej}_{30\%}$, as well as other

models. Note that only the ParT model approaches the performance of our model in some metrics, but MIParT still maintains an overall lead over ParT. In comparison with other models, such as PFN [32], ABCNet [47], and PCT [46], MIParT demonstrates a substantial lead, with metrics extracted from published results. For the fine-tuned MIParT-L model pre-trained on the 100M JetClass dataset, a 6% enhancement in background rejection performance is achieved, outperforming fine-tuned ParT. Detailed comparison results on the quark-gluon dataset are presented in Table 3. MIParT achieves the best performance across all evaluation metrics, improving background rejection power by approximately 3% compared to ParT. Simultaneously, the background rejection of the fine-tuned MIParT-L model is improved by approximately 2% compared to fine-tuned ParT.

Given that MIParT shares many components with ParT and differs only in the addition of the MIA blocks, the comparative results between these two models highlight the effectiveness of the MIA block. Specifically, MIParT consists of five MIA blocks, five particle attention blocks, and two class attention blocks, whereas ParT consists of eight particle attention blocks and two class attention blocks. Thus, from the results tested on the top tagging and quark-gluon datasets, it is evident that MIParT outperforms ParT, stressing the significant role played by the MIA block. Furthermore, the effectiveness of the particle attention blocks was already established in the ParT paper [36], and the impact of the class attention blocks was tested in the CaiT framework [38].

Table 2. Performance comparison of various models on the top tagging dataset. This table lists the results for the MIParT model alongside those of other prominent models such as Particle Flow Network (PFN), Particle-level Convolutional Neural Network (P-CNN), Point Cloud Transformer (PCT) [46], Clifford Group Equivariant Neural Networks (CGENN) [43], permutation equivariant and Lorentz invariant or covariant aggregator network (PELICAN) [44], Lorentz-Equivariant Geometric Algebra Transformers (L-GATr) [45], LorentzNet [42], ParticleNet [34], and ParT [36]. Metrics of other models are extracted from published results. The results for the fine-tuned version of our model, MIParT-L f.t., are shown at the bottom of the table for comparison with those of the fine-tuned ParT model, that is, ParT f.t.

	Accuracy	AUC	$\text{Rej}_{50\%}$	$\text{Rej}_{30\%}$
PFN	—	0.9819	247±3	888±17
P-CNN	0.930	0.9803	201±4	759±24
PCT	0.940	0.9855	392±7	1533±101
CGENN	0.942	0.9869	500	2172
PELICAN	0.9426	0.9870	—	—
L-GATr	0.9417	0.9868	548±26	2148±106
LorentzNet	0.942	0.9868	498±18	2195±173
ParticleNet	0.940	0.9858	397±7	1615±93
ParT	0.940	0.9858	413±16	1602±81
MIParT (ours)	0.942	0.9868	505±8	2010±97
ParT f.t.	0.944	0.9877	691±15	2766±130
MIParT-L f.t. (ours)	0.944	0.9878	640±10	2789±133

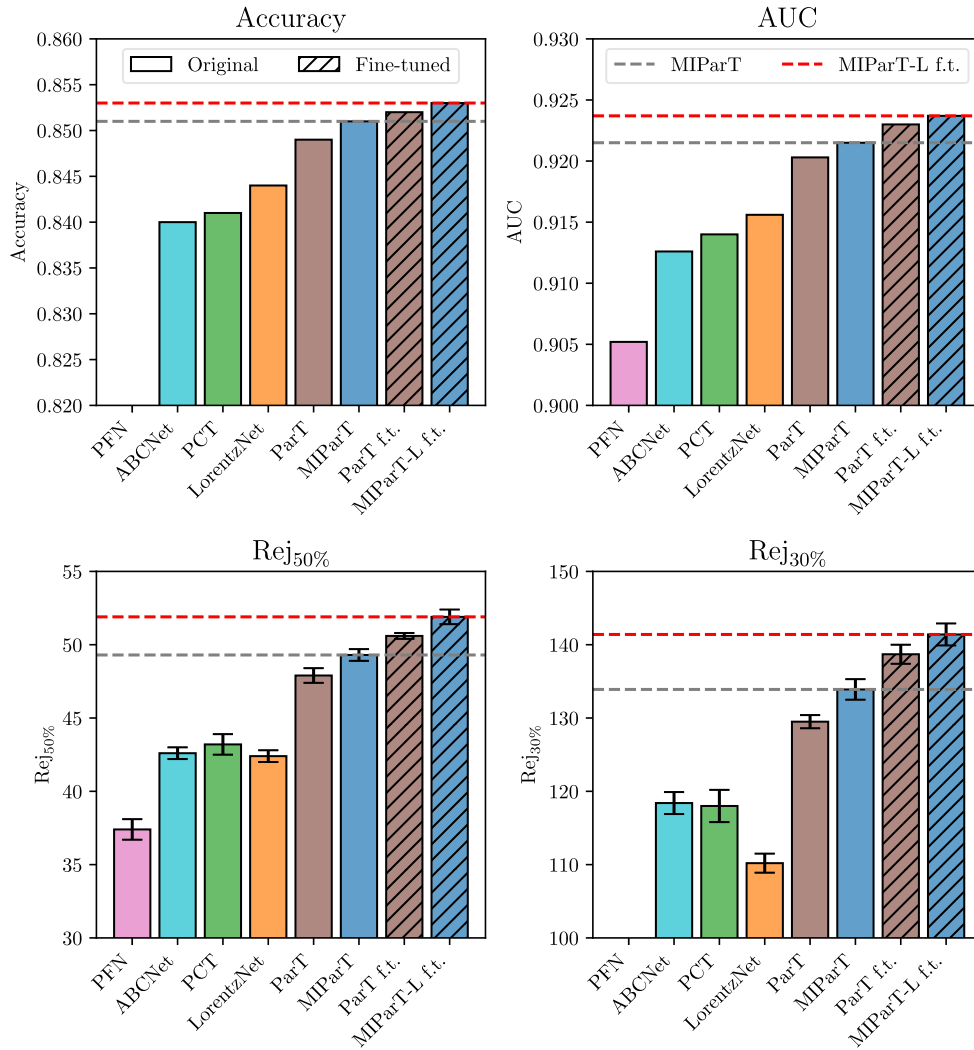


Fig. 6. (color online) Comparison of MIParT performance metrics with those of other models on the quark-gluon dataset. This figure shows Accuracy, AUC, Rej_{50%}, and Rej_{30%} metrics for the MIParT model alongside Particle Flow Network (PFN), attention-based Cloud Net (ABCNet) [47], Point Cloud Transformer (PCT) [46], LorentzNet [42], and ParT [36]. Metrics of other models are extracted from published results. Detailed outcomes are provided in Table 2. Bars without slashes indicate original models without fine-tuning, while bars with slashes indicate models with fine-tuning. The gray dashed line indicates the results for MIParT, whereas the red dashed line shows the results for fine-tuned MIParT-L (MIParT-L f.t.).

Regarding the impact of hyperparameter choices on model performance, we found that MIParT is not overly sensitive to hyperparameter settings, but is more influenced by the overall network architecture. In particular, increasing the number of MIA blocks and particle attention blocks generally leads to better performance at the cost of increased complexity. Architectural modifications show that placing MIA blocks before particle attention blocks is optimal. Placing MIA blocks after particle attention blocks or alternating them significantly reduces effectiveness, sometimes to the point of performing worse than ParT. We consider that MIA blocks function similarly to embeddings, allowing better integration of interaction information into the jets for improved information fusion and classification.

Table 4 presents the parameters, FLOPs (Floating Point Operations Per Second), and accuracy of various models on the top tagging and quark-gluon datasets. Parameters denote the number of trainable elements within a model, which indicates its capacity to learn. Having more parameters generally increases the complexity of the model. FLOPs is a metric that measures the computational complexity required to process data through the model. Reducing the number of parameters typically reduces FLOPs, simplifying the model and making it more computationally efficient.

However, reducing the number of parameters to reduce FLOPs usually results in lower accuracy. In contrast, our MIParT model has only 30% of the parameters and 53% of the FLOPs of the ParT model, significantly

Table 3. Performance comparison of various models on the quark-gluon dataset. This table lists the results for the MIParT model along with other significant models, including Particle Flow Network (PFN), attention-based Cloud Net (ABCNet) [47], Point Cloud Transformer (PCT) [46], LorentzNet [42], and ParT [36]. Metrics of other models are extracted from published results. The fine-tuned version of our model, MIParT-L f.t., is shown at the bottom of the table for comparison with the fine-tuned ParT model, that is, ParT f.t.

	Accuracy	AUC	Rej _{50%}	Rej _{30%}
PFN	—	0.9052	37.4±0.7	—
ABCNet	0.840	0.9126	42.6±0.4	118.4±1.5
PCT	0.841	0.9140	43.2±0.7	118.0±2.2
LorentzNet	0.844	0.9156	42.4±0.4	110.2±1.3
ParT	0.849	0.9203	47.9±0.5	129.5±0.9
MIParT (ours)	0.851	0.9215	49.3±0.4	133.9±1.4
ParT f.t.	0.852	0.9230	50.6±0.2	138.7±1.3
MIParT-L f.t. (ours)	0.853	0.9237	51.9±0.5	141.4±1.5

Table 4. Parameters, FLOPs, and accuracy for various models on the top tagging (TOP) and quark-gluon (QG) datasets. Parameters refer to the number of trainable elements within a model, while FLOPs (Floating Point Operations Per Second) is a metric that measures the computational complexity involved in processing data through the model.

	TOP	QG	Params	FLOPs
PFN	—	—	86.1k	4.62M
P-CNN	0.930	—	354k	15.5M
ParticleNet	0.940	—	370k	540M
ParT	0.940	0.849	2.14M	340M
MIParT (ours)	0.942	0.851	720.9k	180M
MIParT-L f.t. (ours)	0.944	0.853	2.38M	368M

reducing model complexity. Despite this reduction, there is no compromise in accuracy; in fact, accuracy improves on both top tagging and quark-gluon datasets. For the fine-tuned version of MIParT-L, the parameters and FLOPs are comparable to those of the ParT model, but with a slight improvement in accuracy.

Table 5 presents a performance comparison of various models on different sizes of the JetClass dataset. We displaced the results for the MIParT-L model alongside ParticleNet [34] and ParT [36] across the 2M, 10M, and 100M JetClass datasets. Note that as the dataset size increases, the performance of the models improves. Specifically, MIParT-L and ParT exhibit nearly identical effectiveness on very large datasets, outperforming ParticleNet. In addition, our evaluation of models on the JetClass dataset serves to test the ability of MIParT to generalize across different classification tasks. The JetClass datasets represent a more complex classification challenge, aiming at identifying Higgs boson decays to charm quarks. Our MIParT model shows remarkable stability on this task, highlighting its generalization capabilities.

Here, we discuss the improvements attributed to the pre-training performed on the JetClass dataset, with sub-

sequent performance improvements observed on the top tagging and quark-gluon datasets. These three jet tagging tasks differ in their objectives: the JetClass dataset focuses on identifying Lorentz boosted W , Z , Higgs bosons and top quarks, the top tagging dataset aims to identify top quarks, and the quark-gluon dataset aims to distinguish between quark and gluon jets. The improvements across such diverse tasks suggest that MIParT has learned more generalized jet properties during the pre-training phase. These characteristics are effectively transferable to other tasks, demonstrating the robustness of the model and its adaptability to different jet identification challenges. This capability highlights the potential of pre-trained models to improve performance in a wide range of applications by capturing and exploiting general features applicable to multiple scenarios.

Regarding the interpretability of MIParT, it is important to acknowledge that, as a model based on a transformer-based neural network architecture, its interpretability remains limited, similar to many neural networks currently in use. Despite these interpretability challenges, the CMS collaboration has successfully used the graph neural network ParticleNet [34], another model that lacks full interpretability, to search for Higgs boson decay to

Table 5. Performance comparison of various models on different sizes of the JetClass dataset. This table lists the results for the MIParT-L model alongside ParticleNet [34] and ParT [36] across 2M, 10M, and 100M JetClass datasets. Metrics of other models are extracted from published results. Models trained using the full 100M training dataset are highlighted in bold text.

	All classes		$H \rightarrow b\bar{b}$	$H \rightarrow c\bar{c}$	$H \rightarrow gg$	$H \rightarrow 4q$	$H \rightarrow \ell\nu qq'$	$t \rightarrow bq q'$	$t \rightarrow b\ell\nu$	$W \rightarrow qq'$	$Z \rightarrow qq'$
	Accuracy	AUC	Rej _{50%}	Rej _{50%}	Rej _{50%}	Rej _{50%}	Rej _{99%}	Rej _{50%}	Rej _{99.5%}	Rej _{50%}	Rej _{50%}
ParticleNet (2 M)	0.828	0.9820	5540	1681	90	662	1654	4049	4673	260	215
ParticleNet (10 M)	0.837	0.9837	5848	2070	96	770	2350	5495	6803	307	253
ParticleNet (100 M)	0.844	0.9849	7634	2475	104	954	3339	10526	11173	347	283
ParT (2 M)	0.836	0.9834	5587	1982	93	761	1609	6061	4474	307	236
ParT (10 M)	0.850	0.9860	8734	3040	110	1274	3257	12579	8969	431	324
ParT (100 M)	0.861	0.9877	10638	4149	123	1864	5479	32787	15873	543	402
MIParT-L (2 M)	0.837	0.9836	5495	1940	95	819	1778	6192	4515	311	242
MIParT-L (10 M)	0.850	0.9861	8000	3003	112	1281	3650	16529	9852	440	336
MIParT-L (100 M)	0.861	0.9878	10753	4202	123	1927	5450	31250	16807	542	402

charm quarks [48]. This success underscores that the lack of interpretability does not prevent the use of neural network models in particle physics experiments. In fact, ParticleNet, which functions as a non-interpretable "black box" model, is playing a significant role in particle experiments, demonstrating that the non-interpretable nature of these models should not be a barrier to their use in advancing scientific discovery.

IV. CONCLUSIONS

In this paper, we propose a novel deep-learning approach for jet tagging called MIParT that increases the dimensionality of particle interaction embeddings through More-Interaction Attention (MIA) to better exploit particle interaction inputs. We tested our model on two popular datasets and compared it with other models:

- **On the Top Tagging Dataset:** The MIParT model achieves accuracy and AUC metrics nearly identical to those of LorentzNet, and its Rej_{50%} and Rej_{30%} metrics are comparable within the error range to LorentzNet. Moreover, a series of Lorentz-equivariant methods demonstrate similar performance to that of LorentzNet. The MIParT model significantly outperforms ParT in the top tagging benchmark, achieving approximately 25% better background rejection at a 30% signal efficiency. Among the models evaluated, MIParT, along with LorentzNet and other Lorentz-equivariant-based models, ranks at the top tier, consistently delivering top-tier performance and robustness. For the fine-tuned MIParT-L model pre-trained on the 100M JetClass dataset, a 39% enhancement in background rejection performance is achieved, comparable to that of fine-tuned ParT.

- **On the Quark-gluon Dataset:** The MIParT model

significantly outperforms LorentzNet across all metrics, including accuracy, AUC, Rej_{50%}, and Rej_{30%}, as well as other models. MIParT achieves the best performance across all evaluation metrics, improving background rejection power by approximately 3% compared to ParT. For the fine-tuned MIParT-L model, background rejection performance improves by 6%, outperforming fine-tuned ParT. Specifically, the background rejection of fine-tuned MIParT-L improves by an additional 2% compared to that of fine-tuned ParT.

Overall, MIParT outperforms ParT on both the top and quark-gluon tagging tasks while also exhibiting lower computational complexity and fewer parameters. Previously, it was generally assumed that transformer-based models required large-scale dataset pre-training to achieve optimal results. Our MIParT model demonstrates that with higher-dimensional particle interaction embeddings, top-tier performance can be achieved without pre-training on large datasets, even outperforming ParT.

Furthermore, given that pre-training ParT on the larger multi-class JetClass dataset and subsequently fine-tuning it on the top tagging dataset can enhance performance, we applied this approach to MIParT-L in this study. We found that MIParT-L can further capitalize on the knowledge from large datasets, showing superior capabilities after fine-tuning. Specifically, it performs better on the quark-gluon dataset than fine-tuned ParT. Finding more efficient approaches to fine-tune a base transformer model will be especially helpful for future experiments when generic and foundation models are deployed, and downstream application tasks are varied. Moreover, MIParT is not limited to jet tagging but can also be applied to event identification, which could be notably helpful in the search for new physics signals.

References

- [1] A. J. Larkoski, I. Moulton, and B. Nachman, *Phys. Rept.* **841**, 1 (2020), arXiv: 1709.04464[hep-ph]
- [2] M. Feickert and B. Nachman, arXiv: 2102.02770[hep-ph]
- [3] J. Gallicchio and M. D. Schwartz, *Phys. Rev. Lett.* **107**, 172001 (2011), arXiv: 1106.3076[hep-ph]
- [4] J. Gallicchio and M. D. Schwartz, *JHEP* **04**, 090 (2013), arXiv: 1211.7038[hep-ph]
- [5] A. J. Larkoski, J. Thaler, and W. J. Waalewijn, *JHEP* **11**, 129 (2014), arXiv: 1408.3122[hep-ph]
- [6] B. Bhattacharjee, S. Mukhopadhyay, M. M. Nojiri *et al.*, *JHEP* **04**, 131 (2015), arXiv: 1501.04794[hep-ph]
- [7] D. Ferreira de Lima, P. Petrov, D. Soper *et al.*, *Phys. Rev. D* **95**, 034001 (2017), arXiv: 1607.06031[hep-ph]
- [8] P. Gras, S. Höche, D. Kar *et al.*, *JHEP* **07**, 091 (2017), arXiv: 1704.03878[hep-ph]
- [9] C. Frye, A. J. Larkoski, J. Thaler *et al.*, *JHEP* **09**, 083 (2017), arXiv: 1704.06266[hep-ph]
- [10] J. Cogan, M. Kagan, E. Strauss *et al.*, *JHEP* **02**, 118 (2015), arXiv: 1407.5675[hep-ph]
- [11] L. G. Almeida, M. Backović, M. Cliche *et al.*, *JHEP* **07**, 086 (2015), arXiv: 1501.05968[hep-ph]
- [12] L. de Oliveira, M. Kagan, L. Mackey *et al.*, *JHEP* **07**, 069 (2016), arXiv: 1511.05190[hep-ph]
- [13] P. T. Komiske, E. M. Metodiev, and M. D. Schwartz, *JHEP* **01**, 110 (2017), arXiv: 1612.01551[hep-ph]
- [14] G. Kasieczka, T. Plehn, M. Russell *et al.*, *JHEP* **05**, 006 (2017), arXiv: 1701.08784[hep-ph]
- [15] S. Macaluso and D. Shih, *JHEP* **10**, 121 (2018), arXiv: 1803.00107[hep-ph]
- [16] A. Butter, *et al.*, *SciPost Phys.* **7**, 014 (2019), arXiv: 1902.09914[hep-ph]
- [17] M. Kagan, (2020), arXiv: 2012.09719[physics.data-an]
- [18] R. T. de Lima, arXiv: 2102.06128[physics.data-an]
- [19] H. Kheddar, Y. Himeur, A. Amira *et al.*, arXiv: 2403.11934[hep-ph]
- [20] S. Mondal and L. Mastrolorenzo, arXiv: 2404.01071[hep-ex]
- [21] J. Pumplin, *Phys. Rev. D* **44**, 2025 (1991)
- [22] J. Lin, M. Freytsis, I. Moulton, and B. Nachman, *JHEP* **10**, 101 (2018), arXiv: 1807.10768[hep-ph]
- [23] D. Guest, J. Collado, P. Baldi *et al.*, *Phys. Rev. D* **94**, 112002 (2016), arXiv: 1607.08633[hep-ex]
- [24] E. Bols, J. Kieseler, M. Verzetti *et al.*, *JINST* **15**, P12012 (2020), arXiv: 2008.10519[hep-ex]
- [25] G. Louppe, K. Cho, C. Becot *et al.*, *JHEP* **01**, 057 (2019), arXiv: 1702.00748[hep-ph]
- [26] T. Cheng, *Comput. Softw. Big Sci.* **2**, 3 (2018), arXiv: 1711.02633[hep-ph]
- [27] I. Henrion, K. Cranmer, J. Bruna *et al.*, *Proceedings of the Deep Learning for Physical Sciences Workshop at NIPS (2017)*, (2017).
- [28] M. Abdughani, J. Ren, L. Wu, and J. M. Yang, *JHEP* **08**, 055 (2019), arXiv: 1807.09088[hep-ph]
- [29] J. Arjona Martínez, O. Cerri, M. Pierini *et al.*, *Eur. Phys. J. Plus* **134**, 333 (2019), arXiv: 1810.07988[hep-ph]
- [30] J. Ren, L. Wu, and J. M. Yang, *Phys. Lett. B* **802**, 135198 (2020), arXiv: 1901.05627[hep-ph]
- [31] X. Ju *et al.*, *33rd Annual Conference on Neural Information Processing Systems* (2020), arXiv: 2003.11603[physics.ins-det]
- [32] P. T. Komiske, E. M. Metodiev, and J. Thaler, *JHEP* **01**, 121 (2019), arXiv: 1810.05165[hep-ph]
- [33] M. Zaheer, S. Kottur, S. Ravanbakhsh *et al.*, arXiv: 1703.06114[cs.LG]
- [34] H. Qu and L. Gouskos, *Phys. Rev. D* **101**, 056019 (2020), arXiv: 1902.08570[hep-ph]
- [35] Y. Wang, Y. Sun, Z. Liu *et al.*, arXiv: 1801.07829[cs.CV]
- [36] H. Qu, C. Li, and S. Qian, *Proceedings of the 39th International Conference on Machine Learning*, 18281 (2022), arXiv: 2202.03772[hep-ph]
- [37] A. Vaswani, N. Shazeer, N. Parmar *et al.*, *31st International Conference on Neural Information Processing Systems (2017)*, arXiv: 1706.03762[cs.CL]
- [38] H. Touvron, M. Cord, A. Sablayrolles *et al.*, *IEEE/CVF International Conference on Computer Vision (ICCV)*, 32 (2021), arXiv: 2103.17239[cs.CV]
- [39] S. Shleifer, J. Weston, and M. Ott, arXiv: 2110.09456[cs.CL]
- [40] A. Paszke, S. Gross, F. Massa *et al.*, *Advances in neural information processing systems* **32**, (2019), arXiv: 1912.01703[cs.LG]
- [41] F. A. Dreyer and H. Qu, *JHEP* **03**, 052 (2021), arXiv: 2012.08526[hep-ph]
- [42] S. Gong, Q. Meng, J. Zhang *et al.*, *JHEP* **07**, 030 (2022), arXiv: 2201.08187[hep-ph]
- [43] D. Ruhe, J. Brandstetter, and P. Forré, arXiv: 2305.11141[cs.LG]
- [44] A. Bogatskiy, T. Hoffman, D. W. Miller *et al.*, arXiv: 2211.00454[hep-ph]
- [45] J. Spinner, V. Bresó, P. de Haan *et al.*, arXiv: 2405.14806[physics.data-an]
- [46] V. Mikuni and F. Canelli, *Mach. Learn. Sci. Tech.* **2**, 035027 (2021), arXiv: 2102.05073[physics.data-an]
- [47] V. Mikuni and F. Canelli, *Eur. Phys. J. Plus* **135**, 463 (2020), arXiv: 2001.05311[physics.data-an]
- [48] A. Tumasyan, *et al.* (CMS), *Phys. Rev. Lett.* **131**, 061801 (2023), arXiv: 2205.05550[hep-ex]