
Interpreting Transformers for Jet Tagging

Aaron Wang
University of Illinois Chicago
Chicago, IL 60607
aaronw5@uic.edu

Abhijith Gandrakota Jennifer Ngadiuba
Fermi National Accelerator Laboratory
Batavia, IL 60510
{abhijith,ngadiuba}@fnal.gov

Vivekanand Sahu Priyansh Bhatnagar Elham E Khoda Javier Duarte
University of California San Diego
La Jolla, CA 92093
{vsahu,prbhatnagar,ekhoda,jduarte}@ucsd.edu

Abstract

Machine learning (ML) algorithms, particularly attention-based transformer models, have become indispensable for analyzing the vast data generated by particle physics experiments like ATLAS and CMS at the CERN LHC. Particle Transformer (ParT), a state-of-the-art model, leverages particle-level attention to improve jet-tagging tasks, which are critical for identifying particles resulting from proton collisions. This study focuses on interpreting ParT by analyzing attention heat maps and particle-pair correlations on the η - ϕ plane, revealing a binary attention pattern where each particle attends to at most one other particle. At the same time, we observe that ParT shows varying focus on important particles and subsets depending on decay, indicating that the model learns traditional jet substructure observables. These insights enhance our understanding of the model’s internal workings and learning process, offering potential avenues for improving the efficiency of transformer architectures in future high-energy physics applications.

1 Introduction

Machine learning (ML) algorithms are becoming crucial for effectively analyzing the enormous data produced by particle physics experiments like ATLAS and CMS at the CERN LHC [1]. Attention-based transformer models [2], with their ability to effectively capture and weigh the relative importance of different elements in input data, have revolutionized various domains, including natural language processing and computer vision [3–5]. The ability of the attention mechanism to discern intricate correlations in the collimated spray of particles jets initiated from massive particle decays has proven invaluable in searches for new physics and standard model measurements [6].

High-energy collisions of protons at LHC can create new unstable particles which then decay and produce sprays of outgoing particles, called jets, which are observed by these experiments[7]. Jet tagging, or the process of identifying the particle that initiates this spray, is a critical step in data analysis at the LHC [8, 9]. One such state-of-the-art model for jet tagging is Particle Transformer (ParT) [10], which is a transformer variant that leverages physics-inspired particle-particle pairwise features (or “interactions”) to augment the attention mechanism. Although this is one of the best performing jet-tagging models, there is effort required to interpret the underlying workflow and why the model performs so well.

The attention-based architecture of transformers may offer deeper insight into the inner workings of neural networks [11]. In the case of a jet tagger, the attention matrix captures the particle-

to-particle correlations potentially making these ML models more interpretable. These attention scores between particles highlight the most important connections between particles necessary for classification. This helps us to understand whether the neural network is learning the physics we know. In addition, looking at which parts of the input the neural network focuses on can help us improve the computational performance by pruning unnecessary parts of the neural network architecture, resulting in a more efficient jet tagging model.

Related Work This paper builds on prior research that has applied machine learning, specifically graph neural networks (GNNs), to particle physics problems at the CERN LHC [12, 13]. Reference [14] attempts to explain how ParticleNet [8], a GNN designed for jet tagging learns to classify three-prong hadronic top quark decays using edge relevancy (R) graphs based on layerwise relevance propagation. Similarly, we visualize the attention scores between pairs of particles in order to investigate how ParT classifies jets.

2 Methods

This study uses the pre-trained ParT model, trained using the JETCLASS dataset [10, 15], which consists of 100 million examples of jets for training, 5 million for validation, and 20 million for testing, containing 10 classes of jets with different decay modes of quarks, gluons, W and Z bosons, Higgs bosons, and top quarks. The classes are q/g , $H \rightarrow b\bar{b}$, $H \rightarrow c\bar{c}$, $H \rightarrow gg$, $H \rightarrow 4q$, $H \rightarrow \ell\nu qq'$, $t \rightarrow bq q'$, $t \rightarrow b\ell\nu$, $W \rightarrow qq'$, and $Z \rightarrow q\bar{q}$, indicating the decays to combination of light quarks (q), charm quark (c), bottom quark (b), and lepton (ℓ). There are 17 input features divided into 3 categories: kinematics, particle ID, and trajectory displacement. Each jet is composed of a maximum of 128 particles, ordered by p_T , with an average of 30–50 particles per jet.

ParT introduces a modified attention mechanism, called particle multihead attention (P-MHA). Given a particle-level representation of a jet with N constituents with each constituent having d features, $x \in \mathbb{R}^{N \times d}$, the attention is computed as

$$\begin{aligned} \text{P-MHA}(x) &= \text{concat}(\text{head}_1, \dots, \text{head}_h)W^O \\ \text{head}_i &= \text{softmax}\left(\frac{xW_i^W(xW_i^K)^\top}{\sqrt{d_k}} + U_i\right)xW_i^V, \end{aligned} \quad (1)$$

where $U_i \in \mathbb{R}^{N \times N}$ is a single feature for each pair of particles, which can be concatenated into a larger structure, known as the pairwise interaction matrix $U \in \mathbb{R}^{N \times N \times d'}$.

In ParT, the number of pairwise features d' must be equal to the number of heads h , which is chosen to be 8. In addition, $d_k = 16$ so that the overall output dimension is the same as the input dimension $d = hd_k = 128$. The pairwise interaction matrix U is learned by applying a 4-layer pointwise MLP with (64, 64, 64, 8) channels and GELU nonlinearity from the 4 features motivated by Ref. [16], $(\ln \Delta, \ln k_T, \ln z, \ln m^2)$, where

$$\begin{aligned} \Delta &= \sqrt{(y_a - y_b)^2 + (\phi_a - \phi_b)^2}, \\ k_T &= \min(p_{T,a}, p_{T,b})\Delta, \\ z &= \min(p_{T,a}, p_{T,b}) / (p_{T,a} + p_{T,b}), \\ m^2 &= (E_a + E_b)^2 - \|\mathbf{p}_a + \mathbf{p}_b\|^2, \end{aligned} \quad (2)$$

y_i is the rapidity, ϕ_i is the azimuthal angle, $p_{T,i} = (p_{x,i}^2 + p_{y,i}^2)^{1/2}$ is the transverse momentum, and $\mathbf{p}_i = (p_{x,i}, p_{y,i}, p_{z,i})$ is the momentum 3-vector and $\|\cdot\|$ is the norm, for $i = a, b$. We use the exact pre-trained model weights provided from the Particle Transformer repository [10].

3 Results

Distribution of Attention Scores We visualize the attention scores across all heads in a heat map and observe that the ParT model’s attention scores exhibit a nearly binary (on/off) pattern. As shown in Fig. 1, the distribution of attention scores reveals that most values are either close to one or zero. This binary characteristic indicates that particles generally attend to, at most one other particle. This behavior raises the intriguing question of whether the model is capable of capturing underlying physical laws and how we might interpret the specific physics the model is learning from the data. The binary nature of the attention weights also suggests potential for more efficient transformer




Figure 1: Distribution of attention scores (left) illustrating the binary nature. Heat map of attention scores for example jets in the $t \rightarrow b\ell\nu$ (center) and $t \rightarrow bqq'$ classes (right) between the constituent particles.

models by reducing the computational complexity of attention mechanisms, focusing only on key interactions between particles.

Particle Attention Graphs We decluster each jet into a specific number of subjets using the k_T algorithm [17] implemented in the FastJet package [18] with Python bindings [19]. We cluster into two subjets for leptonic top quark decays ($t \rightarrow b\ell\nu$), three subjets for hadronic top decays ($t \rightarrow bqq'$), and four subjets for Higgs decays into four quarks ($H \rightarrow 4q$). For each process, we visualize the attention scores by representing each particle as a point in η - ϕ space in Fig. 2. We compare these results with those obtained with randomly initialized attention weights. We focus on the heads in the final layer of P-MHA.

In Fig. 2, we observe that compared to an untrained, randomly initialized model, ParT demonstrates a significantly stronger tendency to attend to the lepton in the $t \rightarrow b\ell\nu$ process. This behavior indicates that ParT recognizes the importance of the lepton in classifying the underlying resonance and decay mode. Additionally, in $t \rightarrow bqq'$ and $H \rightarrow 4q$ P-MHA captures both inter-subjet and intra-subjet interactions. Specifically, ParT exhibits attention between different subjets and also within individual subjets. For example, in Fig. 2, P-MHA connects all different subjets but does not form connections within the subjets. In contrast, P-MHA focuses its attention on interactions between subjets 0 and 1 while neglecting connections with subjct 2, instead of forming connections within subjct 2.

In order to quantify how often ParT exhibits these behaviors, we display the distribution of attention values for different jet classes. In particular, we find the sum of all attention values to a lepton in $t \rightarrow b\ell\nu$ jets, and divide it by the sum of all attention values in the head. Since the sum of all attention values in a head should be equal to the number of particles, this effectively is the proportion of attention weights that attend to the lepton. Similarly, for $t \rightarrow bqq'$ and $H \rightarrow 4q$ jets, we display the sum of attention weights between particles in different subjets divided by the sum of all attention weights. In Fig. 3, we observe that the trained ParT model is more likely to have a head that consists of only inter-subjet connections or only intra-subjet connections than an untrained, randomly initialized model. ParT attending to only leptons or to no leptons for $t \rightarrow b\ell\nu$ jets, and only intra-subjet or only inter-subjet for $t \rightarrow bqq'$ and $H \rightarrow 4q$ jets is some evidence that ParT is learning relevant physical properties of the jet data.

Limitations Our analysis is constrained to the final attention layer, limiting insights into how earlier layers contribute to the representations. Additionally, the observed attention patterns may depend on the specific subjet clustering algorithm used, which could influence the model's interpretability. Exploring intermediate layers and varying subjet algorithms in future studies could provide a more comprehensive understanding of ParT's attention mechanisms.

4 Optimizing Attention Mechanism for Efficiency

Leveraging the binary nature of P-MHA, we constrain the total number of particles used in the attention mechanism, and explore the impact of the number of particles attended to in each attention head on the overall performance. For only one particle, we directly replace attention with softmax. For other numbers of particles, we keep the top- k highest attention particles, setting the rest of the

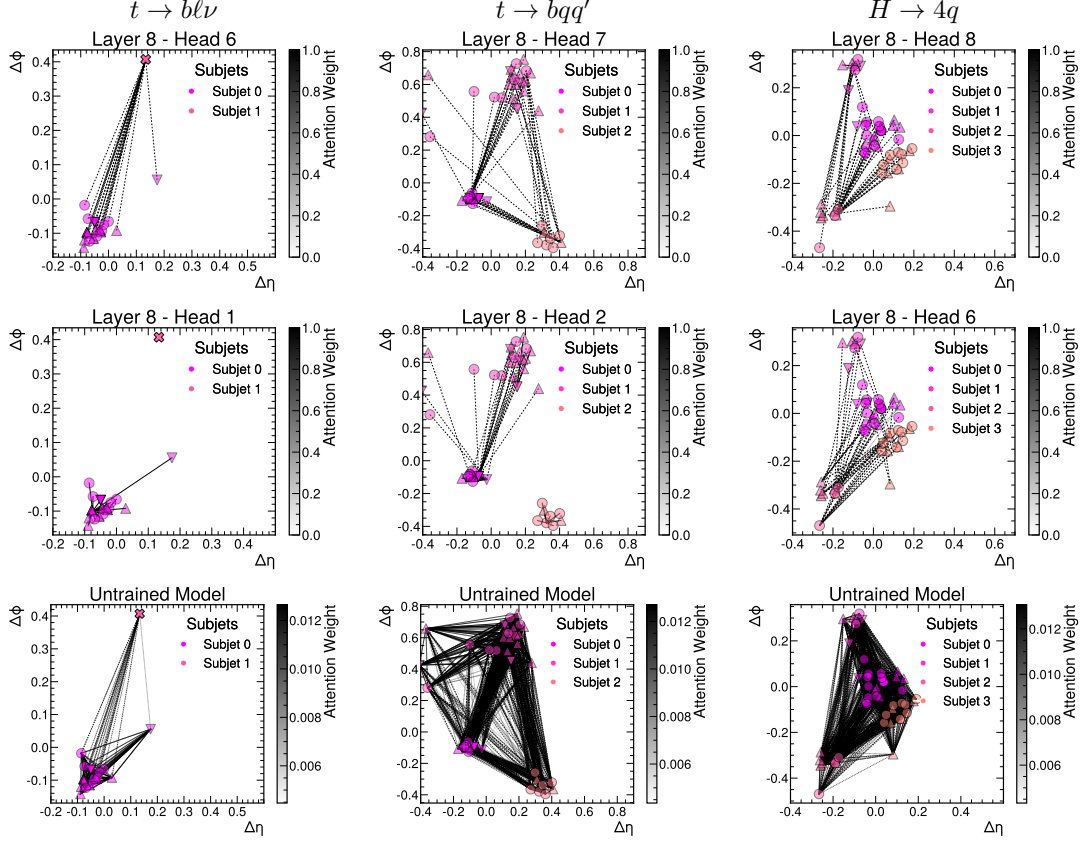


Figure 2: Visualization of attention values in the η - ϕ plane for pairs of particles (\times : muon, \blacktriangle : charged hadron, \blacktriangledown : neutral hadron, \bullet : photon, and \oplus : electron) within jets. The transparency of each point is proportional to the transverse momentum of the particle (p_T). The intensity of the connecting lines represents the magnitude of the attention scores. Solid lines depict intra-subject connections, while dotted lines represent inter-subject connections. The plots in the first column correspond to $t \rightarrow b\ell\nu$, the second column to $t \rightarrow bq\bar{q}'$, and the third column to $H \rightarrow 4q$ processes.

Table 1: Performance of the ParT model with each attention head constrained to use a maximum of k particles.

Max particles	Accuracy	AUC	$H \rightarrow b\bar{b}$ Rej _{50%}	$H \rightarrow c\bar{c}$ Rej _{50%}	$H \rightarrow g\bar{g}$ Rej _{50%}	$H \rightarrow 4q$ Rej _{50%}	$H \rightarrow \ell\nu q\bar{q}'$ Rej _{99%}	$t \rightarrow bq\bar{q}'$ Rej _{50%}	$t \rightarrow b\ell\nu$ Rej _{99.5%}	$W \rightarrow qq'$ Rej _{50%}	$Z \rightarrow q\bar{q}$ Rej _{50%}
1	0.770	0.9754	4396	1325	58	246	1944	2384	6601	217	175
2	0.798	0.9795	5780	1862	67	390	2894	4219	10050	286	231
3	0.814	0.9817	6757	2268	73	515	3460	5917	11111	341	272
4	0.825	0.9831	7326	2558	79	661	3884	7813	12422	379	299
6	0.838	0.9849	8621	3063	89	938	4474	12821	13889	438	339
10	0.851	0.9865	9852	3670	103	1352	5102	20619	15748	502	375
20	0.859	0.9875	11111	4237	118	1802	5263	29412	16000	539	399
30	0.860	0.9877	10989	4193	122	1859	5362	33898	16129	543	402
128	0.861	0.9877	10638	4149	123	1864	5479	32787	15873	543	402

particles to zero, and then applying softmax over those k particles (out of 128 maximum). We find that with just 1 particle, some performance is maintained and with 30 particles, the performance of ParT is recovered as shown in Table 1. By filtering the particles that P-MHA can consider, we can potentially reduce the number of calculations and increase computational efficiency.

5 Summary and Outlook

The ParT model demonstrates a unique binary attention pattern, in which each particle typically attends to at most one other particle. This focused attention mechanism contrasts with other transformer models, such as vision transformers, which spread attention across many elements. The




Figure 3: Distribution of the proportion of attention values that attend to leptons in $t \rightarrow b\ell\nu$ (left). Distribution of the proportion of attention in between subjects in the $t \rightarrow bqq'$ (center) and $H \rightarrow 4q$ classes (right)

sparse and selective attention in ParT prioritizes key relationships, potentially uncovering important substructures in jet tagging, and enhancing model interpretability in high-energy physics.

This binary pattern also opens up opportunities for optimization. This could lead to more efficient architectures without compromising performance, particularly in tasks such as jet tagging, where the identification of essential interactions is vital. Future work could explore these optimizations to further improve both the interpretability and efficiency of transformers in physics applications. This work will be helpful in designing a more efficient physics-based attention mechanisms.

In addition, the importance of the interaction matrix for the performance of ParT has not been explored fully. In Ref. [20], many P-MHA layers are replaced with the interaction matrix and find improved performance and efficiency. Understanding what the interaction matrix is learning to augment P-MHA could lead to a more efficient architecture.

Broader Impact Studies that interpret machine learning algorithms are important to increase confidence in model predictions and also to improve the performance of future machine learning models. This study uses an xAI method in order to explain the performance of the state-of-the-art transformer model for the LHC, by demonstrating that the model is indeed learning physics. This study opens doors for future studies to extract and discover physics information in machine learning models that are currently not being used for prediction.

The code used for this analysis is publicly available at: <https://github.com/aaronw5/Interpreting-Transformers-for-Jet-Tagging>

Acknowledgments and Disclosure of Funding

This work is supported by the DOE Office of Science, Award No. DE-SC0023524, Fermi Research Alliance, LLC under Contract No. DE-AC02-07CH11359 with the DOE, Office of Science, Office of High Energy Physics, LDRD L2024-066-1, DOE Office of Science, Office of High Energy Physics “Designing efficient edge AI with physics phenomena” Project (DE-FOA-0002705), DOE Office of Science, Office of Advanced Scientific Computing Research under the “Real-time Data Reduction Codesign at the Extreme Edge for Science” Project (DE-FOA-0002501), and DOE Early Career Research Program Award No. DE-SC0021187. The work is also supported in part by RCSA Grant #CS-CSA-2023-109, Sloan Foundation Grant #FG-2023-20452, NSF awards CNS-1730158, ACI-1540112, ACI-1541349, OAC-1826967, OAC-2112167, CNS-2100237, CNS-2120019, the University of California Office of the President, and the University of California San Diego’s California Institute for Telecommunications and Information Technology/Qualcomm Institute, AI2050 program at Schmidt Futures (Grant G-23-64934), and the NSF HDR Institute A3D3 (PHY-2117997).

References

- [1] P. Harris et al., “Physics Community Needs, Tools, and Resources for Machine Learning”, in *Snowmass 2021*. 2022. arXiv:2203.16255.

- [2] A. Vaswani et al., “Attention is all you need”, in *Advances in Neural Information Processing Systems*, I. Guyon et al., eds., volume 30. Curran Associates, Inc., 2017.
- [3] A. Ramesh et al., “Hierarchical text-conditional image generation with clip latents”, 2022. arXiv:2204.06125.
- [4] Gemini Team, “Gemini: A family of highly capable multimodal models”, 2024. arXiv:2312.11805.
- [5] T. B. Brown et al., “Language models are few-shot learners”, 2020. arXiv:2005.14165.
- [6] CMS Collaboration, “Search for highly energetic double Higgs boson production in the two bottom quark and two vector boson all-hadronic final state”, CMS Physics Analysis Summary CMS-PAS-HIG-23-012, 2024.
- [7] B. Andersson, “The Lund Model”, *Nucl. Phys. A* **461** (1987) 513C, doi:10.1016/0375-9474(87)90510-0.
- [8] H. Qu and L. Gouskos, “Jet tagging via particle clouds”, *Phys. Rev. D* **101** (2020) doi:10.1103/physrevd.101.056019, arXiv:1902.08570.
- [9] R. Liu et al., “Efficient and Robust Jet Tagging at the LHC with Knowledge Distillation”, in *Machine Learning and the Physical Science Workshop at the 37th Conference on Neural Information Processing Systems*. 2023. arXiv:2311.14160.
- [10] H. Qu, C. Li, and S. Qian, “Particle Transformer for jet tagging”, in *Proceedings of the 39th International Conference on Machine Learning*, p. 18281. 2022. arXiv:2202.03772.
- [11] H. Chefer, S. Gur, and L. Wolf, “Transformer interpretability beyond attention visualization”, in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, p. 782. 2021.
- [12] J. Shlomi, P. Battaglia, and J.-R. Vlimant, “Graph Neural Networks in Particle Physics”, *Mach. Learn.: Sci Technol.* **2** (2021) 021001, doi:10.1088/2632-2153/abbf9a, arXiv:2007.13681.
- [13] J. Duarte and J.-R. Vlimant, “Graph neural networks for particle tracking and reconstruction”, in *Artificial Intelligence for High Energy Physics*, P. Calafiura, D. Rousseau, and K. Terao, eds., p. 387. World Scientific, 2022. arXiv:2012.01249. doi:10.1142/9789811234033_0012.
- [14] F. Mokhtar, R. Kansal, and J. Duarte, “Do graph neural networks learn traditional jet substructure?”, in *Machine Learning and the Physical Sciences Workshop at the 36th Conference on Neural Information Processing Systems*. 2022. arXiv:2211.09912.
- [15] H. Qu, C. Li, and S. Qian, “JETCLASS: A large-scale dataset for deep learning in jet physics”, 2022. doi:10.5281/zenodo.6619768.
- [16] F. A. Dreyer and H. Qu, “Jet tagging in the Lund plane with graph networks”, *JHEP* **03** (2021) 052, doi:10.1007/JHEP03(2021)052, arXiv:2012.08526.
- [17] S. Catani, Y. L. Dokshitzer, M. H. Seymour, and B. R. Webber, “Longitudinally invariant K_t clustering algorithms for hadron hadron collisions”, *Nucl. Phys. B* **406** (1993) 187, doi:10.1016/0550-3213(93)90166-M.
- [18] M. Cacciari, G. P. Salam, and G. Soyez, “FastJet User Manual”, *Eur. Phys. J. C* **72** (2012) 1896, doi:10.1140/epjc/s10052-012-1896-2, arXiv:1111.6097.
- [19] A. Roy, J. Pivarski, and C. W. Freer, “An array-oriented Python interface for FastJet”, *J. Phys. Conf. Ser.* **2438** (2023) 012011, doi:10.1088/1742-6596/2438/1/012011, arXiv:2202.03911.
- [20] Y. Wu et al., “Jet Tagging with More-Interaction Particle Transformer”, 2024. arXiv:2407.08682. Accepted by *Chin. Phys. C*.