PANOPTAG: Simultaneously Tagging All Jets in a Particle Collision Event

Umar Sohail Qureshi[1,2] **Brendon Bullard**[2] **Ariel Schwartzman**[2]

[1]Department of Physics, Stanford University, Stanford, CA 94305, USA
[2]SLAC National Accelerator Laboratory, Menlo Park, CA 94025, USA

uqureshi@cern.ch bbullard@slac.stanford.edu sch@slac.stanford.edu

Abstract

Jet tagging, identifying the origin of jets produced in particle collisions, is a critical classification task in high-energy physics. Despite the revolutionary impact of deep learning on jet tagging, the paradigm has remained unchanged: jets are classified independently, one at a time. This approach ignores correlations, overlaps, and wider event context between jets. We introduce PANOPTAG, which simultaneously tags all jets in an event using an encoder-decoder architecture where jet kinematics serve as queries that cross-attend to particle flow object embeddings. We evaluate PANOPTAG on heavy-flavor (b/c)-tagging and demonstrate significant performance improvements over state-of-the-art single-jet baselines, while maintaining the topology independence required for conventional calibration techniques.

1 Introduction

Particle colliders, such as the Large Hadron Collider (LHC) at CERN, accelerate beams of protons to nearly the speed of light and bring them to collide tens of millions of times per second. Each collision event produces complex sprays of outgoing particles that traverse a layered detector system [3, 12] used to reconstruct particle trajectories, momenta, energies, and identities.

A central object in this reconstruction chain is a jet [42, 39], a collimated spray of particles originating from quarks or gluons. Jets are ubiquitous in collider analyses and as such determining the origin of each jet, dubbed jet tagging, is a key classification task that underpins both precision measurements and searches for new physics. A jet is naturally represented as a variable-size set of its reconstructed constituent particles, called particle flow objects (PFOs). Modern deep learning methods [16, 36] that operate directly on low-level constituent PFOs have become standard, with the prevailing state-of-the-art treating jets as unordered sets or "particle clouds" [34].

Within this paradigm, Particle Flow Networks [21] adapted the Deep Sets [44] formulation, ParticleNet [32] introduced dynamic graph construction with EdgeConv from DGCNN [43], and the Particle Transformer (ParT) [34] uses attention augmented with pairwise kinematic information. Symmetry-aware models like LGATr [41, 7] incorporate Lorentz equivariance within the Geometric Algebra Transformer [8] framework. Alongside these, unsupervised [19], self-supervised [37], and fully-supervised pretraining methods [28, 14, 24, 6] have been explored for representation learning. In addition, both CMS and ATLAS collaborations at the LHC have deployed increasingly sophisticated deep-learning-based taggers [29, 1, 4], underscoring that continued progress in jet tagging is a key enabler for the LHC physics program.

In this work, we introduce PANOPTAG, a new paradigm that reframes jet tagging as an event-level, multi-jet inference problem. Instead of classifying each jet in isolation, PANOPTAG processes the full set of reconstructed jets and particles once with a DETR-inspired [11] encoder–decoder architecture, assigning labels to all jets simultaneously. A permutation equivariant backbone generates embeddings

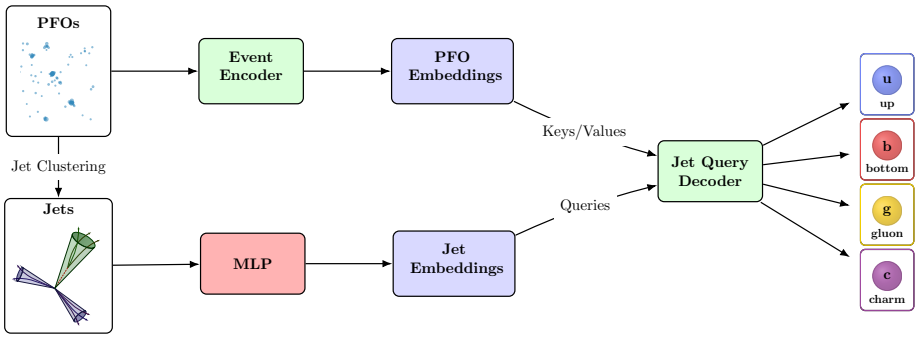

Figure 1: A high-level overview of the PANOPTAG architecture. PFOs are processed through the Event Encoder to produce embeddings, which serve as keys and values in a cross-attention mechanism. Simultaneously, jet kinematics are passed through an MLP to produce jet embeddings that act as queries, with the outputs after classification heads predicting class labels.

for all PFOs in the event. Jet kinematics are then used as queries in a decoder that cross-attends to the particle embeddings, producing one latent representation per jet, which is mapped to label probabilities (Fig. 1). This design allows information to be shared coherently across jets, while letting each jet focus on the subset of particles most informative for its identity. Evaluated on heavy-flavor identification, PANOPTAG achieves remarkably higher performance than existing state-of-the-art single-jet baselines, with gains that generalize beyond topologies included in training.

2 Model Architecture

The model consists of an Event Encoder (EE) block that processes all PFOs to compute contextual embeddings, and a Jet Query Decoder (JQD) block that cross-attends per-jet kinematic embeddings to these PFO embeddings to produce per-jet classification logits.

2.1 Event Encoder Block

The EE combines local and global feature extraction. PFO input features $\mathbf{X} \in \mathbb{R}^{B\times N\times d_{\text{in}}}$ are first projected to the model dimension d_h via an input MLP. Local geometric structure is captured through sequential EdgeConv [43] layers. The first layer constructs k-nearest neighbor graphs in the PFO coordinate space (pseudorapidity η and azimuthal angle ϕ). EdgeConv computes edge features as $\mathbf{e}_{ij} = [\mathbf{x}_j - \mathbf{x}_i, \mathbf{x}_i]$ for all $j \in \mathcal{N}_i$, passed through an MLP and aggregated via max pooling:

$$\mathbf{x}_i^{\text{EdgeConv}} = \max_{j\in\mathcal{N}_i} \left(\text{MLP}(\mathbf{e}_{ij})\right). \tag{1}$$

Subsequent EdgeConv layers use Euclidean distances in the learned feature space. This hierarchical design is motivated by [22, 28], where incorporating local features is shown to improve performance. The addition of EdgeConv layers creates a latent representation that is better aware of distances and topology, which is important since PANOPTAG needs to both learn geometric correlations between particles and associate individual PFOs to jets.

Following local feature extraction, global context is captured through Induced Self-Attention Blocks (ISABs) from Set Transformers [23]. Each ISAB introduces learnable inducing points $\mathbf{I} \in \mathbb{R}^{M\times d_h}$, first attending to the full PFO set as $\mathbf{H} = \text{MAB}(\mathbf{I}, \mathbf{X}^{\text{local}})$, then PFOs attend back as $\mathbf{X}^{\text{ISAB}} = \text{MAB}(\mathbf{X}^{\text{local}}, \mathbf{H})$. This reduces complexity from $\mathcal{O}(N^2)$ to $\mathcal{O}(MN)$, which is important since collision events can contain $\mathcal{O}(1000)$ PFOs. Each MAB implements scaled dot-product attention without positional encoding to ensure permutation equivariance, with residual connections and layer normalization [5]. Both local and global features are accumulated with skip connections.

2.2 Jet Query Decoder Block

The JQD generates per-jet predictions by cross-attending jet-specific query embeddings to the contextual PFO embeddings. Each jet is represented by its kinematic properties (p_{T}, η, ϕ, m),

	All Classes		b-jets		c-jets	
	Accuracy	AUC	$\text{Rej}_{50\%}$	$\text{Rej}_{80\%}$	$\text{Rej}_{50\%}$	$\text{Rej}_{80\%}$
ParticleNet	0.910	0.9742	10,003	194	157	18
ParT	0.913	0.9744	13,757	232	176	18
PANOPTAG	**0.9350 $\pm$ 0.0009**	**0.9787 $\pm$ 0.0007**	**18,087**	**383**	**252**	**24**

Table 1: Jet tagging performance on the heavy-flavor tagging dataset for all processes. Higher is better. Best results in **boldface**.

embedded via an MLP: $\mathbf{Z}_0 = \text{MLP}_{\text{jet}}(\mathbf{J})$, where $\mathbf{J} \in \mathbb{R}^{B \times M_{\text{jets}} \times 4}$. The decoder consists of L_{dec} cross-attention layers computing $\mathbf{Z}_{l+1}, \mathbf{A}_{l+1} = \text{MAB}(\mathbf{Z}_l, \mathbf{H}_c)$, with mask-aware attention to prevent attention to padded PFOs. A final classification head projects jet embeddings to per-class logits.

3 Heavy-Flavor Tagging Experiment

Existing jet-tagging benchmarks [18, 20, 33] provide data in a jet-centric format with no access to full event context. Since PANOPTAG operates at the event level, we simulate dedicated datasets.

3.1 Dataset

Samples are simulated with PYTHIA [40] for event generation, parton showering, and hadronization, and DELPHES [13] for detector simulation using the default ATLAS configuration. Jets are clustered from reconstructed PFOs via the anti-k_{T} algorithm [10] with $R = 0.4$ and $p_{\text{T}} > 20$ GeV using FASTJET [9, 38]. Unlike the ATLAS and CMS heavy-flavor tagging studies which consider only $t\bar{t}$ production, we also include single-top, tW, ZH, and $t\bar{t}H$ to ensure PANOPTAG does not memorize features of a particular topology. We train only on single-top, $t\bar{t}$, and $t\bar{t}H$, reserving ZH and tW for out-of-distribution validation. For each topology, we simulate 1M events, yielding $\sim$30M jets in three classes: light jets (55%), b-jets (33%), and c-jets (12%). The leading 128 PFOs and 10 jets are kept per event. Each PFO uses 14 input features spanning kinematics, particle identification, and trajectory displacement. The baselines use the same jets but with PFO constituents restricted to those used in jet clustering.

3.2 Setup

A training-validation-testing split of 750k-100k-150k events is used, with additional 150k ZH and tW events for testing. The nominal model contains 2 EdgeConv layers and 3 ISAB layers, with $d_h = 256$, $m = 48$ inducing points, $n_h = 32$ attention heads, $k = 20$ nearest neighbors, GELU [15] activations, and 3.42M learnable parameters. We use ADAMW [26] with weight decay 10^{-3} to minimize the weighted focal cross-entropy [25] loss with $\gamma = 1$. The learning rate warms up linearly over one epoch from 10^{-6} to $2{\times}10^{-4}$, followed by cosine annealing [27] with warm restarts every 15 epochs, training for 30 epochs with batch size 512. The model is implemented in PyTorch [30].

3.3 Baselines

We compare against ParticleNet [32] and ParT [34], using existing implementations from the WEAVER [31] framework. All models are trained end-to-end on the same dataset, with the same input features, optimizer, loss, and LR schedule. The key difference is that baselines input jets one-at-a-time with relative coordinates $(\Delta\eta, \Delta\phi)$.

3.4 Results

We report classification accuracy, AUC, and the background rejection $\text{Rej}_{X\%} \equiv 1/\text{FPR}$ at TPR $= X\%$, reported separately for each signal (b/c) jet category at operating points of 50% and 80%. Background rejection directly maps onto discovery sensitivity at the LHC: doubling it corresponds to an approximately 40% gain in discovery potential.

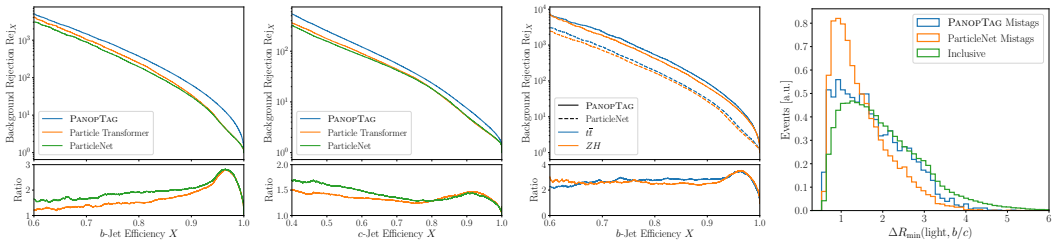

Figure 2: Background rejection versus b-jet (left) and c-jet (center-left) tagging efficiency for PanopTag, ParT, and ParticleNet, with ratio panels below. (Center-right) In-distribution ($t\bar{t}$) versus out-of-distribution (ZH) comparison. (Right) ΔR to nearest b/c jet for mistagged light jets compared to the inclusive distribution.

Table 1 and Fig. 2 summarize the results. PanopTag delivers the best performance on every metric, outperforming ParT by a wide margin. Overall accuracy increases by 2.2% and 2.5% compared to ParT and ParticleNet, respectively. For background rejection, PanopTag improves upon ParticleNet and ParT by a factor of 1.8 (2.0) and 1.7 (1.6) respectively for b-jets, and 1.6 (1.3) and 1.4 (1.3) for c-jets at 50% (80%) signal efficiency. The performance gains persist when evaluating on topologies not included during training: PanopTag maintains its improvement on the ZH process, demonstrating generalization beyond the training distribution.

To investigate the source of these gains, we study the $\Delta R \equiv \sqrt{\Delta\eta^2 + \Delta\phi^2}$ distance distribution of light jets to the nearest heavy-flavor jet. Relative to the inclusive spectrum, ParticleNet mistags are strongly biased toward smaller ΔR, indicating ambiguities when a light jet lies close to a true b/c jet and constituent overlap makes classification challenging. PanopTag shows a much smaller bias, following the inclusive distribution more closely, suggesting it suppresses proximity-driven mistags.

3.5 Calibration

Although jet classification is performed jointly within each event, compatibility with standard per-jet efficiency calibrations—typically derived from $t\bar{t}$ control samples—requires that per-jet flavour probabilities remain statistically independent. To validate this, we perform a perturbative test: for each event, a randomly selected b-jet and its associated PFOs are replaced by a light-flavour jet drawn from a cross-event cache, and model inference is re-evaluated. The resulting accuracy of 0.9348 and AUC of 0.9780 are consistent with Table 1 within training-seed uncertainties. Further tests of conditional independence of b-jet probabilities and compatibility with conventional calibration strategies is left to future work. This result is a natural consequence of the topology independence.

4 Discussion and Conclusion

We introduce PanopTag, a new paradigm presenting an event-level formulation of jet tagging. All jets in an event are simultaneously tagged by using each jet's kinematics as a learnable query that cross-attends to particle embeddings. On heavy-flavor tagging, PanopTag yields performance gains over state-of-the-art single-jet baselines while being topology-agnostic—enabling the standard approach of efficiency measurement and calibration using $t\bar{t}$ or $Z(\to \mu\mu)$ + Jets control samples.

These results support a broader conclusion: treating jets as a set of correlated objects and enforcing a shared event representation is a productive modeling direction, motivated by the physical reality that jets within an event are constrained by global energy-momentum conservation, characteristic production topologies, and even partial overlap. The PanopTag formulation is a general recipe that can be repurposed for pileup-jet identification [2], jet regression [17], background subtraction [35], and more, making it a candidate for a high-energy physics foundation model. The event-level perspective is also attractive from a fast-inference standpoint, as the particle representation is computed once per event and reused across all jets. Finally, PanopTag is compatible with alternative backbones (e.g. Lorentz-equivariant layers) and is a natural target for self-supervised pretraining with masking, denoising, or contrastive objectives; PanopTag improves upon state-of-the-art performance today, while providing a template for extending collider machine learning toward unified, event-level solutions for a broad class of jet analyses tomorrow.

References

[1] "A unified approach for jet tagging in Run 3 at $\sqrt{s}$=13.6 TeV in CMS". In: (2024). URL: https://cds.cern.ch/record/2904702.

[2] Morad Aaboud et al. "Identification and rejection of pile-up jets at high pseudorapidity with the ATLAS detector". In: *Eur. Phys. J. C* 77.9 (2017). [Erratum: Eur.Phys.J.C 77, 712 (2017)], p. 580. DOI: 10.1140/epjc/s10052-017-5081-5. arXiv: 1705.02211 [hep-ex].

[3] G. Aad et al. "The ATLAS Experiment at the CERN Large Hadron Collider". In: *JINST* 3 (2008), S08003. DOI: 10.1088/1748-0221/3/08/S08003.

[4] Georges Aad et al. "Transforming jet flavour tagging at ATLAS". In: (May 2025). arXiv: 2505.19689 [hep-ex].

[5] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E. Hinton. *Layer Normalization*. 2016. arXiv: 1607.06450 [stat.ML]. URL: https://arxiv.org/abs/1607.06450.

[6] Wahid Bhimji et al. "OmniLearned: A Foundation Model Framework for All Tasks Involving Jet Physics". In: (Oct. 2025). arXiv: 2510.24066 [hep-ph].

[7] Johann Brehmer et al. "A Lorentz-Equivariant Transformer for All of the LHC". In: (Nov. 2024). arXiv: 2411.00446 [hep-ph].

[8] Johann Brehmer et al. "Geometric Algebra Transformer". In: *Advances in Neural Information Processing Systems*. Vol. 36. 2023. eprint: 2305.18415. URL: https://arxiv.org/abs/2305.18415.

[9] Matteo Cacciari, Gavin P. Salam, and Gregory Soyez. "FastJet User Manual". In: *Eur. Phys. J. C* 72 (2012), p. 1896. DOI: 10.1140/epjc/s10052-012-1896-2. arXiv: 1111.6097 [hep-ph].

[10] Matteo Cacciari, Gavin P. Salam, and Gregory Soyez. "The anti-k_t jet clustering algorithm". In: *JHEP* 04 (2008), p. 063. DOI: 10.1088/1126-6708/2008/04/063. arXiv: 0802.1189 [hep-ph].

[11] Nicolas Carion et al. "End-to-End Object Detection with Transformers". In: *Computer Vision – ECCV 2020*. Ed. by Andrea Vedaldi et al. Cham: Springer International Publishing, 2020, pp. 213–229. ISBN: 978-3-030-58452-8.

[12] S. Chatrchyan et al. "The CMS Experiment at the CERN LHC". In: *JINST* 3 (2008), S08004. DOI: 10.1088/1748-0221/3/08/S08004.

[13] J. de Favereau et al. "DELPHES 3, A modular framework for fast simulation of a generic collider experiment". In: *JHEP* 02 (2014), p. 057. DOI: 10.1007/JHEP02(2014)057. arXiv: 1307.6346 [hep-ex].

[14] Tobias Golling et al. "Masked particle modeling on sets: towards self-supervised high energy physics foundation models". In: *Mach. Learn. Sci. Tech.* 5.3 (2024), p. 035074. DOI: 10.1088/2632-2153/ad64a8. arXiv: 2401.13537 [hep-ph].

[15] Dan Hendrycks and Kevin Gimpel. *Gaussian Error Linear Units (GELUs)*. 2023. arXiv: 1606.08415 [cs.LG]. URL: https://arxiv.org/abs/1606.08415.

[16] HEP ML Community. *A Living Review of Machine Learning for Particle Physics*. URL: https://iml-wg.github.io/HEPML-LivingReview/.

[17] "Jet energy scale and resolution of jets with ParticleNet p_{T} regression using Run3 data collected by the CMS experiment in 2022 and 2023 at 13.6 TeV". In: (2024). URL: https://cds.cern.ch/record/2904700.

[18] Gregor Kasieczka et al. *Top Quark Tagging Reference Dataset*. Version v0 (2018_03_27). Zenodo, Mar. 2019. DOI: 10.5281/zenodo.2603256. URL: https://doi.org/10.5281/zenodo.2603256.

[19] Subash Katel et al. "Learning Symmetry-Independent Jet Representations via Jet-Based Joint Embedding Predictive Architecture". In: *Postponed: Machine Learning and the Physical Sciences: Workshop at NeurIPS 2024*. Dec. 2024. arXiv: 2412.05333 [hep-ph].

[20] Patrick Komiske, Eric Metodiev, and Jesse Thaler. *Pythia8 Quark and Gluon Jets for Energy Flow*. Version v1. Zenodo, May 2019. DOI: 10.5281/zenodo.3164691. URL: https://doi.org/10.5281/zenodo.3164691.

[21] Patrick T. Komiske, Eric M. Metodiev, and Jesse Thaler. "Energy Flow Networks: Deep Sets for Particle Jets". In: *JHEP* 01 (2019), p. 121. DOI: 10.1007/JHEP01(2019)121. arXiv: 1810.05165 [hep-ph].

[22] Abihith Kothapalli et al. "Equivariant vs. Invariant Layers: A Comparison of Backbone and Pooling for Point Cloud Classification". In: *Proceedings of the Geometry-grounded Representation Learning and Generative Modeling Workshop (GRaM)*. Ed. by Sharvaree Vadgama et al. Vol. 251. Proceedings of Machine Learning Research. PMLR, 29 Jul 2024, pp. 364–380. URL: https://proceedings.mlr.press/v251/kothapalli24a.html.

[23] Juho Lee et al. "Set Transformer: A Framework for Attention-based Permutation-Invariant Neural Networks". In: *Proceedings of the 36th International Conference on Machine Learning*. Ed. by Kamalika Chaudhuri and Ruslan Salakhutdinov. Vol. 97. Proceedings of Machine Learning Research. PMLR, Sept. 2019, pp. 3744–3753. URL: https://proceedings.mlr.press/v97/lee19d.html.

[24] Matthew Leigh et al. "Is Tokenization Needed for Masked Particle Modelling?" In: (Sept. 2024). arXiv: 2409.12589 [hep-ph].

[25] Tsung-Yi Lin et al. "Focal Loss for Dense Object Detection". In: *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*. Oct. 2017.

[26] Ilya Loshchilov and Frank Hutter. "Decoupled Weight Decay Regularization". In: *International Conference on Learning Representations*. 2019. URL: https://openreview.net/forum?id=Bkg6RiCqY7.

[27] Ilya Loshchilov and Frank Hutter. "SGDR: Stochastic Gradient Descent with Warm Restarts". In: *International Conference on Learning Representations*. 2017. URL: https://openreview.net/forum?id=Skq89Scxx.

[28] Vinicius Mikuni and Benjamin Nachman. "Solving key challenges in collider physics with foundation models". In: *Phys. Rev. D* 111.5 (2025), p. L051504. DOI: 10.1103/PhysRevD.111.L051504. arXiv: 2404.16091 [hep-ph].

[29] Spandan Mondal and Luca Mastrolorenzo. "Machine learning in high energy physics: a review of heavy-flavor jet tagging at the LHC". In: *Eur. Phys. J. ST* 233.15-16 (2024), pp. 2657–2686. DOI: 10.1140/epjs/s11734-024-01234-y. arXiv: 2404.01071 [hep-ex].

[30] Adam Paszke et al. *PyTorch: An Imperative Style, High-Performance Deep Learning Library*. 2019. arXiv: 1912.01703 [cs.LG]. URL: https://arxiv.org/abs/1912.01703.

[31] Huilin Qu. *weaver-core: Streamlined neural network training*. Version v0.4.17. GitHub repository (accessed 2025-12-16). Aug. 6, 2024. URL: https://github.com/hqucms/weaver-core.

[32] Huilin Qu and Loukas Gouskos. "ParticleNet: Jet Tagging via Particle Clouds". In: *Phys. Rev. D* 101.5 (2020), p. 056019. DOI: 10.1103/PhysRevD.101.056019. arXiv: 1902.08570 [hep-ph].

[33] Huilin Qu, Congqiao Li, and Sitian Qian. *JetClass: A Large-Scale Dataset for Deep Learning in Jet Physics*. Version 1.0.0. Zenodo, June 2022. DOI: 10.5281/zenodo.6619768. URL: https://doi.org/10.5281/zenodo.6619768.

[34] Huilin Qu, Congqiao Li, and Sitian Qian. "Particle Transformer for Jet Tagging". In: *Proceedings of the 39th International Conference on Machine Learning*. 2022, pp. 18281–18292. arXiv: 2202.03772 [hep-ph].

[35] Umar Sohail Qureshi and Raghav Kunnawalkam Elayavalli. "Deep image reconstruction for background subtraction in heavy-ion collisions". In: *Phys. Rev. C* 112.6 (2025), p. 064906. DOI: 10.1103/dqnx-1fqk. arXiv: 2507.14036 [hep-ph].

[36] Alexander Radovic et al. "Machine learning at the energy and intensity frontiers of particle physics". In: *Nature* 560.7716 (2018), pp. 41–48. DOI: 10.1038/s41586-018-0361-2. URL: https://doi.org/10.1038/s41586-018-0361-2.

[37] Patrick Rieck et al. "Self-supervised learning strategies for jet physics". In: *Mach. Learn. Sci. Tech.* 6.4 (2025), p. 045015. DOI: 10.1088/2632-2153/ae1100. arXiv: 2503.11632 [hep-ph].

[38] Aryan Roy, Jim Pivarski, and Chad Wells Freer. "An array-oriented Python interface for FastJet". In: *Journal of Physics: Conference Series* 2438.1 (Feb. 2023), p. 012011. ISSN: 1742-6596. DOI: 10.1088/1742-6596/2438/1/012011. URL: http://dx.doi.org/10.1088/1742-6596/2438/1/012011.

[39] Sebastian Sapeta. "QCD and Jets at Hadron Colliders". In: *Prog. Part. Nucl. Phys.* 89 (2016), pp. 1–55. DOI: 10.1016/j.ppnp.2016.02.002. arXiv: 1511.09336 [hep-ph].

[40] Torbjörn Sjöstrand et al. "An introduction to PYTHIA 8.2". In: *Comput. Phys. Commun.* 191 (2015), pp. 159–177. DOI: 10.1016/j.cpc.2015.01.024. arXiv: 1410.3012 [hep-ph].

[41] Jonas Spinner et al. "Lorentz-Equivariant Geometric Algebra Transformers for High-Energy Physics". In: *Advances in Neural Information Processing Systems*. Vol. 37. 2024. eprint: 2405.14806. URL: https://arxiv.org/abs/2405.14806.

[42] George Sterman and Steven Weinberg. "Jets from Quantum Chromodynamics". In: *Phys. Rev. Lett.* 39 (23 Dec. 1977), pp. 1436–1439. DOI: 10.1103/PhysRevLett.39.1436. URL: https://link.aps.org/doi/10.1103/PhysRevLett.39.1436.

[43] Yue Wang et al. "Dynamic Graph CNN for Learning on Point Clouds". In: *ACM Transactions on Graphics (TOG)* (2019).

[44] Manzil Zaheer et al. "Deep Sets". In: *Advances in Neural Information Processing Systems*. Ed. by I. Guyon et al. Vol. 30. Curran Associates, Inc., 2017. URL: https://proceedings.neurips.cc/paper_files/paper/2017/file/f22e4747da1aa27e363d86d40ff442fe-Paper.pdf.