PROCEEDINGS OF SCIENCE

Deep Sets and Event-Level Maximum-Likelihood Estimation for Fast Pile-Up Jet Rejection in ATLAS

Mohammed Aboelela* **on behalf of the ATLAS Collaboration**

Southern Methodist University,
Dallas, TX, U.S.A

E-mail: mo.aboelela@cern.ch

Multiple proton–proton collisions occur at every bunch crossing at the LHC, with the mean number of interactions reaching about 64 during Run 3 and is expected to rise to around 200 at the High-Luminosity LHC. As a direct consequence, events with multijet signatures will appear at increasingly high rates. To cope with the higher luminosity, efficiently grouping jets according to their origin along the beamline is crucial, particularly at the trigger level. In this work, a novel uncertainty-aware jet regression model based on a Deep Sets architecture, DIPz, is introduced to predict the longitudinal origin of jets along the beamline using the charged-particle tracks associated with each jet. An event-level discriminant, the Maximum Log Product of Likelihoods (MLPL), is then constructed by combining the DIPz per-jet predictions. MLPL is optimized to select events compatible with targeted multijet signatures. This combined approach provides a robust and computationally efficient method for pile-up jet rejection in multijet final states, suitable for real-time event selections at the ATLAS High-Level Trigger.

The European Physical Society Conference on High Energy Physics (EPS-HEP 2025)
06 July – 11 July 2025
Marseille, France

*Speaker

© Copyright owned by the author(s) under the terms of the Creative Commons
Attribution-NonCommercial-NoDerivatives 4.0 International License (CC BY-NC-ND 4.0).
All rights for text and data mining, AI training, and similar technologies for commercial purposes, are reserved.
ISSN 1824-8039. Published by SISSA Medialab.

PoS(EPS-HEP2025)574

1. Introduction

At the Large Hadron Collider (LHC), proton–proton collisions occur at a rate of about 40 MHz, producing multiple interactions per bunch crossing, known as pile-up. During Run 3, the average number of interactions per crossing is about 64, and it is expected to reach values near 200 at the High-Luminosity LHC (HL-LHC). These overlapping interactions each produce their own jets, leading to dense event topologies that complicate the reconstruction of multijet final states and challenge real-time event selection at the trigger level.

The ATLAS experiment [1] employs a two-level trigger system that reduces the event rate from 40 MHz to about 3 kHz for permanent storage. Within the High-Level Trigger (HLT), where refined reconstruction algorithms within strict latency constraints are executed, charged-particle tracking provides essential information for jet and vertex reconstruction. Performing track finding over the entire inner detector region is essential for precise jet and vertex reconstruction, but is highly CPU-intensive. To mitigate this cost, a pre-selection stage is introduced, in which regional tracking is performed in localized super-Regions of Interest (super-ROIs) around calorimeter jets. Used primarily for fast b-tagging [2], this super-ROI tracking information can also be exploited to identify and reject pile-up jets early in the trigger selection, improving the efficiency of true multi-jet signal event reconstruction.

This work presents a novel approach for fast and efficient pile-up jet rejection in ATLAS. An uncertainty-aware Deep Sets-based neural network, named **DIPz**, regresses the longitudinal production vertex (z-position) of each jet from the charged-particle tracks reconstructed within the corresponding jet super-ROI. Per-jet likelihoods from DIPz are subsequently combined into an event-level discriminant, the **Maximum Log Product of Likelihoods** (**MLPL**), designed to efficiently reject events dominated by pile-up jets. Although broadly applicable, this combined DIPz and MLPL approach is particularly well-suited for optimizing HLT capabilities for high jet-multiplicity physics signatures. It provides a fast rejection mechanism analogous to fast b-tagging strategies, but designed explicitly for flavor-neutral multijet signatures, addressing an important gap in current trigger preselection options.

2. DIPz for uncertainty-aware regression

The DIPz algorithm follows the Deep Sets framework [3], which is well suited for variable-length unordered inputs such as tracks within a jet. Each track is processed through shared feed-forward layers, producing latent representations that are summed to ensure permutation invariance. These are merged with a small set of jet-level features (jet p_{T} and η), and the combined representation is passed to a final network that predicts the mean (μ_z) and standard deviation (σ_z) of a Gaussian describing the jet-vertex position along the beamline. As illustrated in Figure 1, the complete model comprises three sub-networks: a track-level Deep Sets encoder, a jet-level encoder, and a final post-merge network producing the regression outputs. The training target corresponds to the true z-position of the collision vertex from which the jet originated.

The model is trained using simulated $t\bar{t}$ events at $\sqrt{s} = 13.6$ TeV, generated with Powheg+Pythia 8 [4] [5], with pile-up overlaid to emulate Run-3 conditions [6]. The loss function is the Gaussian

PoS(EPS-HEP2025)574

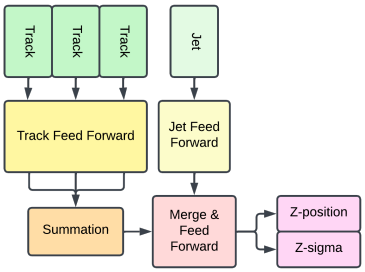

Figure 1: Schematic of the DIPz neural network architecture

negative log-likelihood,

$$L = -\log\left[\frac{1}{\sqrt{2\pi}\sigma_z}\exp\left(-\frac{1}{2}\left(\frac{z_{\text{true}}-\mu_z}{\sigma_z}\right)^2\right)\right], \tag{1}$$

allowing the model to learn both the central value and its associated uncertainty. The model has about 19,650 trainable weights and is implemented using the Keras framework [7] with a TensorFlow backend [8] and trained using an Adam optimizer [9] with early stopping. Training is performed on a dataset of 3 million jets, with validation on a separate set of 600,000 jets.

The DIPz output has been validated on independent $t\bar{t}$, $HH \to 4b$, and multijet samples. Figure 2 shows the distributions of the significance of the prediction error δ of the z-position of jets for each of the three samples. The distributions are well-behaved, with a center around zero and width near unity, indicating good calibration of the per-jet likelihoods. This confirms that the model's uncertainty estimates are statistically consistent with the empirical spread of prediction errors. Such results demonstrate that DIPz functions as an accurate, well-calibrated regressor, capable of not only making precise predictions for the jet vertex z-position but also providing reliable per-instance uncertainty estimates. This robustness across domains further supports the use of DIPz in high-throughput environments such as the ATLAS HLT, where both predictive accuracy and uncertainty quantification are critical.

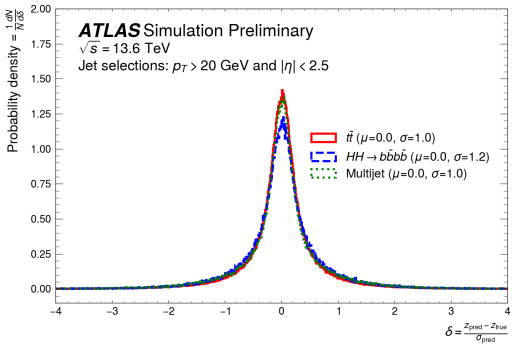

Figure 2: Distributions of the significance of the prediction error δ for the DIPz jet regression model, defined as $(z_{pred} - z_{true})/\sigma_{pred}$, where z_{pred} and σ_{pred} represent the predicted mean and standard deviation of the jet's vertex position along the beamline, respectively [10].

PoS(EPS-HEP2025)574

3. MLPL: the fast rejection algorithm

This approach relies on constructing an event-level discriminant based on the per-jet vertex predictions provided by DIPz. For a given n-jet final state selection, a cut is optimized and applied to this discriminant to reject pile-up-dominated events efficiently, enabling fast decision-making within tight latency constraints faced at the ATLAS trigger system.

The goal of this approach is to identify events in which at least n jets originate from the same primary collision vertex, using only the DIPz neural network outputs. For that, each jet is assigned a Gaussian likelihood function, characterized by its predicted mean μ_z and standard deviation σ_z along the beamline:

$$\mathscr{L}_j(z;\mu_{z_j},\sigma_{z_j}) = \frac{1}{\sqrt{2\pi}\sigma_{z_j}} \exp\left(-\frac{(z-\mu_{z_j})^2}{2\sigma_{z_j}^2}\right). \tag{2}$$

To assess the compatibility of multiple jets with a common vertex, the individual jet likelihoods are multiplied to form a product likelihood for a group of n jets J_n:

$$L(z;J_n) = \prod_{j\in J_n} \mathscr{L}_j(z;\mu_z,\sigma_z) \tag{3}$$

Maximizing $L(z;J_n)$ with respect to z finds the most probable common vertex $\hat{z}$ from which this group of jets J_n can be originating. As a result, this gives a measure of how likely this combination of jets are originating from one vertex when evaluating $L(z;J_n)$ at that most probable $\hat{z}$. It is expected that a group of jets actually originating from one point will give a much larger value of $L(\hat{z};J_n)$ than a group of jets in which the jets are coming from different vertices.

In an n-jet final state candidate event with an arbitrary number of jets m (where $m \geq n$), C_n^m distinct combinations of jets $J_n^{(m)}$ can be formed. When evaluating $L(\hat{z};J_n^{(m)})$ for each group of jets $J_n^{(m)}$, the jet group corresponding to the group of n-jets coming from the primary vertex (collision of interest) is expected to have the largest value compared to the other n-jet groups. Following this logic, the **Maximum Log Product of Likelihoods (MLPL)** discriminant is introduced as:

$$\mathrm{MLPL}(n,m) = \max_{J_n^{(m)}} \left(\max_z \log L(z;J_n^{(m)})\right), \tag{4}$$

where:

- m is the number of the highest-p_T jets considered for combinations in the event (i.e. the size of the pool from which J_n is constructed). It can take any integer value between n (the size of the jet combination) and n_{all} (the total number of jets in the event)

The inclusion of the parameter m controls the computational complexity, which otherwise grows combinatorially with the number of jets. By appropriately choosing n and m, a balance between performance and computational efficiency is achieved.

To evaluate the performance of the MLPL discriminant, a validation sample of simulated $t\bar{t}$ events is used. Targeting optimization for a four-jet signature, the signal region is defined as events containing at least four selected jets, all matched to the same truth collision vertex. The background

PoS(EPS-HEP2025)574

region consists of events containing at least four selected jets where no group of four jets are matched to the same truth collision vertex.

Running the algorithm over the entire signal and background datasets ($\sim 25k$ events each), MLPL distributions are obtained for multiple MLPL(n, m) configurations. Then, Receiver Operating Characteristic (ROC) curves are constructed by scanning thresholds on the discriminant and computing the corresponding signal efficiency and background rejection (Figure 3). The comparison between configurations—varying the number of jets entering the likelihood ($n = 2 - 4$) and the number of highest-p_{T} jets considered ($m = 5, 6, \text{all}$)—shows that larger n values impose stricter topological requirements, yielding stronger signal–background separation, while reducing m limits the combinatorial complexity with only a modest impact on performance. Results show that for a four-jet trigger selection, the pre-selection rejection provided for a signal efficiency of 80% allows to reduce the input rate to the later CPU-intensive selection steps by a factor of four, reducing CPU needs.

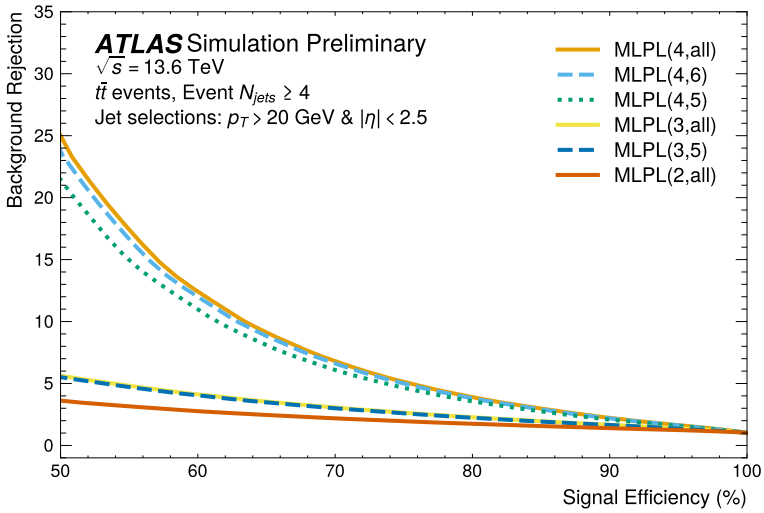

Figure 3: ROC curves showing performance comparison of multiple MLPL(n,m) configurations when targeting optimization for four-jet signature selection [10].

4. Conclusion

Pile-up mitigation is a persistent challenge for multi-jet final states at the LHC. The presented approach combines an uncertainty-aware Deep Sets-based regression model, DIPz, for jet-vertex prediction with an event-level likelihood discriminant, MLPL(n,m), to enable fast and robust pile-up jet rejection. Both components are fully implemented in the ATLAS Athena framework [11] and integrated into the High-Level Trigger, where they run with negligible additional CPU cost compared to existing tracking stages, effectively providing a near cost-free enhancement. Deployed online at Point 1, the DIPz + MLPL algorithm is already used in ATLAS collision-data taking, optimizing trigger chains targeting multijet signatures in 2025 and offering a scalable, low-latency solution for future HL-LHC operations.

PoS(EPS-HEP2025)574

References

[1] ATLAS Collaboration. The ATLAS Experiment at the CERN Large Hadron Collider. *JINST*, 3:S08003, 2008.

[2] ATLAS Collaboration and G. Aad et al. Fast b-tagging at the high-level trigger of the ATLAS experiment in LHC Run 3. *JINST*, 18(11):P11006, 2023. ArXiv preprint.

[3] Patrick T. Komiske, Eric M. Metodiev, and Jesse Thaler. Energy flow networks: deep sets for particle jets. *JHEP*, 01(121), 2019. arXiv preprint.

[4] Stefano Frixione, Paolo Nason, and Carlo Oleari. Matching NLO QCD computations with parton shower simulations: The POWHEG method. *JHEP*, 11:070, 2007.

[5] Torbjörn Sjöstrand, Stephen Mrenna, and Peter Skands. A brief introduction to PYTHIA 8.1. *Comput. Phys. Commun.*, 178:852–867, 2008.

[6] ATLAS Collaboration. Emulating the impact of additional proton–proton interactions in the ATLAS simulation by presampling sets of inelastic Monte Carlo events. *Comput. Softw. Big Sci.*, 6:3, 2022.

[7] François Chollet et al. Keras. https://keras.io, 2015.

[8] Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S. Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Ian Goodfellow, Andrew Harp, Geoffrey Irving, Michael Isard, Yangqing Jia, Rafal Jozefowicz, Lukasz Kaiser, Manjunath Kudlur, Josh Levenberg, Dandelion Mané, Rajat Monga, Sherry Moore, Derek Murray, Chris Olah, Mike Schuster, Jonathon Shlens, Benoit Steiner, Ilya Sutskever, Kunal Talwar, Paul Tucker, Vincent Vanhoucke, Vijay Vasudevan, Fernanda Viégas, Oriol Vinyals, Pete Warden, Martin Wattenberg, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015. Software available from tensorflow.org.

[9] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization, 2017.

[10] ATLAS Collaboration. Deep sets and Event-level Maximum (Jet Trigger Public Results). https://twiki.cern.ch/twiki/bin/view/AtlasPublic/JetTriggerPublicResults#Deep_sets_and_Event_level_Maximu. [Online; accessed 15 Oct 2025].

[11] ATLAS Collaboration. Athena: Software framework of the ATLAS experiment. https://doi.org/10.5281/zenodo.4772550, 2021. Version 21.0.127.

PoS(EPS-HEP2025)574