
Explaining machine-learned particle-flow reconstruction

Farouk Mokhtar, Raghav Kansal, Daniel Diaz, Javier Duarte

University of California San Diego
La Jolla, CA 92093, USA

Joosep Pata

National Institute of Chemical Physics and Biophysics (NICPB)
Rävala pst 10, 10143 Tallinn, Estonia

Maurizio Pierini

European Organization for Nuclear Research (CERN)
CH-1211 Geneva 23, Switzerland

Jean-Roch Vlimant

California Institute of Technology
Pasadena, CA 91125, USA

Abstract

The particle-flow (PF) algorithm is used in general-purpose particle detectors to reconstruct a comprehensive particle-level view of the collision by combining information from different subdetectors. A graph neural network (GNN) model, known as the machine-learned particle-flow (MLPF) algorithm, has been developed to substitute the rule-based PF algorithm. However, understanding the model’s decision making is not straightforward, especially given the complexity of the set-to-set prediction task, dynamic graph building, and message-passing steps. In this paper, we adapt the layerwise-relevance propagation technique for GNNs and apply it to the MLPF algorithm to gauge the relevant nodes and features for its predictions. Through this process, we gain insight into the model’s decision-making.

1 Introduction

Machine learning (ML) has been used in many facets of Large Hadron Collider (LHC) computing and physics analyses. It can provide a faster and more accurate paradigm for designing reconstruction algorithms, compared to hand-coded, rule-based, imperative algorithms. Graph neural networks (GNNs), which are well-suited to non-Euclidean structured data such as detector data, have been shown to be performant for high energy physics (HEP) tasks [1, 2] ranging from charged particle tracking [3–8], to jet classification [9–12], clustering [13, 14], and simulation [15].

An important core software algorithm at the LHC is the particle-flow (PF) algorithm [16–27], which takes heterogeneous detector information, namely charged particle tracks [28–30] and calorimeter energy clusters [31–33], and returns a list of final-state particles known as PF candidates.

In a previous effort, the machine-learned particle-flow (MLPF) [34] algorithm was developed as an end-to-end trainable substitute based on parallelizable, computationally efficient, and scalable GNNs. Those results show that the MLPF algorithm gives comparable performance to a rule-based PF algorithm. In this paper, we expand on this by applying layerwise-relevance propagation (LRP) [35, 36], an explainable artificial intelligence (XAI) [37] technique, to MLPF.

ML models, which operate non-trivially on high-dimensional representations of input features often suffer from a lack of interpretability. Understanding the decision-making behind the PF-candidate class and feature predictions is extremely valuable in order to increase confidence, ensure robustness under changing conditions, and enable us to potentially glean new insights about the detector performance or reconstruction not currently utilized by the rule-based PF algorithm. The term “black box” has long been associated with ML algorithms as they operate on, and learn directly from, data in ways that are not obvious to humans. In recent years, the interpretability of ML models have received increasing attention by both the academic research community and industry under the umbrella of XAI [37].

This paper is organized as follows. Section 2 reviews related work, while section 3 describes the modifications we made to the LRP technique in order to apply it to GNNs, in general, and specifically to the MLPF algorithm. Section 4 describes the dataset and the results in terms of the relevancy scores of different features. Finally, section 5 summarizes our findings.

2 Related Work

XAI in HEP One application of LRP in HEP is a framework presented to extract and understand decision-making information from a DNN classifier of jet substructure tagging techniques [38]. The LRP technique was combined with expert augmented variables to not only find the relevant network information, but also rank features, potentially highlighting a reduced set that capture the network performance. Another application of XAI in HEP is a white box AI [39] approach which demonstrated that a generative adversarial network (GAN) was able not only to learn the final distribution of partons in an event but also extract the underlying physics using full event information registered in the detectors.

XAI on Graphs One attempt to develop a consistent LRP application on graphs for a graph classification task is the GNN-LRP method [40] which proposes explanations in the form of scored sequences of edges on the input graph. Another attempt is the graph layerwise-relevance propagation (GLRP) [41] method which extends the procedure of LRP to make it available for Graph-CNN [42] and test its applicability on a large breast cancer dataset. This was done for a graph signal classification task, which is different from the problem at hand. Another XAI method developed to operate on graphs is the GNNExplainer [43] which attempts to highlight a compact subgraph structure and a small subset of node features that are most relevant for a GNN’s prediction. It attempts to maximize the mutual information between a GNN’s prediction and distribution of possible subgraph structures. We build upon existing LRP methods to adapt them to the GravNet layer in MLPF.

3 LRP Applied to the MLPF Algorithm

The MLPF algorithm is based on the GravNet [13] layer, which learns a high-level embedding of the node positions and performs dynamic graph construction by connecting the $k = 16$ nearest neighbors in this embedding space. It receives as input a set of detector elements $X = \{x_i\}$ per event (up to 5,000 in LHC Run 2 conditions) and attempts to predict the corresponding set of PF candidates $Y = \{y_i\}$. The set X can be best represented as a heterogeneous point cloud in feature space due to the different elements—from different subdetectors—and inherently non-Euclidean structure of the detector. Detector elements and target PF candidates are specified by

$$x_i = [\text{type}, p_T, \eta, \phi, \eta_{\text{out}}, \phi_{\text{out}}, E_{\text{ECAL}}, E_{\text{HCAL}}, \text{charge}, \text{is_gen_el}, \text{is_gen_mu}] \quad (1)$$

$$y_i = [\text{PID}, p_T, E, \eta, \phi, \text{charge}] \quad (2)$$

where $\text{type} \in \{\text{cluster}, \text{track}\}$, $\text{PID} \in \{\text{null}, \text{charged hadron}, \text{neutral hadron}, \text{photon}, \text{electron}, \text{muon}\}$, E_{ECAL} (E_{HCAL}) is the energy in the electromagnetic (hadron) calorimeter, and η_{out} and ϕ_{out} are the extrapolations of η (pseudorapidity) and ϕ (azimuthal angle) to the tracker edge. For input tracks, only the type, p_T (transverse momentum), η , ϕ , η_{out} , ϕ_{out} , and charge features are filled. Similarly, for input clusters, only the type, E_{ECAL} , E_{HCAL} , η , ϕ , entries are filled. As is done in Ref. [34], the target set Y is zero-padded such that $|Y| = |X|$, and Y is then arranged such that if a target PF candidate can be associated geometrically to a detector element input, it is arranged to be in the same location in the sequence.

The tasks at hand are both the classification of PF candidates and the regression of their kinematics. The MLPF model is implemented and trained using PyTorch Geometric [44], which accommodates variable-sized input graphs.

LRP Rules MLPF largely consists of multilayer perceptron (MLP) transformations, to which the standard LRP rules [35, 36] are applicable, with the additional operation of aggregating messages from neighboring nodes in a graph. In this step, as shown in Figure 1, we distribute the relevance scores per node post-aggregation across the nodes pre-aggregation according to the following formula:

$$\mathbf{R}_j^{(l)} = \sum_k \frac{x_j A_{jk}}{\sum_m x_m A_{mk}} \mathbf{R}_k^{(l+1)} \quad (3)$$

where $\mathbf{R}_j^{(l)}$ represent the R -scores of the features of node j at layer l , while the quantity $x_j A_{jk}$ models the extent to which node j at layer l , with activation x_j , contributes to the relevance of node k at layer $l + 1$, where A is the adjacency matrix.




Figure 1: The flow of R -scores of node 1 across the different layers in MLPF. For MLP layers, the redistribution of R -scores follows the standard LRP rules [35, 36]. For the aggregation step in the message passing layer, the redistribution follows Equation 3. We only show three nodes for simplicity.

4 Results

Dataset We train the MLPF algorithm using a detector-agnostic DELPHES [45] dataset from a previous publication [34] found on the Zenodo platform [46] under CC BY 4.0 license. It contains a particle-level training dataset of 50,000 top quark-antiquark ($t\bar{t}$) events produced in proton-proton collisions at 14 TeV generated with PYTHIA 8 [47, 48] and DELPHES 3 [45] from the HepSim software repository [49], overlaid with minimum bias events corresponding to 200 pileup collisions on average. An additional testing dataset comprises 5,000 events composed uniquely of jets produced through the strong interaction, referred to as quantum chromodynamics (QCD) multijet events, with the same pileup conditions, is used for testing the model’s performance and explainability. Additional details can be found in Ref. [34].

We apply LRP to the fully trained MLPF model¹. We run the LRP evaluation on the full testing dataset for about two days on one Nvidia A100 GPU. For each output neuron of each node, we obtain a tensor of the same dimensions as the input tensor of detector element features, which contains the R -scores of all input features in the event for that particular output neuron. We refer to it as a relevancy-tensor (R -tensor). Explicitly, the R -tensor has dimensions $[|Y|, n_y, |X|, n_x]$, where n_y (n_x) is the number of PF candidate (detector element) features. Given the dynamic graph construction, we note only the $k = 16$ nearest neighbors identified in the latent space will have non-zero R -scores. To examine the general decision-making of the model, we randomly sample up to five R -tensors for each of the five PF-candidate node classes per event, sort the nearest neighbors by relevance score and average over many events to produce relevancy-maps (R -maps).

R -maps Examples of R -maps for two common classes in our dataset, charged and neutral hadrons are presented in Figure 2.

¹Code can be found in Ref. [50]




Figure 2: *R*-maps (with normalized *R*-scores computed using LRP) for elements associated to charged hadrons (top), and neutral hadrons (bottom). The *R*-maps are those of the output neuron corresponding to the classification. We see that charged hadrons use more neighbor information than neutral hadrons.

Histograms Another way of visualizing the results of the *R*-maps is to create histograms for each output neuron for each of the classes. Examples of such histograms for charged and neutral hadrons are given in Figure 3. Similar to the *R*-maps we have shown before, we only focus on the classification output neurons. As expected, we find that the charge feature is significantly more relevant for identifying a charged hadron as compared to a neutral hadron. Although it is unsurprising that LRP reaches this conclusion, this work paves the way for exploring other, potentially less obvious, connections that can be surmised using LRP.




Figure 3: We focus on three input features: η , ϕ , and charge, for clarity in the histograms. The charge feature is significantly more relevant for identifying a charged hadron compared to a neutral hadron.

5 Summary

We implement the LRP technique on the MLPF algorithm, a GNN developed for the task of PF reconstruction. We find that, on average, correctly classified charged hadrons make significantly more use of neighboring information than correctly classified neutral hadrons. We suspect that this is because charged hadrons may give rise to multiple tracks (i.e. a primary associated track and its neighbors), which all need to be considered to correctly identify the charged hadron. However, since our model only contains one message passing layer, it might be difficult, in general, for neighbors to have large relevant contributions. This motivates us to further explore models with several message passing layers and check whether more neighbors will become relevant, for neutral

hadron classification for instance, and how this would impact the model performance. It is also interesting to explore in more detail the distribution of the most relevant input features across the different classes to start obtaining physics or detector insights that might not be obvious from a traditional standpoint. Finally, it would be interesting to explore LRP on the MLPF model when trained on a dataset more realistic than DELPHES, i.e, a full detector simulation including all material and physics effects.

Broader Impact Physicists are usually wary of complex ML algorithms because the idea of a “black box” method may not be conducive to scientific discovery and understanding. This effort attempts to carefully examine how an ML model makes its decisions for a challenging particle physics task. This allows us to potentially better trust the ML model, and even learn from it. Furthermore, LRP can be potentially used as a pruning method to limit the resources required to perform PF-reconstruction. This would be very helpful in view of the planned high-luminosity upgrade of the CERN LHC.

Acknowledgments and Disclosure of Funding

F. M. is supported by an Halicioğlu Data Science Institute (HDSI) fellowship and an Institute for Research and Innovation in Software for High Energy Physics (IRIS-HEP) fellowship through the U.S. National Science Foundation (NSF) under Cooperative Agreement OAC-1836650. F. M., R. K., D. D., and J. D. are supported by the US Department of Energy (DOE), Office of Science, Office of High Energy Physics Early Career Research program under Award No. DE-SC0021187 and by the DOE, Office of Advanced Scientific Computing Research under Award No. DE-SC0021396 (FAIR4HEP). R. K. is also supported by the LHC Physics Center at Fermi National Accelerator Laboratory, managed and operated by Fermi Research Alliance, LLC under Contract No. DE-AC02-07CH11359 with the DOE. J. P. is supported by the Mobilitas Pluss Grant No. MOBTP187 of the Estonian Research Council. M. P. is supported by the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation program (Grant Agreement No. 772369). J.-R. V. is supported by the DOE, Office of Science, Office of High Energy Physics under Award No. DE-SC0011925, DE-SC0019227, and DE-AC02-07CH11359. J.-R. V. is additionally supported by the same ERC grant as M. P. This work was performed using the Pacific Research Platform Nautilus HyperCluster supported by NSF awards CNS-1730158, ACI-1540112, ACI-1541349, OAC-1826967, 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. Thanks to CENIC for the 100 Gbps networks. Funding for cloud credits was supported by NSF Award #1904444 Internet2 Exploring Clouds to Accelerate Science (E-CAS).

References

- [1] 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](https://doi.org/10.1088/2632-2153/abbf9a), [arXiv:2007.13681](https://arxiv.org/abs/2007.13681).
- [2] J. Duarte and J.-R. Vlimant, “Graph neural networks for particle tracking and reconstruction”, in *Artificial Intelligence for Particle Physics*. World Scientific Publishing, 2020. [arXiv:2012.01249](https://arxiv.org/abs/2012.01249). Submitted to *Int. J. Mod. Phys. A*. [doi:10.1142/12200](https://doi.org/10.1142/12200).
- [3] S. Farrell et al., “Novel deep learning methods for track reconstruction”, in *4th International Workshop Connecting the Dots*. 2018. [arXiv:1810.06111](https://arxiv.org/abs/1810.06111).
- [4] X. Ju et al., “Graph neural networks for particle reconstruction in high energy physics detectors”, in *2nd Machine Learning and the Physical Sciences Workshop at the 33rd Conference on Neural Information Processing Systems*. 3, 2020. [arXiv:2003.11603](https://arxiv.org/abs/2003.11603).
- [5] S. Amrouche et al., “The Tracking Machine Learning Challenge : Accuracy Phase”, in *The NeurIPS ’18 Competition*, p. 231. 2020. [arXiv:1904.06778](https://arxiv.org/abs/1904.06778). [doi:10.1007/978-3-030-29135-8_9](https://doi.org/10.1007/978-3-030-29135-8_9).
- [6] S. Amrouche et al., “Similarity hashing for charged particle tracking”, in *IEEE International Conference on Big Data 2019*, p. 1595. 2019. [doi:10.1109/BigData47090.2019.9006316](https://doi.org/10.1109/BigData47090.2019.9006316).

- [7] N. Choma et al., “Track seeding and labelling with embedded-space graph neural networks”, in *6th International Workshop Connecting the Dots*. 2020. [arXiv:2007.00149](#).
- [8] G. Dezoort et al., “Charged particle tracking via edge-classifying interaction networks”, (2021). [arXiv:2103.16701](#). Submitted to *Comput. Softw. Big Sci*.
- [9] E. A. Moreno et al., “JEDI-net: a jet identification algorithm based on interaction networks”, *Eur. Phys. J. C* **80** (2020) 58, [doi:10.1140/epjc/s10052-020-7608-4](#), [arXiv:1908.05318](#).
- [10] E. A. Moreno et al., “Interaction networks for the identification of boosted $H \rightarrow b\bar{b}$ decays”, *Phys. Rev. D* **102** (2020) 012010, [doi:10.1103/PhysRevD.102.012010](#), [arXiv:1909.12285](#).
- [11] H. Qu and L. Gouskos, “ParticleNet: Jet Tagging via Particle Clouds”, *Phys. Rev. D* **101** (2020) 056019, [doi:10.1103/PhysRevD.101.056019](#), [arXiv:1902.08570](#).
- [12] V. Mikuni and F. Canelli, “ABCNet: An attention-based method for particle tagging”, *Eur. Phys. J. Plus* **135** (2020), no. 6, 463, [doi:10.1140/epjp/s13360-020-00497-3](#), [arXiv:2001.05311](#).
- [13] S. R. Qasim, J. Kieseler, Y. Iiyama, and M. Pierini, “Learning representations of irregular particle-detector geometry with distance-weighted graph networks”, *Eur. Phys. J. C* **79** (2019) 608, [doi:10.1140/epjc/s10052-019-7113-9](#), [arXiv:1902.07987](#).
- [14] J. Kieseler, “Object condensation: one-stage grid-free multi-object reconstruction in physics detectors, graph and image data”, *Eur. Phys. J. C* **80** (2020), no. 9, 886, [doi:10.1140/epjc/s10052-020-08461-2](#), [arXiv:2002.03605](#).
- [15] R. Kansal et al., “Particle cloud generation with message passing generative adversarial networks”, (2021). [arXiv:2106.11535](#). Preprint.
- [16] CELLO Collaboration, “An analysis of the charged and neutral energy flow in e^+e^- hadronic annihilation at 34 GeV, and a determination of the QCD effective coupling constant”, *Phys. Lett. B* **113** (1982) 427, [doi:10.1016/0370-2693\(82\)90778-X](#).
- [17] ALEPH Collaboration, “Performance of the ALEPH detector at LEP”, *Nucl. Instrum. Meth. A* **360** (1995) 481, [doi:10.1016/0168-9002\(95\)00138-7](#).
- [18] H1 Collaboration, “Measurement of charged particle multiplicity distributions in DIS at HERA and its implication to entanglement entropy of partons”, [arXiv:2011.01812](#).
- [19] ZEUS Collaboration, “Measurement of the diffractive structure function $F_2(D(4))$ at HERA”, *Eur. Phys. J. C* **1** (1998) 81, [doi:10.1007/s100520050063](#), [arXiv:hep-ex/9709021](#).
- [20] ZEUS Collaboration, “Measurement of the diffractive cross-section in deep inelastic scattering using ZEUS 1994 data”, *Eur. Phys. J. C* **6** (1999) 43, [doi:10.1007/PL00021606](#), [arXiv:hep-ex/9807010](#).
- [21] DELPHI Collaboration, “Performance of the DELPHI detector”, *Nucl. Instrum. Meth. A* **378** (1996) 57, [doi:10.1016/0168-9002\(96\)00463-9](#).
- [22] A. Bocci, S. Lami, S. Kuhlmann, and G. Latino, “Study of jet energy resolution at CDF”, *Int. J. Mod. Phys. A* **16S1A** (2001) 255, [doi:10.1142/S0217751X01006632](#).
- [23] A. L. Connolly, “A Search for Supersymmetric Higgs Bosons in the Di-tau Decay Mode in $p\bar{p}$ Collisions at 1.8 TeV”. PhD thesis, UC Berkeley, 2003. [doi:10.2172/15017134](#).
- [24] CDF Collaboration, “Measurement of $\sigma(p\bar{p} \rightarrow Z) \cdot \text{Br}(Z \rightarrow 2\tau)$ in $p\bar{p}$ collisions at $\sqrt{s} = 1.96$ TeV”, *Phys. Rev. D* **75** (2007) 092004, [doi:10.1103/PhysRevD.75.092004](#).
- [25] D0 Collaboration, “Measurement of $\sigma(p\bar{p} \rightarrow Z + X) \cdot \text{Br}(Z \rightarrow \tau^+\tau^-)$ at $\sqrt{s} = 1.96$ TeV”, *Phys. Lett. B* **670** (2009) 292, [doi:10.1016/j.physletb.2008.11.010](#), [arXiv:0808.1306](#).
- [26] CMS Collaboration, “Particle-flow reconstruction and global event description with the CMS detector”, *JINST* **12** (2017) P10003, [doi:10.1088/1748-0221/12/10/P10003](#), [arXiv:1706.04965](#).
- [27] ATLAS Collaboration, “Jet reconstruction and performance using particle flow with the ATLAS detector”, *Eur. Phys. J. C* **77** (2017) 466, [doi:10.1140/epjc/s10052-017-5031-2](#), [arXiv:1703.10485](#).

- [28] A. Strandlie and R. Frühwirth, “Track and vertex reconstruction: From classical to adaptive methods”, *Rev. Mod. Phys.* **82** (2010) 1419, [doi:10.1103/RevModPhys.82.1419](#).
- [29] CMS Collaboration, “Description and performance of track and primary-vertex reconstruction with the CMS tracker”, *J. Instrum.* (2014), no. 10, P10009, [doi:10.1088/1748-0221/9/10/P10009](#), [arXiv:1405.6569](#).
- [30] ATLAS Collaboration, “Performance of the ATLAS track reconstruction algorithms in dense environments in LHC Run 2”, *Eur. Phys. J. C* **77** (2017) 673, [doi:10.1140/epjc/s10052-017-5225-7](#), [arXiv:1704.07983](#).
- [31] C. W. Fabjan and T. Ludlam, “Calorimetry in high-energy physics”, *Annu. Rev. Nucl. Part. Sci.* **32** (1982) 335, [doi:10.1146/annurev.ns.32.120182.002003](#).
- [32] R. Wigmans, “Calorimetry in high energy physics”, in *Techniques and Concepts of High-Energy Physics VI*, T. Ferbel, ed., p. 325. Springer, Boston, MA, USA, 1991. [doi:10.1007/978-1-4684-6006-3_6](#).
- [33] C. Fabjan and F. Gianotti, “Calorimetry for particle physics”, *Rev. Mod. Phys.* **75** (2003) 1243, [doi:10.1103/RevModPhys.75.1243](#).
- [34] J. Pata et al., “MLPF: Efficient machine-learned particle-flow reconstruction using graph neural networks”, *Eur. Phys. J. C* **81** (2021), no. 5, 381, [doi:10.1140/epjc/s10052-021-09158-w](#), [arXiv:2101.08578](#).
- [35] S. Bach et al., “On pixel-wise explanations for non-linear classifier decisions by layer-wise relevance propagation”, *PLOS ONE* **10** (07, 2015) e0130140, [doi:10.1371/journal.pone.0130140](#).
- [36] G. Montavon et al., “Layer-wise relevance propagation: An overview”, in Samek et. al. [37], p. 193. [doi:10.1007/978-3-030-28954-6_10](#).
- [37] W. Samek et al., eds., “Explainable AI: Interpreting, Explaining and Visualizing Deep Learning”. Springer International Publishing, Cham, Switzerland, (2019). [doi:10.1007/978-3-030-28954-6](#).
- [38] G. Agarwal et al., “Explainable AI for ML jet taggers using expert variables and layerwise relevance propagation”, *JHEP* **05** (2021) 208, [doi:10.1007/JHEP05\(2021\)208](#), [arXiv:2011.13466](#).
- [39] Y. S. Lai, D. Neill, M. Płoskoń, and F. Ringer, “Explainable machine learning of the underlying physics of high-energy particle collisions”, [arXiv:2012.06582](#).
- [40] T. Schnake et al., “Higher-order explanations of graph neural networks via relevant walks”, (2020). [arXiv:2006.03589](#). Preprint.
- [41] H. Chereda et al., “Explaining decisions of graph convolutional neural networks: patient-specific molecular subnetworks responsible for metastasis prediction in breast cancer”, *Genome Medicine* **13** (2021), no. 1, 42, [doi:10.1186/s13073-021-00845-7](#).
- [42] M. Defferrard, X. Bresson, and P. Vandergheynst, “Convolutional neural networks on graphs with fast localized spectral filtering”, in *Advances in Neural Information Processing Systems*, D. Lee et al., eds., volume 29. Curran Associates, Inc., 2016. [arXiv:1606.09375](#).
- [43] Z. Ying et al., “Gnnexplainer: Generating explanations for graph neural networks”, in *Advances in Neural Information Processing Systems*, H. Wallach et al., eds., volume 32. Curran Associates, Inc., 2019. [arXiv:1903.03894](#).
- [44] M. Fey and J. E. Lenssen, “Fast graph representation learning with PyTorch Geometric”, in *ICLR Workshop on Representation Learning on Graphs and Manifolds*. 2019. [arXiv:1903.02428](#).
- [45] DELPHES 3 Collaboration, “DELPHES 3, a modular framework for fast simulation of a generic collider experiment”, *JHEP* **02** (2014) 057, [doi:10.1007/JHEP02\(2014\)057](#), [arXiv:1307.6346](#).
- [46] J. Pata et al., “Simulated particle-level events of ttbar and QCD with PU200 using PYTHIA 8+DELPHES 3 for machine learned particle flow (MLPF)”, February, 2021. [doi:10.5281/zenodo.4559324](#).
- [47] T. Sjöstrand, S. Mrenna, and P. Z. Skands, “PYTHIA 6.4 physics and manual”, *JHEP* **05** (2006) 026, [doi:10.1088/1126-6708/2006/05/026](#), [arXiv:hep-ph/0603175](#).

- [48] T. Sjöstrand, S. Mrenna, and P. Z. Skands, “A brief introduction to PYTHIA 8.1”, *Comput. Phys. Commun.* **178** (2008) 852, [doi:10.1016/j.cpc.2008.01.036](https://doi.org/10.1016/j.cpc.2008.01.036), [arXiv:0710.3820](https://arxiv.org/abs/0710.3820).
- [49] S. Chekanov, “HepSim: a repository with predictions for high-energy physics experiments”, *Adv. High Energy Phys.* **2015** (2015) 136093, [doi:10.1155/2015/136093](https://doi.org/10.1155/2015/136093), [arXiv:1403.1886](https://arxiv.org/abs/1403.1886).
- [50] F. Mokhtar, “particleflow”, 2021. [doi:10.5281/zenodo.5723133](https://doi.org/10.5281/zenodo.5723133), <https://github.com/faroukmokhtar/particleflow>.

Checklist

1. For all authors...
 - (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
 - (b) Did you describe the limitations of your work? [Yes]
 - (c) Did you discuss any potential negative societal impacts of your work? [No]
 - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
 - (a) Did you state the full set of assumptions of all theoretical results? [N/A]
 - (b) Did you include complete proofs of all theoretical results? [N/A]
3. If you ran experiments...
 - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes]
 - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes]
 - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [No] , it would have been too computationally expensive
 - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
 - (a) If your work uses existing assets, did you cite the creators? [Yes]
 - (b) Did you mention the license of the assets? [Yes] , dataset has a license that we mentioned in a footnote, while code doesn't have a license
 - (c) Did you include any new assets either in the supplemental material or as a URL? [N/A]
 - (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [Yes]
 - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [Yes]
5. If you used crowdsourcing or conducted research with human subjects...
 - (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
 - (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
 - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]

A LRP for electrons and muons

An interesting feature of the DELPHES simulation is that it has access to generator level information that otherwise would not be accessible. For example, electrons and muons observed in the tracker will respectively have input features `is_gen_el` and `is_gen_mu`. This information is necessary for the simple rule-based PF algorithm in DELPHES, for example, to distinguish photons from electrons, due to the lack of full, realistic detector information. Hence, with DELPHES simulation, we normally restrict our view to the two classes: charged and neutral hadrons, when it comes to assessing the quality of reconstruction.

Nevertheless, for this effort, it is interesting to look at the R -maps of electrons and muons as a way of verifying our LRP implementation.

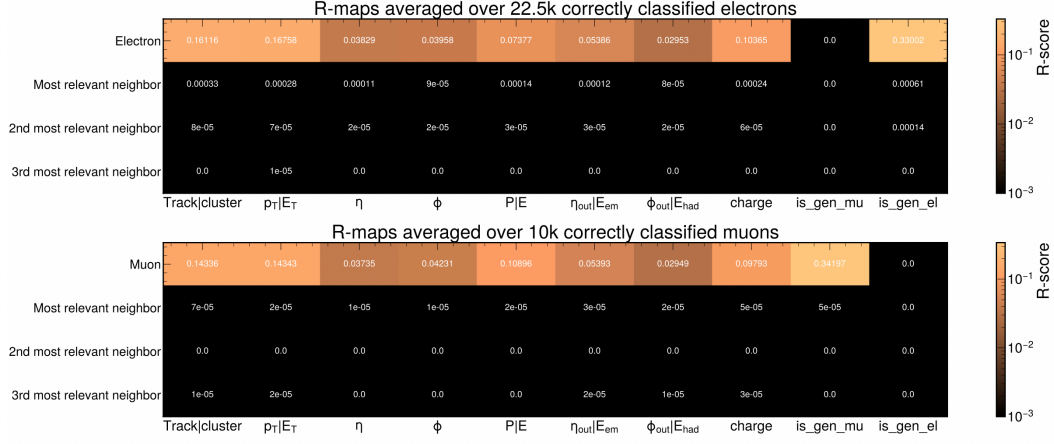


Figure 4: R -maps created for correctly classified electrons (top) and correctly classified muons (bottom). We see that different features are highlighted for each different class, especially `is_gen_el` feature for correctly classified electrons, and `is_gen_mu` feature for correctly classified muons.

Looking at Figure 4, we see that indeed LRP identifies `is_gen_el` and `is_gen_mu` for each of the electrons and muons respectively. Similar to the R -maps we have shown before, we only focus on the classification output neuron for either class.

B Model training and performance

For training the model, we use the same loss function as Ref. [34]. We use the Adam algorithm with a learning rate of 10^{-4} for 300 epochs, training over 4×10^4 events, with 5×10^3 events used for testing. The events are processed in minibatches of four simultaneous events per graphics processing unit (GPU), we train for approximately 12 days using four RTX 1080Ti GPUs simultaneously. The modified architecture achieves comparable performance to the rule-based PF algorithm as shown in Appendix B.

The modified architecture of MLPF achieves comparable performance to the rule-based PF algorithm. Figure 5 shows the classification performance, quantified through a confusion matrix, is similar between the MLPF model and the rule-based PF algorithm. Likewise, Figure 6 shows similar ability of the two algorithms to predict the particle multiplicity for charged and neutral hadrons.

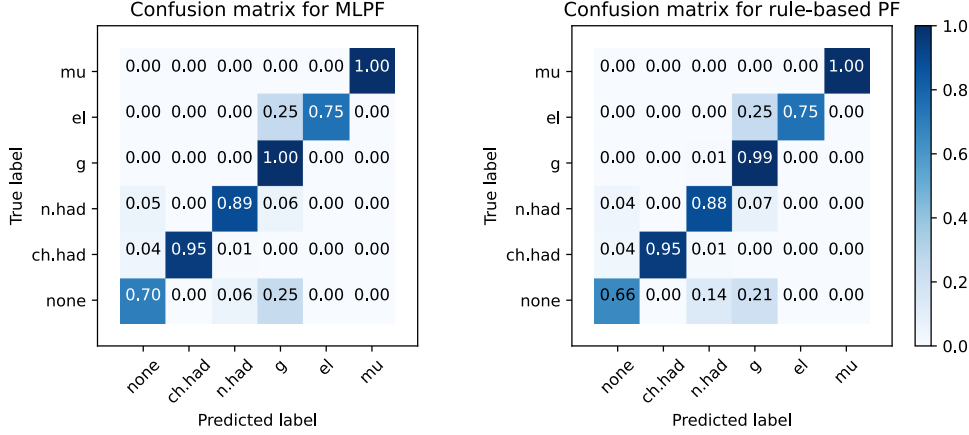


Figure 5: We observe similar classification performance when comparing the MLPF model’s confusion matrix (left) to the rule-based PF algorithm’s confusion matrix (right).




Figure 6: The MLPF model and rule-based PF algorithm perform similarly in predicting the particle multiplicity for charged (left) and neutral hadrons (right).