

# The LHCb ultra-fast simulation option, Lamarr design and validation

Lucio Anderlini<sup>1</sup>, Matteo Barbetti<sup>1,2,\*</sup>, Simone Capelli<sup>3,4</sup>, Gloria Corti<sup>5</sup>, Adam Davis<sup>6</sup>, Denis Derkach<sup>7</sup>, Nikita Kazeev<sup>7</sup>, Artem Maevskiy<sup>7</sup>, Maurizio Martinelli<sup>3,4</sup>, Sergei Mokonenko<sup>7</sup>, Benedetto G. Siddi<sup>8</sup>, and Zehua Xu<sup>9</sup> on behalf of the LHCb Simulation Project

<sup>1</sup>Istituto Nazionale di Fisica Nucleare (INFN), Sezione di Firenze, Italy

<sup>2</sup>Department of Information Engineering (DINFO), University of Firenze, Italy

<sup>3</sup>Istituto Nazionale di Fisica Nucleare (INFN), Sezione di Milano-Bicocca, Italy

<sup>4</sup>Department of Physics, University of Milano-Bicocca, Italy

<sup>5</sup>European Organization for Nuclear Research (CERN), Switzerland

<sup>6</sup>Department of Physics and Astronomy, University of Manchester, United Kingdom

<sup>7</sup>Affiliated with an institute covered by a cooperation agreement with CERN

<sup>8</sup>Department of Physics, University of Ferrara, Italy

<sup>9</sup>Laboratoire de Physique de Clermont (LPC), Université Clermont Auvergne, France

**Abstract.** Detailed detector simulation is the major consumer of CPU resources at LHCb, having used more than 90% of the total computing budget during Run 2 of the Large Hadron Collider at CERN. As data is collected by the upgraded LHCb detector during Run 3 of the LHC, larger requests for simulated data samples are necessary, and will far exceed the pledged resources of the experiment, even with existing fast simulation options. The evolution of technologies and techniques for simulation production is then mandatory to meet the upcoming needs for the analysis of most of the data collected by the LHCb experiment. In this context, we propose LAMARR, a GAUDI-based framework designed to offer the fastest solution for the simulation of the LHCb detector. LAMARR consists of a pipeline of modules parameterizing both the detector response and the reconstruction algorithms of the LHCb experiment. Most of the parameterizations are made of Deep Generative Models and Gradient Boosted Decision Trees trained on simulated samples or alternatively, where possible, on real data. Embedding LAMARR in the general LHCb GAUSS Simulation framework allows combining its execution with any of the available generators in a seamless way. LAMARR has been validated by comparing key reconstructed quantities with Detailed Simulation. Good agreement of the simulated distributions is obtained with two order of magnitude speed-up of the simulation phase.

## 1 Introduction

The LHCb experiment [1] has been originally designed to study rare decays of particles containing  $b$  and  $c$  quarks produced at the Large Hadron Collider (LHC). The LHCb detector is a single-arm forward spectrometer covering the pseudorapidity range of  $2 < \eta < 5$ , that

---

\*e-mail: [matteo.barbetti@cern.ch](mailto:matteo.barbetti@cern.ch)

includes a Tracking system and a Particle Identification (PID) system [2]. The Tracking system provides high-precision measurements of the momentum  $p$  of charged particles and the position of primary vertices. The PID system counts two ring-imaging Cherenkov (RICH) detectors, an electromagnetic calorimeter (ECAL), a hadron calorimeter (HCAL), and a system specialized in the identification of muons called MUON. The separation of charged hadrons mostly relies on the RICH detectors, while the calorimeter system is used to reconstruct photon and electron candidates, and provides additional information for hadrons identification.

Interpreting signal, rejecting background contributions and performing efficiency studies requires a full understanding of the data sample, from the high-energy  $pp$  collisions to the set of physics processes responsible for the detector high-level response, namely the set of variables used by analysts. These studies greatly benefit from the use of simulated samples. At LHCb, the simulation production relies mainly on the GAUSS framework [3], which implements the generation and simulation phases, and is based on the GAUDI processing framework [4]. High-energy collisions and all the decay products (e.g., muons, pions, kaons or protons) are simulated during the *generation phase* using software such as PYTHIA8 [5] and EVTGEN [6]. The radiation-matter interactions between the detector materials and the traversing particles are reproduced during the *simulation phase* that aims to compute the energy deposited in the active volumes and relies on the GEANT4 toolkit [7]. Then, a separate application converts the energy deposits into raw data compatible with the real one collected by LHCb.

The simulation of particle interactions with the detector material is the major consumer of CPU resources at LHCb, having used more than 90% of the total computing budget during LHC Run 2. The upgraded version of the experiment is designed to collect one-order-of-magnitude larger data samples during Run 3. Meeting the upcoming and future requests for simulated samples is not sustainable relying only on the traditional *detailed simulation*. For this reason, the LHCb Collaboration putting a lot of effort in modernizing the simulation software stack through the novel experiment-independent framework GAUSSINO<sup>1</sup> [8, 9] on which a newer version of GAUSS will be built on, and in developing faster simulation options, some of which also powered by machine learning algorithms [10–13].

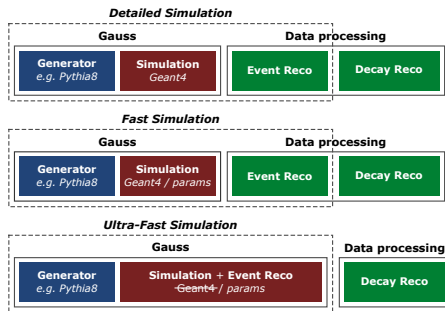
## 2 Fast simulation VS. ultra-fast simulation

Simulating all the decay processes of interest for the LHCb physics program is extremely expensive in terms of computing resources, especially the GEANT4-based step that is the major CPU-consumer. Speeding up the computation of the energy deposits or, more generally, the detector response is mandatory to satisfy the demand for simulations expected for Run 3 and those that will follow. Actually, this is a common problem faced by the High Energy Physics (HEP) community. The latest achievements in Computer Science are therefore exploited and *deep generative models* are typically adapted by various experiments [14–16] to parameterize the detectors low-level response, namely the energy deposited in the active volumes. The literature refers to this kind of strategies with the term *fast simulation*. Fast simulations share their data processing scheme and the reconstruction step with the detailed simulation (as depicted in Figure 1), and have proven to be capable of reducing the computation cost of a simulated sample up to a factor of 20.

To meet the upcoming and future requests for simulated samples, the LHCb Collaboration is also considering a more radical approach based on the so-called *ultra-fast simulation* paradigm. In this case, the aim is to directly reproduce the high-level response of the detector

---

<sup>1</sup>Visit <https://gaussino.docs.cern.ch> for additional details.



**Figure 1.** Schematic representation of the data processing flow in *detailed* (top), *fast* (center) and *ultra-fast* (bottom) simulation paradigms.

relying on a set of parameterizations developed to transform generator-level particles information into reconstructed physics objects, as schematically represented in Figure 1 (bottom). Such parameterizations can be built using generative models, like *Generative Adversarial Networks* (GAN), that have proven to succeed in reproducing the high-level response of the LHCb detector [17] and offering reliable simulated samples [18]. Following pioneering studies on the ultra-fast simulation of the electromagnetic calorimeter based on GANs [19], the CMS Collaboration has recently started developing a full-scope ultra-fast simulation based on *Normalizing Flow*, named FLASHSIM [20].

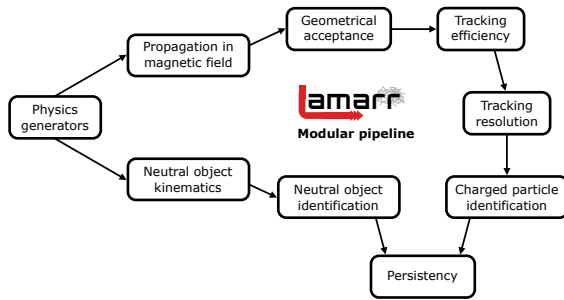
### 3 LAMARR: the LHCb ultra-fast simulation framework

LAMARR [12, 13] is the official ultra-fast simulation framework of the LHCb experiment, which is able to offer the fastest options for simulation. Originating from an attempt of a LHCb-tuned version of DELPHES [21, 22], LAMARR became an independent project retaining only the inspiration of its modular layout from DELPHES. In particular, the LAMARR framework consists of a pipeline of modular parameterizations, most of which are based on machine learning algorithms, designed to take as input the particles generated by the physics generators and provide as output the high-level response of the various LHCb sub-detectors.

The LAMARR pipeline can be logically split in two separated chains according to the charge of the generated particles. We expect that charged particles leave a mark in the Tracking system that LAMARR characterizes in terms of acceptance, efficiency and resolution, as described in Section 3.1. The reconstructed tracking variables are then used to compute the response of the PID system for a set of traversing charged particles (muons, pions, kaons or protons), as detailed in Section 3.2. In case of neutral particles (e.g., photons), the calorimeters play a key role and, since multiple photons can concur to the energy of a single calorimetric cluster, parameterizing particle-to-particle correlation effects is of major relevance. The solutions under investigation are reported in Section 3.3. The LAMARR pipelines described above are shown in Figure 2.

#### 3.1 Tracking system

One of the aims of the LHCb Tracking system is to measure the momentum  $p$  of charged particles (i.e., electrons, muons, pions, kaons and protons), exploiting the deflection of their trajectories due to the dipole magnet located in between the tracking detectors. Hence, the first step of the *charged chain* reported in Figure 2 is the propagation of charged particles provided by physics generators through the magnetic field. LAMARR parameterizes the particle trajectories as two rectilinear segments with a single deflection point (inversely proportional to the transverse momentum  $p_T$ ), implementing the so-called *single  $p_T$  kick* approximation.



**Figure 2.** Scheme of the LAMARR modular pipeline. According to the charge of the particle provided by the physics generator, two sets of parameterizations are defined: the charged particles are passed through the Tracking and PID models, while the neutral ones follow a different path where the calorimeter modeling plays a key role.

The next step requires to select the subset of tracks that fall within the LHCb geometrical acceptance and that have any chance to be reconstructed. To this end, LAMARR uses *Gradient Boosted Decision Trees* (GBDT) trained to learn the fraction of candidates that are in the acceptance of the detector as a function of the kinematic information provided by physics generators. Given a generated track in acceptance, we ask whether it will be reconstructed and, in case of positive answer, which tracking detectors will be involved in the reconstruction procedure. LAMARR statistically infers such information, namely the tracking efficiency, relying on *neural networks* trained to perform a multi-class classification according to the track kinematics. A major effort is ongoing to improve the performance of the efficiency model on the basis of the type of tracks and particle species (i.e., electrons, muons or hadrons).

At this point, LAMARR disposes of the subset of the generated particles that can be considered as reconstructed tracks, but their kinematics and geometry are still identical to those provided by physics generators. Smearing of these features, mimicking the effect of the reconstruction algorithms, is achieved using GANs. Driven by a *binary cross-entropy* loss function and powered by *skip connections* [23], GANs are able to describe the resolution effects only relying on track kinematic information at generator-level as input conditions. A similar GAN-based architecture is used to provide the correlation matrix obtained from the Kalman filter adopted in the reconstruction algorithm to define the position, slope and curvature of each track.

Stacking the parameterizations described above, LAMARR is able to provide the high-level response of the LHCb Tracking system. The resulting reconstructed quantities can be further processed using the LHCb analysis software to combine the parameterized tracks into decay candidates, as depicted by the green slot in Figure 1 (bottom).

### 3.2 Particle identification system

To accomplish the LHCb physics program, a high-performance PID system is crucial since it allows for discriminating various particle species that traverse the detector. LAMARR provides parameterizations for the majority of charged particles for which the PID detectors are relevant (i.e., muons, pions, kaons or protons). Specialized parameterizations for electrons, encoding multiple scattering and Bremsstrahlung emission contributions in the interaction with the detector materials, is planned as a future development.

Identifying these particles involves mainly the RICH and MUON detectors, while the role played by the calorimeters is minor. In general, we expect that the response of the PID system depends only on the specie of the traversing particle, its kinematics, and the detector occupancy. According to these dependencies, LAMARR parameterizes the high-level response of the RICH and MUON detectors using GAN-based models properly conditioned [11, 18]. Given the particle specie from the physics generators, its analysis-level kinematic information

results from the LAMARR Tracking modules, while the detector occupancy is described by the total number of tracks traversing the detector.

In real data, the combination of responses from RICH detectors, calorimeters, MUON system and a binary muon-identification criterion implemented via Field Programmable Gate Array (FPGA) devices and named `isMuon` [24], allows to compute the higher-level response of the PID system, referred to as GlobalPID variables. The parameterization of these variables relies on conditioned GANs (GlobalPIDGAN), adding as input what results from the `RichGAN` and `MuonGAN` models. The binary output of a neural-network-based implementation of `isMuon` is used as additional input feature, while no explicit calorimeters contribution is defined, leaving the missing information problem to the generator *latent space*.

`RichGAN` and `MuonGAN` models, driven by a *Wasserstein distance* loss function and trained using a Lipschitz-constrained discriminator [25], succeed in describing the high-level response of the RICH and MUON systems. Chaining together different GANs, LAMARR allows to provide the higher-level response of the LHCb PID system using the GlobalPIDGAN model, that injects an implicit contribution from the calorimeters adopting the same training strategy of RICH and MUON parameterizations.

### 3.3 Electromagnetic calorimeter

Providing a parameterization for electrons requires describing the response to Bremsstrahlung photons by the LHCb ECAL detector. Since interested in a large number of secondary particles, detailed simulation of the calorimeter system is the most computationally expensive step in the simulation pipeline. It is a common problem in the HEP community, that is investing great efforts in tuning deep generative models to properly parameterize the energy deposited in the calorimeter cells [10, 14–16]. Such studies belong to the fast simulation paradigm that aims to reduce the GEANT4 use by providing models for the low-level response of the various experiments.

The current version of LAMARR provides a simplified parameterization for the LHCb calorimeter, designed for detector studies and based on a fast simulation approach. Retrieving information at the calorimeter cell level requires running reconstruction algorithms that may become rather CPU-expensive for high-multiplicity events. In addition, since non-physical strategies are used to simulate the energy deposits (as in the case of GANs), there is no certainty that the reconstruction software stack can correctly reproduce the expected distributions for the high-level variables [26]. Hence, the LAMARR project is actively working to provide an ultra-fast solution for the ECAL detector.

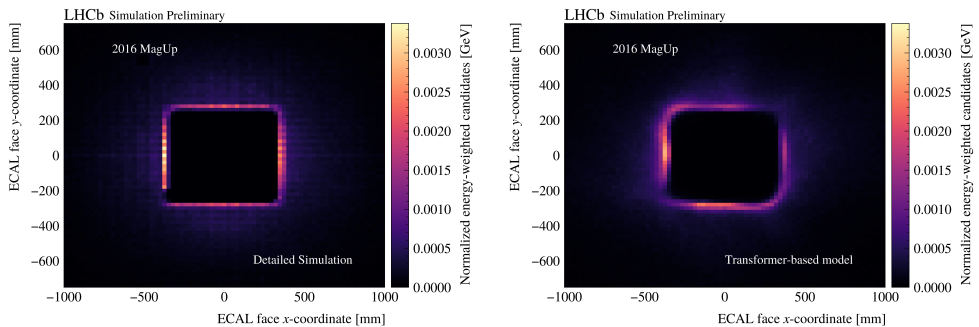
Reproducing the calorimeter high-level response is a non-trivial task since traditional generative models rely on the hypothesis that there is an unambiguous relation between the generated particle and the reconstructed object<sup>2</sup>. Instead, the presence of merged  $\pi^0$  and Bremsstrahlung photons may lead to having  $n$  generated particles responsible for  $m$  reconstructed objects (in general with  $n \neq m$ ). A strategy to face this particle-to-particle correlation problem can be built using techniques designed in the context of Language Modeling, describing the calorimeter simulation as a *translation problem*. For this purpose, *Graph Neural Network* (GNN) [27] and *Transformer* [28] models are currently under investigation.

Both the models are designed to process a sequence of  $n$  generated photons and infer the kinematics of a sequence of  $m$  reconstructed clusters. The non-trivial correlations between any particles of the source sequence (photons) and the target one (clusters) rely on the *attention mechanism* [28, 29]. To improve the quality of the resulting parameterizations, the

---

<sup>2</sup>To a first approximation, the response of the Tracking and PID systems satisfy this condition.

training of both GNN-based and Transformer-based models is driven by an adversarial procedure (similarly to GANs). The discriminator is currently implemented through a *Deep Sets* model [30], while further studies are ongoing to replace it with a second Transformer [31]. Considering the complexity of the problem, the preliminary results are promising, as depicted in Figure 3, where the joint action of Transformer and Deep Sets succeeds in deriving the energy distribution on the ECAL face. The center of the calorimeter has no active material as it is used to host the LHC beam pipe. It should be pointed out that no constraints are applied to the model output to reproduce such conditions, and that the empty space shown in Figure 3 (right) is the result of the adversarial training procedure.




**Figure 3.** Distribution of the  $(x, y)$ -position of the reconstructed clusters on the LHCb ECAL face for a  $2000 \times 1500 \text{ mm}^2$  frame placed around the center. The geometrical information is combined with the energy signature by properly weighting each bin entry. The results obtained from detailed simulation are reported on the left, while the predictions of an adversarial trained Transformer model are shown on the right. The corresponding LHCb-FIGURE is in preparation.

## 4 Validation campaign and timing performance

The ultra-fast philosophy at the base of the LAMARR framework is being validated by comparing distributions obtained from machine-learned models trained on detailed simulation and distributions resulting from standard simulation strategies. In particular, we will briefly discuss the validation studies performed for the charged particles using simulated  $\Lambda_b^0 \rightarrow \Lambda_c^+ \mu^- \bar{\nu}_\mu$  decays with  $\Lambda_c^+ \rightarrow p K^- \pi^+$ . The semileptonic nature of the  $\Lambda_b^0$  decay requires an interface with dedicated generators, in this case EVTGEN. Deeply studied by LHCb, this decay channel includes in its final state the four charged particle species parameterized in the current version of LAMARR, namely muons, pions, kaons and protons.

The validation of the LAMARR Tracking modules is depicted in Figure 4 (left) where the agreement between the  $\Lambda_c^+$  invariant mass distribution resulting from the ultra-fast paradigm and the one obtained from detailed simulation proves that the decay dynamics is well reproduced and the resolution effects correctly parameterized. To demonstrate good performance of the LAMARR PID models, a comparison between the selection efficiencies for a tight requirement on a multivariate proton classifier is shown in Figure 4 (right).

Comparing the CPU time spent per event by GEANT4-based production of  $\Lambda_b^0 \rightarrow \Lambda_c^+ \mu^- \bar{\nu}_\mu$  samples and the one needed by LAMARR, we estimate a CPU reduction of two order of magnitude only for the simulation phase. Interestingly, since the generation of  $b$ -baryons is exceptionally expensive, PYTHIA8 becomes the major consumer of CPU resources in the ultra-fast paradigm. Further speed-up can be reached by reducing the generation time, for example by



**Figure 4.** Validation plots for  $\Lambda_b^0 \rightarrow \Lambda_c^+ \mu^- \bar{\nu}_\mu$  decays with  $\Lambda_c^+ \rightarrow pK^-\pi^+$  simulated with PYTHIA8, EVTGEN and LAMARR (orange markers) and compared to detailed simulation samples relying on PYTHIA8, EVTGEN and GEANT4 (cyan shaded histogram). Reproduced from LHCb-FIGURE-2022-014.

using a *Particle Gun* that simulates directly the signal particles without going through the high-energy collisions which is not needed since LAMARR parameterizes the detector occupancy. Even in these simplified settings, the ultra-fast philosophy succeeds in reproducing distributions obtained from detailed simulation [12].

## 5 Integration with the LHCb simulation framework

To be integrated within the LHCb software stack, the parameterizations provided by LAMARR need to be queried from a C++ application, running in the GAUDI framework. Traditional deployment strategies were found to lead to unacceptably large overheads due to the presence of different multi-threading schedulers and context switching issues. Hence, a custom deployment strategy was preferred: models trained with `scikit-learn` [32] and `Keras` [33] are converted into compatible C code using the `scikinC` toolkit [34], and then distributed through the LHCb Computing Grid via the CERN VM file-system (`cvmfs`) [35].

The modular layout of LAMARR enables a variety of studies and developments on the single parameterizations, providing a unique and shared infrastructure for validation and performance measurements. While crucial for applications within LHCb, the integration with GAUDI and GAUSS makes the adoptions of LAMARR unappealing for researchers outside of the LHCb community. The `SQLamarr` package<sup>3</sup> aims to mitigate this problem, providing a stand-alone ultra-fast simulation framework with minimal dependencies. Based on `SQLite3`, `SQLamarr` provides a set of classes and functions for loading data from physics generators and defining pipelines from compiled models. An integration between `SQLamarr` and `GAUSSINO` is currently under investigation with the goal of providing ultra-fast parameterizations following the experiment-independent philosophy of the newest LHCb simulation framework, called `GAUSS-ON-GAUSSINO`<sup>4</sup> [8, 9].

## 6 Conclusion

Evolution of the LHCb software stack and simulation techniques are mandatory to meet the upcoming and future demand for simulated samples expected for Run 3 and those to follow.

<sup>3</sup>Visit <https://lamarrsim.github.io/SQLamarr> for additional details.

<sup>4</sup>Visit <https://lhcb-gauss.docs.cern.ch/Futurev5> for additional details.

Ultra-fast solutions will play a key role in reducing the pressure on pledged CPU resources without unreasonably compromising the description of the uncertainties introduced in the detection and reconstruction phases. Such techniques, powered by deep generative models, are provided to LHCb via the novel LAMARR framework. Well integrated with physics generators within the GAUSS framework, LAMARR delivers two pipelines according to the charge of the generated particle. The parameterizations for the Tracking and the PID systems have been deployed for charged particles and validated with satisfactory results on  $\Lambda_b^0 \rightarrow \Lambda_c^+ \mu^- \bar{\nu}_\mu$  decays. Several models are currently under investigation for the neutral pipeline, where the translation problem approach offers a viable solution to face the particle-to-particle correlation problem. Further development of the integration between LAMARR and the LHCb simulation framework is one of the major ongoing activities to bring the former into the production and make its parameterizations available to the HEP community.

## Acknowledgements

This work is partially supported by ICSC – Centro Nazionale di Ricerca in High Performance Computing, Big Data and Quantum Computing, funded by European Union – NextGenerationEU.

## References

- [1] A.A. Alves, Jr. et al. (LHCb), *JINST* **3**, S08005 (2008)
- [2] R. Aaij et al. (LHCb), *Int. J. Mod. Phys. A* **30**, 1530022 (2015), 1412.6352
- [3] M. Clemencic et al. (LHCb), *J. Phys. Conf. Ser.* **331**, 032023 (2011)
- [4] G. Barrand et al., *Comput. Phys. Commun.* **140**, 45 (2001)
- [5] T. Sjostrand, S. Mrenna, P.Z. Skands, *Comput. Phys. Commun.* **178**, 852 (2008), 0710.3820
- [6] D.J. Lange, *Nucl. Instrum. Meth. A* **462**, 152 (2001)
- [7] J. Allison et al., *IEEE Trans. Nucl. Sci.* **53**, 270 (2006)
- [8] M. Mazurek, G. Corti, D. Müller, *Comput. Inform.* **40**, 815–832 (2021), 2112.04789
- [9] M. Mazurek, M. Clemencic, G. Corti, *PoS ICHEP2022*, 225 (2023)
- [10] V. Chekalina et al., *EPJ Web Conf.* **214**, 02034 (2019), 1812.01319
- [11] A. Maevskiy et al. (LHCb), *J. Phys. Conf. Ser.* **1525**, 012097 (2020), 1905.11825
- [12] L. Anderlini et al., *PoS ICHEP2022*, 233 (2023)
- [13] M. Barbeti (2023), *ACAT'22*, 2303.11428
- [14] M. Paganini, L. de Oliveira, B. Nachman, *Phys. Rev. D* **97**, 014021 (2018), 1712.10321
- [15] C. Krause, D. Shih, *Phys. Rev. D* **107**, 113003 (2023), 2106.05285
- [16] O. Amram, K. Pedro (2023), 2308.03876
- [17] F. Ratnikov et al., *Nucl. Instrum. Meth. A* **1046**, 167591 (2023)
- [18] L. Anderlini et al. (LHCb), *J. Phys. Conf. Ser.* **2438**, 012130 (2023), 2204.09947
- [19] P. Musella, F. Pandolfi, *Comput. Softw. Big Sci.* **2**, 8 (2018), 1805.00850
- [20] F. Vaselli et al. (CMS), *Tech. rep.* (2023), <https://cds.cern.ch/record/2858890>
- [21] J. de Favereau et al. (DELPHES 3), *JHEP* **02**, 057 (2014), 1307.6346
- [22] B.G. Siddi, *EPJ Web Conf.* **214**, 02024 (2019)
- [23] K. He et al. (2016), *CVPR'16*, 1512.03385
- [24] G. Lanfranchi et al. (LHCb), *Tech. rep.* (2009), <https://cds.cern.ch/record/1202759>



- [25] D. Terjék (2020), ICLR'20, 1907.05681
- [26] A. Rogachev, F. Ratnikov, J. Phys. Conf. Ser. **2438**, 012086 (2023), 2207.06329
- [27] F. Scarselli et al., IEEE Trans. Neural. Netw. **20**, 61 (2009)
- [28] A. Vaswani et al. (2017), NeurIPS'17, 1706.03762
- [29] S. Brody, U. Alon, E. Yahav (2022), ICLR'22, 2105.14491
- [30] M. Zaheer et al. (2017), NeurIPS'17, 1703.06114
- [31] K. Lee et al. (2022), ICLR'22, 2107.04589
- [32] F. Pedregosa et al., Journal of Machine Learning Research **12**, 2825 (2011)
- [33] F. Chollet et al., *Keras* (2015), software available from <https://keras.io>
- [34] L. Anderlini, M. Barbetti, PoS **CompTools2021**, 034 (2022)
- [35] P. Buncic et al., J. Phys. Conf. Ser. **219**, 042003 (2010)