

Refining Jets for CMS Run 3 using Fast Simulation

Açelya Deniz Güngördü^{1,}, Dorukhan Boncukçu^{1,**}, Kevin Pedro^{2,***}, Samuel Bein^{3,****}, Moritz Wolf⁴, Altan Çakır¹, and Mirac Ekim Vural¹ on behalf of the CMS Collaboration*

¹Istanbul Technical University, Science and Letters Faculty, Istanbul, Türkiye

²Fermi National Accelerator Laboratory, Batavia, IL, USA

³Université Catholique de Louvain, Louvain-la-Neuve, Belgium

⁴University of Hamburg, Institute of Experimental Physics, Hamburg, Germany

Abstract. As the LHC moves into its high-luminosity phase, the CMS experiment must handle more complex data collected at much higher rates. While the Geant4-based simulation application (FullSim) provides highly accurate simulation to complement real data, FullSim's intensive consumption of computing resources becomes an increasing liability as the rates increase, while faster tools offer an advantage. The fast MC production application (FastSim) delivers a complete simulation with a factor of 10 speedup over FullSim, but introduces inaccuracies in some observables. A specialized refinement method, Fast Perfect, employs machine learning to improve the accuracy of FastSim. An initial report of this work focused on the refinement of jet flavor tagging observables. This article presents an update on the refinement, focusing on PUPPI jets with Run 3 data-taking conditions. Refinement is extended to include jet transverse momentum as well as its propagation to missing transverse momentum. A grid-based framework and real-time monitoring system have been developed to facilitate optimization and scaling of the refinement to a large number of target variables.

1 Introduction

The CMS experiment [1] at the CERN LHC is preparing to record an order of magnitude more data after introducing new highly complex detectors [2]. In the upcoming High-Luminosity LHC (HL-LHC) runs, efficiently processing and analyzing this large quantity of data will be essential for maintaining the experiment's scientific output and ensuring quality results [3], as will producing sufficiently large and accurate simulated data sets to complement the observed data.

The physics analysis pipeline relies on simulation. To address the increasing computational demands, the CMS experiment employs two distinct simulation chains, called FullSim and FastSim. FullSim [4, 5], based on the GEANT4 toolkit [6, 7], simulates the detector response to collision events at state-of-the-art accuracy, based on realistic transport of particles traversing and interacting with the detector. Although highly accurate, FullSim also has high

*e-mail: acelya.deniz.gungordu@cern.ch

**e-mail: dorukhan.boncukcu@cern.ch

***e-mail: pedrok@cern.ch

****e-mail: samuel.bein@cern.ch

computational intensity, which results in slower processing, posing a challenge for scalability in the high-luminosity era. In contrast, FastSim [8–10] implements a set of approximations that significantly reduce computation time, achieving speeds approximately ten times faster than FullSim. Particles are propagated through a simplified detector geometry comprising infinitesimal surfaces corresponding to detector components, where interactions take place on the layer boundaries rather than in the bulk volume. Particle interactions with the material are modeled with analytic functions and histogram libraries, with calorimeter showering is parametrized with GFlash [11]. Generator-assisted track identification is also implemented to accelerate reconstruction. FastSim brings a significant advantage for statistically limited analyses. However, the acceleration comes at the cost of reduced accuracy in certain observables.

To increase the accuracy of the final outputs of FastSim while maintaining its advantage in throughput, a final step in the production chain is introduced that refines the values of variables and collections. A multivariate regression neural network (NN) is trained to refine the FastSim outputs to be more similar to FullSim. A prototype, based on the Fast Perfekt method [12], was recently introduced [13], focusing on jet flavor observables. This proceeding presents an update on the effort to refine FastSim, extending to the jet momentum and its impact on the missing transverse momentum p_T^{miss} . This strategy aims to ensure that the CMS collaboration can effectively navigate the challenges of the HL-LHC.

2 Methodology

The goal is to train and employ an ML-based regression that modifies the output values of FastSim variables to be closer to FullSim. The regression is conditioned on the generator (truth-level) jet p_T , η , and hadronic flavor, while refining jet flavor tagging variables. Following the methodology of Ref. [13], the training focuses on the simulation of gluino pair production in the framework of the simplified model T1tttt [14] generated using the PYTHIA 8.309 event generator [15] at leading order precision. Each gluino decays into a pair of top quarks and a neutralino, yielding a diverse range of final state objects, including leptons and jets of all flavors over a wide energy range, giving broad support for the phase space of most applications.

Identical sets of generated events are processed with FullSim and FastSim separately, and each generated jet is matched to its reconstructed FullSim and FastSim counterpart, resulting in “jet triplets”. A summary of the jet reconstruction and selection is given in Table 1.

Table 1: Jet reconstruction and selection parameters for the training dataset.

Parameter	Value
Jet algorithm	anti- k_T PUPPI jets [16–19]
Distance parameter (R)	0.4
Jet/jet matching (ΔR)	< 0.2
Overlap removal (ΔR)	< 0.5
p_T	> 15 GeV
$ \eta $	< 5.0

To preserve the association between reconstructed and generator-level jets, the generated jets, FullSim jets, and FastSim jets are required to be separated from neighboring jets (in the same collection) by $\Delta R > 0.5$. The resulting training sample comprises approximately 6×10^6 jet triplets.

2.1 Network training

The task of the NN is to modify the FastSim jet properties in order to render them as similar as possible to their more accurate FullSim counterparts. The NN takes as input $[\vec{x}_{\text{Fast}}, \vec{g}]$ and outputs a refined vector $\vec{x}_{\text{Refined Fast}}$; the concatenation of this output and the generator vector $[\vec{x}_{\text{Refined Fast}}, \vec{g}]$ is compared to the target FullSim vector $([\vec{x}_{\text{Full}}, \vec{g}])$. As shown in Table 2, the process uses four key jet flavor tagging observables derived from the DeepJet algorithm [20]: B for generic b-tagging, and CvB, CvL, and QvG for discrimination of charm vs. bottom, charm vs. light, and quark vs. gluon, respectively. The refinement is guided by auxiliary parameters ($\vec{g}$) including generator-level p_T , η , and true hadron flavor information.

Parameter	Definition
Input: $\vec{x}_{\text{Fast}}^T$	$(p_{T,\text{Fast}}, B_{\text{Fast}}, \text{CvB}_{\text{Fast}}, \text{CvL}_{\text{Fast}}, \text{QvG}_{\text{Fast}})^T$
Auxiliary Parameters ($\vec{g}$)	$p_{T,\text{GEN}}, \eta_{\text{GEN}}, \text{true hadron flavor (b, c, or uds/g)}$
Target: $\vec{x}_{\text{Full}}^T$	$(p_{T,\text{Full}}, B_{\text{Full}}, \text{CvB}_{\text{Full}}, \text{CvL}_{\text{Full}}, \text{QvG}_{\text{Full}})^T$

Table 2: The FastSim input features, FullSim output targets, and auxiliary parameters.

2.2 Network architecture

The refinement network is a ResNet-like NN [21] that estimates the residual differences between FastSim and FullSim, and adds these residuals to the FastSim. The architecture comprises three main components, as illustrated in Fig. 1: a pre-processing layer, a series of four residual blocks, and a post-processing layer, forming an end-to-end refinement pipeline.

The input layer accepts the FastSim feature vector ($\vec{x}_{\text{Fast}}$) and auxiliary parameters ($\vec{g}$), as depicted on the left side of Fig. 1. These inputs undergo pre-processing to normalize and prepare the data for the network’s core processing stages. The backbone of the architecture consists of four consecutive residual blocks, each incorporating skip connections to maintain gradient flow and prevent information loss during training.

The biased maximum mean discrepancy (MMD) is used as the loss function. This loss compares the FastSim and FullSim jet populations by assessing the variation within and between FastSim and FullSim. The MMD is 0 in the limit that the two samples are identical. The MMD is computed over batches of size 2048. The loss is combined with an additional constraint that maintains unitarity among the flavor observables using the modified differential method of multipliers.

3 Results for jet flavor and momentum

Simultaneous refinement of jet flavor and jet p_T has been performed and the results are shown in Fig. 2 (upper left). The refined FastSim exhibits clear improvement in modeling. Good refined modeling of jet p_T is seen throughout the range spanning 30 GeV to 1 TeV, with a marginal improvement seen outside these ranges. Where FastSim’s overly broad jet energy resolution shifts the p_T spectrum to higher values, the refinement brings the spectrum into alignment with FullSim. The results of the b-tagging discriminant are shown in Fig. 2 (upper right), where pronounced discrepancies in FastSim’s modeling are mostly mitigated by the refinement. The FastSim test entry represents an independent sample of refined FatsSim jets, included as a check of overtraining. The impact of the refinement on correlations among the




Figure 1: The architecture of the refinement network, taken from [12]. The diagram shows the complete processing pipeline from FastSim inputs through residual blocks to refined outputs, including the dual loss function implementation. In the presented application, the MSE is not used.

refined and auxiliary variables are studied in terms of Pearson coefficients. These coefficients are displayed in Fig. 2 (lower), along with residuals of the correlations with respect to those present in FullSim. Correlations among the refined FastSim variables are much closer to those of FullSim, with the largest improvement observed the correlation between p_T and the quark-gluon discriminator. A small amount of mismodeling remains after refinement, but the Fast Perfekt approach provides sufficient accuracy for typical analyses.

4 Propagation of corrections to missing momentum

The impact of refinement on the missing transverse momentum is also studied. The so-called hard p_T^{miss} is used, defined as the magnitude of the negative vector sum of the p_T of jets within an event that satisfy $p_T \geq 30 \text{ GeV}$ and $|\eta| \leq 5$.

The effect of refinement on the jet p_T is propagated as a correction to the p_T^{miss} [22]. This correction amounts to replacing the contribution of each jet's momentum to the p_T^{miss} with that of the refined jet.

The result of these corrections is shown in Fig. 3. The distribution of the refined FastSim hard p_T^{miss} aligns closely with that of FullSim, while the original FastSim deviates slightly in the lowest and highest ranges. The event-by-event difference between the FastSim and FullSim hard p_T^{miss} before and after refinement, $\Delta \text{Hard} p_T^{\text{miss}}$ is shown for all events, as well as its profile in bins of FullSim hard p_T^{miss} . The differences corresponding to refined FastSim are narrower and more centered around zero than for the unrefined FastSim, indicating an event-by-event improvement in FastSim modeling. A comparison of per-jet p_T differences to FullSim between FastSim and refined FastSim is also shown, indicating improved modeling of FastSim after refinement, particularly on the lower tail of the jet response.

5 Training and monitoring framework

A framework for launching and monitoring the NN training has been developed. This framework facilitates scaling the refinement to more complex networks and larger numbers of input variables. A schematic of the monitoring system is shown in Fig. 4. The tabulation of the grid search results is illustrated in Fig. 4 (upper), where the progress is viewable as colored bars, and real-time information about the training status is provided through links. Models trained in the grid search are shown under a ‘‘Grouped Trainings’’ tab, where the training variations are organized based on the input variables and the object class. A comprehensive summary




Figure 2: Upper left: Comparison of refined and unrefined jet p_T to FullSim. Upper right: Comparison of refined and unrefined DeepJet b discriminator to FullSim. Lower 6 plots: Correlation matrices for FullSim (left), FastSim (center), and refined FasSim (right), with the corresponding residual differences to FullSim shown below.

table, shown in Fig. 4 (middle), highlights the best-performing models based on predefined metrics, alongside detailed loss values for each model. Finally, real-time loss values of the training are monitored in Fig. 4 (lower).

This framework addresses the growing number of otherwise unconstrained choices in the network and training steps, such as input variable transformation functions, the depth and complexity of the network, and the learning rate.

6 Summary

An update to a machine learning-based approach to refine reconstructed jets from Fast Simulation in the CMS experiment is presented. A ResNet-like network, comprising pre-processing layers, a series of residual blocks, and a post-processing layer, estimates a residual





Figure 3: Upper: per-event difference in p_T^{miss} between FastSim and FullSim, inclusive (left) and profiled in bins of p_T^{miss} (right). Lower left: Distribution of p_T^{miss} for FullSim, FastSim, and refined FastSim and the ratio of each FastSim version to FullSim. Lower right: per-jet ratio of p_T between FastSim and FullSim jets, as well as those between refined FastSim and FullSim jets.

correction to augment the FastSim output variables to render them more like FullSim. The jet transverse momentum p_T has been incorporated in addition to jet flavor tagging discriminators, which were introduced in the first iteration of the study. The jet p_T correction substantially improves the modeling of the jet p_T and also improves the modeling of p_T^{miss} when propagated to that variable. A grid search mechanism for hyperparameter tuning has been developed, along with a submission and monitoring framework that tracks training progress in real time, facilitating the scaling of the refinement methods to a large number of input parameters. The refinement is applicable to standard jets over a comprehensive kinematic range, independent of the underlying physics process. The developed modules have been incorporated into CMS production workflows and are undergoing further validation. Given the computational challenges posed by the upcoming High Luminosity LHC era, regression-based refinement addresses much of the need for efficient yet accurate simulation methods.

References

- [1] CMS Collaboration, The CMS experiment at the CERN LHC, JINST **3**, S08004 (2008).



- [4] D.J. Lange, M. Hildreth, V.N. Ivantchenko, I. Osborne (CMS), Upgrades for the CMS simulation, *J. Phys. Conf. Ser.* **608**, 012056 (2015). [10.1088/1742-6596/608/1/012056](https://doi.org/10.1088/1742-6596/608/1/012056)
- [5] M. Hildreth, V.N. Ivanchenko, D.J. Lange (CMS), Upgrades for the CMS simulation, *J. Phys. Conf. Ser.* **898**, 042040 (2017). [10.1088/1742-6596/898/4/042040](https://doi.org/10.1088/1742-6596/898/4/042040)
- [6] S. Agostinelli et al., GEANT4—a simulation toolkit, *Nucl. Instrum. Meth. A* **506**, 250 (2003). [10.1016/S0168-9002\(03\)01368-8](https://doi.org/10.1016/S0168-9002(03)01368-8)
- [7] J. Allison et al., Recent developments in GEANT4, *Nucl. Instrum. Meth. A* **835**, 186 (2016). [10.1016/j.nima.2016.06.125](https://doi.org/10.1016/j.nima.2016.06.125)
- [8] S. Abdullin, P. Azzi, F. Beaudette, P. Janot, A. Perrotta, The fast simulation of the CMS detector at LHC, *J. Phys. Conf. Ser.* **331**, 032049 (2011). [10.1088/1742-6596/331/3/032049](https://doi.org/10.1088/1742-6596/331/3/032049)
- [9] A. Giammanco, The fast simulation of the CMS experiment, *J. Phys. Conf. Ser.* **513**, 022012 (2014). [10.1088/1742-6596/513/2/022012](https://doi.org/10.1088/1742-6596/513/2/022012)
- [10] S. Sekmen (CMS), Recent developments in CMS fast simulation, *PoS ICHEP2016*, 181 (2016), 1701.03850. [10.22323/1.282.0181](https://doi.org/10.22323/1.282.0181)
- [11] G. Grindhammer, M. Rudowicz, S. Peters, The Fast Simulation of Electromagnetic and Hadronic Showers, *Nucl. Instrum. Meth. A* **290**, 469 (1990). [10.1016/0168-9002\(90\)90566-O](https://doi.org/10.1016/0168-9002(90)90566-O)
- [12] M. Wolf, L.O. Stietz, P.L.S. Connor, P. Schleper, S. Bein, Fast Perfekt: Regression-based refinement of fast simulation, *SciPost Phys. Core* **8**, 021 (2025), 2410.15992. [10.21468/SciPostPhysCore.8.1.021](https://doi.org/10.21468/SciPostPhysCore.8.1.021)
- [13] S. Bein, P. Connor, K. Pedro, P. Schleper, M. Wolf (CMS), Refining fast simulation using machine learning, *EPJ Web Conf.* **295**, 09032 (2024), 2309.12919. [10.1051/epjconf/202429509032](https://doi.org/10.1051/epjconf/202429509032)
- [14] D. Alves (LHC New Physics Working Group), Simplified Models for LHC New Physics Searches, *J. Phys. G* **39**, 105005 (2012), 1105.2838. [10.1088/0954-3899/39/10/105005](https://doi.org/10.1088/0954-3899/39/10/105005)
- [15] C. Bierlich et al., A comprehensive guide to the physics and usage of PYTHIA 8.3, *SciPost Phys. Codeb.* **2022**, 8 (2022), 2203.11601. [10.21468/SciPostPhysCodeb.8](https://doi.org/10.21468/SciPostPhysCodeb.8)
- [16] M. Cacciari, G.P. Salam, G. Soyez, The anti- k_T jet clustering algorithm, *JHEP* **04**, 063 (2008), 0802.1189. [10.1088/1126-6708/2008/04/063](https://doi.org/10.1088/1126-6708/2008/04/063)
- [17] M. Cacciari, G.P. Salam, G. Soyez, FastJet user manual, *Eur. Phys. J. C* **72**, 1896 (2012), 1111.6097. [10.1140/epjc/s10052-012-1896-2](https://doi.org/10.1140/epjc/s10052-012-1896-2)
- [18] D. Bertolini, P. Harris, M. Low, N. Tran, Pileup per particle identification, *JHEP* **10**, 059 (2014), 1407.6013. [10.1007/JHEP10\(2014\)059](https://doi.org/10.1007/JHEP10(2014)059)
- [19] A.M. Sirunyan et al. (CMS), Pileup mitigation at CMS in 13 TeV data, *JINST* **15**, P09018 (2020), 2003.00503. [10.1088/1748-0221/15/09/p09018](https://doi.org/10.1088/1748-0221/15/09/p09018)
- [20] E. Bols, J. Kieseler, M. Verzetti, M. Stoye, A. Stakia, Jet flavour classification using deepjet, *JINST* **15**, P12012 (2020). [10.1088/1748-0221/15/12/P12012](https://doi.org/10.1088/1748-0221/15/12/P12012)
- [21] K. He, X. Zhang, S. Ren, J. Sun, Deep Residual Learning for Image Recognition, in *IEEE CVPR* (2016), p. 770
- [22] A.M. Sirunyan et al. (CMS), Performance of missing transverse momentum reconstruction in proton-proton collisions at $\sqrt{s} = 13$ TeV using the CMS detector, *JINST* **14**, P07004 (2019), 1903.06078. [10.1088/1748-0221/14/07/P07004](https://doi.org/10.1088/1748-0221/14/07/P07004)



License Agreement

In submitting an article to *Web of Conferences*, I grant EDP Sciences a license to publish the article, and identify EDP Sciences as the original publisher. I certify to the Publisher that:

1. I am authorized by my co-authors to enter into these arrangements.
2. I warrant, on behalf of myself and my co-authors, that:
 - the document is original, has not been formally published in any other journal, is not under consideration by any other journal. If parts from copyrighted works owned by third parties are included (included figures, tables etc.), written permission is obtained from the copyright owners for all uses as set forth in the Journal's Instructions for Authors, and credit to the sources is shown in the Manuscript. This permission is obtained for reproduction in a publication in Open Access with a CC-by 4.0 License;
 - I am/we are the sole author(s) of the article and have full authority to enter into this agreement and in granting rights to the Publisher that are not in breach of any other obligation.
 - the document contains nothing that is unlawful, libelous, or which would, if published, constitute a breach of contract or of confidence or of commitment given to secrecy;
 - I/we have taken due care to ensure the integrity of the article. To my/our – and currently accepted scientific – knowledge all statements contained in it purporting to be facts are true and any formula or instruction contained in the article will not, if followed accurately, cause any injury, illness or damage to the user.
3. I agree to the Creative Commons Attribution License <https://creativecommons.org/licenses/by/4.0/>.

Title of the conference

Conference on Computing in High Energy and Nuclear Physics 2024

Title of the article

Refining Jets for CMS Run 3 using Fast Simulation

Author(s)

Açelya Deniz Güngördü, Dorukhan Boncukçu, Kevin Pedro, Samuel Bein, Moritz Wolf, Altan Çakır and Mirac Ekim Vural

Author's signature

Dorukhan Boncukcu

Date

15/05/2025