PROCEEDINGS OF SCIENCE

Deep learning techniques for high-precision neutral meson reconstruction in the LHCf experiment

Eugenio Berti,[a,*] **Andrea Paccagnella**[a] **and Giuseppe Piparo**[b] **for the LHCf collaboration**

[a] *INFN Firenze section,*
Via G. Sansone 1, Sesto Fiorentino (FI), Italy
[b] *INFN Catania section,*
Via S. Sofia 64, Catania (CT), Italy

E-mail: giuseppe.piparo@ct.infn.it

The Large Hadron Collider forward (LHCf) experiment measures neutral particles produced in the very forward region of proton–proton collisions to constrain hadronic interaction models used in ultra high energy cosmic ray simulations. A key physics goal is the reconstruction of neutral mesons at very large pseudorapidity, in particular $K_s^0 \to \pi^0\pi^0 \to 4\gamma$ decays, whose photons are highly collimated and tend to produce overlapping electromagnetic showers in the Small Tower (TS) and Large Tower (TL) of the Arm2 detector. This topology is difficult to handle with traditional methods based on simple energy sharing and centroid estimates.

We present a multimodal deep learning strategy, generally applicable to neutral meson decays, that is specifically optimised for the identification and reconstruction of forward K_s^0 candidates. The pipeline combines calorimetric information and silicon microstrip profiles in a sequence of dedicated models for event selection, background rejection, topological classification of multi-photon patterns, and per-photon energy and position regression. The approach is developed and validated on detailed simulations of the Arm2 response in proton–proton collisions at LHC energies and is designed to provide a clean and well characterised sample of forward K_s^0 decays, improving the experimental sensitivity to neutral meson production for hadronic interaction studies in cosmic ray physics.

The European Physical Society Conference on High Energy Physics (EPS-HEP2025)
7-11 July 2025
Marseille, France

[*] Speaker

PoS(EPS-HEP2025)665

1. Introduction

Ultra-high-energy cosmic rays (UHECRs) probe extreme astrophysical environments, yet their origin and composition remain uncertain. The interpretation of air-shower data depends on hadronic-interaction models that are weakly constrained at the most forward angles and highest energies. The Large Hadron Collider forward (LHCf [1]) experiment addresses this by measuring neutral particles produced at zero degrees in LHC collisions, providing key inputs to tune these models at laboratory-equivalent energies up to $\sim 10^{17}$ eV. LHCf comprises two compact imaging calorimeters (Arm1 and Arm2) installed at the Target Neutral Absorber (TAN), where charged particles are swept away and only neutrals reach the detectors (Fig. 1, top). Each arm hosts two tungsten–scintillator sampling towers instrumented with position-sensitive sensors. In Arm2, used in this work, the Small Tower (TS) and Large Tower (TL) are read out by GSO scintillators and silicon microstrip layers, providing precise X/Y profiles and enabling measurements at very large pseudorapidities [2].

Neutral meson decays at very forward rapidities are among the most sensitive observables accessible to LHCf. Previous studies have investigated π^0 and η reconstruction with deep-learning techniques, highlighting the potential of data-driven methods for complex multi-photon topologies [3]. Here we extend this programme to K_s^0 mesons, which frequently yield four photons via $K_s^0 \to \pi^0\pi^0 \to 4\gamma$ [4]. The resulting, highly collimated electromagnetic showers often overlap within the two compact sampling towers, where traditional reconstruction based on simple energy sharing and centroiding becomes ambiguous (Fig. 1, bottom). To address this challenge, we develop a deep-learning pipeline tailored to K_s^0 candidates and exploiting the multimodal Arm2 response. The chain comprises: (i) a relevant-event selector that identifies configurations compatible with four photon hits, (ii) a neutron filter that separates four-photon final states from hadronic and other backgrounds, (iii) a topological tagger that assigns each selected event to one of the characteristic four-photon patterns across the Small Tower and Large Tower, and (iv) a photon-regression network that estimates per-photon energies and impact positions from the combined calorimetric and silicon information.

2. Deep-learning pipeline

2.1 Data and inputs

The models are trained and validated on Monte Carlo simulations that reproduce the Arm2 geometry and detector response. Proton–proton collisions at $\sqrt{s}$ = 13.6 TeV are generated with QGSJET II-04 [5], including material interactions in tungsten and the response of the scintillator and silicon sensors, perfomed using DPMJET 3.04 [6]. For each event we store the longitudinal energy deposits in the sampling layers, the silicon microstrip X/Y profiles, and a set of compact global summaries (see Table 1 and Fig. 2). Class labels 1TS + 3TL, 3TS + 1TL, 0TS + 4TL and 4TS + 0TL are assigned from the true number of photons impinging on the Small Tower (TS) and the Large Tower (TL). Only these four patterns are used, as they are the configurations with non-negligible acceptance in the Arm2 detector [7]. The simulated samples are split into training (70%) and test (30%) sets.

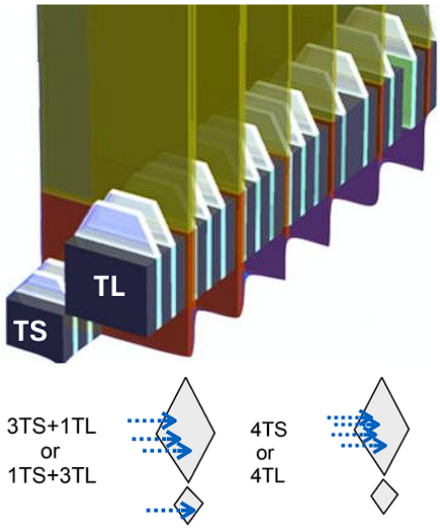

Figure 1: Top: schematic layout of the LHCf Arm2 detector with two tungsten–scintillator towers and silicon microstrips. Bottom: example of a forward multi-photon event topology mapped onto the Small Tower (TS) and Large Tower (TL).

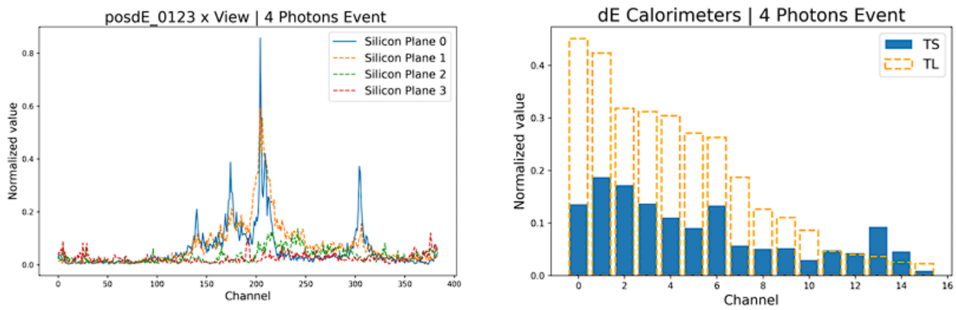

Figure 2: Example simulated event with overlapping electromagnetic showers across the Small Tower (TS) and Large Tower (TL). The inputs fed to the networks include lateral silicon microstrip patterns (left) and longitudinal calorimetric profiles (right) for both towers.

2.2 Multi-stage architecture

To deal with overlapping, highly collimated electromagnetic showers in the Small and Large Tower, we adopt a four-step workflow:

1. **Selection of relevant events (4-hit selector):** binary classifier that separates events with exactly four reconstructed hits in the two towers from all other hit multiplicities.

2. **Neutron filter (4γ vs hadronic):** binary classifier that rejects events dominated by neutral hadrons (mainly neutrons), retaining four-photon final states compatible with $K_s^0 \to \pi^0\pi^0 \to 4\gamma$ decays.

3. **Topological tagger (four-photon pattern):** multiclass classifier that assigns selected four-photon events to one of the tower-hit patterns (1TS+3TL, 3TS+1TL, 0TS+4TL, 4TS+0TL).

Input type	**Shape**	**Description**
Calorimeter signals	$(16, 2)$	Energy deposited in the 16 sampling layers for the two towers (Small Tower TS, Large Tower TL)
Silicon planes	$(384, 4) \times 2$	X/Y profiles from four microstrip layers per tower (384 channels per view)
Global sums	$(10,)$	Aggregated features (per-tower and per-plane sums)

Table 1: Input data structure used across the networks for event selection, classification and regression.

4. **Photon-reconstruction network (under development):** regression model to infer the energy and impact position of each of the four photons from the combined calorimetric and silicon information. Results from this stage will be presented in future work.

The models are applied sequentially: the 4-hit selector acts as a gate for the subsequent steps, such that the neutron filter is evaluated only on events tagged as 4-hit, and the topological tagger is evaluated only on events retained as compatible with a 4γ final state. This staged design reduces class imbalance and focuses each model on the relevant phase space for its task. In all classifiers, silicon microstrip profiles are processed with one-dimensional convolutional blocks [8], calorimeter sequences with shallow convolutional or multilayer-perceptron blocks [9], and the resulting embeddings are concatenated and passed to a fully connected head that outputs the final class scores or regression targets.

2.3 Results

We evaluate the three classification stages on held-out test sets and report precision, recall, F1 score, and support. Operating points are chosen to maximise the F1 score.[1] The consolidated metrics are reported in Table 2 and summarise the performance of each stage separately on the corresponding stage-specific test sample; the overall behaviour of the staged workflow follows from the sequential application of the steps and therefore depends on the chosen operating points and the propagation of upstream decisions. The *relevant-event selector* achieves balanced performance between positive (four-hit) and negative (not-four-hit) classes, providing a stable event preselection. The *neutron filter* prioritises high recall for the four-photon class to retain signal while keeping good precision against hadronic and other backgrounds. The *topological tagger* delivers high accuracy across the four TS/TL patterns, with mild and expected confusion between configurations that are symmetric in the Small Tower and Large Tower occupancy. The confusion matrix in Fig. 3 shows that residual misclassifications of the topological tagger are limited and largely occur between TS/TL patterns that share similar energy-sharing configurations. These results establish a robust input to the ongoing photon-regression stage for per-photon energy and impact-position estimates.

3. Discussion and perspectives

The three-stage classification chain (four-hit selection, neutron filtering, and Small/Large Tower topology tagging) achieves robust performance on realistic simulations (Table 2), enabling a clean

[1]Scores refer to the F1–maximising thresholds on the corresponding validation scans.

Model	Class	Precision	Recall	F1-score	Support
(i) Relevant: 4-hit selector (0 = not-4-hit, 1 = 4-hit)					
Relevant	0	0.95	0.91	0.93	1029
Relevant	1	0.91	0.95	0.93	1069
(ii) Discrimination: neutron filter (0 = hadronic/other, 1 = 4γ)					
Discrimination	0	0.95	0.82	0.88	606
Discrimination	1	0.79	0.94	0.86	443
(iii) Topological tagger (four TS/TL classes)					
Topological	1TS+3TL	0.92	0.86	0.89	64
Topological	3TS+1TL	0.98	0.81	0.89	117
Topological	0TS+4TL	0.93	0.97	0.95	107
Topological	4TS+0TL	0.90	0.99	0.94	176
Macro avg F1 (by model)	**Relevant ≈ 0.93; Discrimination ≈ 0.87; Topological ≈ 0.92**				

Table 2: Test-set performance of the three classifiers at the chosen operating points, reported on a per-stage basis. When applied sequentially, the staged chain provides a clean stream of $K_s^0 \to \pi^0\pi^0 \to 4\gamma$ candidates while preserving high signal efficiency; the overall end-to-end behaviour depends on the selected operating points and the propagation of upstream decisions.

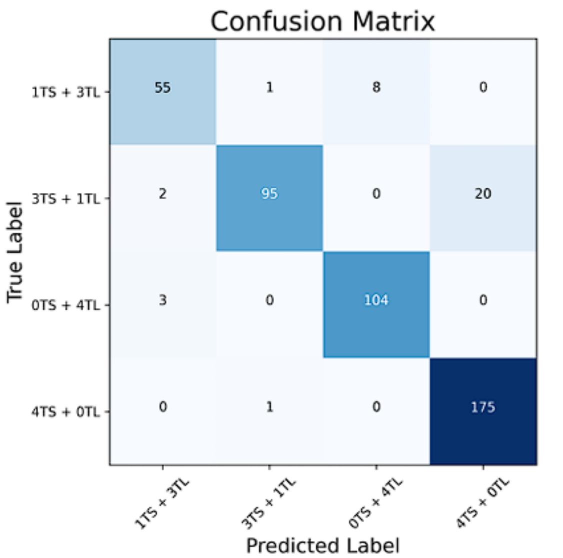

Figure 3: Confusion matrix for the four-class topological tagger, showing the agreement between true and predicted TS/TL four-photon patterns on the test sample.

stream of $K_s^0 \to \pi^0\pi^0 \to 4\gamma$ candidates for downstream reconstruction. In particular, the high F1 scores of the four-hit selector and neutron filter provide effective background suppression, while the topological tagger supplies the event geometry needed to initialise per-photon inference. The final regression stage for photon energy and impact positions is currently under development and will leverage the same multimodal inputs to reach microstrip-limited spatial precision and improved energy estimates. In forthcoming work we will (i) finalise and validate this regression network on proton–proton collisions at $\sqrt{s}$ = 13.6 TeV, (ii) extend training to proton–oxygen and proton–lead simulations, and (iii) explore domain-adaptation strategies and the inclusion of Arm1 information to widen rapidity coverage and reduce systematics. Together, these developments aim to enable precision measurements of forward K_s^0 production for hadronic-model calibration in air-shower physics.

PoS(EPS-HEP2025)665

References

[1] O. Adriani *et al.*, "The LHCf Detector at the CERN Large Hadron Collider," *Journal of Instrumentation*, vol. 3, no. 08, S08006, 2008, doi:10.1088/1748-0221/3/08/S08006.

[2] Y. Makino *et al.*, "Performance Study for the Photon Measurements of the Upgraded LHCf Calorimeters with Gd_2SiO_5 (GSO) Scintillators," *Journal of Instrumentation*, vol. 12, P03023, 2017, doi:10.1088/1748-0221/12/03/P03023.

[3] G. Piparo *et al.*, "Reconstruction of multiple calorimetric clusters in the LHCf experiment with machine learning techniques," *PoS*, vol. ICHEP2024, 1016, 2024, doi:10.22323/1.476.1016.

[4] S. Navas *et al.* (Particle Data Group), "Review of Particle Physics," *Physical Review D*, vol. 110, no. 3, 030001, 2024, doi:10.1103/PhysRevD.110.030001.

[5] S. Ostapchenko, "Monte Carlo treatment of hadronic interactions in enhanced Pomeron scheme: I. QGSJET-II model," *Physical Review D*, vol. 83, 014018, 2011, doi:10.1103/PhysRevD.83.014018.

[6] F. W. Bopp, J. Ranft, R. Engel and S. Roesler, "Antiparticle to particle production ratios in hadron-hadron and d-Au collisions in the DPMJET-III Monte Carlo model," *Physical Review C*, vol. 77, 014904, 2008, doi:10.1103/PhysRevC.77.014904.

[7] O. Adriani *et al.*, "LHCf - Technical Proposal for the LHC Run3," CERN Technical Report CERN-LHCC-2019-008; LHCC-P-014, 2019.

[8] I. Goodfellow, Y. Bengio and A. Courville, *Deep Learning*, MIT Press, 2016, ISBN:9780262035613.

[9] Y. LeCun, L. Bottou, Y. Bengio and P. Haffner, "Gradient-based learning applied to document recognition," *Proceedings of the IEEE*, vol. 86, no. 11, pp. 2278–2324, 1998, doi:10.1109/5.726791.