

Rings of Light, Speed of AI: YOLO for Cherenkov Reconstruction

Martino Borsato^{1,2}, Giovanni Laganá^{1*} and Maurizio Martinelli^{1,2†}

¹ University of Milano-Bicocca

² INFN, Sez. di Milano-Bicocca

* giovannilag66@gmail.com, † maurizio.martinelli@unimib.it



EuCAIF

*The 2nd European AI for Fundamental
Physics Conference (EuCAIFCon2025)
Cagliari, Sardinia, 16-20 June 2025*

Abstract

Cherenkov rings play a crucial role in identifying charged particles in high-energy physics (HEP) experiments. Most Cherenkov ring pattern reconstruction algorithms currently used in HEP experiments rely on a likelihood fit to the photo-detector response, which often consumes a significant portion of the computing budget for event reconstruction. We present a novel approach to Cherenkov ring reconstruction using YOLO, a computer vision algorithm capable of real-time object identification with a single pass through a neural network. We obtain a reconstruction efficiency above 95% and a pion misidentification rate below 5% across a wide momentum range for all particle species.

Copyright attribution to authors.

This work is a submission to SciPost Phys. Proc.

License information to appear upon publication.

Publication information to appear upon publication.

Received Date

Accepted Date

Published Date

Contents

1	Introduction	2
2	Dataset	2
3	YOLO Classification Pipeline	3
3.1	Training and evaluation	3
3.2	Experiments	4
3.3	Comparison with Maximum-Likelihood fit	4
4	Conclusions	5
	References	6

1 Introduction

When a charged particle traverses a medium at a velocity exceeding the phase velocity of light, it emits Cherenkov radiation [1] in the form of conical wakes of light. The size of the light cone depends directly on the mass and momentum of the particle that produced it. In particular, the Cherenkov angle θ_C is related to the particles velocity by:

$$\cos(\theta_C) = \frac{v_p}{v} = \frac{1}{n\beta}, \quad (1)$$

where n is the refractive index of the medium, and $\beta = v/c$. This relation shows that measuring θ_C provides direct access to the particle velocity, which, when combined with the momentum from tracking detectors, allows to estimate the particle's mass, and hence to identify its species.

At the LHCb experiment [2], this principle is exploited by the Ring Imaging Cherenkov (RICH) detectors, which capture the Cherenkov photons as characteristic ring patterns on the photodetector plane. LHCb employs two such detectors (RICH1 and RICH2), optimized for complementary momentum ranges, ensuring efficient particle identification (PID) across the experiments acceptance. Accurate PID is indispensable for suppressing combinatorial background and enhancing event selection, and for this reason the RICH response is integrated directly into the online trigger. However, reconstructing the rings with the traditional maximum-likelihood method is computationally intensive: in the high-level trigger (HLT2), RICH PID consumes about 20–40% of the CPU budget, making it one of the major bottlenecks for real-time operation. At Run 3 luminosities, the LHCb trigger reduces the 30 million protonproton collisions per second to an HLT2 input rate of 0.5–1.5 million events per second [3]. Since each event typically contains on the order of one hundred Cherenkov rings, the RICH PID task requires processing $\mathcal{O}(10^8)$ ring classifications every second. To address those computational challenges we propose a YOLO-based pipeline trained to directly infer particle species from ring images. Similar ML-based approaches have already been explored with promising results [4], showing the potential of deep learning for RICH PID.

2 Dataset

The training and evaluation data are produced with a dedicated synthetic event generator [5] that emulates the response of the RICH1 detector (the procedure transfers straightforwardly to RICH2). Each simulated event consists of a set of Cherenkov rings projected onto the photodetector plane.

For every event, the generator samples a random number of primary particles (between 160 and 180) from LHCb-like momentum distributions. Each particle is assigned a species (pion, kaon, proton, electron, or muon) and a momentum. From these parameters, the Cherenkov angle is computed using Eq. 1, and the corresponding ring radius R on the detector plane is

$$R = \frac{\tan(\theta_c) \cdot R_{\max}}{\tan(\theta_{c_{\max}})}, \quad (2)$$

where $R_{\max}$ is the maximum ring radius (100 mm was used to roughly emulate the optics of RICH1), and $\theta_{c_{\max}}$ is the maximum Cherenkov angle ($\beta = 1$).

Ring centers are then sampled from a distribution reproducing the geometry and acceptance of the RICH1 detector, ensuring a realistic spread across the sensor plane. The expected photon yield for each ring is proportional to $\sin^2(\theta_C)$ and is set to a maximum of 60 hits. Photon hits are radially smeared with a Gaussian resolution of 1.5 mm (0.8 mrad). No Cherenkov hit is generated for the ring centers of particles that have a speed below threshold.




Figure 1: From full RICH1 event to masked polar view. **Left:** full event (selected ring in red); **middle:** single-ring image in Cartesian coordinates with highlighted hits; **right:** annular mask and polar transform (r, θ). In middle and right, all lit pixels are set to 255 for clarity.

For the YOLO models, we generate images centered around each tracked ring (to emulate the LHCb RICH1 tracking efficiency, only $\sim 32\%$ of the rings are retained as tracked in the test set). Before rasterization, an annular mask is applied to restrict the image content to the radial region compatible with possible Cherenkov radii at the given momentum, thereby suppressing irrelevant background. The hits are then transformed into polar coordinates, which linearizes the circular rings into horizontal lines (see figure 1).

Finally, to build a balanced training dataset, generated events are used as realistic backgrounds, on top of which a controlled number of signal rings of specific momentum and particle type are added. This procedure allows us to produce class-balanced samples for each momentum bin while retaining the realism of events.

3 YOLO Classification Pipeline

YOLO (You Only Look Once) [6, 7] is a family of single-pass convolutional models designed for real-time vision tasks. In this work we use the *classification* variant (*YOLO-cl*s) as a fast per-ring classifier within a momentum-binned PID scheme.

As discussed in Sec. 1, the Cherenkov radius is a deterministic function of β and thus of momentum at fixed mass. We partition the full momentum range into bins that fulfill the following requirements: the expected radii of distinct particle species do not overlap, and for each particle type the variation of the radius within the bins momentum range is below a threshold. These conditions ensure that each image corresponds to a unique particle identification, avoiding multiple possibilities. Each bin is then treated as an independent, closed-set classification problem with its own label space. In addition, for each bin, we merge species with very close Cherenkov radius into the same class. We tested both strategies—training with merged classes in these ambiguous regions and training with all particle types kept separate—and compared their performance, finding that the grouping approach yields higher accuracy.

3.1 Training and evaluation

For each momentum bin we train an independent YOLOv11-cl network, initialized from the pretrained Ultralytics weights [6]. The loss function is the standard cross-entropy over the bins label set. To ensure balance, the training sets are constructed with a uniform distribution of particle types and momenta, consisting of 25 000 images per bin split into 80% training and 20% validation.

All models corresponding to the different bins are loaded onto the device. During infer-

ence, each ring image is dispatched to the appropriate model according to its momentum. While effective, this dispatching procedure introduces a non-negligible overhead, as we will discuss later.

The global test set is built from 850 000 rings with realistic momentum and particle-type distributions. Evaluation focuses primarily on the pion misidentification rate (pion misID) versus the kaon identification efficiency (Kaon ID), the standard benchmark in RICH PID.

3.2 Experiments

We conducted three main sets of experiments to evaluate different design choices of the YOLO pipeline and to identify the optimal configuration for RICH PID.

Grouping vs non-grouping. As discussed earlier, at high momenta the Cherenkov radii of different particle species converge to nearly identical values, making them physically indistinguishable. We therefore tested two training strategies: (i) keeping each particle type as a separate class across all bins, and (ii) merging indistinguishable types into a single class, with probabilities later split among the hypotheses at inference time. The results show that class grouping significantly improves stability and reduces misidentification rates at low and high momentum. Consequently, the grouped configuration was adopted for the following experiments.

Model size. YOLOv11 comes in multiple size variants, from *nano* to *extra-large*, trading off accuracy against computational cost. We compared the *nano*, *small*, and *medium* models on identical training and test samples. Surprisingly, the *medium* variant did not outperform the smaller models: in the momentum range 6–60 GeV the *small* and *nano* variants achieved lower pion misidentification rates and similar kaon efficiencies, while also providing faster inference. We attribute this to the increased complexity of the *medium* model, which may require more data and training to converge optimally. Given its balance of accuracy and efficiency, the *small* variant was selected as the baseline.

Image size. We also investigated the effect of input resolution, training models with images of 32×32 , 64×64 , and 128×128 pixels. Larger images capture more detail and generally improve accuracy, with the 128×128 model giving the lowest pion misID rates. However, this resulted in increased inference times and higher memory consumption. The 64×64 configuration provided a good compromise, offering nearly the same accuracy as 128×128 while maintaining fast inference comparable to the 32×32 setup.

Optimal pipeline. From these studies we conclude that the best configuration consists of the YOLOv11-small model trained on 64×64 images, with grouped particle classes in high-momentum bins.

3.3 Comparison with Maximum-Likelihood fit

Figure 2 presents a comparison between the YOLO pipeline and the maximum-likelihood (ML) fitter at a fixed operating point (TPR= 0.99, FPR= 0.01). The separation of kaons and pions as a function of the track momentum is used as a benchmark. In the central momentum region, where the Cherenkov separation between species is maximal and photon yield is high, the two approaches perform almost identically: kaon identification efficiencies remain above 95% while pion misID rates stay at or below the percent level. At lower momenta, close to the pion threshold, the ML fitter reaches its plateau more quickly and maintains lower misID than YOLO, which shows a broader transition region. At higher momenta, where Cherenkov




Figure 2: Comparison of YOLO Object Detection Performance and Likelihood Estimation.

radii saturate and inter-species separation decreases, ML again provides better accuracy, with YOLO showing a gradual loss of efficiency and a corresponding increase in misID. Overall, YOLO provides competitive performance in the momentum range most relevant for RICH1 physics. However, across the full spectrum, the likelihood-based approach achieves higher accuracy. This is because it leverages the complete hit information from the entire detector plane to infer the PID of all particles simultaneously, whereas YOLO relies only on the local subset of hits around the track under consideration.

4 Conclusions

We have presented a YOLO-based pipeline for Cherenkov ring reconstruction in the LHCb RICH detectors. Using the *optimal* configuration (Sec. 3.2), we achieved kaon identification efficiencies above 95% at a pion misID rate of 1% over a broad momentum range, with an average inference time of ~ 20 ms per full event (118 rings) on an NVIDIA A6000 GPU.

These results are very promising in terms of both accuracy and speed. The main limitation we identified is the dispatching mechanism, required by the bin-wise training of models, which introduces a non-negligible overhead during inference. To quantify the dispatching overhead, we ran inference on a single bin without it, obtaining an estimated inference time per event of ~ 4 ms, confirming that dispatching dominates the inference time.

Next steps include modifying the YOLO model so that it takes as input, in addition to the ring image, the particle momentum and the ring center position on the detector. This will make it possible to train a single unified model without binning, thereby removing the dispatching bottleneck and enabling an even faster and more scalable RICH PID pipeline. Furthermore, we plan to train the YOLO model on PID calibration datasets where one of the tracks is identified by fully reconstructing the decay it comes from. Since only a single track per event can be tagged, this procedure is not viable to tune the maximum-likelihood based model and could represent an advantage of the YOLO approach to PID. The capability of identifying the PID of a single track very quickly with YOLO by using only the photon hits around its center could be useful in order to use the PID information to refine the first trigger-level decisions.

References

- [1] P. A. Čerenkov, *Visible radiation produced by electrons moving in a medium with velocities exceeding that of light*, Phys. Rev. **52**, 378 (1937), doi:[10.1103/PhysRev.52.378](https://doi.org/10.1103/PhysRev.52.378).
- [2] R. Aaij *et al.*, *The LHCb Upgrade I*, JINST **19**(05), P05065 (2024), doi:[10.1088/1748-0221/19/05/P05065](https://doi.org/10.1088/1748-0221/19/05/P05065), [2305.10515](https://arxiv.org/abs/2305.10515).
- [3] *RTA and DPA dataflow diagrams for Run 1, Run 2, and the upgraded LHCb detector*, URL <https://cds.cern.ch/record/2730181> (2020).
- [4] M. P. Blago, *Convolutional neural networks and photonic crystals for particle identification at high energy collider experiments*, <https://cds.cern.ch/record/2791645>, Presented 11 Nov 2021 (2021).
- [5] G. Laganà, *RICHGenerator: A lightweight python library for synthesising cherenkov ring data*, <https://github.com/lhcb-bicocca/RICHGenerator> (2025).
- [6] G. Jocher, J. Qiu and A. Chaurasia, *Ultralytics yolo* (2023).
- [7] J. Redmon, S. Divvala, R. Girshick and A. Farhadi, *You only look once: Unified, real-time object detection* (2016), [1506.02640](https://arxiv.org/abs/1506.02640).