



MLOps Pipeline for Continuous Deployment of Machine Learning Algorithms for HEP

Thursday 4 September 2025 17:00 (1 hour)

We present an MLOps-based approach for managing the end-to-end lifecycle of machine learning (ML) algorithms deployed on FPGAs in real-time trigger systems, as used in experiments such as CMS and ATLAS. The primary objective of this pipeline is to enable agile and robust responses to evolving detector and beam conditions by automating the collection of new training data, retraining and optimizing models, validating performance, synthesizing firmware, and deploying updated versions to both online and offline environments. To monitor model stability over time, we incorporate dedicated data streams that bypass trigger selections (e.g., scouting or express streams). These streams allow for continuous monitoring of model outputs and the detection of distributional drifts, enabling us to assess model operational lifetimes and support strategies like continual learning, periodic retraining, or threshold adjustment to ensure consistent performance.

Our pipeline uses existing computing infrastructure, which includes distributed computing resources, container orchestration frameworks like Kubeflow, and CI/CD tools such as GitLab, to provide a scalable and maintainable foundation for real-time ML integration. This architecture supports rapid iteration cycles while promoting long-term sustainability, both of which are essential as ML becomes more central to trigger design and real-time data processing in modern collider experiments.

We invite discussion on shared challenges, solutions, and future directions for managing the ML lifecycle in low-latency HEP environments. This includes topics like model validation, deployment strategies, firmware synthesis workflows, and the role of community-developed tooling across experiments.

Authors: BROWN, Christopher Edward (CERN); XIOTIDIS, Ioannis (CERN); GLOWACKI, Maciej Mikolaj (CERN); KÖPPEL, Marius (ETH Zurich (CH))

Presenters: GLOWACKI, Maciej Mikolaj (CERN); KÖPPEL, Marius (ETH Zurich (CH))

Session Classification: Birds-of-a-Feather