

MLOps for Scientific Applications: Building Reliable, Reproducible, and Transparent AI Workflows

Wednesday, 18 March 2026 16:22 (22 minutes)

Modern scientific workflows increasingly rely on Machine Learning (ML) models whose development, deployment, and validation must meet high standards of reliability, transparency, and reproducibility. Yet many scientific ML pipelines still lack robust engineering practices, making experiments difficult to track, compare, and replicate. In this contribution, we present a structured MLOps methodology tailored for scientific data analysis, integrating data exploration and versioning, experiment tracking, and model lineage within a unified workflow.

We illustrate the methodology with a representative use case inspired by a High-Energy Physics application, showing how a fully instrumented pipeline built on MLOps principles can improve reproducibility, enable transparent comparison across models, and support effective collaboration in multidisciplinary teams.

Our results highlight how well-designed MLOps workflows can bridge the gap between experimental science and modern ML engineering, providing a scalable foundation for deploying reliable AI models in scientific applications.

Primary authors: BONACORSI, Daniele (University of Bologna); CLISSA, Luca (University of Bologna & INFN); PLINI, Leonardo (University La Sapienza)

Presenter: CLISSA, Luca (University of Bologna & INFN)

Session Classification: Artificial Intelligence (AI) - II

Track Classification: Track 10: Artificial Intelligence (AI)