

Cloud MLOps - 16 academic hours

Workshop Overview:

The Cloud MLOps workshop focuses on designing, implementing, and operating Machine Learning workflows using cloud-based platforms and services. The workshop provides a practical understanding of how cloud infrastructure enables scalable, reliable, and automated ML systems.

Throughout the workshop, participants explore MLOps capabilities across leading cloud providers, with a strong focus on AWS services. They learn how cloud-native tools support data storage, model training, deployment, automation, and monitoring, and how these services are integrated into end-to-end Machine Learning pipelines.

What Participants Gain?

- Practical skills in deploying ML models to production using batch and real-time strategies
- Ability to design and operate CI/CD pipelines for ML systems
- Hands-on experience with monitoring model performance, system metrics, and logs
- Applying testing methodologies (unit, integration, performance) in MLOps workflows
- Operating scalable ML pipelines using distributed and streaming technologies (e.g., Spark, Kafka/Flink)
- Managing end-to-end model deployment and monitoring in real production environments

Target Audience

- Data Scientists and ML Engineers working with cloud-based ML systems
- Data Engineers responsible for cloud data pipelines and ML infrastructure
- DevOps and Platform Engineers supporting AI and Machine Learning workloads
- Developers building cloud-based ML applications
- Data and technology professionals seeking to adopt cloud-native MLOps practices

Prerequisites

- Basic familiarity with Machine Learning concepts and workflows
- Experience with at least one programming language, preferably Python
- General understanding of data processing and ML pipelines
- Basic familiarity with cloud or DevOps concepts is an advantage
- Good technical English proficiency

Workshop Syllabus

- **Introduction to Cloud MLOps**
 - Overview - AWS, Azure, and GCP and MLOps services
- **AWS MLOps services**
 - Amazon SageMaker: A fully managed service for building, training, and deploying machine learning models.
 - AWS Lambda: For serverless computing, enabling automation and integration in MLOps pipelines.
 - Amazon S3: For storing and managing data, model artifacts, and pipeline artifacts.
 - Amazon CloudWatch: For monitoring and logging of your MLOps processes and infrastructure
 - Amazon ECR (Elastic Container Registry): To store and manage Docker containers for model deployment.
- **Cloud-based Machine Learning Pipelines**
 - Implementing end-to-end machine learning pipelines in the cloud
 - Integration of various cloud services for data storage, computation, and deployment