

MLOps Infrastructure - 16 academic hours

Workshop Overview:

The MLOps Infrastructure workshop focuses on the core infrastructure components required to operate Machine Learning systems in production environments. The workshop provides a practical and architectural understanding of how ML workloads are deployed, scaled, and optimized using modern infrastructure technologies.

Throughout the workshop, participants learn how containerization enables reproducible and scalable ML systems, how compute resources such as CPUs and GPUs impact model performance, and how Feature Stores support consistent and reliable feature management across training and inference workflows. The workshop emphasizes production-ready infrastructure practices used in modern MLOps environments.

What Participants Gain?

- Practical skills in building and operating production-ready ML infrastructure
- Hands-on experience with containerized ML workloads and orchestration concepts
- Ability to select and optimize CPU and GPU resources for ML performance
- Understanding how to design scalable and efficient ML infrastructure architectures
- Implementing Feature Stores to manage features across training and inference pipelines

Target Audience

- Data Scientists and ML Engineers working with production ML systems
- Data Engineers responsible for ML infrastructure and data pipelines
- DevOps and Platform Engineers involved in AI and Machine Learning platforms
- Developers building ML-enabled services and applications
- Data and technology professionals seeking a deeper understanding of MLOps infrastructure

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

- **Understanding Containers**
 - Overview of containerization platforms like Docker
 - Creating and Managing Containers
 - Container Orchestration - Kubernetes
- **Understanding GPU and CPU Computing**
 - The role of GPUs and CPUs in machine learning computations
 - The differences, advantages, and limitations of GPU and CPU computing
 - Leveraging GPU Acceleration
 - Best practices for optimizing GPU utilization
- **Introduction to Feature Stores**
 - Essential in managing features for machine learning models
 - The benefits of centralizing feature management
 - Using Feature Stores - Feast