

MLOps is the application of DevOps practices and principles to the machine learning (ML) process. In DevOps for AI and Machine Learning, DevOps engineers and testers will learn the foundational knowledge and practical skills necessary to integrate machine learning operations (MLOps) into existing AI Model development workflows.

This workshop covers the entire MLOps lifecycle, focusing on essential concepts, tools, and methodologies required to deploy and maintain machine learning models within DevOps environments.

Key takeaways from this class include:

- Gaining an in-depth understanding of MLOps principles and their integration with DevOps practices.
- Learning to set up automated data engineering pipelines.
- Tracking model experiments effectively.
- Implementing CI/CD pipelines for machine learning models.
- Establishing robust monitoring and retraining strategies.
- Navigating security and compliance considerations in MLOps.

Throughout this workshop, students will gain real-world context through practical hands-on exercises, such as setting up feature stores, implementing CI/CD processes for ML models, and deploying and monitoring models.

Who Should Attend

This workshop is ideal for DevOps engineers, software testers, and operations personnel looking to expand their skill set into MLOps. Professionals involved in software development, deployment, infrastructure management, quality assurance, or operations who wish to understand the unique challenges and best practices in deploying and maintaining machine learning models will benefit. It caters to individuals with a technical background in DevOps practices but with limited exposure to machine learning, aiming to bridge the gap between traditional DevOps workflows and the specialized requirements of MLOps.

Course Outline

Session 1: Introduction to MLOps

- Definition and importance of MLOps
- Cross-disciplinary collaboration
- Key challenges in deploying and maintaining machine learning models
- Machine Learning Process & Roles
- The MLOps lifecycle
- Exercise #1

Session 2: Automated Data Engineering

- Data engineering process and ETL/ELT transformations
- Understanding and managing data and feature sets
- Introduction to feature stores and their role in data-centric ML pipelines
- Exercise #2: Creating datasets and setting up a feature store

Session 3: Model experiment tracking

- Integration testing and model evaluation
- Model registries and best practices in deploying ML models (AB Testing, Canary)
- Exercise #4: Setup CI/CD process and integrate testing

Session 5: Monitoring, Logging, and Retraining

- Setting up monitoring systems for deployed ML models
- Scalability and auto-scaling considerations for models
- Implementation of logging and error-tracking systems
- Retraining strategies when model accuracy deteriorates
- Exercise #5: Deploy and monitor a model

Session 6: Security and Compliance

- Security aspects in MLOps: understanding the threats
- Compliance considerations such as GDPR, HIPAA
- Implementing authentication and authorization
- Creating and using AIBOMs
- Exercise #6: Setting up security measures and AIBOMs

- Model experimentation process
- Testing and validating models with datasets
- Introduction to collaborative notebooks
- Capturing experimentation information
- Deploying notebooks vs. model code
- Exercise #3: Track experiments during model development

Session 4: CI/CD for MLOps

- Introduction to CI/CD pipelines and their role in MLOps
- Managing model and data versioning
- Automating model deployment with CI/CD (models vs. code)

Exercise #7: Putting it all together

Q&A and Wrap-Up

- Summary and wrap-up of the course
- References
- Q&A session to address participant queries