

MLOps- Draft Syllabus

Fall 2025

Course Description/Overview

This course is focused on the development and deployment of Machine Learning (ML) models into a production environment. This includes operational excellence of ML systems, data governance, and ensuring the ROI of deployed models. These are skills that a MLE (Machine Learning Engineer) needs to support all stages of ML development: early data access for analysis and exploration, the ML model development process, preparing models for production, setting up pipelines for automated releases, working within a larger organization to provide models the data they need, ensuring robustness, technical performance (e.g. response time), model performance (e.g. detecting drift and recognizing when something goes wrong), monitoring, observability, and disaster response/recovery.

Learning Objectives

By the end of this course, students should be able to:

- Differentiate between core ML models, what they can be used for, what kind of data they can be used on
- Differentiate between foundational ML models and newer AI/GPT models and their capabilities
- Differentiate between the system architecture needed to support ML vs AI technologies
- Recognize where privacy and compliance concerns must be designed for in the pipeline
- Differentiate among available tools by different providers
- Build an automated ML pipeline that uses a provided ML model to re-train periodically, test, and deploy into a "production" environment that provides metrics and observability

Course Materials

Technology

Students are required to have a computer they can develop Python applications on

This course will utilize:

- A Python development environment
- Github or GitLab (TBD)
- Miscellaneous libraries and tools

Readings

Required readings will be provided by the instructor.

Potentially helpful readings:

- Implementing MLOps in the Enterprise, Haviv and Gift, 2024

Grading

Assignment weight/point distribution

- 2 Exams: 30%
- 3 Programming/Technical Assignments: 30%
- 1 Project: 40%

Course Schedule: Topical Outline

Weekly topics

- Week 1: Intro to MLOps, Basic ML Development
- Week 2: The MLOps Process
- Week 3: A First mini-MLOps Project
- Week 4: Data Governance for MLOps
- Week 5: Development and Deployment of Models and ML Apps, Part I
- Week 6: Development and Deployment of Models and ML Apps, Part II
- Week 7: Building Production Grade MLOps projects
- Week 8: Deep Learning and Large Language Model Projects (versus Foundational Models)
- Week 9: Ethics, Privacy, and Responsibility for MLOps Engineers
- Week 10: The Industry Landscape: Tools and Approaches

Assignments

- Assignment 1: My First MLOps Pipeline
- Assignment 2: Tracking ML Jobs
- Assignment 3: Observability Instrumentation
- Final Project: Building a full ML system; scope and scale TBD

Labs

- Labs will be in-class workshops/hands-on exercises

Exams

- Midterms on Week 4 and Week 8