

Introduction to Machine Learning Operations



Level: Intermediate

Duration: 6 hours

In this course we will train an example Machine Learning model using the scikit-learn library, then use open source tools to explore the steps involved in taking the model to production. Along the way we will cover model versioning using pins, model deployment using FastAPI, and good practices for monitoring model performance as the data evolves over time. We will finish with some considerations for automated reporting of model outputs and standardising your Machine Learning workflows using cloud platforms like AWS and Databricks.



Course Outline

- **Motivation:** When to take the leap from classical statistics to a Machine Learning Operations (MLOps) workflow.
- **Model Selection:** A recap of key considerations for model selection, including scoring metrics for classification and regression, and grid-search cross validation for hyperparameter tuning.
- **Versioning:** An overview of good practices for storing versioned iterations of your model, source code and data.
- **Source Code:** An example folder structure for maintaining Machine Learning projects in Python.
- **Deployment:** Illustrating the key steps for deploying a model, including building and interacting with a model API.
- **Monitoring:** Tracking model performance over time as the data evolves and considering when to retrain.
- **Automated Reporting:** Using Quarto to build report templates for automatic communication of model outputs.
- **Cloud MLOps:** A brief description and comparison of cloud platforms for MLOps, including Databricks and AWS.

Learning Outcomes

Session 1:

By the end of session 1 participants will...

- understand what distinguishes MLOps from a typical statistical workflow.
- be familiar with key terminologies in MLOps:
 - versioning
 - deployment
 - monitoring
 - retraining
- remember how to build and train models using the scikit-learn library.
- be able to perform a grid-search cross validation routine for hyperparameter tuning.
- know how to set up a sensible folder structure for model source code.
- have familiarity with the vetiver package for MLOps.

Session 2:

By the end of session 2 participants will...

- understand how to store and retrieve versioned iterations of a model using pins.
- be able to build and interact with a model API using FastAPI.
- be familiar with the steps involved in deploying a model to the cloud using docker.
- know how to track model scoring metrics over time.
- be familiar with automated reporting using quarto.
- be able to recognise the key features of the most popular cloud platforms for MLOps.

This course does not include:

- An introduction to Machine Learning fundamentals using scikit-learn, see our [Machine Learning with Python](#) course.
- Deep learning modelling techniques, see our [Python and Tensorflow](#) and [PyTorch with Python](#) courses.
- Training and running models on the cloud. We focus primarily on free and open source software in this course, deploying models to the localhost only. We touch on cloud platforms for MLOps very briefly at the end of the course.

Prior Knowledge

While we do provide a brief recap of training models in scikit-learn, it is expected that attendees will already have a good grasp of Machine Learning fundamentals and some practical experience with scikit-learn or a similar library (check out our [Machine Learning with Python](#) course). We do not expect attendees to have any prior knowledge of concepts like versioning, deployment, docker, monitoring or cloud platforms for MLOps.

Attendee Feedback

- "Interactive, very responsive. All of us who did some sort of formal data science training before mentioned that this is very practical and applied knowledge, that a large part of the course is not part of a normal data science education and really helpful for our daily work. Fantastic course, thanks Myles!"

Contact

hello@jumpingrivers.com