KU LEUVEN



# AI Lifecycle Management

## MLOps in Practice

Wednesday 2nd December - Friday 4th December 2026

VIVES - Bruges

PUC | CONTINUE

# AI Lifecycle Management | MLOps in Practice

This three-day deep-dive course provides professionals with a practical and end-to-end understanding of MLOps as a core discipline within AI lifecycle management. Participants will learn how machine learning models are developed, deployed, monitored, and continuously improved in real-world production environments. The course combines conceptual insights with hands-on experience and industry practices, bridging the gap between experimentation and scalable, reliable AI systems.

After completing this course, participants will be able to:

- Understand and apply MLOps principles across the AI lifecycle
- Deploy and manage machine learning applications using containers and CI/CD pipelines
- Track, version, and reproduce data and model experiments
- Monitor models in production and manage drift and performance issues
- Design automated, scalable MLOps workflows based on real-world practices

## Programme

### Day 1 — Foundations of DevOps and Containerisation for MLOps

The first day introduces the core DevOps principles that underpin modern MLOps environments, with a strong focus on container-based software deployment. Participants explore the advantages, limitations, and practical considerations of containerisation, using Docker as the primary container technology and Kubernetes as the leading orchestration platform. Through a hands-on DevOps lab, participants learn how to containerise applications, automate deployments using Git-based CI/CD pipelines, and understand how these tools form the technical backbone of an MLOps setup.
Note: the hands-on sessions require a machine with an x64-compatible CPU (no Snapdragon laptops or ARM-based Macs) and administrator access.
*Lecturers: Sander Borny, Tom Goethals and Bruno Volckaert (UGent)*

### Day 2 — Managing the MLOps Lifecycle: Versioning, Tracking, and Monitoring

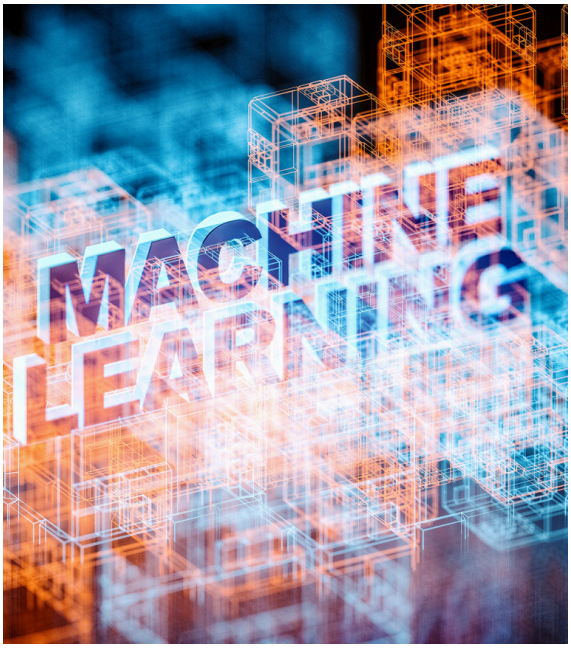
Day two starts with a concise overview of the complete MLOps lifecycle, setting the context for managing machine learning systems beyond initial deployment. The focus then shifts to version control and experiment tracking for both data and models during training, enabling reproducibility and traceability. The day concludes with monitoring machine learning applications in production, covering dashboards, data and model drift detection, and feedback mechanisms that help close the MLOps loop and support continuous improvement.
*Lecturer: Lara Luys (KU Leuven)*

### Day 3 — Model serving, Optimisation, and Automation for Production MLOps

The final day focuses on operationalising machine learning in production. Participants deploy an AI model as a containerised service on a server, expose it via an API, and connect the key building blocks around the service, such as monitoring and a lightweight front-end to create an end-to-end ML application. Next, we look at performance and scalability. You will optimise inference by exporting models to ONNX and running them with ONNX Runtime, and we discuss what professional inference serving looks like in GPU-enabled production environments using an inference server.
After lunch, a guest lecturer from Captic shares how they deploy AI-driven vision and robotics for industrial automation, with a strong focus on robust edge deployments and real-world data collection in demanding production settings, especially food-related environments.
In the afternoon, participants build an automation workflow with a data orchestration framework to make training, packaging and deployment repeatable. The day ends with an overview of ML cloud platforms and services (AWS SageMaker, Azure ML, Vertex AI) and how to position them in an MLOps architecture.
*Lecturer: Alexander D'hoore (Vives)*

## Target audience

This course is intended for professionals involved in the development, deployment, or operation of machine learning systems, including data scientists, machine learning engineers, software engineers, and DevOps engineers. It is also relevant for technical leads, architects, and AI practitioners who are responsible for bringing machine learning models into production and maintaining them over time. Participants are expected to have a basic understanding of machine learning concepts and software development practices.




## Practical

### Date and location

- Wednesday 2 December, Thursday 3 December and Friday 4 December 2026, from 9 AM until 4.30 PM.
- The training is organised at VIVES - Bruges

### Registration

- Register online before 25th November 2026. The fee for the session is 1200 euros.
- Online payment is possible, use 400/0029/00440 as reference.
- Pay by bank transfer to account number IBAN BE31 2850 2133 2955 of PUC - KU Leuven Continue, stating '400/0029/00440 + name of participant(s) and do not receive an invoice.
- If you would like an invoice, please indicate this when registering and pay AFTER receipt of the invoice.

### Partners

This programme is organised by PUC - KU Leuven Continue in cooperation with UGent, VIVES, and VAIA (Vlaamse AI Academie).

**PUC - KU Leuven Continue**

KU Leuven Kulak
E. Sabbelaan 53 bus 7643 - 8500 Kortrijk
+32 56 24 61 84 - puc@kuleuven.be
puc.kuleuven.be

By registering, I agree that the information I provide may be used to contact me in the context of this training course and for any useful follow-up.