

POST-ACADEMIC COURSE

MACHINE LEARNING INTO PRACTICE DEEP DIVE INTO MLOps

16, 17 & 18 DECEMBER 2024



MLOps



GHENT
UNIVERSITY

VIA

MACHINE LEARNING INTO PRACTICE: DEEP DIVE INTO MLOPS

16, 17 & 18 December 2024

This foundational course offers a comprehensive journey through **the various stages of deploying and maintaining machine learning models to applications using the MLOps paradigm.**

MLOps is a paradigm that aims to **deploy and maintain machine learning models in production reliably and efficiently.** The word is a compound of “machine learning” and the continuous delivery practice (CI/CD) of DevOps in the software field.

Through hands-on workshops participants will gain insights into **the core steps of MLOps: Data preparation and versioning, model deployment, monitoring, scaling, and continuous training.** They will understand the significance of having a clear understanding of **what to expect in real-world scenarios when deploying a machine-learning model.**

Additionally, this workshop covers **the specific challenges of deploying LLMs and RAG solutions.** We conclude the workshop with **techniques for downscaling models to edge devices for real-time processing.**

Throughout the course, participants will acquire **practical skills and knowledge essential for navigating ML deployment smoothly, empowering them to face various real-world challenges.**

🔗 TARGET AUDIENCE

⇒ Data scientists working at companies looking to acquire the skills needed to deploy their models to robust and scalable applications.

🔗 LECTURERS

⇒ Gilles Ballegeer (Vintec), Sander Borny (Ghent University), Cedric De Boom (Dataminded), Dimitri De Rocker (Datashift), Sam Leroux

🔗 PROGRAMME

⇒ 16/12 WORKSHOP DOCKER & KUBERNETES (4 X 1,5H)

This workshop will give you practical hands-on experience with using Docker and Kubernetes for MLOps. It will also give you the theoretical foundations to make you comfortable with relying on Docker and Kubernetes for MLOps.

⇒ 16/12 DATA LOGGING FOR INDUSTRIAL EDGE DEVICES (1H)

Having a good data logging setup is crucial for developing industrial machine vision algorithms that run on the edge.

In this sessions we present how we set up our current projects and show several use-cases to demonstrate the effectiveness.

⇒ 17/12 APPLYING DEVOPS TO MACHINE LEARNING (4H)

This workshop introduces MLOps by applying DevOps practices to machine learning. Participants will deploy, scale, and monitor a model in production, covering topics like model deployment, versioning, scaling, and concept drift. Optional subjects like A/B testing and LLMops may be included, along with recommended tools and best practices.

⇒ 17/12 AI ASSISTANTS: FROM POC TO PRODUCTION (3H)

This workshop focuses on setting up AI assistants in production, addressing challenges faced during deployment. Participants will build a RAG-based solution in the cloud and explore tools and frameworks for traceability and automated quality monitoring.

⇒ 18/12 ML MODEL OPTIMIZATIONS FOR EFFICIENT EDGE AI DEPLOYMENT (1H)

This lesson covers optimizing machine learning models for resource-constrained edge devices. It explores ways to reduce computational cost and memory footprint, enabling image recognition and natural language processing models to run on mobile and embedded devices.

⇒ 18/12 FINE-TUNING SMALL LLMs FOR PII DETECTION AND SECURE EDGE DEPLOYMENT (3H)

Learn to fine-tune small language models for PII detection with a focus on privacy and secure edge deployment. You'll generate synthetic data, perform on-device fine-tuning, and implement encryption for secure deployment. By the end, you'll build a privacy-focused AI app that operates offline on edge devices.

⇒ 18/12 INTRODUCTION TO FEDERATED LEARNING WITH PYTHON & FLOWER (2H)

This session introduces Federated Learning, a technique for training models without centralizing data. It covers key principles, available frameworks, and challenges. In the hands-on portion, participants will use the Flower framework and Python to apply the concepts in practice.

🔗 PRACTICAL

⇒ **Fee:** 16/12 -> € 465,- | 17/12 -> € 465,- | 18/12 -> € 400,- | complete -> € 1.200,- (Reduction for multiple subscriptions)

⇒ **Location:** Ghent University, classroom 0.2, building 60, Technologiepark Zwijnaarde. *With the support of VAIA.*