



COURSE DESCRIPTION CARD - SYLLABUS

Course name

Machine Learning Operations (MLOps) [S2SI1E>NUM]

Course

Field of study

Artificial Intelligence

Year/Semester

1/2

Area of study (specialization)

–

Profile of study

general academic

Level of study

second-cycle

Course offered in

English

Form of study

full-time

Requirements

elective

Number of hours

Lecture

15

Laboratory classes

15

Other

0

Tutorials

0

Projects/seminars

0

Number of credit points

3,00

Coordinators

dr hab. inż. Mikołaj Morzy prof. PP
mikolaj.morzy@put.poznan.pl

Lecturers

Prerequisites

Student partaking this course should have basic knowledge in the areas of machine learning and data processing. Student should be able to solve basic problems related to information systems' design and implementation (versioning, testing, code integration). Basic command of Python programming language is required. Student should be able to utilize external APIs and integrate external resources. Student should also understand the need to broaden one's competences and be able to effectively cooperate within a project team. Regarding social competences student should exhibit qualities such as responsibility, endurance, intellectual curiosity, creativity, respect for other people, ability to work in a group.

Course objective

The main objective of the course is to present a rich set of tools required to apply machine learning solutions to information systems in practice. During lectures students learn techniques and tools for statistical model productization, in particular, tools allowing for a seamless integration of machine learning models with the existing computing infrastructure or ecosystem. During laboratories students experiment with the tools using various scenarios. The list of topics covered by the course includes: - versioning of data and statistical models, - workflow management tools for machine learning, - tools for data annotation for machine learning models' training, - monitoring of the learning process, - practical aspects of machine learning models' productization, - tools for managing of machine learning projects.

Course-related learning outcomes

Knowledge Student has advanced and in-depth knowledge of machine learning and of practical aspects of productization of machine learning solutions [K2st_W1].

Student has knowledge of good practices related to the development and practical application of machine learning solutions in information systems, in particular, good practices related to the testing and verification of statistical models [K2st_W2].

Student has detailed knowledge of the collection, annotation, and versioning of data used for machine learning models' training [K2st_W3].

Student is aware of the currently available tools, libraries, and frameworks for development, integration, and maintenance of information systems utilizing machine learning solutions [K2st_W4].

Student understands the full life cycle of information systems utilizing machine learning solutions, is able to evaluate the correctness of the life cycle and identify non-trivial dependencies between subsequent steps of the life cycle (data collection and annotation, training of statistical models, validation and optimization of statistical models, monitoring of statistical models) [K2st_W5]

Student knows basic methods, techniques and tools used to design complex information systems and is able to apply these methods, techniques and tools to the specific case of designing systems utilizing machine learning solutions [K2st_W6].

Skills Student knows how to use various APIs and documentation of complex information systems [K2st_U1].

During the design of an information system utilizing machine learning solutions the student is able to plan and carry out measurement experiments in accordance with scientific method and to interpret the results of conducted experiments [K2st_U3].

Student knows how to use visualization tools to monitor the process of machine learning model training [K2st_U4].

Student is able to use agile programming methodologies to manage a project containing statistical models at its core. Student understands the need to include domain expertise and knowledge in the process of designing machine learning based solutions [K2st_U5].

Student is able to construct information systems with machine learning components using the containerization technique which allows a seamless integration of the existing workflow with new tools and solutions [K2st_U6].

Student is able to evaluate the time required to complete development steps of an information system utilizing statistical models [K2st_U7].

Student understands the nature of the technological debt incurred by machine learning and is able to use available tools to mitigate the negative impacts that potential errors in the statistical model training process can project on the remaining part of an information system [K2st_U8].

During the design of an information system utilizing machine learning solutions the student is able to analyze available data in search for latent biases and correlations. The student is able to analyze the life cycle of data and spot threats to the integrity of the model training process [K2st_U9].

Student is able to design and develop a solution to a given economic, technological, or social problem using machine learning tools [K2st_U10].

Social competencies Student understands the exceptional dynamics of the area of machine learning and is aware of the existence of multiple tools [K2st_K1].

Student understands the need for life-long learning in the area of machine learning tools due to a rapid cycle of tool exchange [K2st_K2].

Student can effectively communicate in a group and collaborate within the ramifications of agile programming techniques [K2st_K3].

Methods for verifying learning outcomes and assessment criteria

Learning outcomes presented above are verified as follows:

The grading method for the course utilizes continuous assessment based on five mini-projects (graded binarily: done/not done), through which students deepen their knowledge in selected areas of MLOps. Students are given full technological freedom in each area, such as the choice of cloud provider (e.g., AWS vs. GCP), and are strongly encouraged to integrate these mini-projects into their master's theses or larger projects within other courses. The final grade for the entire course is common for both the lecture and laboratory parts and is determined based on the successful completion of these mini-projects.

Programme content

General ramifications of machine learning tools are presented during lectures. Laboratories present detailed case studies of tool usage at a given stage of the development of a machine learning-based system.

Topics covered during lectures include:

- technical debt and its consequences
- ML workflows
- universal platforms for ML
- preparation, deployment and monitoring of ML solutions
- best & worst practices of ML operations
- LLM operations

Topics covered during laboratories include:

- brief recap of Pytorch
- Pytorch Lightning; model monitoring; hyper-parameter optimization
- model serving and deployment & dockerization
- model deployment in cloud environments
- cloud serverless deployment
- inference optimization
- LLMs & RAGs

Course topics

The course introduces practical and theoretical aspects of modern machine learning system development, with a focus on the integration of contemporary tools and frameworks. Lectures provide an overview of the general implications and methodologies of machine learning tools, while laboratory sessions involve case studies illustrating their application in various development stages.

The lectures introduce the theoretical and strategic aspects of building and managing modern machine learning (ML) systems in production environments, focusing on methodologies for effective and scalable integration of ML models with broader IT systems. Key topics begin with the concept of technical debt specific to ML, covering the hidden costs and strategies for minimizing the maintenance burden associated with complex systems. We will explore ML workflows, detailing the model lifecycle from data collection to continuous monitoring, emphasizing CI/CD/CT and workflow orchestration tools. Practical aspects of preparation, deployment, and monitoring are covered, including safe deployment strategies (Canary, A/B testing) and the differentiation between operational and ML-specific monitoring metrics like data and model drift. Finally, we discuss the best and worst practices of ML operations, emphasizing reproducibility and Responsible AI (XAI, bias minimization), and conclude with an overview of LLM operations (LLMOps), covering the unique challenges of scaling, monitoring, and adapting Large Language Models in production via techniques like fine-tuning and Prompt Engineering.

The laboratory part of the course covers the recap of fundamentals of PyTorch and its structured extension, PyTorch Lightning, along with methods for model monitoring and hyperparameter optimization. It addresses model serving and deployment, including containerization with Docker and deployment on cloud infrastructure such as AWS EC2 and AWS Lambda. Further topics include inference optimization, the use of large language models and retrieval-augmented generation (RAG)

Teaching methods

Lecture: multimedia presentation, seminar with students' presentations, information retrieval

Laboratory: programming examples, small individual programming exercises

Project: practical problem solution, group work, design thinking, documentation

Bibliography

Basic

1. Treveil, Mark, Nicolas Omont, Clément Stenac, Kenji Lefevre, Du Phan, Joachim Zentici, Adrien Lavoillotte, Makoto Miyazaki, and Lynn Heidmann. *Introducing MLOps*. O'Reilly Media, 2020.
2. Gift, N., & Deza, A. (2021). *Practical MLOps*. O'Reilly Media, Inc.
3. Géron, Aurélien. *Hands-on machine learning with Scikit-Learn, Keras, and TensorFlow: Concepts, tools, and techniques to build intelligent systems*. O'Reilly Media, 2019.
4. Raschka, Sebastian, and Vahid Mirjalili. *Python machine learning*. No. 1. Packt Publishing,, 2019.

Additional

1. Pustejovsky, James, and Amber Stubbs. *Natural Language Annotation for Machine Learning: A guide to corpus-building for applications*. O'Reilly Media, Inc. 2021.
2. Zheng, Alice, and Amanda Casari. *Feature engineering for machine learning: principles and techniques for data scientists*. O'Reilly Media, Inc., 2018.

Breakdown of average student's workload

	Hours	ECTS
Total workload	75	3,00
Classes requiring direct contact with the teacher	30	1,50
Student's own work (literature studies, preparation for laboratory classes/ tutorials, preparation for tests/exam, project preparation)	45	1,50