

Content

Course Code	Course Name	Semester	Theory	Practice	Lab	Credit	ECTS
VM 536	Applications of Data Science	3	0	4	0	3	8
Prerequisites							
Admission Requirements							
Language of Instruction	English						
Course Type	Compulsory						
Course Level	Masters Degree						
Objective	The aim of the course is to teach how to manage, with engineering discipline, the whole life cycle of machine learning models from development to production, and to introduce the construction of systems based on large language models.						
Content	The course covers the engineering practices needed to take data science and machine learning models into production: version control, testing of code and data, feature engineering and selection, MLOps principles and experiment management, pipeline automation with continuous integration, model deployment through APIs and containers, model monitoring in production and continuous delivery. The final part is devoted to large language models: foundations, prompt engineering, advanced techniques, and LLM-based system architectures (LLMOps and agents).						
References	Course page: https://gitlab.com/onayg/vm536 Foundations of Data Science: Avrim Blum, John Hopcroft, and Ravindran Kannan						

Theory Topics

Week	Weekly Contents
1	Git tooling and version control
2	Testing strategies for code and data
3	Feature engineering and selection
4	MLOps principles and experiment management
5	Pipeline automation with continuous integration (CI)
6	Model deployment strategies: APIs and containers
7	Midterm exam
8	Model monitoring in production and continuous delivery (CD)
9	Foundations of large language models (LLM) and prompt engineering
10	Advanced techniques for LLMs
11	LLM-based system architectures: LLMOps and agents