Machine Learning Operations

1. **Course number and name:** 020MLOES5 Machine Learning Operations

2. **Credits and contact hours:** 4 ECTS credits, 2x1:15 contact hours

3. **Name of course coordinator:**

4. **Instructional materials:**

   References:
   - "Building Machine Learning Powered Applications: Going from Idea to Product" by Emmanuel Ameisen
   - "Machine Learning Engineering" by Andriy Burkov

5. **Specific course information**
   a. **Catalog description:**
   This course offers a comprehensive exploration of software engineering principles specifically adapted for artificial intelligence (AI) applications. It covers the full software development lifecycle (SDLC) of AI systems, including requirements engineering, design patterns for machine learning workflows, and software architecture for intelligent systems. Emphasis is placed on modern machine learning operations (MLOps) practices, such as automated training and deployment pipelines, model monitoring and performance evaluation, model versioning, and lifecycle management. The course also addresses responsible AI development, focusing on fairness, bias mitigation, and explainability, equipping students with the tools and methodologies needed to build robust, scalable, and ethical AI-powered software solutions.

   b. **Prerequisites:** 020MLRES4 Machine Learning

   c. **Required** for CCE Artificial Intelligence Option students; **Selected Elective** for students in the CCE Software Engineering and Telecommunication Networks Options.

6. **Educational objectives for the course**
   a. **Specific Course Learning Outcomes:**
      - Describe the software development lifecycle (SDLC) for AI applications, including key differences from traditional software engineering.
      - Apply design patterns and architectural principles specific to machine learning systems.
      - Implement end-to-end machine learning workflows, from data ingestion and preprocessing to model training and deployment.
      - Develop and manage CI/CD pipelines tailored for AI models and data workflows.

- Monitor deployed machine learning models, detect data/model drift, and evaluate performance over time.
- Use version control tools for managing machine learning models, datasets, and experiments.
- Analyze and mitigate ethical risks related to bias, fairness, and explainability in AI systems.
- Integrate responsible AI practices into the development, deployment, and maintenance of ML-enabled software products.

b. PI addressed by the course:

PI	1.1	1.2	1.3	2.1	2.2	2.3	2.4	2.5	4.1	4.2	7.1
Covered	x	x	x	x	x	x	x	x	x	x	x
Assessed	x	x	x	x	x	x	x	x			

7. Brief list of topics to be covered

- Introduction to AI Software Engineering and the MLOps Lifecycle (2 lectures)
- Requirements Engineering and System Design for AI Applications (2 lectures)
- Design Patterns and Architecture for ML Systems (2 lectures)
- Data Engineering: Ingestion, Validation, and Feature Pipelines (2 lectures)
- Model Development and Experiment Tracking (2 lectures)
- Continuous Integration and Continuous Deployment (CI/CD) for ML (2 lectures)
- Model Packaging, Containerization, and Infrastructure as Code (2 lectures)
- Monitoring, Logging, and Model Drift Detection (2 lectures)
- Model Versioning and Lifecycle Management (2 lectures)
- Explainable AI and Model Interpretability Techniques (2 lectures)
- Security, Privacy, and Governance in MLOps (2 lectures)
- Responsible AI and Ethics: Bias, Fairness, and Accountability (2 lectures)