

Deep Learning & GenAI - 52 academic hours

Course Overview:

This course provides an in-depth and practical introduction to modern deep learning and generative AI technologies. Participants will explore the foundations of neural networks and deep learning architectures and gain hands-on experience implementing models using PyTorch. The course covers the core building blocks of generative AI, including transformer and diffusion architectures, embeddings, and the internal mechanics of large language models (LLMs). In the practical part of the course, participants will design, build, and evaluate real-world LLM-based applications, including retrieval-augmented generation (RAG), agentic AI systems, workflow automation, and fine-tuning strategies, with a strong emphasis on best practices and end-to-end solution design.

What Participants Gain?

- Designing and implementing deep learning models using PyTorch
- Applying modern GenAI architectures such as Transformers, Diffusion models, and embeddings
- Working effectively with Large Language Models using Hugging Face
- Building and evaluating LLM-based applications, including RAG and Agentic Workflows
- Applying prompt engineering and fine-tuning techniques in real-world scenarios
- Designing end-to-end, production-ready generative AI solutions

Target Audience

- Data Scientists
- Machine Learning Engineers
- AI Engineers
- Experienced Data Analysts
- Software Engineers involved in AI-driven products
- Technical professionals transitioning into advanced AI roles

Prerequisites

- Prior experience with Python programming
- Familiarity with data analysis concepts
- Basic understanding of machine learning fundamentals (recommended)
- Background in linear algebra, probability, or statistics is an advantage but not mandatory

Workshop Syllabus

- Deep Learning
 - Introduction to Neural Networks
 - Deep Learning with Pytorch
 - Deep Learning Architectures
- Generative AI Fundamentals
 - Transformers + Diffusion architectures + Embeddings
 - HuggingFace
 - Behind the Scenes of an LLM
- Practical LLMs
 - LLM-based Applications Overview
 - Prompt engineering
 - LangeChain + LlamaIndex frameworks
 - Use Case: RAG
 - Agentic AI
 - Fine tuning (why, approaches, tools)
 - Workflow automation tools
 - Evaluation