

Deep Learning

[View PDF](#)

Instructor(s):

Luca Szegletes

Short description of the course:

Artificial Intelligence (AI) has emerged as a key research area over the past decade. Among its various subfields, deep learning stands out for achieving state-of-the-art results across numerous application domains, including image recognition, speech recognition and synthesis, natural language processing, and reinforcement learning. Deep learning models are capable of representation learning and modeling jointly, thus, other machine learning methods (which require feature engineering) can be outperformed in case a large amount of data is available.

This course provides a comprehensive introduction to deep learning, combining foundational theory with practical experience through hands-on sessions, and project work.

The course is divided into three main sections:

- I. **Deep Learning Fundamentals:** Introduction to artificial neural networks and basic architectures.
- II. **AI Agents:** Exploration of natural language processing basics, transformers, and GPT models.
- III. **Generative AI:** Analysis of the capabilities and limitations of generative AI models.

At the end of the course, we will address various practical issues related to the development of such models.

Aim of the course:

The objective of the course is to teach fundamental knowledge of the most important methods in deep learning and to develop practical knowledge, to prepare students to create their own complex deep learning solutions.

Prerequisites:

This is an introductory level deep learning course. The following background is required:

- Basic programming skills (e.g. able to write a small game with a simple interface)
- Basic matrix algebra knowledge (addition, subtraction, multiplication, transpose, L1, L2 norm, inverse, linear independence)
- Basic probability theory knowledge (mean and variance calculation, probability distributions)
- Optional, but useful: basic knowledge of signal processing (Fourier transform, filters, multiresolution analysis)

Grading:

80% Project work: Project work can be done independently or in a group of 2 or 3 students. The instructor will work with the students to ensure that the scope of the project is appropriate for the number of students involved. Group projects require documentation of each member's contributions as peer-assessment of each individual's contributions to the project. The projects' progress and final results must be presented over the course of the semester.

- Project registration: 10%
- Two project milestones: 30% (15% each)
- Final project submission: 30%
- Code documentation: 10%

- Poster presentation: 20%

Syllabus:

Week	Topic	Assignment
1	Course introduction. Requirements. Examples of deep learning and end-of-term projects. Fundamentals of statistical learning, machine learning, and deep learning. Overview of supervised, unsupervised, and reinforcement learning. Complete machine learning pipeline	
2	Introduction to artificial neural networks (ANN). Role of geometry (structure of data: image, sequential, text). Tensor operations and deep learning frameworks. Project Meeting 1.	
3	Introduction to image data and convolutional neural networks (CNNs). Deep learning models for computer vision.	Project registration.
4	Introduction to sequential data, 1D CNNs, and recurrent neural networks (RNNs).	
5	Basics of natural language processing (NLP). Project Meeting 2.	
6	Introduction to Sequence2sequence and word2vec models.	Project milestone 1.
7	Understanding attention mechanism and transformer models.	
8	AI agents, GPT models, and retrieval-augmented generation (RAG)	
9	Introduction to autoencoders (AE). Project Meeting 3.	
10	Basics of generative AI, including Variational Autoencoders (VAE) and Generative Adversarial Networks (GAN)	Project milestone 2.
11	Understanding diffusion models.	

- 12 Project Meeting 4.
- 13 MLOps: production-ready deep learning, explainability (XAI), and safe AI practices. Project final submission.
- 14 Poster session.

Example of project work:

Students will have considerable autonomy in selecting their own course project. Here, we offer an example of one possible project and the steps that would be involved in this project.

Create a deep learning model that predicts the temperature in Budapest

First, the students will find a suitable data source; many public weather websites can be used. Next, the data should be downloaded, cleaned, and preprocessed. Once data preprocessing is complete, a deep learning model should be built, trained, and evaluated. In addition to the regular project updates, the final deliverables for this project comprise well-commented and documented source code and a poster that describes the rationale for design decisions and an analysis of the predictive power of the model.

Textbooks:

Goodfellow, Ian, Yoshua Bengio, and Aaron Courville. Deep learning. MIT press, 2016., online: <https://www.deeplearningbook.org/>

Franchois Chollet, Deep Learning with Python, Manning Publications, 384 pages, 2017, <https://www.manning.com/books/deep-learning-with-python>

Instructors' bio:

Luca Szegletes is an assistant professor at the Budapest University of Technology and Economics, Department of Automation and Applied Informatics. She received her Ph.D. at BME in 2016. She has a strong background in deep learning and advanced signal processing with an interest in interdisciplinary research. As a visiting scholar, she carried out research at EPFL (Switzerland), Stanford (US), and Araya (Japan). She received a Fulbright scholarship, Google's Anita Borg award, and a JPSP fellowship in the past. Her primary research area is geometric deep learning and generative models. She was a co-author of several scientific papers related to deep learning.