

Deep Learning

[View PDF](#)

Instructor(s):

Bálint Gyires-Tóth

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, assignments, and project work. You will explore the essential components of deep neural networks, optimization techniques for model training, and advanced architectures tailored for general classification and regression problems, computer vision, natural language processing, anomaly detection, etc. Additionally, the course covers the deep learning hardware and software ecosystem, and real-world applications throughout the complete course.

Aim of the course:

The objective of the course is to teach a fundamental knowledge of the most important methods in deep learning and to develop a practical knowledge, which helps the students to create complex deep learning solutions (e.g. as a researcher, an engineer, or even an AI startup founder).

Prerequisites:

This is a beginner level deep learning class, basic math and computer science knowledge are required:

- Basic programming skills (e.g. able to write a small game with a simple interface, like invaders)
- Basic matrix algebra knowledge (addition, subtraction, multiplication, transpose, L1, L2 norm, inverse, linear independence)
- Basic probability theory knowledge (mean and variance calculation, probability distributions)

Grading:

50% Assignments: there are deep learning related assignments. A successful proficiency assignment is required for this class. The assignments have to be done individually. Assignments may vary between in-class and homework, determined by the instructor.

40% Project work: students in groups of three have to develop a complete deep learning solution with source code and documentation.

10% Presentations: the project's progress and final results must be presented over the course of the semester.

Syllabus:

Week Topic

Assignments

1	Fundamentals of machine and deep learning. Supervised, unsupervised and reinforcement learning. A complete machine learning pipeline. Deep learning software stack. Python basics.	
2	Training deep neural networks, the backpropagation algorithm	Assessment 1.
3	Fundamentals of Deep Learning part 1/2	Project work registration
4	Fundamentals of Deep Learning part 2/2	Assignment 2.
5	Data preparation, weight initialization, regularization in deep learning. High level deep learning frameworks with GPU acceleration.	Assignment 3.
6	Advanced optimizers, activation functions. Regression and classification. Cost functions. Advanced architectures. Residual, highway, dense networks. Skip connections.	Project work milestone 1.
7	Hyperparameter optimization. Pattern recognition with convolutional neural networks (CNN). 2D CNN for computer vision.	Assignment 4.
8	Advanced deep learning models for computer vision. Training neural networks with few data. Data augmentation. Transfer learning.	Assignment 5.
9	Sequential data. 1D CNN for pattern recognition in sequential data.	Assignment 6.
10	Recurrent neural networks (RNNs) and LSTM (Long Short-Term Memory). Natural Language Processing basics.	Project work milestone 2.
11	Building Transformer-Based Natural Language Processing Applications part 1/2	Assignment 7.
12	Building Transformer-Based Natural Language Processing Applications part 2/2	Assignment 8.
13	Autoencoders. Example application(s) of autoencoders (e.g. clustering, anomaly detection, image generation).	Project work final submission.

Example of assessment:

Create a deep learning model that predicts the temperature in Budapest

First, the students have to find a proper data source. There are many public weather homepages, that can be used. Next, the data is downloaded, cleansed and preprocessed. When the data preprocessing is ready, a deep learning model is built, trained and evaluated. The source codes with detailed comments and the output of the scripts are submitted until the deadline.

Example of project work:

Create a deep learning model that predicts the temperature in Budapest

First, the students have to find a proper data source. There are many public weather homepages, that can be used. Next, the data is downloaded, cleansed and preprocessed. When the data preprocessing is ready, a deep learning model is built, trained and evaluated. The source codes with detailed comments and the output of the scripts are submitted until the deadline.

Example of project work

Design and implement a deep learning solution for webcam based real-time sex and age prediction

First, the data collection is performed: public, easy-to-access datasets with sex and age labels are used for training. The chosen dataset and the source code of preprocessing are submitted until the milestone 1 deadline. Next, to be able to detect faces on images and to make sex and age predictions a pretrained deep neural network is optimized. For milestone 2 the initial approach of the solution is submitted with source code. Last, the video feed of a webcam is streamed to the deep neural network to make real-time sex and age prediction. Furthermore, enhancements are made to increase the speed and accuracy of the deep learning models, and evaluation is carried out. At the end of the semester, the source codes and a 2-4 pages long documentation are submitted, furthermore, an 8 minute long presentation is held in the exams *period*.

The possible topics for project work are announced in the first week.

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>

Instructor's bio

Bálint Gyires-Tóth conducts research on fundamental and applied machine learning since 2007. With his leadership, the first Hungarian hidden Markov-model based Text-To-Speech (TTS) system was introduced in 2008. He obtained his PhD degree from the Budapest University of Technology and Economics with summa cum laude in January 2014. Since then, his primary research field is deep learning. His main research interests are sequential data modeling with deep learning and deep reinforcement learning. He also participates in applied deep learning projects, like time series classification and forecast, image and audio classification and natural language processing. He was involved in various successful research and industrial projects. In 2017 he was certified as NVidia Deep Learning Institute (DLI) Instructor and University Ambassador.

Instructors' bio:

Bálint Gyires-Tóth conducts research on fundamental and applied machine learning since 2007. With his leadership, the first Hungarian hidden Markov-model based Text-To-Speech (TTS) system was introduced in 2008. He obtained his PhD degree from the Budapest University of Technology and Economics with summa cum laude in January 2014. Since then, his primary research field is deep learning. His main research interests are sequential data modeling with deep learning and deep reinforcement learning. He also participates in applied deep learning projects, like time series classification and forecast, image and audio classification and natural language processing. He was involved in various successful research and industrial projects. In 2017 he was certified as NVidia Deep Learning Institute (DLI) Instructor and University Ambassador.