

Deep Learning

[View PDF](#)

Instructor(s):

Bálint Gyires-Tóth

Short description of the course:

Artificial Intelligence (AI) has become a focused research area in the past decade. Among its subfields, deep learning is one of the most important due to the state-of-the-art results achieved in several application scenarios, e.g. 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 teaches not only the fundamentals of deep learning but also helps to gain practical knowledge through hands-on classes, assignments and project work. The course covers the most important building blocks of deep neural networks, the optimization methods that are used for training, advanced architectures for computer vision, speech technologies, natural language processing and self-supervised learning, the deep learning hardware and software stack, and we discuss the possible 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, a programmer, or even a 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 (matrix operations and partial derivative)
- Basic probability theory knowledge (mean and variance calculation, probability distributions)

Grading:

30% Assignments: there is a proficiency assignment of basic programming skills and two deep learning related assignments. A successful proficiency assignment is required for this class. The assignments have to be done individually.

50% Project work: students in groups of two have to develop a complete deep learning solution with source code and documentation.

20% Presentation: the project work has to be presented at the end of the semester.

Syllabus:

Week Lecture

Practice

**Assignments and
deadlines**

1	Fundamentals of machine and deep learning. Supervised, unsupervised and reinforcement learning. A complete machine learning pipeline.	Deep learning software stack. Python basics.	Proficiency assignment (Python code)
2	How to train deep neural networks? The backpropagation algorithm.	Numpy based backpropagation implementation with a toy example.	Deadline of proficiency assignment
3	Regularization methods. Early stopping, L1, L2, dropout, batch normalization.	Introducing regularization techniques to backpropagation.	
4	Advanced activation functions. Regression and classification. Cost functions.	High-level deep learning frameworks with GPU acceleration. Regression and classification toy examples.	Deadline for project work milestone 1.
5	Fully-connected neural networks.	Solving a regression problem with a fully-connected neural network.	Deep Learning Assignment 1.
6	Pattern recognition with convolutional neural networks (CNN). 2D CNN for computer vision.	Image classification with simple 2D CNNs.	
7	Advanced deep learning models for computer vision. Training neural networks with few data. Data augmentation. Transfer learning.	Inference with pretrained CNN for image classification and transfer learning.	Deadline for Assignment 1.
8	1D CNN for pattern recognition in sequential data.	Constructing 1D CNN for natural language processing (NLP) task.	
9	Sequential data modeling – methods and challenges. Recurrent neural networks (RNN), backpropagation through time.	Low-level implementation of RNN. Applying RNN for NLP task.	Deadline for project work milestone 2.
10	Vanishing and exploding gradients. Long Short-Term Memory (LSTM).	Low-level implementation of LSTM. GPU optimized LSTM implementation. Handling inputs with different sizes.	Deep Learning Assignment 2.

11	Advanced architectures. Residual, highway, dense networks. Skip connections.	Constructing and displaying advanced deep neural networks.	
12	Hyperparameter optimization. Grid, random, Bayesian and TPE search.	Hyperparameter search for an arbitrary deep learning model.	Deadline for Assignment 2.
13	Generative models, self-supervised learning 1: Auto encoders.	Example application(s) of auto encoders (e.g. clustering, anomaly detection, image generation).	
14	Generative models, self-supervised learning 2: Generative Adversarial Networks (GANs). Final thoughts.	Example application(s) of GANs (e.g. image generation, anomaly detection).	Deadline for project work submission.

The theoretical parts are introduced in the lectures, while the corresponding hands-on practices take place afterwards, before the next lecture.

Example of assignment:

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 minutes long presentation is held in the exams period.

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:

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.