

Deep Learning with Python - from tabular to multimedia

(taught by Dr. Christian Kauth)

Deep learning with neural networks is a fascinating field, especially on non-tabular data (like images and text). The mixture of faster hardware, new techniques, highly optimized open source libraries and large datasets allow very large networks to be created with frightening ease. Deep neural networks have repeatedly proven impressively skillful on a range of problems.

This course is a guide to deep learning in *Python*. You will discover the *Keras* Python library for deep learning and how to use it to develop and evaluate deep learning models. You will discover the techniques and develop the skills in deep learning that you can then bring to your own machine learning projects.

After familiarizing with *Keras*, we will illustrate the skill of deep learning on some well-understood case study machine learning problems from the UCI Machine learning repository (<http://archive.ics.uci.edu/ml/index.php>) and compare the performance to the classical machine learning approaches used in the course “Machine Learning with Python – from Zero to Hero”. Next we introduce convolutional layers to the networks and use them to classify handwritten digits (e.g. MNIST dataset <http://yann.lecun.com/exdb/mnist/>) and real-world objects (e.g. CIFAR-10 <https://www.cs.toronto.edu/~kriz/cifar.html>).

We then shift our focus from images to text data, and learn how to solve common Natural Language Processing (NLP) tasks such as text classification, token classification, summarization and translation. To that purpose, we’ll use state-of-the-art transformer-based pre-trained models and see how to fine-tune them on our own dataset (transfer-learning).

Lastly you’ll learn how to integrate the temporal dimension into your machine learning projects to make forecasts on time-series. We’ll compare the performance of several techniques, reaching from tabular representation to state-of-the-art causal transformer architectures.

Fun fact: Deep learning is taking AI performance in compute vision, natural language and time series analysis from deceptive to disruptive, and the Attention mechanism plays a crucial role in this success story. If you’re curious, have a look at the paper “Attention is all you need” <https://arxiv.org/abs/1706.03762>

Objectives

- To understand the structure and working principles of neural networks and transformers
- To gain insights into some layer types of feed-forward neural networks (dense, convolutional, dropout, attention, embedding) and how they are trained.
- To learn how to classify images with neural networks
- To learn how to solve common NLP tasks with transformers
- To learn how to frame time-series tasks for machine and deep learning
- To gain hands-on experience with Python and the deep learning library *Keras*
- To be able to leverage and fine-tune state-of-the-art models (from Microsoft, OpenAI, Google, DeepMind, HuggingFace)

Content

- How to classify images (day 1), how to analyze text (day 2), how to predict time-series (day 3)
- How to prepare your data for neural networks (structure, training methods, data preparation)
- How to train models (Keras, model saving, checkpointing)
- How to leverage pre-trained models (models, spaces, fine-tuning)
- How to compare deep- and machine learning models

Preconditions

- Fluency in the programming language “Python”, as e.g. provided in the course “Machine Learning with Python – from Zero to Hero”.
- **Google account to access Google Colab.**

Duration

3 days on Feb 14th-16th (roughly 7*45 minutes)

Evaluation

take home exam: project work to be solved in *Python* on image, text and time-series data

ECTS

1.5