

Course description

This is an advanced course on modern Machine Learning methods. It is mostly focused on Deep Learning basic topics (backpropagation, training tricks) and applications for visual and text data. We will also cover some advanced Machine learning topics like recommender systems, learning to rank and other complex problems. All topics will be covered with practical homework assignments on Python.

Course requirements, grading, and attendance policies

Class attendance and participation are encouraged, but not required. The course grade will be based on homework assignments (70% of the grade) and the final exam (30%).

Course Contents

1. Fully connected neural networks. Nonlinearities. Backpropagation and gradient-based training.
2. Neural networks for images. Convolutions. Receptive field. Padding. Typical convolutional architecture.
3. Training neural networks. Optimization methods. Regularization. Dropout. Batch normalization. Weight initialization.
4. Popular CNN architectures. VGG, Inception, ResNet. Embeddings from pre-trained CNNs. Fine-tuning.
5. Computer vision applications. Style transfer and perceptual loss. Image segmentation. Metric learning.
6. Shallow approaches for texts. Word embeddings. Self-supervision.
7. Recurrent neural networks. Backpropagation through time. LSTM. Seq2seq architectures.
8. Attention. Transformers. BERT.
9. Generative models. Variational autoencoders. Generative adversarial networks.
10. Recommender systems and ranking.
11. Up-lift modeling.

Course materials

There is no required textbook, but selected chapters from Goodfellow-Bengio-Courville “Deep Learning” will be useful.

Academic integrity policy

Cheating, plagiarism, and any other violations of academic ethics at NES are not tolerated.