

CEng 501 Deep Learning

Department of Computer Engineering @ METU – Fall 2023

Instructor: Emre Akbas; Office B-208; emre@ceng.metu.edu.tr; Office hours: <https://dar.vin/ea-oh>.

Lectures: Monday 09:40-12:30 @ BMB4.

Web: <https://dar.vin/501-F231> and <https://odtuclass.metu.edu.tr/>.

Description: This course aims to teach the fundamentals of deep learning. We will study the three major types of deep neural networks, namely, Multi-layer Perceptrons, Convolutional Neural Networks, and Recurrent Neural Networks, and take an in-depth look at their use in supervised learning, unsupervised learning, generative modeling and reinforcement learning. We will also explore the most recent developments in the field and state of the art applications of deep neural networks in computer vision and natural language processing. Weekly tentative schedule is as follows.

Date	Topic	Activities	
1	Oct 2	Course logistics. High-level introduction to deep learning.	Hw1 given Oct 2
2	Oct 9	Machine learning background and basics. <i>An overview of supervised learning</i>	Hw1 due Oct 11
3	Oct 16	Artificial neurons. Multi-layer Perceptrons <i>Biological neuron, artificial neuron, Perceptron, Multilayer Perceptrons, Activation Functions, Loss Functions, Backpropagation, Stochastic Gradient Descent</i>	
4	Oct 23	Convolutional neural networks (CNNs) <i>Connectivity types, Convolution, Convolutional neural networks, Pooling, AlexNet, Data augmentation, Dropout, Batch and group norm.</i>	
5	Oct 30	CNNs <i>More loss functions, Initialization, Implementing backpropagation in a modular way, Adaptive learning rate methods, Deconvolution</i>	
6	Nov 6	Applications of CNNs Practical aspects of training deep models	Hw2 given Nov 6
7	Nov 13	Recurrent neural networks (RNNs) <i>Recurrent neural networks, Backpropagation through time, Long short-term memory networks, Gated recurrent units, Encoder-decoder architectures</i>	Hw2 due Nov 19
8	Nov 20	Applications of RNNs	
9	Nov 27	Self-attention, Transformers	
10	Dec 4	Applications of Transformers	Hw3 given Dec 4
11	Dec 11	Deep generative models <i>Boltzmann machines, Deep belief networks, Auto-encoders, Variational autoencoders, Generative Adversarial Networks, Diffusion</i>	Hw3 due Dec 17
12	Dec 18	Deep reinforcement learning (RL) <i>Intro to RL, Deep Q-Learning, Deep policy gradient, Applications of RL</i>	
13	Dec 25	Misc. topics – latest trends, limitations, open issues	
14	Jan 1	No class – New year’s day	

Textbook: There is no official textbook for the class. We will follow the state of the art mainly with papers and by using parts of the “Deep Learning” book by Goodfellow et al., which is available online.

Grading: Homework assignments 55% (15+20+20); Final exam 40%; Participation 5%

Prerequisites: Foundational knowledge in machine learning, calculus, linear algebra. Proficiency in Python.

¹Full url: <http://user.ceng.metu.edu.tr/~emre/Fall12023-DeepLearning.html>