

# EECS E6691 Advanced Deep Learning

## (TOPICS: DATA-DRIVEN ANALYSIS & COMPUTATION)

Spring 2023, 3 credits

**Instructor:** Dr. Mehmet Kerem Turkcan      mkt2126 (at) columbia.edu

### Description

This is an advanced-level course in which the students study topics in deep learning. It is required that students had previously taken a first-course in deep learning. The course consists of: (i) lectures on state-of-the-art architectural and modeling concepts, (ii) assignments, (iii) exam, and a (iv) final project. The course will address topics beyond material covered in the first course on Deep Learning (such as ECBM E4040), with applications of interest to students. In 2023, the main subject of the lectures will be object detection.

Students entering the course have to have prior experience with deep learning and neural network architectures including Convolutional Neural Nets (CNNs), Recurrent Neural Networks (RNNs), Long Short Term Memories (LSTMs), and autoencoders. They need to have working knowledge of coding in Python, Python libraries, Jupyter notebook, TensorFlow both on local machines and on Google Cloud, and of GitHub or similar code hosting tools. The framework and associated tools which will be the focus of this course are PyTorch and Google Cloud. Students have to be self-sufficient learners and to take an active role during classroom activities.

There will be a few (3-4) assignments throughout the semester focusing on coding. In the second half of the course, there will be a midterm exam comprised of multiple-choice questions.

Final projects need to be documented in a conference-style report, with code deposited in a GitHub repository. The code needs to be documented and instrumented such that the instructor can run it after a download from the repository. A Google Slides presentation of the project suitable for a poster presentation is required.

### Prerequisites

- (i) Machine Learning (taken previously, or in parallel with this course).
- (ii) ECBM E4040 Neural Networks and Deep Learning, or an equivalent neural network/DL university course taken for academic credit.
- (iii) The course requires an excellent theoretical background in probability and statistics, and linear algebra.

Students are strongly advised to drop the class if they do not have an adequate theoretical background and/or previous experience with programming of deep learning models. It is strongly advised (the instructor's requirement) that students **take no more than 12 credits of any coursework (including this course and project courses) during the semester while this course is being taken.**

### Registration

The enrollment is limited to several dozen students. The instructor's permission is required to register. Students interested in the course need to populate the SSOL waitlist, and **MUST** also populate the [questionnaire](#). The instructor will move the students off of the SSOL waitlist after reviewing the questionnaire.

### Grading for the course (2023 Spring)

1. Assignments: 30%
2. Midterm Exam (Delivered at Week 11): 30%
3. Project (Final report & Code Repository): 40%