ECBM E4040: Neural Networks and Deep Learning

Fall 2023, 3 credits

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

Description

The course covers theoretical underpinnings and practical aspects of Neural Networks and Deep Learning. Students will learn about and implement a range of different deep learning architectures, including convolutional and recurrent neural networks. The focus of the course is on applications and projects. The main resource for the course is the book entitled "Deep Learning" by Ian Goodfellow, Yoshua Bengio, and Aaron Courville, MIT Press, http://www.deeplearningbook.org, 2016.

Prerequisites

(i) Machine Learning (taken previously, or in parallel with this course).

(ii) The course requires background in probability and statistics, and in linear algebra.

(Tentative) Grading for the course (2023 Fall)

1. **Assignments (40%):** The course usually has 4 assignments (homeworks). All assignments are graded. Each student has to complete their own coding tasks and questions, using Python, TensorFlow, Jupyter notebooks, and code management tools.
2. **Exam (at Week 11, 25%):** Students will take an in-class exam covering the deep learning theory, architecture and design, as well as mastery of prerequisite topics in linear algebra, probability and statistics, calculus, and basic machine learning. The exam will be after the break.
3. **Final Projects (35%):** Students will work on one project, based on contemporary papers or original ideas, in groups of no more than 3. Projects will have to be documented in code, a report, presentation/poster/slides and (possibly) a website.

   Late homeworks (Assignments) - Slip Days: A student is entitled to 4 late days without penalty. For all homeworks together, a student can divide those four days in any fashion needed. Examples: (i) Homework 2 is late 4 days, in which case no other homework can be late for any amount of time; (ii) Homework 1 is late 1 day, homework 2 is late 2 days, in which case the student still has one more late-day for future assignments. The unit of delay can not be divided into less than a full day (like hours). Requests for additional extensions will not be granted: if the budget of 4 days is blown, the student will be given 0 credit for homework(s) for which their submission is late. Late policy does not apply to project submissions.

(Tentative) Weekly Syllabus (2023 Fall)

- **Week 1:** Introduction to the Course, Introduction to Deep Learning, Introduction to Computing Resources for Deep Learning
  Course Textbook: https://www.deeplearningbook.org/
  Other Instructional Resources: https://d2l.ai/
- **Week 2:** Introduction to the Course (Continued), Review of Machine Learning
- **Week 3:** Deep Forward Networks, Introduction to Backpropagation

- **Week 4:** Backpropagation, t-SNE, Universal Approximation Theorem
- **Week 5:** Optimization for Deep Learning
- **Week 6:** Convolutional Neural Networks
- **Week 7:** Regularization for Deep Learning
- **Week 8:** Practical Methodology
- **Week 9:** Modern CNN Architectures
- **Week 10:** RNNs and Exam Preparation
- **Week 11:** Exam
- **Week 12:** Autoencoders
- **Week 13:** GANs, Variational Autoencoders, Trends in Deep Learning
- **Week 14:** GANs, Variational Autoencoders, Trends in Deep Learning (Continued)