ECE-GY 7123 / CS-GY 6953 / Deep Learning

Fall 2021

Objectives

Upon successful completion of this course, you will be able to:

- grasp the mathematical basics of deep learning,
- formulate practical machine learning problems arising in a variety of practical applications,
- implement software prototypes of deep learning-based solutions to these problems,
- and test these prototypes on real-world datasets.

Pre-requisites

- CS-GY 6923, ECE-GY 6143, or equivalent graduate-level machine learning course
- Mathematical maturity (esp., linear algebra and optimization)
- Expertise in Python programming

Outline of lectures (tentative)

Week	Lecture topic
1	ML basics
2	Deep neural nets
3	Autodiff
4	Convnets
5	Object detectors
6	Recurrent architectures
7	Attention mechanisms
8	Deep learning in NLP
9	Reinforcement learning
10	Deep RL
11	Generative adversarial nets
12	Self-supervision
13	Meta learning
14	Current and future trends

Reading material

The primary text will be the course lecture notes (posted online before each week). We will be loosely following the structure listed in the (excellent) online textbook "Dive Into Deep Learning" available at d2l.ai.

Grading policy

- 60% - Homework assignments (expect one assignment every three weeks)
- 40% - Course project

Expect to receive 4 homework assignments. Homework will consist of a mix of theory and programming assignments. Projects will be semester-long, will be student-defined, and will involve solving a real-world problem using tools from deep learning.

Contact info

Chinmay Hegde (chinmay.h@nyu.edu)