

Columbia University
Department of Statistics
GR5245 Python for Deep Learning
Fall 2024

Instructor

Ka Yi NG (email: kn15@columbia.edu)
Office hours: by appointment

Teaching Assistant

Office hours: TBD

Class schedule

Section 001: Thurs 5:35pm – 6:35pm
Section 002: Thurs 6:45pm – 7:45pm

Course Syllabus

Description

This course is an optional companion lab course for GR5242 Advanced Machine Learning. The aim of this course is to help students acquire the basic computational skills in a python-based Deep Learning library (such as Torch, TensorFlow) to implement deep learning models. Lab class materials will be aligned closely with the topics covered in GR5242. Google Colab will be used as the main tools for the hands-on lab exercises.

Open to GR5242 students only.

Prerequisites

Some familiarity with python is assumed, but we will begin the class with a tutorial on 'Python for machine learning'.

Textbooks

Textbook to be determined. Additional materials depend on which Deep Learning library will be covered in class.

Grading Policy

There will be 5-6 homework assignments in total. Student with an average grade of 70% or higher will receive a Pass in this course.

Academic integrity

Students are expected to truthfully represent their course participation and submit their own work. They should adhere to the [GSAS policy on Academic integrity and Responsible Conduct of Research](#).

Course outline (subject to change)

Students are expected to use their own laptops to run their own Python code during the classes.

- Python overview, data generation and visualization (numpy, pandas, matplotlib)
- Tensors and auto differentiation
- Building blocks for neural networks
- Data preprocessing
- Convolutional neural networks
- Natural language processing
- Transfer learning
- Representation learning
- Reinforcement learning