Deep Learning CS 236606 Syllabus   למידה עמוקה

The course introduces techniques and principles of "deep learning" — a branch of machine learning concerned with the research, development and application of deep multi-layered artificial neural networks. Deep neural networks implement high capacity models capable of automatically extracting hierarchical representations of data, and enable state-of-the-art performance in many interesting applications such as image recognition, speech recognition, language translation, and autonomous driving. The course will address a range of topics that cover the principles of basic neural networks, convolutional and recurrent network structures, unsupervised learning, and deep reinforcement learning. The course will also introduce the main ideas behind various current applications in machine vision, text processing, image captioning and game playing.
Homework will include both dry and wet exercises. Wet exercises will require Python proficiency.

Lecturer: Professor Ran El-Yaniv
Teaching Assistants: Yonatan Geifman, Yair Feldman
Credit points: 3 (2 hour lecture and 1 hour tutorial each week)

Prerequisite courses (enforced):
- Introduction to Machine Learning 236756 or
- Introduction to Machine Learning 046195

    Python programming proficiency: the course will require substantial proficiency in Python programming.

Conflicts: Due to some overlap, it will not be possible to receive credit to both this course and 236605.

Grading: The final grade will be based on a combination of homework assignments (both dry and wet) and a final project. The precise grading scheme will be announced in class.

Time/Location: to be announced