FRE-GY-7871,
Advanced Topics in Deep Learning

Instructor Information

- Ken Perry
- Office hours on request (Zoom preferred)
- ken.perry@nyu.edu

Course Information

- FRE-GY-7871
- Advanced Topics in Deep Learning
- Course description: survey of advanced and recent topics in Deep Learning, along with the practical skills to use them
- Prerequisites: Familiarity and competence with Deep Learning and associated tools (Keras, Jupyter). This can be satisfied either by having taken the FRE-GY-7773 section taught by Professor Perry OR by reviewing the notebooks on Deep Learning that were part of the latest iteration of that course. You should also be comfortable reading research papers as some of the material we will learn is recently published.
- Face-to-face class meeting
- Time and classroom varies by semester

Course Overview and Goals

A Deep Learning Neural Network is a combination of parameterized components. The parameters are “learned” by a training process that minimizes a Loss function.

In the Deep Learning section of my Introduction to Machine Learning course there were several assumptions

1. Networks (models) were small and organized as a sequence of layers
2. Training datasets were also small and fit in memory
3. The Loss functions were chosen from a small, pre-defined set

These simplifications are increasingly not representative of the state of the art in Deep Learning

A. Model architecture is fully general: the Functional (versus Sequential) models of Keras
B. Modern models can be large: hundreds of millions (or billions) of parameters
C. Datasets need to be large as well: too big to fit in memory
D. Loss functions are hand-engineered to capture the semantics of the task at hand

The objective of this course is to study state of the art Deep Learning. We will gain an understanding of the theory and techniques that enable recent advances that have captured the imagination of the public (such as Text to Image) as well as acquire the technical skills that allow us to participate in the advances.

Upon completion of this course, students will:

- be able to use advanced programming techniques for constructing Neural Networks
- understand fundamental techniques in advanced Deep Learning networks
- understand recent breakthroughs in the field of Deep Learning that have contributed to the current rapid innovations
- be able to jumpstart projects by using Transfer Learning and Model Hubs

Course Requirements

Class Participation

Students are strongly encouraged to participate by raising questions, offering insights and participating in class discussion. The best way to be prepared to participate is by reviewing the weekly notebook and reading material in advance of class.

Assignments

There will be a Final Project that will constitute the bulk of the Course Grade.

The Project will require that you submit a Jupyter notebook that solves a Deep Learning task that will be chosen for you from among recent innovations. You may need to use some of the advanced techniques that you will have learned in the early lectures.

Tests & Quizzes

None

Assigned Readings

Readings for each week will be posted online as part of the course schedule. It is suggested that students read the material prior to class so as to enable them to participate in discussion.

Grading of Assignments

The Final Project will be assigned a numeric grade. There will be a number Bmax assigned as the maximum number of “base points” that may be earned by successfully completing a required task and a number Emax assigned as the maximum number of “extra points” awarded as extra credit for succeeding at optional tasks or for exceptional approaches to the required tasks.

The grade for the project will be a sum of the “base points” earned (B) and the “extra points” earned (E) with no distinction between the two.

The Course Letter Grade will be determined by computing normalized points earned for the Project. The points for the project will be normalized by dividing the points earned (B+E) by the maximum number of base points (Bmax)

$$(B + E)/Bmax$$

The grade for this course will be determined according to the following formula:

Assignments/Activities	% of Final Grade
Class participation	5%
Final project	95%

Letter Grades

The Letter Grade will be almost totally (95%) based on the normalized points earned on the Final Project.

The threshold for earning each Letter Grade will depend on the distribution (across students) of the normalized points earned on the Final Project. The threshold for earning the highest letter grade is usually in excess of 100% (so earning extra credit is highly encouraged) but the threshold is not predetermined.

Class participation (5%) will be used to shift border-line cases to a higher Letter Grade.

View Grades

Grades will be posted on Brightspace

Course Schedule

The course schedule will be available online (Brightspace). The posted schedule is approximate and will be adjusted based on actual progress.

My goal is to keep slightly ahead of schedule so we will typically complete the current lecture's plan and make a head-start on the following lecture. This will give us a buffer to spend more time on subjects that the students find challenging or particularly interesting.

The exact content and material for each week's lecture (and the following week's lecture, so that we can get a head-start) will be posted at least one day in advance of the course meeting.

Course Materials

The course will be presented as a collection of Jupyter notebooks, available on Github, and released at least one day in advance of the weekly meeting.

The notebooks will also contain a variety of readings (mostly Research papers, historic and recent). These will be identified as follows

- <u>Suggested Reading</u>: you are expected to have at least reviewed this material in advance of the class meeting. This will enable you to be part of the Class Participation.

- Further Reading: Available for students who want a deeper understanding of the topics, but not required.

Required Textbooks & Materials

- Jupyter notebooks distributed via Github will be provided weekly.

Resources

- **Access your course materials**: NYU Brightspace
- **Databases, journal articles, and more**: Bern Dibner Library (library.nyu.edu)
  NYU Virtual Business Library (guides.nyu.edu/vbl)
- **Obtain 24/7 technology assistance:** Tandon IT Help Desk (soehelpdesk@nyu.edu, 646.997.3123)
  NYU IT Service Desk (AskIT@nyu.edu, 212-998-3333)

Policies

Academic Misconduct

E. Introduction: The School of Engineering encourages academic excellence in an environment that promotes honesty, integrity, and fairness, and students at the School of Engineering are expected to exhibit those qualities in their academic work. It is through the process of submitting their own work and receiving honest feedback on that work that students may progress academically. Any act of academic dishonesty is seen as an attack upon the School and will not be tolerated. Furthermore, those who breach the School's rules on academic integrity will be sanctioned under this Policy. Students are responsible for familiarizing themselves with the School's Policy on Academic Misconduct.

F. Definition: Academic dishonesty may include misrepresentation, deception, dishonesty, or any act of falsification committed by a student to influence a grade or other academic evaluation. Academic dishonesty also includes intentionally damaging the academic work of others or assisting other students in acts of dishonesty. Common examples of academically dishonest behavior include, but are not limited to, the following:

   4. Cheating: intentionally using or attempting to use unauthorized notes, books, electronic media, or electronic communications in an exam; talking with fellow students or looking at another person's work during an exam; submitting work prepared in advance for an in-class examination; having someone take an exam for you or taking an exam for someone else; violating other rules governing the administration of examinations.
   5. Fabrication: including but not limited to, falsifying experimental data and/or citations.
   6. Plagiarism: Intentionally or knowingly representing the words or ideas of another as one's own in any academic exercise; failure to attribute direct quotations, paraphrases, or borrowed facts or information.

7. Unauthorized collaboration: working together on work that was meant to be done individually.
8. Duplicating work: presenting for grading the same work for more than one project or in more than one class, unless express and prior permission have been received from the course instructor(s) or research adviser involved.
9. Forgery: altering any academic document, including, but not limited to, academic records, admissions materials, or medical excuses.

Disability Disclosure Statement

Academic accommodations are available for students with disabilities. Please contact the **Moses Center for Students with Disabilities** (212-998-4980 or mosescsd@nyu.edu) for further information. Students who are requesting academic accommodations are advised to reach out to the Moses Center as early as possible in the semester for assistance.

Inclusion Statement

The NYU Tandon School values an inclusive and equitable environment for all our students. I hope to foster a sense of community in this class and consider it a place where individuals of all backgrounds, beliefs, ethnicities, national origins, gender identities, sexual orientations, religious and political affiliations, and abilities will be treated with respect. It is my intent that all students' learning needs be addressed both in and out of class, and that the diversity that students bring to this class be viewed as a resource, strength and benefit. If this standard is not being upheld, please feel free to speak with me.