FRE 7871

Advanced Topics in Deep Learning

Instructor Information

- Ken Perry, Adjunct Professor
- ken.perry@nyu.edu

Course Information

- Title: Advanced Topics in Deep Learning
- Description: Study and apply cutting edge techniques in Deep Learning
- Prerequisites:
    - FRE 7773 Introduction to Machine Learning
    - Fluency with the Machine Learning tools
        - Jupyter
        - Keras
    - In addition: the student
        - Should be comfortable reading the dense and stylized language of research papers
        - Be able to explore and understand complex code
- There will be one lecture per week, presented as Jupyter notebooks.
- Since much of what we learn is state of the art: there will be assigned readings from recently published (as well as historically important) research papers.

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

- Networks (models) were small and organized as a sequence of layers
- Training datasets were also small and fit in memory
- 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

- Model architecture is fully general: the Functional (versus Sequential) models of Keras
- Modern models can be large: hundreds of millions (or billions) of parameters
- Datasets need to be large as well: too big to fit in memory
- 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 the course, the students will be able to

- Understand and use advanced model types (e.g., Transformer)
- Understand cutting edge developments in Deep Learning
- Be able to jumpstart projects by using Transfer Learning and Model Hubs

Course Requirements

Class Participation

Students are expected to read and understand research papers and be able to discuss them in class.

Assignments

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

Tests & Quizzes

None.

Assigned Readings

TBD

Grading of Assignments

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

Assignments/Activities	% of Final Grade
Class participation	10%
Final project: Computer code plus description of process, presented as a Jupyter notebook.	90%

Course Schedule (tentative and subject to change)

Topics and Assignments

Week/Date	Topic	Reading	Assignment Due
Week 1	Basics Review of key concepts in Deep Learning; introduction to enabling technologies (Functional Models, Dataset API).		
Week 2	Generative Deep Learning Part 1 Synthetic data; Autoencoders (vanilla, variational); GANs		
Week 3	Generative Deep Learning Part 2 Time Series GANs. Evaluating quality of synthetic data.		
Week 4	Advanced NLP Part 1 Attention, Transformers, Large Language Models.		
Week 5	Advanced NLP Part 2 Zero-shot learning (learning to learn), Text to image (DALL-E, Stable Diffusion)		
Week 6	Model Hubs and Transfer Learning		
Week 7	Wrap-up		Final Project.

			Student presentations (tentative)

Course Materials

Required Textbooks & Materials

- Links will be given to online material

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

A. 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.

B. 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:

    1. 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.

2. Fabrication: including but not limited to, falsifying experimental data and/or citations.
3. 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.
4. Unauthorized collaboration: working together on work that was meant to be done individually.
5. 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.
6. 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.