Boston University Metropolitan College

Advanced Machine Learning and Neural Networks C1

CS 767S

Instructor:
Name: Chertushkin Mikhail
Email: mikhaild@bu.edu

TA:
Name: Zeyu Cai
Email: zeyucai@bu.edu

Course Description

Theories and methods for learning from data. The course covers a variety of approaches, including Supervised and Unsupervised Learning, Neural Nets and Deep Learning. The course is using Pytorch as a main framework for training neural networks.

Books

Prince, J. D. (2023). Understanding Deep Learning. Retrieved from: https://www.amazon.com/Understanding-Deep-Learning-Simon-Prince/dp/0262048647/ - Required

Goodfellow, I (2016). Deep Learning. Retrieved from: https://www.amazon.com/Deep-Learning-Adaptive-Computation-Machine/dp/0262035618 - Recommended

We are going to use 2 books: one is required, and one is recommended. In syllabus the default book is the Required book (Prince), if not stated explicitly. For example, the RNN topic is covered better in the Recommended book (Goodfellow)

Courseware

Blackboard

Class Policies

1) **Attendance & Absences** – this course emphasizes a lot on practice and requires full attendance on lectures. Working laptops with full charge batteries are necessary as they are needed for passing the in-lecture submissions. During all lectures, students will implement at least one of the tasks covered by theoretical material. The lectures will consist of 50% theory and 50% practice and will be organized as “Reverse Seminars” - this means that students must read the chapters of the coursebook before coming into class. They first are presented with an machine learning problem and then they try to

solve it. After trial-and-error students get familiar with necessary theoretical concepts and submit the solution to the grading system after the lecture.

2) **Assignment Completion & Late Work** – every week students will have to solve one homework assignment and 1 Lab assignment, which will usually have 3-4 tasks. The time for submission of homework is next Monday 11:59 PM, the day before lecture. Late submissions are not possible. The time for submission of Lab assignment is Friday 11:59 PM the same week.
3) **Academic Conduct Code** – Cheating and plagiarism will not be tolerated in any Metropolitan College course. They will result in no credit for the assignment or examination and may lead to disciplinary actions. Please take the time to review the Student Academic Conduct Code:

http://www.bu.edu/met/metropolitan_college_people/student/resources/conduct/code.html. This should not be understood as a discouragement for discussing the material or your particular approach to a problem with other students in the class. On the contrary – you should share your thoughts, questions and solutions. Naturally, if you choose to work in a group, you will be expected to come up with more than one and highly original solutions rather than the same mistakes."

Important information

Course is organized as "Reverse Classroom" (Flipped Classroom) - https://bokcenter.harvard.edu/flipped-classrooms . As part of a **flipped classroom** format (Harvard's Flipped Classroom Model), here are a few important points:

1. **Preparation**
   You *must* read the relevant textbook chapters before each lecture. For our **first lecture**, please read **Chapters 1 and 2** from our textbook, *Understanding Deep Learning* by Simon Prince (Amazon link).
   a. To know which chapters to read for each lecture, consult the **Syllabus**. You'll see that Lecture 1 requires Chapters 1 and 2, and so forth.
2. **Attendance**
   Attendance accounts for **12%** of your overall grade. If you need to miss a class due to a serious or personal issue, please let me know in advance, and you won't be penalized. Otherwise, missing classes could significantly affect your final grade.
3. **Presentations and Problem-Solving**
   During each class, we will solve problems from the textbook, and everyone will present at least once. Although it may be challenging, this hands-on experience will greatly enhance your understanding.

Grading Criteria

Grades are calculated as a weighted combination of five pieces:

- Homeworks (100 points each)
- Labs (100 points each)
- Presenting Lab (500 points, done twice per course for each student)
- Kaggle Competitions (500 points, 3 competitions)
- Attendance (1% of total grade for each lecture, total 12% of total grade)

Each Homework and Lab costs 100 points.

Kaggle competition costs 500 points if you secure the first place. If not, you will receive a penalty, based on your position on the leaderboard.

There will be 3 Kaggle Competitions:

- Classical Machine Learning methods (predict who wins Dota 2 game based on first 5 minutes of the game)
- Natural Language Processing methods (generate the answer based on the question – competition involves building Retriever Model and Generator model, this is a little glance into Gen AI and RAG)
- Computer Vision methods (predict malignant, benign or no tumor based on ultrasound image of breast, you will need to build the Semantic Segmentation model)

Presenting lab costs 500 points and must be done by every student at least twice. If you do not present lab in class, you will fail the course. Every class we will solve 3 tasks, so every class 3 different students will present their solution. We will have ~12 lectures, so this gives us 36 presentations. Presenting Lab is a challenging and rewarding experience. If you volunteer to present lab yourself – the grading of presenting will be more generous.

Attendance costs 12% (1% of total grade per one lecture) and is needed to get a perfect score. For example, if you don't attend any lecture in semester, but do everything else perfectly, the maximum grade you will be able to get is 100-12=88.

Class Meetings, Lectures & Assignments

There will be lectures every week for the following set of topics. We will examine various topics starting from basics such as principles of Machine Learning and then advancing to Deep Learning.

Lectures, Readings, and Assignments subject to change, and will be announced in class as applicable within a reasonable time frame.

Date	Topic	Readings Due	Assignments Due
May 20	Introduction to Machine Learning, Deep Learning and AI. Classical ML Algorithms. Supervised and unsupervised learning.	Ch. 1, 2	n/a
May 27	Classical ML Algorithms.	Ch. 3, 4, 5	Assignments of first week

	Ensemble Methods.		
June 3	Fully connected Networks	Ch. 6, 7	Assignments of second week
June 10	Fully connected Networks with Classification	Ch. 8, 9	Assignments of third week Kaggle Competition #1
June 17	Computer Vision. CNNs and types of computer vision tasks	Ch. 10	Assignments of fourth week
June 24	Computer Vision. ResNets. Advanced Types of tasks.	Ch. 11	Assignments of fifth week
July 1	Natural Language Processing. Recurrent Neural Networks	Goodfellow: Ch. 10	Assignments of sixth week Kaggle Competition #2
July 8	Language Models. Seq2Seq and Attention	Goodfellow: Ch 12.4 "Applications: Natural Language Processing"	Assignments of seventh week
July 15	Transformers. Large Language Models.	Ch 10	Assignments of eighth week
July 22	Transfer Learning. BERT & GPT	https://www.ruder.io/transfer-learning/	Assignments of ninth week
July 29	ChatGPT. RLHF. Mistral and Mixtral.	Ch 19	Assignments of tenth week
August 5	Generative Adversarial Networks. Variational Autoencoders. Diffusion Models. Conclusion of the course	Ch 15, 17 Ch 18, Ch 20	Assignments of eleventh week Kaggle Competition #3


BOSTON
UNIVERSITY