Boston University Metropolitan College

Advanced Machine Learning and Neural Networks A1

CS 767

Name: Chertushkin Mikhail
Email: mikhaild@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."

Grading Criteria

Grades are calculated as a weighted combination of four pieces:

- Homeworks
- Labs
- Presenting Lab
- Kaggle Competitions

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 (predict the toxicity/politeness of the text)
- Computer Vision methods (predict lung disease based on the X-ray images of lungs)

Presenting lab costs 500 points and must be done by every student at least once. 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. Presenting Lab is challenging and rewarding experience. If you volunteer to present lab yourself – the grading of presenting will be generous.

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

January 21	Introduction to Machine Learning, Deep Learning and AI. Classical ML Algorithms. Supervised and unsupervised learning.	Ch. 1, 2	n/a
January 28	Classical ML Algorithms. Ensemble Methods.	Ch. 3, 4	Assignments of first week
February 4	Fully connected Networks	Ch. 5, 6	Assignments of second week
February 11	Fully connected Networks with Classification	Ch. 9	Assignments of third week Kaggle Competition #1
February 18	No Lecture (Substitute Monday class)	No Lecture (Substitute Monday class)	Assignments of fourth week
February 25	Computer Vision. CNNs and types of computer vision tasks	Ch. 10	N/A
March 4	Computer Vision. Advanced Types of tasks.		Assignments of fifth week
March 11	No Lecture (Spring Recess)	No Lecture (Spring Recess)	Assignments of sixth week
March 18	Natural Language Processing. Recurrent Neural Networks	Goodfellow: Ch. 10	Assignments of seventh week Kaggle Competition #2
March 25	Seq2Seq and Attention	Ch 12	Assignments of eighth week
April 1	Transformers	Ch 17, 20	Assignments of ninth week
April 8	Autoencoders & Diffusion Models	Goodfellow: Ch. 15, 16	Assignments of tenth week

April 15	Self-supervised Learning & Semi-supervised Learning	Ch 14	Assignments of eleventh week Kaggle Competition #3
April 22	Reinforcement Learning	Ch 19	Assignments of twelve week
April 29	Graph Neural Algorithms	Ch 13	Assignments of thirteen week