Boston University Metropolitan College

Advanced Machine Learning and Neural Networks SC1

CS 767

Name: Chertushkin Mikhail
Email: mikhaild@bu.edu
LinkedIn: https://www.linkedin.com/in/michael-chertushkin-4315b4100/
Office hours: by appointment

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

Bishop, C. M. (2024). Deep Learning: Foundations and Concepts. Retrieved from: https://www.amazon.com/Deep-Learning-Foundations-Christopher-Bishop/dp/3031454677 - 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 (Bishop), 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 first are presented with an algorithmic 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, which will usually have 3-4 tasks. The time for submission of homework is next Tuesday 11:59 PM, the day before lecture. Late submissions are not possible.

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
- Volunteering in class
- Kaggle Competitions

Every lecture there will be an opportunity for volunteering – showing the code to your classmates and explaining how it was solved. This can be done several times by one student, but lecturer will have an ability to restrict the number of volunteering opportunities in rare cases. Given that we have only 12 lectures, there will be only 12 volunteering opportunities. Volunteering opportunity costs 100 points and is highly recommended for securing A or A+ in the class.

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 22	Introduction to Machine Learning, Deep Learning and AI. Classical ML Algorithms. Supervised and unsupervised learning.	Ch. 1, 2, 3, 4	n/a
May 29	Neural Networks and Deep Learning. MLP.	Ch. 5, 6, 7, 8, 9	Assignments of first week
June 5	Convolutional Neural Networks	Ch. 10	Assignments of second week

June 12	Recurrent Neural Networks	Goodfellow: Ch. 10	Assignments of third week
June 19	Attention and Transformers	Ch. 12	Assignments of fourth week
June 26	Graph Neural Networks	Ch. 11, 13	Assignments of fifth week
July 3	Sampling & Latent Variables	Ch 14, 15, 16	Assignments of sixth week
July 10	GANs	Ch 17	Assignments of seventh week
July 17	Autoencoders & Diffusion Models	Ch 19, 20	Assignments of eight week
July 24	Self-supervised Learning & Semi-supervised Learning	Goodfellow: Ch. 15, 16	Assignments of ninth week
July 31	Bayesian Optimization	NA	Assignments of tenth week
Aug 7	Advanced Topics	Goodfellow: Ch. 20	Assignments of eleventh week