Boston University Metropolitan College

Advanced Machine Learning and Neural Networks

MET CS767 A1, Fall 2024

Learn from Anywhere Course Format, Offered Simultaneously on Campus and Remote

Zoran B. Djordjević, PhD, Instructor
Chandra, Rajesh, Teaching Assistant
zdjordje@bu.edu
Office hours: by appointment

Course meets in the classroom, CAS 324, 685-725 Commonwealth Avenue, every Thursdays between 6:00 and 8:45 PM, starting September 5th, 2023. The last lecture on December 5th.
The last class meeting for the final project presentations on December 19th, 2024.
All lectures will be available online on Zoom. All lectures will be recorded and recordings posted on the class BlackBoard site.

Course Description

In this course, students will master the most important neural network and deep learning concepts and techniques and start applying them productively in modern AI workplace. Deep learning is the primary technique for data analysis and the solution for many complex problems in natural sciences, linguistics, and engineering. We use deep learning for image classification, manipulation and generation, speech recognition and synthesis, natural language translation, sound and music manipulation and generation, navigation of self-driving cars, and many other activities. Students will master several key deep learning architectures, such as convolutional neural networks (CNNs), long short-term memory networks (LSTMs), autoencoders (AEs) and variational autoencoders (VAEs), generative adversarial networks (GANs), transformers with attention, and graph neural networks (GNNs). Students will master the most essential skills for the efficient use of Large Language Model (LLM) based applications such as ChatGPT and DALL*E. The course starts with a review of the theoretical foundations of neural networks approach to machine learning including auto-differentiation and backpropagation. The emphasis is on practical applications of deep learning APIs: Keras 3 (a package within TensorFlow 2.x framework) and PyTorch.

Prerequisites

MET CS 521 and either MET CS 622, MET CS 673 or MET CS 682.
MET CS 677 is strongly recommended.
Or instructor's consent.

Learning Objectives

Students will accomplish the following.

(1) Understand the goals and applications of Machine Learning and Deep Learning
(2) Master practical neural deep learning concepts, techniques and architectures.
(3) Master practical details of Keras 3 and PyTorch on CNNs, Auto-Encoders and VAE, GANs, Transformers, and Large Language Models (LLMs).

Reference Books

We will use parts of:

- ***Deep Learning with Python****, 2nd Edition,* by François Chollet, Manning Publishing, 2021
- ***Hands-on Machine Learning with Scikit-Learn & TensorFlow****, 3rd edition,* by Aurelie Geron, O'Relly 2022
- ***Machine Learning with Pytorch and Scikit-Learn*** by Sebastian Raschka, Yuxi (Hayden) Liu and Vahid Mirjalili, Packt Publishing Ltd, 2022
- ***Deep Learning*** by Ian Goodfellow, Yoshua Bengio and Aaron Courville, MIT Press 2017. We will provide a free PDF copy of this book.

Courseware

Blackboard content areas will provide copies of texts, technical and scientific articles and web links with scientific and other materials relevant for the class.

Fall 2024 COVID-19 Policies

Compliance: Below are the current University guidelines. Please remember that the best ways to protect yourself and others against severe COVID-19 infections are:

- Stay up-to-date with COVID-19 vaccines.
- Wear a mask in crowded, indoor spaces.
- Test if you are not feeling well.

For more information about COVID-19, students can visit the Student Health Services COVID-19 website, and faculty and staff can visit the Occupational Health Center COVID-19 website.

Class Policies

1) **Assignment Completion & Late Work** – all the assignments have to be submitted on the courseware site. If there is a special justification, assignment could be submitted by an email. No late work will be acceptable.
2) **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.

Grading Criteria

The course grade will be based on

- Active class participation measured by the number of useful responses on the discussion board. (5%)
- Assignments (80%)
- Final project (15%) (Details will be provided in the final project assignment)

Solutions for all assignments must be submitted by the respective due date. Late submissions will not be accepted. If, under special circumstances, late submissions are permitted, they will be penalized.

Tentative Class Syllabus

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

Tentative List of Class and Lab Topics:

	Date	Lecture Topic	Lab Topic
1	09/05/24	**Introduction to Neural Networks and Deep Learning,** Brief review of several key examples from physiology and neurology motivate basic concepts and patterns in Neural Networks (NNs).	a) Review of Python, NumPy & Pandas b) Installing TF and Keras c) Course Logistics
2	09/12/24	**Gradient Descent and Back Propagation and advanced TensorFlow 2.x** Concepts and mechanisms behind back propagation and auto-differentiation, algorithms which made large scale NNs feasible.	a) BackPropagation b) Autodifferention
3	09/19/24	**Keras, Loss Functions and Regularization. Sequential and Functional Models** Detailed architecture and patterns for building models in Keras, a high-level API for efficient construction of Deep Learning (DL) networks. Features of the Sequential and Functional Model and prescription on how to construct and format input data sets. Loss functions and the regularization terms for improving the training of the network.	a) Keras, Regression, Callback, etc.
4	09/26/24	**Convolutional Neural Networks (CNNs)** Starting from a physiological model and an analysis of the computational efficiency of fully connected networks we will introduce CNNs. We will also learn how to determine parameters of various layers in CNN Models	CNNs for Classification and other tasks. Determining numbers of model parameters. Sizing the layers.
5	10/03/24	**Visualizing Feature Maps of CNN Layers, Locating Objects in Images.** We will learn how to monitor the evolution of layer's (network's) parameters as NN evolves through the optimization and observe the distributions of values of those parameters in optimized NNs. Those observations will provide a deep insight of the nature of NN data processing.	Examples for Visualization of Feature Maps in CNNs
6	10/10/24	**Transfer Learning, Fine Tuning, Augmentation** Transfer Learning, Fine Tuning, and Data Augmentation are three most frequently used techniques for training networks on restricted data sets and the increasing the precision of Deep Learning analysis.	Examples for Transfer Learning & Augmentation
7	10/17/24	**Autoencoders and Variational Autoencoders** We will learn that NNs behave as if they are searching for a minimal representation of any objects. Using that idea we will "discover" embedded vector representations of words in texts.	Autoencoders Generation of embedded vectors for Word2Vec. Generation

		We will also see an extension of the autoencoders which allows generation of "higher quality" objects.	of high-resolution images.
8	10/24/24	**Sequence Modeling** and **Recursive Neural Networks** (RNNs and LSTMs). Large classes of physical and computational analyses require long range (time) memory of past (and future) events. We will learn how to approach such problems using RNNs and LSTMs.	Use of LSTMs for weather and time series forecasting
9	10/31/24	**Natural Language Processing (NLP) with Doc2Vec like APIs and Large Language Models (LLMs).** DL and LLMs provide sophisticated tools for analysis and representation of text. We will introduce and demonstrate the most important applications of DL techniques in NLP.	Extraction of NLP features, **With this lecture we will start shifting to PyTorch**
10	11/07/24	**Structure of Transformers, Seq2Seq Models and Machine Translation.** Transformers are the core of LLMs. Using LSMT and LLMs we will demonstrate ability to perform machine (automated) translations of texts from one to another language**.**	Machine Translation of Texts from one Language to another language
11	11/14/24	**Analysis and Transcription of Speech** We will learn how to encode human speech and present it as a time series of codes. Subsequently, we will learn how to use such representations to train DL networks to transcribe speech into text.	Capturing speech characteristics. Speech to text.
12	11/21/24	**Generative Adversarial Networks (GANs)** GANs are DL networks of special architecture that could act as generators of objects such as speech, text, images and others. We will learn how to construct such networks and use them in various creative tasks in art and technology.	Examples of generation of visual art and music.
	11/28/24	**No class, Thanksgiving Recess**	
13	12/05/24	**Graphs Neural Networks.** We will extend the convolution ideas to the **Graph Domain**. Graph and **Graph Convolution**. We introduce **spectral graph convolutional neural networks** and discuss how to perform spectral convolution.	Graph Convolutional Networks (GCNs)
14	12/19/24	**Final Project Presentations**	