CS 767

Boston University MET College 2025
Department of Computer Science

Advanced Machine Learning and Neural Networks (CS 767)

General Information

Time & Location: Tue 9 AM - 11:45 AM, MET 101 (1010 Commonwealth Avenue)

Instructor: Avi Mohan (avimohan@bu.edu)
Office hours (tentative): CS 302 (1010 Commonwealth Ave.) by appointment.
Subject line: **[CS 767 Spring 2025] i**ncluding square brackets.

Teaching Assistants: Rajesh Chandra (rj21@bu.edu)
TA Office hours: TBD

Course Web Site Blackboard Learn: https://learn.bu.edu

Required Textbook:

[G22] Géron, A., 2022. "***Hands-on machine learning with Scikit-Learn, Keras, and TensorFlow.***" O'Reilly Media, Inc.
ISBN-13: 978-1098104030

Additional Material

1. [GBC16] Ian Goodfellow, Yoshua Bengio, and Aaron Courville, "***Deep learning,***" The MIT Press, 2016.
   ISBN- 13: 978-0262035613

2. [ZLLS21] Zhang, A., Lipton, Z.C., Li, M. and Smola, A.J., 2021. "***Dive Into Deep Learning***." *arXiv preprint arXiv:2106.11342*.
   Link: https://d2l.ai/index.html

3. [BB22] Bishop, C.M. and Bishop, H., 2024, "***Deep Learning: Foundations and Concepts***." Springer Nature.
   Online version available for free at: https://www.bishopbook.com/
   ISBN- 13: 978-3031454677

4. [DFO20] Deisenroth, Marc Peter, A. Aldo Faisal, and Cheng Soon Ong. ***Mathematics for Machine Learning.*** Cambridge University Press, 2020.
   ISBN- 13: 978-1108455145

5. [J24] James, Gareth, Daniela Witten, Trevor Hastie, Robert Tibshirani, and Jonathan Taylor. ***An Introduction to Statistical Learning: With Applications in Python***. Springer Nature, 2024.
   ISBN- 13: 978-3031387463

6. [NFI21] Navlani, Avinash, Armando Fandango, and Ivan Idris. ***Python Data Analysis: Perform data collection, data processing, wrangling, visualization, and model building using Python***. Packt Publishing Ltd, 2021.
   ISBN- 13: 978-1787127487

Prerequisites:

1. (CS 544 or CS 555), AND (CS 688 or CS 699 or CS 677). This course involves a lot of theory, and students will be expected to learn Scikit-learn, TensorFlow, and PyTorch on their own. Classes will be dedicated to explanation of theory and homework assignments to test theoretical understanding and Python implementation.

Grading (tentative):

- Class attendance and participation (8%)
- Attending research seminar on March 28th 2025 (2%)
- Homework (50%)
  - Students are allowed to collaborate on homeworks (see section on **grading policies** for more details).
- One mid-term and a final test (20% each)
  - Students are allowed to bring 1 A4 size sheet to both tests. Apart from that, the tests are **closed book**.

Collaborating with anyone on tests is strictly prohibited.

Grading Policies:

1. Late homework submissions will not be accepted.
3. It is forbidden to use any human resource or large language model outside of class (including web-based help services, outside tutors, ChatGPT, Google Bard etc.) in doing your homework.
4. Collaboration in solving homework assignments is acceptable. However:
   a. Names of collaborator(s) must be provided on top of the submission.
   b. Each student must provide detailed explanations of his/her solutions expressed in his/her own words. Students must be able to explain their solutions to the instructor, if requested.
   c. Copying solutions from other students or other sources is strictly unacceptable. Plagiarized solutions will be heavily sanctioned.

Grading Scale:

The final grade, denoted by G, will be a number between 0 and 100. Relative grading will be used. In the worst case, this numerical grade will be converted into a letter grade using the following scale:

Numerical Grade	**Letter Grade**
G >= 90	A
85 <= G < 90	A-
80 <= G < 85	B+
75 <= G < 80	B
70 <= G < 75	B-
65 <= G < 70	C+
60 <= G < 65	C
55 <= G < 60	C-
50 <= G < 55	D
G < 50	F

Laboratories and Homework:

1. Written homework will cover topics such as data science libraries (Pandas, MatplotLib, Seaborn, Scikit-Learn, TensorFlow, and the Keras API), basic data analysis, regression and classification using different machine learning models, feed-forward networks, CNNs, RNNs, LSTMs, etc.
2. The first homework tests students' knowledge of Machine Learning basics. Please take this homework very seriously.
3. Students are strongly encouraged to familiarize themselves with using Jupyter Notebooks and bring their machines to class. An introductory tutorial on installing the Anaconda navigator can be found on the Anaconda docs page (Windows) and (Mac OS).

Submitting Assignments:

1. All homework (except code) should be submitted via Blackboard as a **pdf file.**
2. To get full points, all uploaded code (.py or .ipynb files) must compile with **no modifications** and produce the desired output.
3. Perfunctory or insufficient commenting, vague variable, function, class names etc. will result in loss of points [if you're defining a function to add two numbers, using `def f1(n1,n2):` is bad nomenclature – the grader should be able to surmise what `f1` does from its name, so a better option would be `def add_float(feature_1,weights_2):`].
4. Students are encouraged to type their assignments. Scanned answers will be accepted only if they are of sufficiently high quality.
5. After uploading an assignment on Blackboard, please remember to submit the assignment. Otherwise, your assignment will not be available to the grader.

Academic Misconduct:

BU takes academic integrity *very seriously*. Academic misconduct is conduct by which a student misrepresents his or her academic accomplishments or impedes other students' opportunities of being judged fairly for their academic work. Knowingly allowing others to represent your work as their own is as serious an offense as submitting another's work as your own. More information on BU's Academic Conduct Code, with examples, can be found at
https://www.bu.edu/academics/policies/academic-conduct-code/

Computing Ethics:

BU lab and computing facilities must be used responsibly; misuse by even a few individuals has the potential to disrupt University business or the work of others. You are therefore required to exercise responsible, ethical behavior when using the University's computing facilities.
More information on BU's policy on computing ethics, with examples, can be found at
https://www.bu.edu/dos/policies/lifebook/computing-ethics/

Classroom Conduct:

- Attend class regularly and on time. If arriving late, use back door of classroom.
- Actively participate. If some materials require clarification or writing on blackboard is difficult to read, raise your hand.
- Use your laptop only to take notes or participate in class activities. Do not text or surf the Internet during class (students may be requested to leave or move to a back row if this happens).