

COMP/EECE 7745

Machine Learning

Spring 2024

Instructor: Md Zahangir Alom

Time and Location: TR, 6:00-7:25, International Center 115

Email: malom1@memphis.edu

Office Hours: Please send me an email to set up an appointment.

TA: TBA

TA Office Hours: TBA

Learning Objectives

The aim of this course is to study, analyze and apply fundamental techniques in Machine Learning. Both theoretical and empirical aspects of Machine learning will be emphasized. At the end of this course, students should have the necessary background to build practical machine learning systems and be able to apply machine learning in their own research.

Topics

Concept of Learning and data preparation

Supervised Learning

- Linear and logistic regression
- K-Nearest Neighbors (KNN) Algorithm
- Support Vector Machine (SVM) and Kernelization tricks
- Principal Component Analysis (PCA)
- Decision tree, Random Forest (RF), and XGBoost

Neural Network (NN)

- Multilayer perceptron (MLP), and training and optimization methods
- Variants: Convolutional Neural Network (CNN), Autoencoder (AE), Deep Belief Network (DBN), Recurrent Neural Networks (RNN), and Transformer Architecture.

Graph Neural Networks (GNN)

Transfer Learning, self-supervised and contrastive learning

Multi-modality Learning

Unsupervised Learning: K-means, Generative models, tSNE, and UMAP

Reinforcement Learning

Explainable ML (XML)

Regulations and Ethics in ML

Course Text

- a) Machine Learning, Tom Mitchell, McGraw Hill, 1997
- b) Neural Networks: A systematic Introduction, Raul Rojas. Springer-Verlag, 1996.
- c) Deep Learning by Ian Goodfellow, Yoshua Bengio, and Aaron Courville (available at <http://www.deeplearningbook.org/>).

Pre-Requisites

Discrete Mathematics, Probability Theory, Algorithm Analysis and Data Structure

Evaluation

Your final grade for this course will be determined by the following averaging procedure (subject to change):

Assignments	≈ 40 %
Examinations	≈ 30 %
Progress Reports	≈ 10 %
Term project	≈ 20 %

For 8000-level Ph.D. students: Research paper presentation (10%)

- Select a paper from one of the following conferences (preferably published within the last 3 years): NIPS, ICML, KDD, CVPR, ICLR, and AAAI
- Send me the paper that you wish to present by October 15. I will break ties if there are any and try to make sure we have a good overall schedule
- Presentations in the last two weeks of class. Schedule will be announced later.
- Note: If you want to present your own work related to machine learning by tying it to your research, it is highly encouraged. However, it should be current, i.e., you should perform the research during this semester and not previous work. In this case, by October 15th, please send me a proposal of your intended research.

Grading

Grading: A+ ≥ 95%, A ≥ 92%, A- ≥ 87%, B+ ≥ 82%, B ≥ 77%, B- ≥ 72%, C ≥ 65%

Note: A modified curve may be used for determining the grades at the discretion of the instructor.

Policies

1. Both midterm and final exams are closed book and closed notes
2. No late homeworks will be accepted unless well-documented reasons are presented
3. All homeworks must be individual work. Plagiarizing assignments or code sharing is not permitted.
4. Regular class attendance is mandatory. There is a strong correlation between regular attendance and obtaining a good grade. Instructor reserves the right to lower grades for lack of attendance. Students are responsible for any material and contents of missed lectures.
5. No early or late exams will be given unless under extreme situations
6. Any grading errors in assignments should be notified within a week to the TA
7. By taking this course, you agree that any assignment turned in may undergo a review process and that the assignment may be included as a source document in Turnitin.com's restricted access database solely for the purpose of detecting plagiarism in such documents. Any assignment not submitted according to the procedures given by the instructor may be penalized or may not be accepted at all.
8. If plagiarism or cheating occurs, the student will receive a failing grade on the assignment and (at the instructor's discretion) a failing grade in the course. The course instructor may also decide to forward the incident to the University Judicial Affairs Office for further disciplinary action. For further information on U of M code of student conduct and academic discipline procedures, please refer to <http://www.people.memphis.edu/~jaffairs/>

References:

- Pattern Classification, David G. Stork, Peter E. Hart, and Richard O. Duda, Wiley, 2000
- Pattern Recognition and Machine Learning, Chris Bishop, Springer-Verlag, 2006
- Elements of Statistical Learning, Trevor Hastie, Robert Tibshirani and Jerome Friedman, Springer, 2009
- Machine Learning: A Probabilistic Perspective, Kevin Murphy, MIT press