CSCE 5215.001 and 4930.002 Introduction to Machine Learning

Theory and practice of machine learning. Linear regression, logistic regression, decision trees, neural network learning, support vector machines, kernel methods, bagging, boosting, random forests, ensemble learning, deep learning, unsupervised learning including *k*-means and hierarchical agglomerative clustering, semi-supervised learning, active learning, and reinforcement learning. Practical applications of machine learning algorithms. Topics in experimental design and computational learning theory.

TOPICS (tentative), in approximate chronological order:

1. ML Overview
2. Designing a Learning System & ML Algorithms Overview

SUPERVISED LEARNING

3. Linear Regression
4. Perceptron and Classification
5. Logistic Regression and Classification
6. Decision Trees
7. Experimental Design
8. Evaluation
9. Artificial Neural Networks
10. Instance-based learning
11. Bias and Variance
12. Recurrent and Convolutional Neural Networks
13. Deep Learning
14. Deep Neural Networks
15. Support Vector Machines
16. Computational Learning Theory
17. Kernel methods
18. Ensemble methods and Bagging
19. Random Forests and Boosting

UNSUPERVISED LEARNING

20. K-means clustering
21. Hierarchical Agglomerative clustering
22. Expectation Maximization
23. Dimensionality Reduction, PCA

OTHER TOPICS

24. Active and Semi-Supervised Learning
25. Learning from large datasets
26. Reinforcement Learning Overview

Learning Objectives:

Given a problem statement, have the skills to be able to:

- Recognize whether and where machine learning techniques are applicable,
- Determine which set of the ML algorithms covered in the course are applicable,
- Design a machine learning application component,
- Design and conduct appropriate machine learning experiments,
- Analyze and properly describe the outcome of ML experiments,
- Recognize the relevance of computational learning theory, and be able to extract key information from literature.

Major Assignments:

Major Project: You will need to find an appropriate significant problem, determine which ML algorithm(s) are well suited to apply to that problem, design and implement an ML application that contributes to the solution to the problem, design and conduct a statistically valid experiment to compare the results of your application to the results of another strong alternative system (others' or your own) and or to appropriate baseline systems, and analyze and properly describe the outcome of your experiment.

Midterm Exam: A midterm exam will assess your competency with regard to the learning objectives and topics covered in roughly the first half of the semester.

Final Exam: A final exam will assess your competency with regard to the learning objectives and topics covered throughout the semester.

Required Reading: Deep Learning. Ian Goodfellow, Yoshua Bengio, and Aaron Courville

Recommended Reading: Machine Learning. Tom Mitchell

Grading:

15%	Class participation (asking and answering thought provoking questions)
20%	Homework assignments
15%	Midterm Exam
20%	Final Exam
20%	Project
10%	Significant constructive feedback on peer projects

Under extraordinary circumstances, late assignments might be accepted for partial credit if negotiated in advance with the instructor.

Attendance is required and will be reflected as a component of the class participation grade.