CS 767 Machine Learning

Summary

Theories and methods for automating and representing knowledge with an emphasis on learning from input/output data. The course covers a wide variety of approaches, including Supervised Learning, Neural Nets and Deep Learning, Reinforcement Learning, Expert Systems, Bayesian Learning, Fuzzy Rules, and Genetic Algorithms. Each student focuses on two of these approaches and creates a term project.

Prerequisites

MET CS 566 or instructor's consent.

Syllabus

The order of the following topics is 90% certain. They are subject to change according to student background and rapidly emerging technology.

1. January 24: Introduction I (The Introduction will enable you to pick a tentative term topic)

2. January 31: Introduction II

3. February 7: Supervised Learning
   Learning with Neural Nets I
4. February 14: k-means; Introduction to TensorFlow

5. February 21: Learning with Neural Nets II

6. February 28: Deep Learning

7. (Note) March 14: Expert Systems

8. March 21: Fuzzy Rules

10. March 28: Genetic Algorithms I

11. April 4: Genetic Algorithms II; Decision Trees

12. April 11: Bayesian Learning; Dealing with Uncertainty

13. (Note) April 25: tbd
14. May 2: Presentations

Textbook and Other Source Book

We'll use parts of "Machine Learning" by Marsland (2nd edition) ISBN-13: 978-1466583283.

A good book on deep learning is "Fundamentals of Deep Learning" by Buduma (O'Reilly). This is recommended if you select a deep learning project, and especially if you use Tensorflow.

Learning Objectives

Students will accomplish the following.

(1) Understand the goals of Machine Learning and Expert Systems

(2) Understand the main ML technologies

(3) Implement more than one of these techniques in a significant manner

Assignments and Evaluation

Attendance at class: Class discussions and in-class group work is an important part of learning. The percentages below are predicated on virtually full attendance.

Labs: 10%

Project Proposal: 15%

Project Design: 30%

Project Implementation: 45%

There is also the possibility of a joint research project with the instructor. Ask about this option if you are interested in research.

Evaluation of Student Work

Work will be evaluated according to an evaluation matrix. Unless a matrix particular to an assignment is given, the default matrix here will be used. Make sure that your work conforms well to each of the criteria.
The main goal of grading for the course is for each student to improve via feedback.

The average grade of MET graduate students is expected to be very good--B+.
The project phases are graded according to the attached evaluation matrix. These are averaged using A+=97, A=95, A-=93, B+=87, B=85 etc.

To get an A grade for the course, your weighted average should be

>93 for an A

>=90 for an A-

>=87 for a B+

>83 for a B

>=80 for a B- etc.

The lab grades are Pass (1), Fail (0), or Neither (1/2). The grade at the end of the semester are:

90% or more passed: A

80-89% passed: B

70-79% passed: C

60-69% passed: D

Less than 60% passed: F

Late homework or lab will not be accepted unless there is cause why it was not reasonably possible to perform the work in the time. Extraordinary workloads, illness and emergency conditions will be accepted and documentation will be required. If the reason is acceptable, missing work may be graded on a pass/fail basis.