

MET CS767 A1 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 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
- (2) Apply the principal ML technologies
- (3) Implement more than one of these techniques in a significant manner

Syllabus

1. Introduction I (The Introduction will enable you to pick a tentative term topic)
2. Introduction II
3. Introduction to TensorFlow
4. Learning from Data; Unsupervised Learning; k-means
5. Learning with Neural Nets I
6. Learning with Neural Nets II
7. Learning with Neural Nets III and Deep Learning +
8. Genetic Algorithms I
9. Genetic Algorithms II
10. Bayesian Networks
11. Learning Rules
12. Tba or presentations
13. Presentations

Readings

1. =====

Introduction I: Textbook Chapter 3 pp 39-43 and Chapter 8 p169

2. =====

Introduction II: Textbook Chapter 1 and Chapter 10 pp 211-214

3. and 4. ===

Unsupervised Learning: 281

k-means: pp 282-287

Learning with Neural Nets I : pp 15-20, 43-49

5. =====

TensorFlow: https://www.tensorflow.org/get_started/get_started_for_beginners

Neural Nets: pp 73-85

9,10. =====

Primary: Chapters 10.1 and 10.2

Secondary: Chapter 10.3

11 =====

Chapters 2.3 and 16.1

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.

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: Acceptably on track (1), Not yet acceptably on track (0), and Neither (0.5).

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