COMP SCI 760 section 001 Syllabus

Machine Learning

COURSE INFORMATION

Course Information:

Machine Learning
COMP SCI 760 001 (3 Credits)
2023 Fall [1242]

Description
Computational approaches to learning: including inductive inference, explanation-based learning, analogical learning, connectionism, and formal models. What it means to learn. Algorithms for learning. Comparison and evaluation of learning algorithms. Cognitive modeling and relevant psychological results.

Prerequisite(s)
Graduate/professional standing

Instruction Mode
Classroom Instruction

Section Level Com B
False

Department: Computer Sciences
College: Letters and Science

Meeting Time and Location:

2023 Fall [1242]
Term Start Date: Monday, 21-Aug-2023 **Term End Date:** Friday, 29-Dec-2023

ADD TO CALENDAR

Location and Schedule: Mosse Humanities Building 1111 TR 9:30 AM-10:45 AM
CRN: 168961242

Instructor Information:

Instructors

Josiah HANNA
JPHANNA@CS.WISC.EDU

Hadi Khader
HKHADER@WISC.EDU

Adam Labiosa
LABIOSA@WISC.EDU

Instructor Availability and Preferred Contact:

Detailed office hours for the instructor and TAs are available on the course website: https://pages.cs.wisc.edu/~jphanna/teaching/2023fall_cs760/office_hours.html

TA Availability and Preferred Contact :

Detailed office hours for the instructor and TAs are available on the course website: https://pages.cs.wisc.edu/~jphanna/teaching/2023fall_cs760/office_hours.html

How Credit Hours are Met:

How Credit Hours are Met

This class meets for two 75-minute class periods each week over the semester and carries the expectation that students will work on course learning activities (reading, writing, problem sets, studying, etc.) for about 3 hours out of classroom for every class period. The syllabus includes more information about meeting times and expectations for student work.

Course Learning Outcomes (CLOs):

- Identify various types of regularization techniques and their properties
- Implement and analyze a variety of supervised models for classification and regression, including decision trees, instance-based models, naive Bayes, support vector machines, a variety of neural networks, linear and logistic regression, and others
- Identify different aspects of machine learning, including supervised learning, unsupervised learning, and reinforcement learning
- Implement and analyze neural network models, starting with the perceptron, and continuing to multilayer perceptrons, convolutional neural networks, recurrent neural networks, along with deep generative models
- Analyze unsupervised learning techniques for clustering, dimensionality reduction, and latent models

- Implement optimization techniques used in modern machine learning, including gradient descent and stochastic gradient descent
- Apply various concepts and metrics involved in evaluating models: accuracy, F measures, ROC, and precision/recall curves, and implement cross-validation
- Identify classical and modern techniques to improve models or deal with dearth of data: ensemble methods, semi-supervised learning, weak supervision
- Identify and understand how to implement major classes of reinforcement learning algorithms such as value-based, policy gradient, and actor-critic algorithms.
- Be able to analyze the effect of various design decisions (e.g., reward design) and problem features on the performance of reinforcement learning algorithms.

INSTRUCTOR to STUDENT COMMUNICATION

Course Overview:

All course details are available on the course website: https://pages.cs.wisc.edu/~jphanna/teaching/2023fall_cs760/index.html.

Course Website and Digital Instructional Tools:

We will use four main websites for materials, communication, and assignments:

1. Course website: https://pages.cs.wisc.edu/~jphanna/teaching/2023fall_cs760/
2. Gradescope: https://www.gradescope.com/courses/587829
3. Canvas: https://canvas.wisc.edu/courses/360792
    - Do not share materials on canvas outside of class.
4. Piazza: https://piazza.com/class/llcq5n8vttb5xc
    - Access Code: mlfall23
    - Preferred for questions! Sometimes your peers might be able to better answer your questions than the instructor/TA.

Discussion Sessions:

We will use Piazza for asynchronous questions and discussions. Best Piazza practices are as follows:

1. Piazza is preferred for course communication!

2. Please check if someone has posted the same / similar question before you; it's much easier if we build on the thread.
3. Use an informative "Summary" line to help others.
4. You may post and reply anonymously, however, the instructor reserves the right to disable anonymous posting in the case of discourteous behavior.

Required Textbook, Software and Other Course Materials:

Optional and helpful readings will be posted on the course schedule webpage. The following textbooks are optional but may be useful:

- *Machine Learning*, Tom Mitchell.
- *Pattern Recognition and Machine Learning*, Chris Bishop.
- *Machine Learning: A Probabilistic Perspective*, Kevin Murphy.
- *Understanding Machine Learning: From Theory to Algorithms*, Shalev-Shwartz, Ben-David.
- *Deep Learning*, Ian Goodfellow, Yoshua Bengio and Aaron Courville.
- *Foundations of Machine Learning*, Mehryar Mohri, Afshin Rostamizadeh, and Ameet Talwalkar.
- Reinforcement Learning: An Introduction, Richard Sutton and Andy Barto

These textbooks should be available online or at campus libraries with digital access.

Homework and Other Assignments:

Homework assignments will include written problems and sometimes programming. Homework is typically due the minute before class starts on the due date (9:29AM).

Late submissions will not be accepted. You may ask, but it is very unlikely that an extension will be granted.

Assignment grading questions must be raised with the TAs within 72 hours after it is returned. Regrading request for a part of a homework question may trigger the grader to regrade the entire homework and could potentially take points off. Regrading will be done on the original submitted work, no changes allowed.

We will drop the (single) lowest homework score from your final homework average calculation. The drop is meant for emergency situations. We do not provide additional drops, late days, or homework extensions.

All homework assignments must be done individually. Cheating and plagiarism will be dealt with in accordance with university procedures (see the Academic Misconduct Guide for Students). For example, code for programming assignments must not be developed in groups, nor should code be shared. You are encouraged to discuss with your peers, the TA, or the instructor ideas, approaches and techniques broadly, but not at a level of detail where specific implementation issues are described by anyone. If you have any questions on this, please ask the instructor before you act.

If you latex your solutions *entirely*, extra 5 points per homework.

EXAMS, QUIZZES, PAPERS, COURSE SCHEDULE and GRADING

Exams, Quizzes, Papers and Other Major Graded Work:

Exam format: There will be a midterm exam and a final exam. All exams will be conducted in-person.

Midterm Exam: Planned for an evening during the week of October 16th though exact date is subject to midterm exam room scheduling.
Final Exam: Dec 18 from 2:45 PM - 4:45 PM (Room TBD)

Please plan for exams at these times and let us know about any exam conflicts during the first two weeks of the semester. If an emergency arises that conflicts with the exam times, email us as soon as possible. Emergency exam conflicts will be handled on a case-by-case basis. Alternate times will be granted on the basis of how prompt notification is and how unavoidable the conflict is. Please note that final exam times are already available so you should already know if our final exam time conflicts with another one of your courses.

Exam grading questions must be raised with the instructor within 72 hours after it is returned. If a regrade request is submitted for a part of a question on the exam, the grader reserves the right to regrade the entire exam and could potentially take points off.

Course Schedule/Calendar:

See the course webpage: https://pages.cs.wisc.edu/~jphanna/teaching/2023fall_cs760/schedule.html

Note that this schedule (and homework dates) is tentative and subject to change.

Lecture slides are posted before each lecture. Slides currently posted are from the spring semester and will be updated as the semester progresses. Slides should be considered tentative until lecture time.

Grading:

The following weights are used:

- Midterm Exam: 25%
- Final Exam: 25%
- Homework Assignments: 50%

The weights placed on the assignments will be strictly enforced.

ACADEMIC POLICIES and STATEMENTS

Academic Policies:

Syllabus Statements

https://guide.wisc.edu/courses/#SyllabusStatements