COMPSCI 390A Introduction to Machine Learning

Number of Credits: 3
Type of Course and Format: Lecture
Book: None. Readings will be provided using open access digital resources.
Prerequisites: COMPSCI 220 (or COMPSCI 230), COMPSCI 240 (or STAT 515), and Math 233. A grade of C or better is required for all prerequisites.

Course Description

The course provides an introduction to machine learning algorithms and applications. Machine learning algorithms answer the question: "How can a computer improve its performance based on data and from its own experience?" The course is roughly divided into thirds: supervised learning (learning from labeled data), reinforcement learning (learning via trial and error), and real-world considerations like ethics, safety, and fairness. Specific topics include linear and non-linear regression, (stochastic) gradient descent, neural networks, backpropagation, classification, Markov decision processes, state-value and action-value functions, temporal difference learning, actor-critic algorithms, the reward prediction error hypothesis for dopamine, connectionism for philosophy of mind, and ethics, safety, and fairness considerations when applying machine learning to real-world problems.

Learning Objectives

To understand the mathematical representation of machine learning problems and techniques for solving them; to be capable of applying machine learning algorithms *responsibly* to real problems, accounting for issues of safety and fairness; to be prepared for ethical considerations that arise with the use of machine learning; to understand key machine learning concepts including regression, classification, neural networks, and reinforcement learning, among others.

Grading

- Homework Assignments (60%)
- Midterm (15%)

- Final Exam (25%)

The following table will be used for converting numerical grades at the end of the course to letter grades. Minimum values are inclusive, while maximum values are exclusive (except for 100). For example, 90 corresponds to an A, not an A-.

Minimum	**Maximum**	**Letter**
0	50	F
50	55	D
55	60	C–
60	65	C
65	70	C+
70	75	B–
75	80	B
80	85	B+
85	90	A–
90	100	A

Schedule

The course will be divided in to thirds. The list below indicates planned topics, in order. This is subject to alteration, depending on pacing and student abilities.

Part I: Supervised Learning

1. Introduction, gradient descent, and linear regression. [Week 1]
2. Non-Linear regression and stochastic gradient descent. [Week 2]
3. Neural Networks, deep learning, and backpropagation. [Week 3]
4. Classification and over-fitting. [Week 4]

Part II: Reinforcement Learning

1. Introduction and Markov decision processes. [Week 5]
2. State-value functions, action-value functions, and the Bellman equation. [Week 6]
3. Temporal difference learning. [Week 7]
4. Actor-critic algorithms and their relation to stochastic gradient descent. [Week 8]

Part III: Ethics, Safety, Fairness, and Connections to other Areas

1. Connections to psychology and neuroscience. [Week 9]
2. Connections to philosophy. [Week 10]

3. Issues of safety, fairness, accountability, and transparency when applying machine learning to make decisions that impact people. [Week 11]

4. Ethical considerations when applying machine learning. [Week 12]

Part IV: Conclusion This fourth part of the course is not listed in the discussion of the course being broken into thirds because it will likely be a single concluding lecture (perhaps two).

1. A high-level survey of other topics in machine learning, including natural language processing, computer vision, robotics, intelligent tutoring systems, and ecological and social science applications. This includes pointers to other courses for continued learning at UMass. [Week 13]

Accommodation Statement

The University of Massachusetts Amherst is committed to providing an equal educational opportunity for all students. If you have a documented physical, psychological, or learning disability on file with *Disability Services* (DS), you may be eligible for reasonable academic accommodations to help you succeed in this course. If you have a documented disability that requires an accommodation, please notify me within the first two weeks of the semester so that we may make appropriate arrangements.

Class Policies

Attendance: There is no required attendance policy; students who cannot attend class are responsible for any material covered during their absence. Late arrivals must enter the classroom quietly and discreetly.
Exams: The midterm and final exams will be open notes, though no electronic devices will be allowed (except perhaps a calculator).
Collaboration on Assignments: Instructions regarding allowed collaboration and consultation of outside sources will be included with each assignment, and may vary from assignment to assignment.
Late Submissions: Homework submissions will be accepted up to 48 hours after the due time with a 10% penalty if turned in within the first 24 hours and an additional 10% penalty if turned in within the second 24 hours. Submissions will not be accepted later without instructor permission.

Academic Honesty

Since the integrity of the academic enterprise of any institution of higher education requires honesty in scholarship and research, academic honesty is required of all students at the University of Massachusetts. Academic dishonesty is prohibited in all programs of the University. Academic dishonesty includes but is not limited to: cheating, fabrication, plagiarism, and facilitating dishonesty. Appropriate sanctions may be imposed on any student who has committed an act of academic dishonesty. Instructors should take reasonable steps to address academic misconduct. Any person who has reason to believe that a student has committed academic dishonesty should bring such information to the attention of the appropriate course instructor as soon as possible. Instances of academic dishonesty not related to a specific course should be brought to the attention of the appropriate department Head or Chair. Since students are expected to be familiar with this policy and the commonly accepted standards of academic integrity, ignorance of such standards is not normally sufficient evidence of lack of intent. Further details can he found here: `https://www.umass.edu/honesty/sites/default/files/academic_honesty_policy_rev_sen_doc_no16-038a.pdf`.