

# COMPSCI 389 Introduction to Machine Learning

Spring 2024

**Lecture Time and Place:** Tuesdays and Thursdays, 2:30–3:45 in Morrill 2, Room 222.

**Instructor:** Philip Thomas (pthomas@cs.umass.edu)

**Teaching Assistant:** Alexandra Burushkina (aburushkina@umass.edu).

**Number of Credits:** 3

**Type of Course and Format:** Lecture

**Course Webpage:** [https://people.cs.umass.edu/~pthomas/courses/COMPSCI\\_389\\_Spring2024.html](https://people.cs.umass.edu/~pthomas/courses/COMPSCI_389_Spring2024.html)

**Book:** None. Slides and code notebooks will be provided on the course webpage.

**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.

## Office Hours

The teaching assistant (TA) will provide office hours **Tuesdays from 12:00-1:00** and **Thursdays from 4:00-5:00** in room **LGRT T220**.

## Grading

- Homework Assignments (80%).
- Test (20%)

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	60	D
60	63	C-
63	70	C
70	73	C+
73	80	B-
80	83	B
83	90	B+
90	93	A-
93	100	A

## Schedule

The course will be divided *roughly* into thirds. The list below indicates planned topics, in order. This is subject to alteration, depending on pacing and student abilities. Each number corresponds to a week (two lectures).

### Part I: Supervised Learning

1. Introduction, regression, nonparametric methods ( $k$ -nearest neighbors).
2. Parametric methods, linear regression, gradient descent part 1.
3. Gradient descent part 2, basis functions, feature normalization.
4. Perceptrons, artificial neural networks.
5. Backpropagation, classification, over-fitting, and other topics.

### Part II: Reinforcement Learning

1. Introduction, notation, problem formulation
2. Episodes, policy representations, softmax and likely the midterm.
3. MENACE, value functions, temporal difference error.
4. Actor-critics (and their relation to gradient descent), options, and off-policy evaluation.

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

1. Issues of safety, fairness, accountability, and transparency when applying machine learning to make decisions that impact people.
2. Ethical considerations when applying machine learning.
3. Connections to psychology, neuroscience, and philosophy of mind.

**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.

## Inclusivity Statement

We celebrate the diversity in our community and actively seek to include and listen to voices that are often silenced in the computing world. We welcome all individuals regardless of age, background, citizenship, disability, sex, education, ethnicity, family status, gender, gender identity, geographical origin, language, military experience, political views, race, religion, sexual orientation, socioeconomic status, and work experience.

This course may require you to work in groups. As such, we expect that you will observe social decorum at all times when interacting with peers. Please make sure to read the UMass Guidelines for Classroom Civility and Respect: [http://www.umass.edu/dean\\_students/campus-policies/classroom](http://www.umass.edu/dean_students/campus-policies/classroom).

## 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:** Attendance is not required, but students who cannot attend class are responsible for any material covered during their absence.

**Late Submissions:** I trust your judgement when determining whether you need an extension on an assignment. You may take a 7 day extension on any and every assignment with no penalty. Assignments submitted more than 7 days late will not be accepted unless there are extenuating circumstances.

## Academic Honesty

**General Statement:** 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 be found here: [https://www.umass.edu/honesty/sites/default/files/academic\\_honesty\\_policy\\_rev\\_sen\\_doc\\_no16-038a.pdf](https://www.umass.edu/honesty/sites/default/files/academic_honesty_policy_rev_sen_doc_no16-038a.pdf).

**Course-Specific Statement:** For this course the penalty for cheating on a homework assignment, the midterm, or the final exam will be a failing letter grade for the **course**.