COMPSCI 389 Introduction to Machine Learning

Spring 2025

Lecture Time and Place: Tuesdays and Thursdays, 2:30–3:45 in Ag. Engineering Building, Room 119. Instructor: Philip Thomas (pthomas@cs.umass.edu)

Teaching Assistants: Blossom Metevier (bmetevier@umass.edu) and Zhiheng "Andy" Wang (zhihengwang@umass.edu) Number of Credits: 3

Type of Course and Format: Lecture

Course Webpage: https://people.cs.umass.edu/~pthomas/courses/COMPSCI_389_Spring2025.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 actionvalue 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 realworld 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 assistants (TAs) and a UCA will run office hours at the following times. When office hours are canceled, they will attempt to notify you as soon as possible, and will typically provide make-up office hours at a different time.

- Tuesday: 10–12, Andy, Location LGRT T222.
- Wednesday: 3:30–5:30, Blossom, Room LGRT T220.
- Thursday: 11:30–1:30, Andy, Location LGRT T222.
- Thursday (Evening): 7:30pm-8:30pm, Victor, Zoom link: https://umass-amherst.zoom.us/j/8464845765.
- Friday: 10–12, Blossom, Room LGRT T220.

Prof. Thomas strongly encourages questions during lecture, and will provide significant time during each lecture for questions. If you have a question for Prof. Thomas that you do not believe would be appropriate for lecture (e.g., asking about disability accommodations or for an extension on a homework assignment beyond the standard 7 days), you can email Prof. Thomas at pthomas@cs.umass.edu.

Piazza

We will not be using Piazza or any similar system for this course. You are strongly encouraged to attend lecture and ask questions during the time allotted for questions during each lecture. If your questions require more of a back-and-forth discussion, we encourage you to attend office hours.

Grading

- Homework Assignments (75%).
- Test (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, 93 corresponds to an A, not an A-.

Minimum	Maximum	Letter
0	50	F
50	60	D
60	63	C–
63	70	С
70	73	C+
73	80	B–
80	83	В
83	90	B+
90	93	A–
93	100	А

Gradescope

We will use Gradescope for assignments (submission, grading, and regrade requests). You will receive an invitation to Gradescope during the first week or two of classes. Regrade requests for homework assignments will only be considered if they are submitted within 1 week of the grade for the assignment being announced.

Schedule

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

Part I: Supervised Learning

- 1. Introduction, supervised learning, data sets, nonparametric methods (k-nearest neighbors).
- 2. Model and algorithm evaluation, validation sets.
- 3. Linear regression, objective functions, and optimization
- 4. Gradient descent, data cleaning, neural networks, automatic differentiation
- 5. PyTorch, overfitting, classification
- 6. Generative AI, supervised learning review

Part II: Reinforcement Learning

- 1. Introduction, notation, Markov decision processes (MDPs)
- 2. MENACE and REINFORCE
- 3. Value Functions, temporal difference learning, and actor-critics
- 4. Actor-critics, survey of advanced reinforcement learning topics

Test: The test will cover supervised learning and reinforcement learning, and will likely take place near the end of the part of the course focusing on reinforcement learning.

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 Time-permitting, a lecture near the end of the course will provide 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.

Class Policies

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

Lecture Recording: Lectures will not be recorded.

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: For this course the penalty for cheating on any graded material (e.g., homework assignment or test) will be a failing letter grade for the **course**.

University policies regarding Accommodations, Academic Honesty, and Title IX

University policies regarding accommodations, academic honesty, and Title IX apply to all courses. The policies can be found at: https://www.umass.edu/senate/book/required-syllabus-statements. Note that I am a "Responsible Employee" under the Title IX definition.