

Syllabus

Course Number: COMPSCI 389

Course Title: Introduction to Machine Learning

Term: Spring 2026

Number of Credits: 3

Number of Instructional Contact Hours Per Week: 2.5

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.

Instructor: Professor Philip S. Thomas, (pthomas@cs.umass.edu)

Teaching Assistants: Renhao “Norman” Zhang (renhaozhang@umass.edu)

Undergraduate Course Assistants: TBD (TBD@umass.edu) and TBD (TBD@umass.edu)

Office Hours:

- **Monday:** TBD
- **Tuesday:** TBD
- **Wednesday:** TBD
- **Thursday:** TBD
- **Friday:** TBD

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.

Student 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.

Delivery Mode: In-person, Tuesdays and Thursdays, 1:00–2:15 in the Computer Science Building, Room 142.

Course Materials:

- **Course Webpage:** https://people.cs.umass.edu/~pthomas/courses/COMPSCI_389_Spring2026.html
- **Book:** None. Slides and code notebooks will be provided on the course webpage.
- **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.
- **Piazza:** We will not be using Piazza or a similar system. We encourage you to ask questions during lecture and office hours.
- **Lecture Recordings:** Lectures will not be recorded.

- **Practice test:** A practice test will not be provided. **However**, at least 80% of the points on the test will be for slight variations of questions asked on homework assignments, and so studying the homework assignments provides effective test preparation.

Individual Course Requirements and Grade Weights:

- **Participation:** Not required.
- **Attendance:** Not required.
- **Grade Weights:**
 - Homework Assignments (70%).
 - Test (30%)

- **Grading Scale:** 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	C
70	73	C+
73	80	B-
80	83	B
83	90	B+
90	93	A-
93	100	A

Course 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 (April 23):** The test will cover supervised learning and reinforcement learning. It will be closed-book, closed-notes, and electronic devices will not be permitted. At least 80% of the points on the test will be for slight variations of questions asked on homework assignments, and so studying the homework assignments provides effective test preparation.

- **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.

Attendance Policy: Attendance is not required.

Late or Make-Up Work Policies:

- **Late Homework Assignment Policy:** I trust your judgment 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.
- **Regrade Request Policy:** Regrade requests for homework assignments will only be considered if they are submitted within 1 week of the grade for the assignment being announced.
- **Make-Up Test Policy:** If you have a conflict with the date of the test, email Prof. Thomas at least 1 week prior to the test, and we will work out an alternative.

Required Syllabus Statements: University policies regarding Accommodations, Academic Honesty, and Title IX, apply to all courses. Note that these statements differ depending on whether the instructor (Prof. Thomas in this case) is a “Responsible Employee” (i.e., mandatory reporter). Prof. Thomas is a responsible employee, and so the appropriate statements can be found here: <https://www.umass.edu/senate/book/responsible-employee-required-syllabus-statements>.

Course-Specific Academic Honesty Policy: For this course the penalty for cheating on any graded material (e.g., homework assignment or test) will be a grade of zero for all graded material impacted by the cheating (e.g., if one cheats on a homework assignment, that assignment will be graded as a zero).

Optional Syllabus Statements:

- **Pronoun statement:** My pronouns are he/him.
- **Land acknowledgment:** I will not be including the standard land acknowledgment statement. In my view, as long as the university bears the name of Jeffrey Amherst—a figure associated with advocating genocide against Native peoples—such a statement would be disingenuous. For historical context, see: https://people.umass.edu/derrico/amherst/lord_jeff.html
- **Generative AI Policy:** The use of generative AI is *not* allowed on part of the homework assignments, but is *encouraged* on other parts. See the instructions for each homework assignment for details. The use of electronic devices (including generative AI systems) is *not* allowed on the test.