



COMPSCI 389

Introduction to Machine Learning

Course Introduction

Prof. Philip S. Thomas (pthomas@umass.edu)

“Phil”

Course Materials

- Course Website:

https://people.cs.umass.edu/~pthomas/courses/COMPSCI_389_Spring2026.html

or

<https://people.cs.umass.edu/~pthomas/> → Teaching → COMPSCI 389 → Spring 2026

- The website contains links to the syllabus, slides, code notebooks, and assignments.

Assignments

- We will have *approximately 5* homework assignments.
- Assignments will be a mix of written questions (derivations, short answer, and essays) and programming assignments.
- There are (almost) no restrictions on collaboration or the use of tools (online or offline)!
 - More details on this will be provided in the assignment instructions.

Test

- There will be one in-class test on April 23.
- The test will **not** be open notes.
- The questions on the test will be *mostly* questions from the homework assignments with numbers and values changed.
- $\leq 20\%$ of the points will be new questions.

Late Submission Policy

- Assignments will have a specific due date and time, X.
- If there are special circumstances that justify your having an extension, then you may submit the assignment by $(X + 7 \text{ days})$ with no penalty.
 - You do not need to request this extension – I trust your judgement.
- I encourage you to still submit by date/time X.
 - You will only have one assignment at a time, but the next assignment may be given any time after the due date X (e.g., before $X + 7 \text{ days}$).
 - The assignments are timed to be related to the material that was discussed recently. It will be easier for you to complete the assignments when they are assigned rather than on day X-1 or X+6.
- We aim to post grades for assignments by $(X + 14 \text{ days})$.
- If there are special circumstances that justify your having an extension beyond $(X + 7 \text{ days})$, please e-mail me.
 - Note that your grade for the assignment may not be posted until the end of the semester.
 - Without prior approval, assignments submitted after $X + 7 \text{ days}$ will receive a grade of zero.

Attendance Policy

- Attendance is optional.

Recording and Slides Policy

- Lectures will **not** be recorded.
- Slides (and code notebooks) will be posted.
 - If you miss a lecture, review the slides, talk with a peer, and ask questions during office hours.

Grading

- Let H be your normalized homework assignment grade (0.0–1.0)
 - Extra-credit on the homework assignments cannot cause $H > 1.0$.
- Let T be your normalized grade on the test (0.0–1.0)
 - Extra-credit on the test cannot cause $T > 1.0$.
- Your complete number grade is $C = 0.7H + 0.3T$
- Your letter grade is computed from C using this table:

A	$C \geq 0.93$
A-	$0.93 > C \geq 0.9$
B+	$0.9 > C \geq 0.83$
B	$0.83 > C \geq 0.8$
B-	$0.8 > C \geq 0.73$

C+	$0.73 > C \geq 0.7$
C	$0.7 > C \geq 0.63$
C-	$0.63 > C \geq 0.6$
D	$0.6 > C \geq 0.5$
F	$0.5 > C$

Piazza

- This course will **not** use Piazza.

Office Hours

- Teaching Assistants (TAs):
 - Renhao “Norman” Zhang (renhaozhang@umass.edu)
- Undergraduate Course Assistants (UCAs):
 - Olive Do (ydo@umass.edu) and Rohit Joshi (rjoshi@umass.edu)
- Office Hours
 - **Monday:** TBD
 - **Tuesday:** TBD
 - **Wednesday:** TBD
 - **Thursday:** TBD
 - **Friday:** TBD

Cheating Policy

- The penalty for cheating will be a grade of zero for all graded material impacted by the cheating.

Disability Services

- 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.

Course Content

- **Part I: Supervised Learning (~70%)**
 - Regression and classification, parametric and nonparametric methods, train/test splits, cross-validation (k-fold), loss functions, gradient-based methods, data processing, overfitting, generalization, epistemic and aleatoric uncertainty, generative methods, language models, and other topics.
- **Part II: Reinforcement Learning (~20%)**
 - Problem formulations, MENACE, value functions, temporal difference error, actor-critics, options, and off-policy evaluation.
- **Part III: Ethics, Safety, Fairness, and Connections to other Areas (~10%)**
 - Ethics, safety, fairness, accountability, and transparency when applying machine learning to make decisions that impact people.
 - Connections to psychology, neuroscience, and philosophy.
- **Part IV: Conclusion (~rounding error)**
 - Survey of other topics

Target Audience

- Students with no prior knowledge about machine learning.
 - COMPSCI 589 is intended for students with some background in ML.
- Students who are familiar with basic concepts from calculus (derivatives, gradients, direction of steepest ascent, etc.)
 - Unlike COMPSCI 589, this course assumes no background in linear algebra.
- Students who have experience programming in python.
- Students wondering “is this a topic I might find interesting?”
 - Students who already know they want to focus on machine learning might consider COMPSCI 589 instead.

Language and IDE

- This course will use Python and Jupyter Notebooks.
- You can find instructions for installing Python and setting Visual Studio Code up for Jupyter Notebooks on the course webpage in the file:

Jupyter Notebooks Setup.pdf

- You can find a brief introduction to Python and Jupyter Notebooks on the course webpage in the file:

Jupyter Notebook Introduction.ipynb

- We will *briefly* review these documents now.