COMPSCI 589: Machine Learning
Syllabus

Spring 2026

Instructor:	Bruno Castro da Silva	**Time:**	Tue & Thur from 10:00 to 11:15 am
Email:	bsilva@cs.umass.edu	**Place:**	***In person students*** Computer Science Labs E110
			Students in the online section Please visit our online platforms

1 About the course

This course will introduce core machine learning (ML) models and algorithms for classification, regression, clustering, and dimensionality reduction. On the theory side, the course will cover the mathematical foundations underlying the most commonly used ML algorithms, focusing on understanding models and their relationships. On the applied side, the course will focus on effectively using ML methods to solve real-world problems, with an emphasis on model selection, regularization, experiment design, and presentation and interpretation of results. Assignments will consist of mathematical problems and implementation tasks.

Broad topics covered in this course will include classification algorithms in general, decision trees, random forests, probabilistic models, Naive Bayes methods, ensemble algorithms, gradient-based techniques, linear regression, logistic regression, neural networks, convolutional neural networks and deep learning, unsupervised learning, clustering algorithms, and dimensionality reduction techniques.

Students in the in-person section will meet twice a week. Students in the online section will follow the course asynchronously, through short videos to be watched each week. In both sections, the course includes hands-on projects (done both in groups and individually) and online discussions to solidify understanding.

In this course, each voice in the classroom has something of value to contribute. Please take care to respect the different experiences, beliefs, and values expressed by students and staff. My colleagues and I support UMass' commitment to diversity, and welcome individuals regardless of age, background, citizenship, disability, sex, gender, gender identity, sexual orientation, education, ethnicity, family status, geographical origin, language, military experience, political views, race, religion, socioeconomic status, and work experience.

2 Learning Objectives

1. Distinguish between the main classes of machine learning algorithms and identify the most appropriate one for solving a particular problem.
2. Explain how key algorithms in the field work and under what conditions.
3. Learn to implement and apply key machine-learning algorithms to solve real-world problems.
4. Learn to solve issues that may arise when deploying particular algorithms to tackle real-world problems.
5. Identify the changes that may be required for an algorithm to be applicable in novel settings and conditions.

3 Course Platform

This course's syllabus will be hosted here. Homework assignments and other material will be posted on Canvas:

- Canvas for students enrolled in the in-person section.
- Canvas for students enrolled in the online section.

In-person lectures will be recorded using Echo360. Students in the online section will have dedicated pre-recorded videos covering each week's topics. All recordings and videos, along with PDF slides, will be available on Canvas.

4 Discussion Board

We will use CampusWire *(access code: 0947)* as our interactive discussion board. This will be a dynamic platform for ongoing conversations throughout the course. On Campuswire, you can reach out to instructors with questions, participate in discussions, and engage with your peers. This will be the main platform through which students can collaborate, share insights, and debate different strategies for solving real-world problems with machine learning. It is important that you actively participate in the discussions, as such continuous interaction is key to enhancing your learning experience, allowing you to explore various approaches and deepen your understanding of course concepts as you work on assignments together.

5 Class

In-person classes will be held on Tuesdays and Thursdays from 10:00 am to 11:15 am in Computer Science Labs E110. Students enrolled in the online section of the course should refer to our online platforms.

6 Textbooks

The course has no mandatory textbook. However, you may find the following (freely available) books helpful:

- Machine Learning: a Probabilistic Perspective, by Kevin Patrick Murphy.
- The Elements of Statistical Learning, by Hastie, Tibshirani and Friedman.
- An Introduction to Statistical Learning, by James, Witten, Hastie and Tibshirani.

7 Required background

While this course has an applied focus, it still requires an appropriate mathematical background in probability and statistics, multivariate calculus, linear algebra, and programming. The official prerequisites for undergraduate students can be found here. Graduate students can check the descriptions for these courses to verify that they have a sufficient mathematical background for COMPSCI 589. The following references can provide a useful review:

- Probability Theory.
- Linear Algebra and Matrix Calculus.
- Matrix Cookbook.
- Optimization: Any calculus textbook.

8 Course Schedule

- **Week 1**. Introduction to machine learning and supervised learning. k-nearest neighbors
- **Week 2**. Decision Trees
- **Week 3**. Probabilistic classifiers. Naive Bayes
- **Week 4**. Document classification. Bernoulli and Multinomial Naive Bayes
- **Week 5**. Evaluation metrics
- **Week 6**. Ensemble methods. Random forests. ROC analysis
- **Week 7**. Bias-variance trade-off. Bagging and boosting
- **Week 8**. Gradient descent methods. Linear regression
- **Week 9**. Logistic regression
- **Week 10**. Introduction to neural networks. Backpropagation. Multiple outputs. Regularization
- **Week 11**. Vectorization. Mini-batch training. Heuristics to accelerate gradient descent
- **Week 12**. Convolutional neural networks. Reinforcement learning and Markov decision processes
- **Week 13**. Q-Learning and deep reinforcement learning. Philosophy of AI. Future and challenges

9 Attendance policy

Students taking this course in person should make every effort to attend and actively participate in each class meeting. However, you are free to choose not to attend "regular" class meetings (i.e., classes where no assignments or exams are given) *if the circumstances warrant*. We do understand that all students are capable of autonomously deciding how to best allocate their time. However, notice that you are responsible for the classes you miss. In the event that you miss class, it is your responsibility to obtain the missed notes from a classmate and for picking up any handouts distributed in class by the instructor. The instructor will not provide class notes or condense the entire contents of a missed class period in an email message or a five-minute summary. *Important: please see Section 14 of this document for more details on this course's late policy regarding late homework submissions and missed exams.*

10 Teaching Assistants

The teaching assistants (TAs) this semester will be:

- Ke Xiao (kexiao@umass.edu)
- Sunjae Kwon (sunjaekwon@umass.edu)
- Nana Akua Agyemang Sereboo (nagyemangser@umass.edu)
- Dhawal Gupta (dgupta@umass.edu)
- John Raisbeck (jraisbeck@umass.edu)

11 Office Hours

Office hours will be held according to the following schedule (starting on 02/05), except *(1)* on holidays; *(2)* if the UMass official schedule follows a different day of the week; or *(3)* when noted otherwise.

→ TBD. Office hours details will be posted early in the second week of classes.

	Monday	**Tuesday**	**Wednesday**	**Thursday**	**Friday**
8:50am–9:50am					
9:00am–10:00am					
11:15am–12:30pm		Bruno		Bruno	
12:00pm–1:00pm					
1:00pm–2:00pm					
2:30pm–3:30pm					

If one of the TAs or the instructor needs to reschedule office hours, an official announcement will be made on Canvas.

▷ For students enrolled in the in-person section of the course, **all office hours will take place in person unless noted otherwise**.

- If an in-person student cannot attend in person for a specific reason (for example, illness or travel), they should arrange with the TA *in advance* to schedule a time to receive help via Zoom. The TA will then provide a Zoom link for that session.

▷ For students enrolled in the remote section of the course, **all office hours will take place on Zoom at the times listed in the table above**.

- Students in the online section of the course can join office hours via this Zoom link.

Locations:

- TA's in-person office hours:
  - **TBD**.
  - **TBD**.
- Prof. da Silva's office hours will be held immediately after each lecture. He will stay in the classroom for 1 hour and 15 minutes after class to answer questions and help students as needed.

12 Grading

Your grade will have three components:

1. **Homework Assignments:** There will be frequent homework assignments, both written and programming. All assignments will have equal weight.
   - In-person section: **50**% of final grade.
   - Online section: **70**% of final grade.
2. **Midterm exam**.
   - In-person section: **30**% of final grade.
   - Online section: *Students in the online section will not have a midterm exam.*
3. **Project**: Given the wide range of real-world applications of machine learning, it is important to learn how to implement, fine-tune, and deploy these algorithms in practice. More details will be shared when the project is assigned, after we have covered the most relevant course material.
   - In-person section: **20**% of final grade.
   - Online section: **30**% of final grade.

Letter grades will be assigned using the following scale:

- $[90\%, 100\%]$: **A- or A**
- $[75\%, 90\%)$: **B-, B, or B+**
- $[65\%, 75\%)$: **C-, C, or C+**
- $[55\%, 65\%)$: **D or D+**

Course grades will be curved only in students' favor (that is, the thresholds above may be lowered, but a cumulative grade of 90% will still receive at least an A-). Some extra-credit opportunities may be offered. Your grade may be reduced at the instructor's discretion in cases of inappropriate behavior, including academic dishonesty.

13 Re-grading policy

Errors in grading assignments and exams can occur despite the best efforts of the course staff. If you believe you have found a grading error, complete an online regrade request form via Gradescope. Regrade requests must be submitted no later than one week after the assignment is returned. Note that regrading may result in your original grade being *increased* or *decreased* as appropriate.

14 Late Policy

- Deadlines in this course are **strict**. A submission one minute after the deadline will receive zero credit. You are strongly encouraged to submit hours before any deadline.
- That said, to allow some flexibility to complete assignments (*homeworks only*) given other constraints, you have a total of *seven* free late days that you can choose to use when submitting a homework. You will be charged one late day for handing in an assignment within 24 hours after it is due, two late days for handing in an assignment within 48 hours after it is due, etc. Your assignment is considered late if either the written or code portions are submitted late. The late homework clock stops when both the written and code portions are submitted. After you have used up your late days, late homework will not count for credit except in special circumstances (e.g., illness documented by a doctor's note).
- Extensions may be granted, for example, if you have a medical emergency and bring proof of such to the instructor before final grades for the given assignments are computed. In any other case (unless those covered by the University's Academic Regulations), missing a deadline will result in a zero for the assignment.
- All exams must be taken at the scheduled time unless *(1)* there is a documented conflict and arrangements have been made with the instructor before the exam; or *(2)* you have a medical emergency and you bring proof of such to the instructor before final grades for the given exam are computed. In any other case (unless those covered by the University's Academic Regulations), missing an exam will result in a grade of "F" for that exam.

15 Cheating

- Cheating will not be tolerated. Assignments may include instructions about what forms of collaboration are allowed, if/when relevant.
- Copying answers or code from external sources (books, web pages, etc.), from other students, or from solutions to assignments from previous years is *always* considered cheating. Note that, according to the new UMass Academic Honesty Policy, the use of AI text generators (such as **ChatGPT**) is **prohibited**. To emphasize: no detectable copying is acceptable, even, e.g., copying a single sentence from an outside source. Sharing your code or solutions with other students is also considered cheating.
- The College of Information and Computer Sciences explicitly forbids any redistribution of CICS course materials, including posting them publicly online and sharing student solutions to homework assignments, projects, or exams, without the express written consent of the instructor. Violations of this policy will be deemed instances of "facilitating dishonesty" (since a student making use of such materials would be guilty of plagiarism) and therefore may result in charges under the Academic Honesty Policy.
- 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.
- **All instances of cheating will be reported to the university's Academic Honesty Board.** Any detected cheating will result either *(i)* in a grade of -100% on the assignment for all students involved (negative credit); or *(ii)* a grade of "F" in the course. The instructor will decide at their discretion which of these possible resolutions is more appropriate.

16 Pass/Fail & SAT/Fail

- If you are an undergraduate student, Pass/Fail is requested through the university.
- If you are a graduate student, at some time near the end of the semester (likely around the last day of class), you will be given the option to take the class SAT/Fail rather than for a letter grade. If you plan to take the course SAT/Fail, keep an eye out for an email (or a message on Canvas) from me around the end of the semester with instructions for requesting SAT/Fail. If you elect SAT/Fail, you will earn a SAT grade if your letter grade would have been a C or higher, and you will receive an F if your letter grade would have been lower.
- The above conditions *do not* hold for students with an academic honesty violation. In these cases, the requests described in this section are disallowed and/or un-approved.

17 Disability Services

The University of Massachusetts 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, you may be eligible for academic accommodations to help you succeed in this course. If you would like to register with Disability Services, please visit their website or their office (161 Whitmore Administration Building; phone (413) 545–0892). Finally, if you have a documented disability that requires an accommodation, please notify me within the first two weeks of the semester so that we can make appropriate arrangements.

18 Title IX

In accordance with Title IX of the Education Amendments of 1972, which prohibits gender-based discrimination in educational settings that receive federal funds, the University of Massachusetts Amherst is committed to providing a safe learning environment for all students, free from all forms of discrimination, including sexual assault, sexual harassment, domestic violence, dating violence, stalking, and retaliation.

- This includes interactions in person or online through digital platforms and social media. Title IX also protects against discrimination on the basis of pregnancy, childbirth, false pregnancy, miscarriage, abortion, or related conditions, including recovery.
- There are resources on campus to support you. For a summary of confidential and non-confidential options, please see the UMass Title IX resources page.
- You can seek support regardless of whether you decide to make a formal report. If you need immediate support, you are not alone. Free and confidential support is available 24 hours a day, 7 days a week, 365 days a year at the SASA Hotline: (413) 545-0800.