COMPSCI 589 – Machine Learning
Syllabus

Spring 2025

Instructor:	Bruno Castro da Silva	**Course Platform:**	Canvas
Email:	bsilva@cs.umass.edu	**Online / Asynchronous**	

1 About the Course

This course will introduce core machine learning 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 machine learning algorithms, focusing on understanding models and their relationships. On the applied side, the course will focus on effectively using machine learning 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. The course includes asynchronous lectures in the form of short videos students should watch each week, hands-on projects (to be done in groups and individually), and online discussions to solidify understanding.

Broad topics covered in this course will include classification algorithms in general, decision trees, random forests, probabilistic models, Naive Bayes methods, various ensemble meta-algorithms (such as bagging and boosting), gradient-based techniques, linear regression, logistic regression, neural networks, convolutional neural networks and deep learning, unsupervised learning and clustering algorithms, k-means, hierarchical clustering, and dimensionality reduction techniques.

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 involved in this course. 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

All course content, including instructional videos, homework assignments, and reading materials (lecture notes and slides), will be hosted on Canvas. Log in to Canvas using your UMass ID and follow the instructions to find the learning resources for this course.

4 Discussion Board

We will use CampusWire (*access code 8326*) 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 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.

6 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 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.

7 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.

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 Teaching Assistants

- **To be announced.**

10 Office Hours

- **To be announced.**

11 Grading

Your grade will have two components:

1. **Homework Assignments** (70%): There will be frequent homework assignments, both written and programming. All assignments will have an equal weight.

2. **Project** (30%): Given the wide range of possible real-life applications of machine learning, it is critical that we study how to implement, fine-tune, and deploy these algorithms in practice. Further details will be available when the project is assigned, after the most relevant course material has been covered.

A cumulative grade in $[90\% - 100\%]$ will be an A- or A, $[75\%, 90\%)$ will be a B-, B, or B+, $[65\%, 75\%)$ will be a C-, C, or C+, and $[55\% - 65\%)$ will be a D or D+. Course grades will be curved only in students' favor (that is, these thresholds may be lowered, but a grade of 90% will not be lower than an A-). **Some extra credit opportunities may be given. Your grade may be reduced by any amount at the instructor's discretion due to inappropriate behavior, such as academic dishonesty.**

12 Re-grading policy

Errors in grading assignments and exams can occur despite the best efforts of the course staff. If you believe you've found a grading error, complete an online re-grade request form via Gradescope. Re-grade requests must be submitted no later than one week after the assignment is returned. Note that re-grading may result in your original grade increasing or decreasing as appropriate.

13 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.
- Having said that, 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 if you have a medical emergency and you 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.

14 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 unapproved.

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 (including publicly available posting on an internet site) of any CICS course materials (including student solutions

to course assignments, projects, exams, etc.) without the express written consent of the instructor of the course from which the materials come. 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.