COMPSCI H589: Honors Colloquium for Machine Learning

Syllabus

Spring 2026

Instructor:	Bruno Castro da Silva	**Time:**	Fridays from 11:15 am to 12:05 pm
Email:	bsilva@cs.umass.edu	**Place:**	LGRC, Room A310

1 About the course

This colloquium will complement and extend the primary Machine Learning course (COMPSCI 589) by exploring more advanced topics, with students reading, presenting, and discussing foundational and recent research papers from the *reinforcement learning* literature. Over the course of the semester, students will lead short seminars and group discussions, and will also have the opportunity to implement and experiment with key reinforcement learning algorithms.

Broad topics covered in this course will include Markov decision processes and a range of reinforcement learning algorithms, including model-free, model-based, batch and online methods, value function-based approaches, actor-critic methods, and policy gradient algorithms.

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 Course Platform

You will find assignments, slides, and other material in Canvas. This course's syllabus is also hosted here.

3 Class

Classes will be held on Fridays from 11:15 am to 12:05 pm in the Lederle Graduate Research Center (LGRC Lowrise), room A310.

4 Textbook

This course will be primarily based on the *first* edition of Sutton and Barto's book, *Reinforcement Learning: An Introduction*. It is available on Amazon here and can also be accessed for free online here. For a more in-depth treatment of some topics in reinforcement learning, students may refer to the *second* edition of the same book. It is available on Amazon here and freely available online here.

Although these books provide an excellent introduction to the topic (and I encourage purchasing a copy if you plan to study reinforcement learning), owning them is not required.

5 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 following references can provide a useful review:

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

6 Attendance policy

This course is designed to be interactive, and it depends on active participation from all members of the class. Your presence is essential to the quality of the course experience and to your individual success. Therefore, regular attendance is expected and considered mandatory. *Important: please see Section 9 of this document for details on the course's late assignment policy.*

7 Grading

Your grade will have three components:

1. **Student-led seminars and group discussions** (50%): Each student will be expected to lead short in-class seminars throughout the semester. Each seminar will consist of a presentation lasting approximately 30–35 minutes, during which the student will present and discuss topics covered in the assigned sections of the main textbook.
2. **Programming assignments** (45%): As reinforcement learning continues to evolve from an academic subject into a practical tool used in real-world applications, it is essential to understand how to implement, fine-tune, and evaluate these algorithms. Students will implement a few key RL algorithms and assess their performance on standard benchmark problems. We anticipate a total of 3–4 programming assignments, each carrying equal weight.
3. **Participation** (5%): Students are expected to attend class each week, stay for the full session, and actively participate in discussions and activities.

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 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.**

8 Office Hours

Given the nature of this colloquium, office hours should be scheduled by students on an as-needed basis by directly contacting the instructor via email. The instructor will also hold brief office hours immediately after each lecture.

9 Late Policy

- To allow some flexibility in completing programming assignments, you have a total of *five* free late days that you can 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 it in within 48 hours, and so on. An assignment is considered late if either the written or code portion is submitted late. The late clock stops when both the written and code portions have been 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).
- All presentations, talks, and group discussions led by a student must take place at the scheduled time unless *(1)* there is a documented conflict and arrangements have been made with the instructor in advance; or *(2)* you have a medical emergency and provide proof to the instructor. In any other case (unless covered by the University's Academic Regulations), missing an activity that the student is responsible for leading will result in a grade of "F" for that assignment.

10 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.

11 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.

12 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.

13 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.