

COMPSCI H589: HONORS COLLOQUIUM FOR MACHINE LEARNING SYLLABUS

Spring 2022

| | | | |
|--------------------|--|---------------|---------------------------------|
| Instructor: | Bruno Castro da Silva | Time: | Thursdays from 1:00 to 1:50pm |
| Email: | bsilva@cs.umass.edu | Place: | Computer Science Bldg., Rm. 142 |

1 About the course

This colloquium will enrich the primary course (COMPSCI 589) by focusing on reading, presenting, and discussing foundational and recent research papers from the *reinforcement learning* literature. Students will lead short seminars and group discussions over the course of the semester. They 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 various types of reinforcement learning algorithms: model-free, model-based, batch/online, value function-based, actor-critics, and policy gradient algorithms.

2 Website

You will find assignments, slides, and other material on [Moodle](#). This course's syllabus is also hosted [here](#).

3 Class

Classes will be held on Thursdays from 1:00 to 1:15pm in the Computer Science building, room 142.

4 Book

This course will be primarily based on the *first* edition of Sutton and Barto's book, *Reinforcement Learning: An Introduction*. It can be found on Amazon [here](#). It is also available for free online [here](#). A more in-depth discussion of some RL topics can be found on the *second* edition of Sutton and Barto's book, *Reinforcement Learning: An Introduction*. It can be found on Amazon [here](#). It is also available for free online [here](#).

Although these books are a fantastic introduction to the topic (and I encourage purchasing a copy if you plan to study reinforcement learning), owning them is not a requirement.

5 Required background

While this course has an applied focus, it still requires 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 assumes an interactive approach in its structure and presentation, which requires engaged participation from all members of the class. Your presence is essential to the liveliness of this course and to your individual success in it. Therefore, regular attendance is expected and considered mandatory. *Important: please see Section 10 of this document for more details on this course's late policy regarding late assignment submissions.*

7 Grading

Your grade will have three components:

1. **Student-led seminars and group discussions (65%):** Each student will be expected to lead short seminars, in class, throughout the semester. Each seminar will consist in a presentation lasting approximately 30-35 minutes, where the student will present and discuss topics covered in assigned sections of the main textbook.
2. **Programming assignments (25%):** As reinforcement learning transitions from an academic curiosity to practical tools that you may use in your professional lives, it is critical that we study how to implement, fine-tune, and deploy these algorithms in practice. Students will have to implement a few key RL algorithms and evaluate their performances when tested on standard benchmark problems. We anticipate a total of 3-4 programming assignments. All assignments will have an equal weight.
3. **Participation (10%):** Students are expected to come to class each week, stay for the entire class session, and to be engaged in the class activities and discussions.

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.**

8 Office Hours

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

9 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 Moodle) 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.

10 Late Policy

- To allow some flexibility to complete programming assignments given other constraints, you have a total of six free late days. 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 assignment clock stops when both the written and code portions are submitted. After you have used up your late days, late assignments will not count for credit except in special circumstances (i.e., illness documented by a doctor's note). *If you do not hand in an assignment at all, this will count as using all six late days.*
- All presentations, talks, and group discussions being 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 before the event; or (2) you have a medical emergency and you bring proof of such to the instructor. In any other case (unless those covered by the [University's Academic Regulations](#)), missing an activity that the student should be leading will result in a grade of "F" for that assignment.

11 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.

12 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. 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.

13 COVID-19 and Face Covering Policy

Students and course staff in COMPSCI H589 are expected to do their part to slow the spread of COVID-19 and minimize risk of illness for all community members. It is important to remember that community members may be vulnerable or live with vulnerable individuals.

Students and course staff in COMPSCI H589 must comply with all [university policies](#) regarding COVID-19, including self-isolating (don't come to class) if you test positive, have COVID-19 symptoms, or are in contact with someone who tests positive and meet the criteria for self-isolation (see [FAQs](#)). If you are in any doubt, please do not come to class—you can request an excused absence from any course meeting by sending the instructor an email; you will not be penalized for any missed quizzes, discussion exercises, etc.

As of the start of the Spring 2022 semester, everyone must wear a high-grade mask (such as KN95, KF94, or N95) to all indoor course meetings, including lectures, discussion sections, exams, and office hours. For more details, see the [face covering FAQ for the start of the Spring 2022 semester](#). The only exceptions to the face-covering policy are:

- Employees who receive a medical exemption through the Accessible Workplace Office (see [face-covering FAQ](#)).
- Students who receive a medical exemption through the Disability Services Office (see [face-covering FAQ](#)). If you have a valid medical exemption, you must notify the instructors and present documentation prior to attending a meeting without a face covering.

If a student in COMPSCI H589 does not comply with public health protocols (e.g. does not wear a mask), the course staff will first remind them of the protocols and their importance, and ask them to comply.

If a student refuses to wear a mask and does not comply after you have reminded them, the staff will ask them to leave the classroom. The instructor will then submit a referral to the [Student Conduct and Community Standards Office](#) (SCCS) in the Dean of Students Office. The SCCS will follow up with the student as appropriate and according to policies outlined in the Classroom Civility and Respect policy as well as the Code of Student Conduct. If a student refuses to leave the classroom, the course staff may elect to end the course meeting to allow other students and staff to leave, and then resume the course meeting on Zoom after a short delay (if logistics permit).