Decision Making & Reinforcement Learning

CS 695 | Fall 2023 | Professor Gregory J. Stein

Course Information

CS 695
3 Credits
Office Hours: TBD, probably Friday 9:30–11AM

Instructor Information

Prof. Gregory J. Stein
He, Him, His
gjstein@gmu.edu

Course Description

This course is a survey of planning algorithms and the fundamentals of reinforcement learning. It focuses on applications of particular relevance to robot decision-making and game playing. Topics will include heuristic and constraint-based search, Markov Decision Processes, game playing (including Alpha-Beta pruning and Monte Carlo Tree Search), task planning, and navigation under uncertainty and both core elements of Reinforcement Learning (e.g., TD methods) and modern progress and challenges, emphasizing the impact of deep learning and foundation models. When possible, the course will also include applications of machine learning to planning, which has facilitated significant recent advances in the state-of-the-art. Other topics pertaining the complexities of real-world agents (e.g., filtering for noisy perception) may also be discussed. By the end of the term students will understand both recent advances and open problem in this challenging application domain.

This course is at the introductory graduate level and aims to introduce students to the various topics to which this course pertains (as relevant to all manner of robotics and decision-making systems) through hands-on programming projects. By the end of the course, students should be familiar with the standard tools used by the community and be able to comprehend research papers in the field.

See the detailed class schedule below for more details on what will be covered during the course.

Course Structure

Ensuring that the students are engaging with the material and with one another can be difficult. As such, most lectures will be broken into subparts: each class will consist of two ~1 hour 10 minute lecture “modules,” occasionally including a 10 minute breakout session during which students will collaborate to solve short programming assignments related to the course material. Students should bring a laptop to class if possible so that they can follow along with group or collaborative

coding sessions. Remaining class time, if it exists, will consist of a shorter lecture covering a related topic or research area.

Prerequisite Knowledge Areas:

- Algorithms and Data Structures
- Artificial Intelligence
- Linear Algebra and Calculus
- [Deep Learning is fairly helpful but not required]

In addition, students are expected to have familiarity with the Python programming language (including numpy) in which all programming projects and assignments will be done. There may be some very light programming in C++ or a lisp-like language called PDDL, though I will provide you with starter code when necessary; I do not expect students to be fluent in either.

Assignments & Grading

The grading in this course will come from three main components:

- 5 programming assignments (50%) [10% each]
- 5 take home "quizzes" (20%) [5% each, since lowest quiz dropped]
- 1 final project (30%)
- Participation (0%)

All assignments will be turned in via Blackboard.

Programming Assignments

In each programming assignment, students will be expected to implement some of the algorithms we will discuss during class. Assignments will be given in the form of Jupyter Notebooks and will often include some partially written code for students to complete.

Students will be expected to write up a report for each programming assignment in LaTeX. As this is an introductory graduate course, some creativity on the solutions may be expected; some of the problems will have open-ended prompts and students should expect to explore the parameter space of the algorithms they implement and report on their findings.

Students may collaborate on the programming assignments in small groups (no more than 2 or 3 students), but solutions must be written up independently.

Quizzes

To supplement the programming assignments, which test mostly practical knowledge, the quizzes will be shorter and ask two or three conceptual questions related to the course material. Solutions are encouraged to be typed using LaTeX. Students may collaborate on the quizzes in small groups (no more than 2 or 3 students), but solutions must be written up independently.

There will be a total of 5 take home quizzes assigned during the semester. At the end of the term the lowest quiz score (including missed quizzes) will be dropped. (Programming assignments cannot be dropped.)

Participation & Lectures

Lectures will occasionally include *breakout sessions* during which students will be expected to work collaboratively to tackle some small problem related to the lecture. However, I have chosen *not* to include participation as a part of the overall grade. You are still strongly encouraged to come to lecture, as the breakout sessions are an integral part of the course experience.

Course Resources

Textbook & Readings

There will be no course textbook to purchase, though there will be periodic supplemental readings that will be given, mostly as optional additional content. The lecture slides are intended to be self-contained. I will also occasionally recommend supplemental readings.

Lecture Slides

I will be providing lecture slides as PDF documents after each lecture via Blackboard.

Discussion

We will be using Piazza for class discussion. I strongly encourage students to message me on Piazza rather than email me.

Detailed Course Schedule

The following schedule is *tentative* and subject to change. Deadlines are typically end-of-day Saturday following class on the week the assignment is due. The deadline for each assignment will be listed on Blackboard.

LEC	DATE	TOPICS	ASSIGNMENTS
1	08/22	Introduction	P1 Out

		Fundamentals Bug Algorithms	
2	08/29	State Representations State Space Search (A*, D* Lite)	Q1 Out
3	09/05	Constraint-based Planning Mixed-Integer Linear Programs	P2 Out
4	09/12	Task Planning Planning with PDDL Task Planning	Q2 Out
5	09/19	The Markov Decision Process (MDPs) Tabular Methods for Solving MDPs Introduction to (Tabular) Reinforcement Learning	P3 Out
6	09/26	Bandit Algorithms Game Playing (and review of alpha-beta pruning) Monte Carlo Tree Search	Q3 Out
7	10/03	*[Guest Lecturer TBD: Prof. Stein @ IROS]* Temporal-Difference Learning Methods TD(0), SARSA, Q-Learning	P4 Out
—	10/10	[No Tuesday Classes due to Fall Break]	
8	10/17	Approximate Reinforcement Learning Machine Learning and Deep Learning Crash Course Deep Q-Learning [AlphaGo and AlphaZero]	Q4 Out
9	10/24	Inverse Reinforcement Learning Imitation Learning [Maybe: Skills and Options]	P5 Out
10	10/31	Policy Optimization Policy Gradient The REINFORCE Algorithm (“Trick”) Actor-Critic Methods	
11	11/07	**[Guest Lecture (TBD): Prof. Stein @ CoRL]**	
12	11/14	Planning Under Uncertainty Noisy Perception: Filtering and State Estimation Partially Observable MDPs (POMDPs) Belief-space planning	
13	11/21	Modern Deep RL The Impact of “Foundation Models”	
14	11/28	Ethical Considerations Applications (cont.) & Open Problems Course Wrap Up	

Lateness Policy

Every assignment (except The Final Project) can be turned in 3 days late without penalty. I ask that you make a best effort to turn in the assignments on time, and use the additional three days if you get stuck or if “life happens” (which it

seems to do a lot these days). Any additional late days will result in a 10%/day penalty for all assignments, up to a maximum of 1 week after the *original deadline*, after which the assignment will not be accepted so that we can post solutions. If you feel you need an extension in addition to the 3 free late days, please try to ask me *before the original deadline*, as I would like to avoid emails asking for extensions with only hours before late penalties start to accumulate. **If you turn in the assignments by the assigned deadline, you will get a 2% bonus on that assignment's grade (up to the maximum).**

Assignments will be released on Tuesdays (after class) and then turned in the Wednesday two weeks later, with the option to turn them as late as that Saturday without penalty if needed. Be warned though: I typically do not work on Friday evenings and Saturdays, so messages to me at the last minute will likely go unanswered. I strongly recommend you get started on the assignments in advance of the deadline.

Inclusion & Integrity

I stand by Mason's commitment to diversity and inclusion and hope to foster an inclusive environment in which all feel welcome in my class.

> *True diversity is defined not only as differences in individual backgrounds, personal identities, intellectual approaches, and demographics; it is also the removal of barriers and the creation of space that allow individuals to fully engage in the life of the university.*

Every student in this class is exactly where they belong and it is our honor to welcome each of you to join us in learning throughout this semester. Every student in this class, regardless of background, sex, gender, race, ethnicity, class, political affiliation, physical or mental ability, veteran status, nationality, or any other identity category, is an equal member of our class.

You have the right to be called by whatever name you wish, to be referred to by whatever pronoun you identify, and to adjust these at any point. If you feel uncomfortable in any aspect of our instruction that results in any barrier to your inclusion in this course, please contact me (your professor) directly.

Honor Code Statement

The GMU Honor Code is in effect at all times. In addition, the CS Department has further honor code policies regarding programming projects, which are detailed here. Any deviation from the GMU or the CS department Honor Code is considered an Honor Code violation.

Disability Accommodation

If you have a documented learning disability or other condition which may affect academic performance, make sure this documentation is on file with the Office of Disability Services and then discuss with the professor about accommodations. Submitting the paperwork at the deadline for a project or quiz is far too late! Even if you don't know if you plan on utilizing the accommodations ahead of time, it's in your best interest to prepare them ahead of time.

Mental Wellness

Graduate School can be a stressful environment. My door is always open; if you are struggling with the course work or would like someone to talk to, feel free to reach out to me. GMU also provides many mental health resources that I encourage you to look at.

Sexual Harassment and Interpersonal Violence

As a faculty member and designated "Responsible Employee," I am required to report all disclosures of sexual assault, interpersonal violence, and stalking to Mason's Title IX Coordinator per university policy 1412. If you wish to speak with someone confidentially, please contact the Student Support and Advocacy Center (703-380-1434), Counseling and Psychological Services (703-993-2380), Student Health Services, or Mason's Title IX Coordinator (703-993-8730; cde@gmu.edu).

Privacy and Email

Students must use their Mason email account to receive important University information, including communications related to this class. I cannot respond to messages sent from or send messages to a non-Mason email address.

To protect your privacy, I also cannot list your Mason email address on any public forum or provide it to any other students. You may, of course, give your email address to any other students.

Any video recordings of class meetings that are shared only with the instructors and students officially enrolled in a class do not violate FERPA or any other privacy expectation.

All course materials posted to Blackboard or other course site are private; by federal law, any materials that identify specific students (via their name, voice, or image) must not be shared with anyone not enrolled in this class.