

MAE 242: Robot Motion Planning. Spring '26

Course web address: `canvas.ucsd.edu`

Instructor: Prof. Sonia Martínez, FAH 3302, 858-822-4243, `soniamd at ucsd dot edu`

Teaching Assistant: Nirabhra Mandal, `nmandal at ucsd dot edu`

Lecture Time and Place: Mon/Wed/Fri 10:00am - 10:50am, EBU-II 105

Type of lectures and etiquette: This course will be taught in person at the location indicated above. (Though there might be some occasional lectures that are pre-recorded via zoom.) The classroom does not have podcast technology, so no podcasts will be made available. It is requested that no mobile phones be used during the lecture to minimize distraction. The use of tablets during the lecture is allowed, but the use of personal computers is not.

Office and Tutorial Hours: Sonia, Wednesdays, 3:30 to 4:30pm, EBU-I 1603
Nirabhra, TBA, TBA, place TBA

Texts: The following books provide the main background for the course (will be complemented with slides and papers):

- S. M. LaValle. *Planning algorithms*. 2006. This text is available at *Amazon.com*, and also freely available at the website <http://planning.cs.uiuc.edu/>.
- D. P. Bertsekas and J. Tsitsiklis. *Neuro-Dynamic Programming*. Athena Scientific, 1996
- R. S. Sutton and A. G. Barto. *Reinforcement Learning: An Introduction*. 2008, second edition available at the website <http://incompleteideas.net/book/RLbook2018.pdf>
- C. Szepesvári. *Algorithms for Reinforcement Learning*. Morgan Claypool Publishers, July 2010, available online at the website www.ualberta.ca/~szepesva/papers/RLAlgsInMDPs.pdf

Course Scope and Audience: This course is an introduction to planning, decision-making and learning algorithms for robots, where, interchangeably, robot is a synonym of “agent” or “decision maker”. Algorithms for planning and learning find application in a number of technologies and disciplines such as advanced manufacturing, logistics, power systems, computer-aided design, computer graphics, the theory of artificial intelligence, and broad robotics applications. In particular, learning algorithms in robotics play today an important role in the field, as all robotic systems are being augmented with AI. Because of this, this course will cover reinforcement learning algorithms.

A main emphasis of the course is on the theory of advanced planning and learning algorithms and their properties with the end-goal of providing a good fundamental knowledge in this area. There will be roughly three main parts in this course: Part (I) on discrete, dynamic programming, part (II) learning for single agents, part (III) multi-agent learning (time permitting.)

Prerequisites: This course is *math intensive and formal* and the approach taken to teach the class is that of *mathematical rigor*. Because of its formality, it may be a challenging course for first-year graduate students. Background knowledge that is required (the more you know the better):

- Probability theory: discrete and continuous RVs, mean, variance, conditional probability;
- Linear algebra and systems: basic (matrix) linear algebra operations, solutions to a set of linear equations, state-space representation of dynamic systems, the concept of linear feedback control;

- Mathematical analysis or calculus or nonlinear systems: will be helpful when we talk about algorithm convergence;
- (Convex) optimization: again, the more you know the better, but at least you should be familiar with what an optimization problem is, have an idea of types of optimization problems, and algorithms to solve them.

Programming knowledge: basic programming assignments (only for homework) will be required in Python, and some programs/tutorials will use this language. The nature of the assignments will mostly require functional-oriented programming (they will not use involved or sophisticated tools).

Topics: Tentatively, we will cover topics among the following (time permitting):

- Intro to the course. Review of searches over graphs: BFS, DFS, BF, A*, Dijkstra's algorithm. An overview of the D* algorithm.
- Markov processes, Markov processes with rewards, Markov decision processes. Value functions and policies.
- Solution methods: value iteration, policy iteration, contraction theory and convergence properties of iterations. The LQR problem (this is tentative.)
- Model-free prediction methods. Monte Carlo learning. Temporal difference learning. TD(λ).
- Model-free control methods. On policy learning methods (mainly Monte Carlo and SARSA, some comments about SARSA(λ)), off-policy learning methods (Q -learning).
- Approximation methods. Value function approximation via differentiable parametric functions (linear and NN approximations). Action-value function approximation.
- Policy-gradient methods. Monte Carlo Policy Gradient (reinforce), actor-value, and actor-critic algorithms (TD, Q actor-critic methods).
- (Partially, time permitting:) Cooperative MARL. Homogeneous agent case. Classic game theory: multi-agent games and NE solution concept. Markov Games and competitive MARL.

Assignments: Homework will be assigned throughout the course so that you can better grasp the course material. The exercises will mainly focus on theoretical analysis exercises and basic programming (in Python). They will be posted in the canvas course website and should be submitted via gradescope.

While all solutions of homework assignments will be made available, the correction of homework assignments by the TA and myself is not possible. Because of this, we will make public the solutions to homework with a rubric, so that students can self-assign marks. The corrected solutions (both programming and analytical) will be then submitted to gradescope, and the instructional team will revise that these marks have been applied faithfully and correctly over the solutions. Instructions on how to do this will be shared promptly. It is very important that all students exercise this task with full honesty and do the self corrections themselves (not with an AI). Any dishonestly reported correction can result in an automatic F for the course. It is very important that each student takes full responsibility to understand how to solve homework problems, otherwise the performance in the exams will not be good.

An in class midterm and final exam will be part of the course grade. These are individual assignments, and collaboration is not possible to do them.

Late homework policy: Only one late homework is allowed with no penalty if turned within 24 hours after the deadline. Other than this, penalty of a late homework is of 15%.

Comprehensive exam: This course offers a Comprehensive Exam component. The component will take the form of an exercise immediately after the Final exam period.

Collaboration policy: You are encouraged to work with other students on your assignments, and to help other students complete their assignments, provided that you comply with the following conditions:

1. **Honest representation:** The material you turn in for course credit must be a fair representation of your work. You are responsible for understanding and being able to explain and duplicate the work you submit. Group submissions are not allowed in this course, and each student should submit their own individual assignment, written in their own words. The same happens with programming exercises: please do not submit exact copies of programming solutions, the autocorrection tool in Gradescope checks for plagiarism. Any student can be asked to report and explain a submitted solution.
2. **Active involvement:** You must ensure that you are an active participant in all collaborations, and are not merely dividing up the work or following along while another student does the work. For example, copying another student's work without actively being involved in deriving the solution is strictly prohibited. To avoid misunderstandings, please turn in solutions written in your own words, not an exact copy of what someone else submits.
3. **Work individually or in small groups:** Working in groups of more than **three** people is discouraged because it limits the amount of participation by each member of the group. In your homework solutions please indicate the names of the people you collaborated with.
4. **Give help appropriately:** When helping someone, it is important not to simply give them a solution, because then they may not understand it fully and will not be able to solve a similar problem next time. It's always important to take the time to help someone think through the problem and develop the solution. Often, this can be accomplished by asking them a series of leading questions.
5. **If in doubt, ask your instructor:** Be sure to ask in advance if you have any doubts about whether a certain type of collaboration is acceptable

Important Dates: The following are tentative dates for homework (expect to have a hw assignment per week, either programming or of analytical type or both, extra credit assignments may be made available over time, and the dates may definitely shift)

Title	Issued	Due
Homework 1 (Analytical)	Apr 3	Apr 10 (Fri)
Homework 1 (Program)	Apr 6	Apr 17 (Fri)
Homework 2 (Analytical)	Apr 10	Apr 17 (Fri)
Homework 2 (Program)	April 20	May 8 (Fri)
Homework 3 (Analytical)	Apr 17	Apr 24 (Fri)
Homework 4 (Analytical)	May 8	May 15 (Fri)
Homework 5 (Analytical)	May 15	May 22 (Fri)
Homework 5 (Program part I)	May 11	May 22 (Fri)
Homework 5 (Program part II)	May 25	Jun 5 (Fri)
Homework 6 (Analytical)	May 29	Jun 5 (Fri)
Midterm	May 1, in class	
Final	June 8, in class	

Grading: The grade of the course will be tentatively calculated as with the following code:

```
Define  $X = \max(\text{midterm}(35\%) + \text{final}(65\%), \text{midterm}(0\%) + \text{final}(100\%))$ 
```

```
if  $X \geq 60\%$ :
```

```
     $G = \max(X, X(85\%) + \text{homework}(15\%))$ 
```

```
else:
```

```
     $G = X$ 
```

Participation in Piazza forums will be rewarded with an extra credit of 2.5% of the final grade. To get an S in the course, you should get a 65% of the grade.

Code of conduct:

All participants in the course are bound by the UCSD Code of Conduct, and Academic Integrity. In this course, the use of ChatGPT or other GenAI tools to solve homework problems or write the programming assignments on your behalf is not allowed and constitutes cheating. Copying from previous homework solutions is also considered cheating. Misrepresenting your own solutions in a homework self-correction (or using AI to do the work for you) is also considered cheating. More information can be found at: <https://academicintegrity.ucsd.edu>. A dishonest behavior is taken very seriously and can result into receiving an F for the course.