



CSCI 599: Autonomous Decision-Making

Units: 4.0

Spring 2024 – TuThu – 5:00-6:50pm

Location: KAP 166

Instructor: Thomy Phan

Office: PHE 324

Office Hours: Every Monday at 11am-12pm

Contact Info: thomy.phan@usc.edu (replies expectable within 48 hours).

Course Description

Recent advances in artificial intelligence (AI) have enabled the development of autonomous agents that are able to solve complex problems like mastering a variety of challenging games, predicting protein structures, managing logistics scenarios, or most recently holding conversations with humans. These agents have the ability to perceive their current situation and make adequate decisions based on prior knowledge and current observation. Autonomous agents can adjust their behavior at runtime depending on changing circumstances.

This course will give an introduction to the techniques used to realize such agents ranging from *automated planning* to *reinforcement learning*. The lecture will first cover autonomous decision-making for *single-agent systems* and then extend the concepts to *multi-agent systems*. Students will have the opportunity to implement planning and reinforcement learning algorithms and test them in some toy problems to gain a deeper understanding of the challenges and concepts. This course might be particularly interesting for students who plan to do AI research.

Learning Objectives and Outcomes

This course will enable students to:

- Understand the fundamentals of autonomous decision-making in AI systems
- Understand the challenges of autonomous decision-making in single- and multi-agent systems
- Understand the concepts of advanced decision-making techniques based on automated planning and reinforcement learning
- Solve problems using the taught decision-making techniques

Recommended Preparation

- Basic knowledge of artificial intelligence (at the level of CSCI 360 or 561) and machine learning (at the level of CSCI 467 or 567)
- Python programming skills are strongly encouraged (classes that use Python depending on instructor: CSCI 353, CSCI 360, CSCI 445, EE 250, EE 364, also ITP 115 as introduction to Python). All coding assignments will be based on Python so students should either already have Python skills or be prepared to learn on their own in parallel with the course. A basic understanding of statistics and decision theory is helpful.

Course Notes

Lecture slides and other class information will be posted on USC Blackboard. A Discord server will be provided for Q&A and open discussions among students.

Technological Proficiency and Hardware/Software Required

Course homework and projects require Python 3 (best installed via Anaconda) and supplementary packages like PyTorch and Gymnasium.

Required Readings and Supplementary Materials

Recommended open-access material as follows (the PDFs will also be provided by the instructor):

Books

- [RLAI] R. Sutton & A. Barto, "*Reinforcement Learning: An Introduction*"
- [MAS] Y. Shoham & K. Leyton-Brown, "*Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations*"
- [CIDP] F. Oliehoek & C. Amato, "*A Concise Introduction to Decentralized POMDPs*"

Research Papers

- [1] L. Kocsis & C. Szepesvári, “Bandit Based Monte-Carlo Planning”, European Conference on Machine Learning, 2006.
- [2] D. Silver & J. Veness, “Monte-Carlo Planning in Large POMDPs”, Advances in Neural Information Processing Systems 23, 2010.
- [3] D. Silver et al., “Mastering the Game of Go without Human Knowledge”, Nature, 2017.
- [4] J. Schrittwieser et al., “Mastering Atari, Go, Chess and Shogi by Planning with a Learned Model”, Nature, 2020.
- [5] V. Mnih et al., “Human-Level Control Through Deep Reinforcement Learning”, Nature, 2015.
- [6] V. Mnih et al., “Asynchronous Methods for Deep Reinforcement Learning”, International Conference on Machine Learning, 2016.
- [7] M. Jaderberg et al., “Human-Level Performance in 3D Multiplayer Games with Population-Based Reinforcement Learning”, Science, 2019.
- [8] O. Vinyals et al., “Grandmaster Level in StarCraft II using Multi-Agent Reinforcement Learning” Nature, 2019.
- [9] G. Sharon et al., “Conflict-based Search for Optimal Multi-Agent Pathfinding”, Artificial Intelligence 2015
- [10] J. Li et al., “Anytime Multi-Agent Path Finding via Large Neighborhood Search”, IJCAI 2021
- [11] M. Tan, “Multi-Agent Reinforcement Learning: Independent vs. Cooperative Learners”, ICML 1993
- [12] M. Littman, “Markov Games as a Framework for Multi-Agent Reinforcement Learning”, ICML 1994
- [13] J. Foerster et al., “Counterfactual Multi-Agent Policy Gradients”, AAI 2018
- [14] T. Rashid et al., “QMIX: Monotonic Value Function Factorisation for Deep Multi-Agent Reinforcement Learning”, ICML 2018

Description and Assessment of Assignments

There will be 4 homework assignments and 2 quizzes throughout the semester.

The final project will be carried out by small groups of 2-3 students to implement and evaluate decision-making algorithms to solve a toy problem. Students need to:

1. Set up a collaborative research and development environment
2. Study the toy problem and identify the main challenges of the problem
3. Identify and formulate adequate measures to evaluate algorithms objectively
4. Implement algorithms to solve the toy problem
5. Evaluate and present the results in written reports and presentations

The written reports and intermediate presentations are due on a regular basis. A final report (+code and documentation) is due at the end of the semester.

Grading Breakdown

Assignment	Points	% of Grade
Homework assignments	50	50
Quizzes	10	10
Final Project	40	40
TOTAL	100	100

Grading Scale

Course final grades will be determined using the following scale

A	95-100
A-	90-94
B+	87-89
B	83-86
B-	80-82
C+	77-79
C	73-76
C-	70-72
D+	67-69
D	63-66
D-	60-62
F	59 and below

Assignment Rubrics

Homework assignments, a project proposal, and project reports on a regular basis.

Assignment Submission Policy

All homework assignments (including code) and project reports must be submitted electronically using the USC Blackboard system.

Grading Timeline

Assignments will be graded and returned within a week. The project includes a proposal, written reports, and intermediate presentations on a regular basis, as well as a final report (+code and documentation). The final project will be graded within a week after the final report.

Use of Generative AI in this Course

Generative AI is not permitted: Since creating, analytical, and critical thinking skills are part of the learning outcomes of this course, all assignments should be prepared by the student working individually or in groups as described on each assignment. Students may not have another person or entity complete any portion of the assignment. Developing strong competencies in these areas will prepare you for a competitive workplace. Therefore, using AI-generated tools is prohibited in this course, will be identified as plagiarism, and will be reported to the Office of Academic Integrity.

Additional Policies

No late submissions are accepted. All deadlines are final.

Course Schedule: A Weekly Breakdown

	Topics/Daily Activities	Readings	Deliverable/ Due Dates
Week 1	Introduction, Decision Making Fundamentals	[RLAI] Chapter 1 and 2	Assigned: Homework 1
Week 2	Markov Decision Processes	[RLAI] Chapter 3	
Week 3	Dynamic Programming	[RLAI] Chapter 4	
Week 4	Reinforcement Learning 1	[RLAI] Chapter 5-6	Due: Homework 1 Assigned: Homework 2
Week 5	Reinforcement Learning 2	Research Papers [3,4]	
Week 6	Search Algorithms	-	Due: Homework 2 Assigned: Homework 3
Week 7	Monte Carlo Planning	Research Papers [1,2]	Assigned: Summary Homework
Week 8	Deep Reinforcement Learning 1	[RLAI] Chapter 9-10	Due: Homework 3 Assigned: Homework 4
Week 9	Applications	Research Papers [5,6]	Due: Project Proposal Due: Summary Homework
Week 10	Deep Reinforcement Learning 2	Research Papers [7,8]	Due: Homework 4
Week 11	Multi-Agent Systems, Project Kickoff	[MAS] Chapter 1-2	Due: Project Report 1
Week 12	Multi-Agent Planning 1	Research Papers [9,10]	
Week 13	Multi-Agent Planning 2	[CIDP] Chapter 1-3	Due: Project Report 2
Week 14	Multi-Agent Learning	Research Papers [11,12]	
Week 15	Applications, Challenges, and Outlook	-	Due: Final Project (Code, Documentation, Presentation), Quiz Final Project
FINAL	-	-	Due on the university-scheduled date of the final exam.

Statement on Academic Conduct and Support Systems

Academic Integrity:

The University of Southern California is a learning community committed to developing successful scholars and researchers dedicated to the pursuit of knowledge and the dissemination of ideas. Academic misconduct, which includes any act of dishonesty in the production or submission of academic work, comprises the integrity of the person who commits the act and can impugn the perceived integrity of the entire university community. It stands in opposition to the university's mission to research, educate, and contribute productively to our community and the world.

All students are expected to submit assignments that represent their own original work, and that have been prepared specifically for the course or section for which they have been submitted. You may not submit work written by others or "recycle" work prepared for other courses without obtaining written permission from the instructor(s).

Other violations of academic integrity include, but are not limited to, cheating, plagiarism, fabrication (e.g., falsifying data), collusion, knowingly assisting others in acts of academic dishonesty, and any act that gains or is intended to gain an unfair academic advantage.

The impact of academic dishonesty is far-reaching and is considered a serious offense against the university. All incidences of academic misconduct will be reported to the Office of Academic Integrity and could result in outcomes such as failure on the assignment, failure in the course, suspension, or even expulsion from the university.

For more information about academic integrity see [the student handbook](#) or the [Office of Academic Integrity's website](#), and university policies on [Research and Scholarship Misconduct](#).

Please ask your instructor if you are unsure what constitutes unauthorized assistance on an exam or assignment, or what information requires citation and/or attribution.

Course Content Distribution and Synchronous Session Recordings Policies

USC has policies that prohibit recording and distribution of any synchronous and asynchronous course content outside of the learning environment.

Recording a university class without the express permission of the instructor and announcement to the class, or unless conducted pursuant to an Office of Student Accessibility Services (OSAS) accommodation. Recording can inhibit free discussion in the future, and thus infringe on the academic freedom of other students as well as the instructor. ([Living our Unifying Values: The USC Student Handbook](#), page 13).

Distribution or use of notes, recordings, exams, or other intellectual property, based on university classes or lectures without the express permission of the instructor for purposes other than individual or group study. This includes but is not limited to providing materials for distribution by services publishing course materials. This restriction on unauthorized use also applies to all information, which had been distributed to students or in any way had been displayed for use in relationship to the class, whether obtained in class, via email, on the internet, or via any other media. ([Living our Unifying Values: The USC Student Handbook](#), page 13).

Students and Disability Accommodations:

USC welcomes students with disabilities into all of the University's educational programs. [The Office of Student Accessibility Services](#) (OSAS) is responsible for the determination of appropriate accommodations for students who encounter disability-related barriers. Once a student has completed the OSAS process (registration, initial appointment, and submitted documentation) and accommodations are determined to be reasonable and appropriate, a Letter of Accommodation (LOA) will be available to generate for each

course. The LOA must be given to each course instructor by the student and followed up with a discussion. This should be done as early in the semester as possible as accommodations are not retroactive. More information can be found at osas.usc.edu. You may contact OSAS at (213) 740-0776 or via email at osasfrontdesk@usc.edu.

Support Systems:

[Counseling and Mental Health](#) - (213) 740-9355 – 24/7 on call

Free and confidential mental health treatment for students, including short-term psychotherapy, group counseling, stress fitness workshops, and crisis intervention.

[988 Suicide and Crisis Lifeline](#) - 988 for both calls and text messages – 24/7 on call

The 988 Suicide and Crisis Lifeline (formerly known as the National Suicide Prevention Lifeline) provides free and confidential emotional support to people in suicidal crisis or emotional distress 24 hours a day, 7 days a week, across the United States. The Lifeline is comprised of a national network of over 200 local crisis centers, combining custom local care and resources with national standards and best practices. The new, shorter phone number makes it easier for people to remember and access mental health crisis services (though the previous 1 (800) 273-8255 number will continue to function indefinitely) and represents a continued commitment to those in crisis.

[Relationship and Sexual Violence Prevention Services \(RSVP\)](#) - (213) 740-9355(WELL) – 24/7 on call

Free and confidential therapy services, workshops, and training for situations related to gender- and power-based harm (including sexual assault, intimate partner violence, and stalking).

[Office for Equity, Equal Opportunity, and Title IX \(EEO-TIX\)](#) - (213) 740-5086

Information about how to get help or help someone affected by harassment or discrimination, rights of protected classes, reporting options, and additional resources for students, faculty, staff, visitors, and applicants.

[Reporting Incidents of Bias or Harassment](#) - (213) 740-5086 or (213) 821-8298

Avenue to report incidents of bias, hate crimes, and microaggressions to the Office for Equity, Equal Opportunity, and Title for appropriate investigation, supportive measures, and response.

[The Office of Student Accessibility Services \(OSAS\)](#) - (213) 740-0776

OSAS ensures equal access for students with disabilities through providing academic accommodations and auxiliary aids in accordance with federal laws and university policy.

[USC Campus Support and Intervention](#) - (213) 740-0411

Assists students and families in resolving complex personal, financial, and academic issues adversely affecting their success as a student.

[Diversity, Equity and Inclusion](#) - (213) 740-2101

Information on events, programs and training, the Provost's Diversity and Inclusion Council, Diversity Liaisons for each academic school, chronology, participation, and various resources for students.

[USC Emergency](#) - UPC: (213) 740-4321, HSC: (323) 442-1000 – 24/7 on call

Emergency assistance and avenue to report a crime. Latest updates regarding safety, including ways in which instruction will be continued if an officially declared emergency makes travel to campus infeasible.

[USC Department of Public Safety](#) - UPC: (213) 740-6000, HSC: (323) 442-1200 – 24/7 on call

Non-emergency assistance or information.

[Office of the Ombuds](#) - (213) 821-9556 (UPC) / (323-442-0382 (HSC)

A safe and confidential place to share your USC-related issues with a University Ombuds who will work with you to explore options or paths to manage your concern.

[Occupational Therapy Faculty Practice](#) - (323) 442-2850 or otfp@med.usc.edu

Confidential Lifestyle Redesign services for USC students to support health promoting habits and routines that enhance quality of life and academic performance.