

[CSCI 360] Introduction to Artificial Intelligence

Spring 2019

	Section 30000D	Section 30282R
Lectures:	Mon/Wed 2:00pm-3:50pm	Tue/Thu 2:00pm-3:50pm
Classroom:	GFS 116	MHP 101
Instructor:	Chao Wang	Andy Nealen
Office:	SAL 334	SCI 201R
Office Hours:	TBA	TBA
E-mail:	wang626@usc.edu	nealen@usc.edu

TAs: Elizabeth Bondi (bondi@usc.edu), Jeremy Hsu (chincheh@usc.edu), Brendan Kennedy (btkenned@usc.edu)

Syllabus: Artificial Intelligence (AI) seeks to understand the mechanisms underlying thought and intelligent behavior, with a particular focus on their embodiment in machines. Core topics include the integrating perspective of intelligent agents and how such systems can engage in: search and problem solving; symbolic and probabilistic knowledge representation and reasoning; planning; and machine learning. The course introduces both basic concepts and algorithms, and explores how to apply them in the construction of systems that can interact intelligently with complex environments.

Target Audience: The course is intended for undergraduate students in computer science or closely related disciplines, usually in the junior year. Graduate students should take CS561 rather than CS360.

Prerequisites: The courses CSCI 104L ("Data Structures and Object-Oriented Design") and CSCI 170 ("Discrete Methods in Computer Science") are necessary prerequisites, which will not be waived. Overall, the prerequisites of CSCI 360 include a solid understanding of data structures and algorithms since you will have to be able to understand algorithms and read pseudo code. You should also know the basics of probability theory, calculus (especially derivatives) and discrete mathematics. Finally, you should know how to program since the projects will be programming related. Do not take this class if you cannot program yet.

The most important prerequisite of all, however, is your interest in the class, motivation, and commitment to learning. If you are not sure whether this class is for you, come and talk to us.

Readings: Most readings will be chosen from the (required) textbook, which is readily available from many standard online retailers:

Stuart Russell and Peter Norvig, **Artificial Intelligence: A Modern Approach** (most recent edition - currently: 3rd edition), Pearson/Prentice Hall (ISBN 978-0-13-604259-4)

The authors made extensive revisions from one edition to the next one. We therefore suggest that you buy the latest edition. Definitely do not use the first edition.

We will not cover all of the chapters and, from time to time, cover topics not contained in the book.

Additional material will be provided as necessary.

Lectures: Lecture attendance is required; please do not take the course if you cannot attend the lectures. In general, the lectures are meant to summarize the readings and stress the important points. Thus, we expect you to read the corresponding part of the textbook. If you miss a class, it is your responsibility to find out what we discussed in class, including which announcements we made in class. If there is something that you do not understand, feel free to interrupt with questions. Your active participation in class is crucial in making the class successful.

Use your fellow students as a resource, e.g., by forming study groups or posting questions on the discussion forum on Piazza that the TAs and CPs/Graders monitor on a daily basis. We encourage you to participate actively on the discussion forum, by both asking and answering questions. If you need additional help, please feel free to come to our office hours. The TAs and CPs/Graders are experienced and will be able to answer all of your technical questions, including about the textbook, lectures, homework assignments, and projects.

Assignments: There will be two graded, textbook-style homework assignments, all of which are mandatory and must be done individually. They will ensure your understanding of the material mentioned during the lectures, and help you prepare for the kind of problems encountered during the exams. You are required to submit your solutions in time.

Projects: There will be three graded, programming projects, all of which must be done individually. You are required to cite all resources you relied on, including people, web pages, publications and other write-ups. You are not allowed to use code or code snippets of others (that is, that you did not write yourself), or discuss with others on how to solve the projects.

Please start working on your homework assignments and projects early and hand them in early. There is a grace period of 24 hours during which time your submission will be accepted, but with a 20% penalty. After 24 hours, your submission will not be accepted.

Discussions: The discussion session will allow you to practice the application of the material taught in the lectures via a mix of the TAs giving lectures and answering your questions, you solving textbook-style problems in small groups, and the TAs discussing projects, assignments, and their solutions. There will be no exercise session during the first week of classes.

Exams: There will be one midterm and one final, all of which are mandatory. The exam dates are listed on the schedule; please do not take the course if you cannot attend the exams. No makeups will be given. All exams will be comprehensive but with a focus on material not yet tested in a previous exam.

Grades: Homework assignments, projects and exams have the following weights:

- Quizzes: 5%
- Homework 1: 10%
- Homework 2: 10%
- Project 1: 10%
- Project 2: 10%
- Project 3: 15%
- Midterm Exam: 20%
- Final Exam: 20%

The instructors reserve the right to adjust the grading scale. There will always be some students who are very close to grade boundaries. There is nothing that we will do about that. Grades are based on performance, not need or personal circumstances, and the instructor does not negotiate grades. Thus, do not take CS360 (or take it at your own risk) if you need a certain grade, for example, because you are graduating or because you have been conditionally admitted.

To receive a good grade, you will need to perform well in the exams **and** the projects, homeworks, and quizzes. Please check the correctness of the grading and the posted scores immediately after we announce the availability of the scores. **You will need to let us know about any grading issue within 7 days of us posting the score.** After that time, we will no longer entertain your requests for changes to your score. If you have a grading issue, you will need to discuss the issue first with the TAs. If you cannot reach consensus, you can appeal the grading issue to the instructors.

Academic Integrity: USC seeks to maintain an optimal learning environment. General principles of academic honesty include the concept of respect for the intellectual property of others, the expectation that individual work will be submitted unless otherwise allowed by an instructor, and the obligations both to protect one's own academic work from misuse by others as well as to avoid using another's work as one's own. All students are expected to understand and abide by these principles. *We will strictly enforce the student conduct code and refer students to the Office of Student Judicial Affairs and Community Standards for further review, should there be any suspicion of academic dishonesty, and suggest that they follow the recommended sanctions in case they should find that there was academic dishonesty.* We typically suggest an F as overall class grade as penalty, if asked. Scampus, the Student Guidebook, contains the student conduct code and the academic review process: <https://policy.usc.edu/scampus-part-b/>.

Problems and Concerns: At some point, you will have questions. For example, you might not be able to get code to run, there might be something in the textbook that you do not understand, and so on. In this case, we encourage you to post the question on the discussion forum and see whether someone can help you. If this approach does not generate the desired result, then the TAs and CPs/Graders will be happy to help you in person during their office hours. TAs do answer email but, unfortunately, often will not manage to answer it on the same day. (Sometimes, they will be out of town and it will take them even longer. Also, they are typically overloaded with questions on exam days or directly before.)

It is very important to us that you voice your concerns about any aspect of the class as soon as they arise. Please send an e-mail to the instructors or talk to us in person.

Artificial Intelligence is a fun topic, and we hope that all of us will have lots of fun!

Tentative Schedule: The instructors may adjust this schedule during the semester

Week	30000D	30282R	Topics	Chapters
1	1/7 1/9	1/8 1/10	Intelligent Agents Problem Solving and Search	[Ch 1.1-1.4 and 2.1-2.4] [Ch 3.1-3.3]
2	1/14 1/16	1/15 1/17	Uninformed Search Heuristic Search (A*)	[Ch 3.3-3.4] [Ch 3.5]
3	1/21 1/23	1/22 1/24	Heuristic Functions Local Search	[Ch 3.6] [Ch 4.1-4.2]
	1/25		Project 1 Out	
4	1/28 1/30	1/29 1/31	Adversarial Search Knowledge Based Agents	[Ch 5.1-5.3] [Ch 7.1-7.3]
5	2/4 2/6	2/5 2/7	Propositional Logic Inference First-Order Logic	[Ch 7.4-7.5] [Ch 8.1-8.4]
	2/8 2/8		Project 1 Due Homework 1 Out	
6	2/11 2/13	2/12 2/14	Rule-Based Systems Search-Based Planning	[Ch 9.3-9.4] [Ch 10.1-10.3]
	2/15		Homework 1 Due	
7	2/18 2/20	2/19 2/21	SAT-Based Planning Knowledge Representation	[Ch 10.4] [Ch 12.1-12.5]
8	2/25 2/27	2/26 2/28	Midterm Review Midterm Exam	
	3/1		Project 2 Out	
9	3/4 3/6	3/5 3/7	Quantifying Uncertainty Bayesian Networks	[Ch 13.1-13.6] [Ch 14.1-14.2]
10	3/11 3/13	3/12 3/14	<i>(spring break, no class)</i> <i>(spring break, no class)</i>	
11	3/18 3/20	3/19 3/21	Advanced Topic (Andy Traveling) Advanced Topic (Andy Traveling)	
	3/23		Project 2 Due	
12	3/25 3/27	3/26 3/28	Inference in Bayesian Networks Decision Theory	[Ch 14.3-14.4] [Ch 16.1-16.3 and 16.5]
	3/29		Homework 2 Out	
13	4/1 4/3	4/2 4/4	Markov Decision Processes Decision Tree Learning	[Ch 17.1-17.2] [Ch 18.1-18.3]
	4/5 4/5		Homework 2 Due Project 3 Out	
14	4/8 4/10	4/9 4/11	Perceptron Learning Neural Network Learning	[Ch 18.7.1-18.7.2] [Ch 18.7.3-18.7.4]
15	4/15 4/17	4/16 4/18	Statistical Learning Reinforcement Learning	[Ch 20.2.1-20.2.2] [Ch 21.1-21.2]
16	4/22 4/24	4/23 4/25	Artificial Intelligence Ethics Wrap-Up and Final Review	
	4/26		Project 3 Due	
	5/3	5/2	Final Exam (2pm-4pm)	

More on Project and Exam Topics: Given the subject ordering below, the three projects will be loosely related to

- search (project 1: A* search with perfect information),
- adversarial search (project 2: Minimax with alpha-beta pruning), and
- decision making under uncertainty (project 3: Markov Decision Process).

The midterm exam will cover search, logical reasoning, and planning, while the final exam will cover uncertainty, decision theory, and learning.