San José State University
School/Department
CS 156-04, Introduction to Artificial Intelligence, Fall, 2020

Course and Contact Information

Instructor:	Dr. Sanjoy Paul
Office Location:	[TBD]
Telephone:	(832) 805-4877
Email:	paul.sanjoy@sjsu.edu
Office Hours:	[TBD]

Class Days/Time:	TuTh: 16:30-17:45 PST
Classroom:	Zoom
Prerequisites:	CS 146 and either CS 151 or CMPE 135 with a grade of C- or better in each

Course Format

Technology Intensive, Hybrid, and Online Courses

This course will be taught online. You need Internet connectivity and zoom installed on your a computer to participate in the classroom activities and/or submit assignments. You need to have a Python software development environment installed on your computer to do the projects.

Course Description

Basic concepts and techniques of artificial intelligence: problem solving, search, deduction, intelligent agents, knowledge representation. Topics chosen from logic programming, game playing, planning, machine learning, natural language, neural nets, robotics.

Course Learning Outcomes

Upon successful completion of this course, students will be able to:

1. Understand what web search, speech recognition, face recognition, machine translation, autonomous driving, and automatic scheduling have in common.
2. Use artificial intelligence (AI) to tackle complex real-world problems with programming & mathematics.
3. Internalize the foundational principles that drive complex real-world applications. Specific topics include machine learning, search, game playing, Markov decision processes, constraint satisfaction, graphical models, and logic.
4. Leverage the latest AI tools to tackle new AI problems one might encounter in life.
5. Implement some of the core AI algorithms taught in class.

Recommended Textbook

Artificial Intelligence: A Modern Approach. 3rd Edition. Stuart Russell and Peter Norvig ISBN: 9780136042594

Software

Python 3

https://www.python.org/downloads/release/python-363/

PyCharm Professional or Community Edition – **Recommended IDE**

Other technology requirements / equipment / material

We will use Zoom for our online sessions

Installing Zoom

https://www.youtube.com/watch?v=fVu9BILRkww

Course Requirements and Assignments

Homework Assignments:

Homework assignments will be posted and submitted on Canvas. For full credit, they must be submitted by the posted due date.

Weekly Quizzes:

We will have a weekly quiz that consists of a single question to check your understanding of the previous week's material. I will count the 10 best scores out of the 12 total quizzes in the semester. You must be in the online classroom to take the quiz. Missed quizzes cannot be made up.

Midterm Exam:

The midterm exam will take place in the classroom during class time on Thursday October 8.

Final Exam:

The final exam will take place on Thursday Dec 10.

Grading Information

The final grade in the course will be calculated based on the following percentages:

Homework Assignments: 40%

Weekly Quizzes: 10%

Midterm: 20%

Final Exam: 30%

Late Work:

Late assignments will not be accepted.

Grade Scale:

The letter grade will be determined based on the following scale:

A+ = 96% - 100%	A = 91% - 95%	A- = 86% - 90%
B+ = 81% - 85%	B = 76% - 80%	B- = 71% - 75%
C+ = 66% - 70%	C = 61% - 65%	C- = 56% - 60%
D = 51% - 55%		
F = below 50		

Classroom Protocol

Regular attendance is an integral part of the learning process. Please arrive on time for the classes.

University Policies

Per University Policy S16-9 *(http://www.sjsu.edu/senate/docs/S16-9.pdf)*, relevant information to all courses, such as academic integrity, accommodations, dropping and adding, consent for recording of class, etc. is available on Office of Graduate and Undergraduate Programs' Syllabus Information web page at http://www.sjsu.edu/gup/syllabusinfo/". Make sure to visit this page, review and be familiar with these university policies and resources.

CS156 Introduction to Artificial Intelligence, Fall 2020, Course Schedule

Please note that this schedule is subject to change with fair notice. Any changes will be announced in class and posted on the Canvas course site.

Course Schedule

Week	Date	Topics	Readings AIMA	HW Due date
1	Aug 20	Course Logistics – What is AI? Why is it important? Overview of the Course	Chapter 1-2	
2	Aug 25	Python - numpy, pandas, matplotlib etc.	https://jakevdp.github.io/PythonDataScienceHandbook/02.02-the-basics-of-numpy-arrays.html	HW1 Sep 2
2	Aug 27	Python– numpy, pandas, matplotlib etc.	https://cloudxlab.com/blog/numpy-pandas-introduction/ https://realpython.com/python-matplotlib-guide/	
3	Sep 1	Machine Learning #1 (Supervised) – Regression-1 – Linear, Multi-linear, Backward Elimination	Chapter 19-20	Quiz#1
3	Sep 3	Machine Learning #2 (Supervised) – Regression-2 – Polynomial, Decision Tree, Decision Forest	Chapter 19-20	HW2 Sep 16
4	Sep 8	Machine Learning #3 (Supervised) – Classification-1 – Linear, Support Vector Machines, Kernel SVM, K Nearest Neighbor	Chapter 19-20	Quiz#2
4	Sep 10	Machine Learning #3 (Supervised) – Classification-2 –Decision Tree, Decision Forest, Naïve Bayes	Chapter 19-20	HW3 Sep 23
5	Sep 15	Machine Learning #5 (Unsupervised) – Clustering – K-means, Generalization	Chapter 19-20	Quiz#3
5	Sep 17	Search #1 (state based) – Uninformed Search – DFS, BFS, Dynamic Programing, Uniform Cost Search	Chapter 3-4	HW4 Sep 30
6	Sep 22	Search #2 (state based) – Informed Search – A* algorithm, Admissibility and Consistency, Relaxation	Chapter 3-4	Quiz#4 HW5 Oct 14
6	Sep 24	Search #3 (state based) – Search under uncertainty – Markov Decision Processes #1, Policy evaluation, Policy iteration		
7	Sep 29	Search #4 (state based) – Search under uncertainty – Markov Decision Processes #2, Value iteration, Reinforcement Learning	Chapter 22	Quiz#5 HW6 Oct 28
7	Oct 1	Search #5 (state based) – Adversarial Search / Gaming - Expectimax, Minimax, Evaluation Functions, Alpha-beta pruning	Chapter 5	
8	Oct 6	Search #6 (state based) – Adversarial Search / Gaming – TD Learning, Game theory	Chapter 5	Quiz#6
8	Oct 8	Midterm		

9	Oct 13	Constraint Satisfaction Problem #1 (variable based) – Factor graphs, Backtracking search	Chapter 6	HW7 Nov 11
9	Oct 15	Constraint Satisfaction Problem #2 (variable based) – Dynamic ordering, Arc consistency	Chapter 6	
10	Oct 20	Constraint Satisfaction Problem #3 (variable based) – Beam search, Local search, Conditional independence, Variable elimination	Chapter 6	Quiz#7
10	Oct 22	Bayesian Networks #1 – Bayesian inference, Hidden Markov Models (HMMs)	Chapter 12-17	HW8 Nov 25
11	Oct 27	Bayesian Networks #2 – Particle sampling, Gibbs sampling, Introduction to Learning	Chapter 12-17	Quiz#8
11	Oct 29	THANKSGIVING HOLIDAY		
12	Nov 3	Bayesian Networks #3 – Learning Bayesian networks, Laplace smoothing, Expectation maximization	Chapter 12-17	
12	Nov 5	Logic #1 – Syntax versus semantics, Propositional logic, Horn clauses	Chapter 7-8	Quiz#9 HW9 Dec 1
13	Nov 10	Logic #2 – First-order logic, Resolution	Chapter 8-9	
13	Nov 12	Exam Review – Practice Problems		Quiz#10
14	Nov 17	Perceptron, Artificial Neural Networks, Stochastic Gradient Descent, Deep Learning	Chapters 18, 21	HW10 Dec 8
14	Nov 19	Recurrent Neural Networks (RNN), Natural Language Processing	Chapter 23-24	Quiz#11
15	Nov 24	Convolutional Neural Networks (CNN), Image Analytics, Computer Vision, Robotics	Chapter 25-26	
15	Nov 26	Higher-order logic, Markov logic, Semantic Parsing	Chapter 12	Quiz#12
16	Dec 1	Ethics, Fairness, Safety, Future of AI, Summary	Chapter 27-28	
16	Dec 3	Review		
Final	Dec 10	FINAL EXAM		