San José State University
College of Science, Department of Computer Science
CS 256, Topics in Artificial Intelligence, Section 1, Fall 2023

Course and Contact Information

Instructor(s): Vidya Rangasayee
Office Location: Duncan 282
Email: vidya.rangasayee@sjsu.edu
Office Hours: MW 10:30 - 11:30 am
Class Days/Time: MW 09:00- 10:15am
Classroom: SCI 311
Prerequisites: CS 156 and Graduate standing. Programming in python, discrete math, probability.
Allowed Declared Major: Computer Science, Bioinformatics, Data Science. Or instructor consent.

Course Description

Introduction to topics in artificial intelligence such as problem solving methods, game playing, understanding natural languages, pattern recognition, computer vision, and the general problem of representing knowledge. Students will be expected to use LISP.

Course Format

Technology Intensive, Hybrid, and Online Courses

All students are required to have access to a wireless laptop (running OSX, Windows, or some version of UNIX), upon which you can install required software. Technology used will include Canvas, programming in Java, and an IDE (Integrated Development Environment).

Faculty Web Page and MYSJSU Messaging

Course materials such as syllabus, handouts, notes, assignment instructions, etc. can be found on *Canvas Learning Management System course login website*. You are responsible for regularly checking with the messaging system through *MySJSU* on *Spartan App Portal* (or other communication system as indicated by the instructor) to learn of any updates.

Course Goal

The goal of artificial intelligence (AI) is to tackle complex real-world problems with rigorous mathematical tools. In this course, you will learn the foundational principles and practice implementing various AI systems. Specific topics include machine learning,

search, Markov decision processes, game playing, constraint satisfaction, graphical models, and logic

Course Learning Outcomes (CLO)

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

1. Understand the Foundations of AI, including its historical context and current significance in solving real-world problems.
2. Apply Machine Learning Techniques: Gain proficiency in machine learning algorithms and methodologies, enabling you to analyze and implement supervised and unsupervised learning models for data-driven decision-making.
3. Master Search Algorithms: Acquire the skills to design, implement, and evaluate search algorithms to solve a variety of complex problems, including pathfinding and optimization tasks.
4. Model Decision-Making with Markov Decision Processes: Learn how to model and analyze decision problems using Markov decision processes, enabling you to make optimal choices in uncertain environments.
5. Strategize Game Playing: Explore the strategies and techniques employed in game playing AI, and implement algorithms for making intelligent decisions in competitive and adversarial game scenarios.
6. Solve Constraint Satisfaction Problems: Develop the ability to model and solve constraint satisfaction problems, essential for a wide range of real-world applications such as scheduling, planning, and resource allocation.
7. Hands-on Implementation: Gain practical experience by implementing AI systems and algorithms in programming assignments and projects, honing your skills in building AI solutions.

Texts/Readings

There is no prescribed textbook. There are no required textbooks for this class, and you should be able to learn everything from the lecture notes and homeworks.

Here are some additional resources if you are interested to learn more or get different perspectives

- Russell and Norvig. Artificial Intelligence: A Modern Approach. A comprehensive reference for all the AI topics that we will cover.
- Koller and Friedman. Probabilistic Graphical Models. Covers factor graphs and Bayesian networks
- Sutton and Barto. Reinforcement Learning: An Introduction. Covers Markov decision processes and reinforcement learning (free online).
- Hastie, Tibshirani, and Friedman. The Elements of Statistical Learning. Covers machine learning from a rigorous statistical perspective (free online).
- Tsang. Foundations of Constraint Satisfaction. Covers constraint satisfaction problems (free online).

Additional Readings

A list of additional readings will be provided on the Canvas page associated with this class under http://sjsu.instructure.com.

Other technology requirements / equipment / material

You will be required to have a wireless-network ready laptop computer to participate in the class. You will also need to use your own laptop with wireless access to submit your assignment inside the SJSU campus. Your laptop needs to have wireless capability and you need to register a free wireless account at https://one.sjsu.edu/.The instructor is not responsible for providing either laptops or alternatives.

Library Liaison

Anamika Megwalu, email: *anamika.megwalu@sjsu.edu*, website: https://libguides.sjsu.edu

Course Requirements and Assignments

Success in this course is based on the expectation that students will spend, for each unit of credit, a minimum of 45 hours over the length of the course (normally three hours per unit per week) for instruction, preparation/studying, or course related activities, including but not limited to internships, labs, and clinical practica. Other course structures will have equivalent workload expectations as described in the syllabus.
a. **Project**: A final team project will be provided for you to practice AI principles. Self-selected teams of 3-4 people will work together to solve some selected problems discussed in the course. This team project will be a collaborative group project. You are free to choose your own partners but you cannot change your partners in the middle of the project. Progressive design and implementation of the term project will be done through assignments as part of the learning objectives.
b. **Exams**: There will be one midterm and one final exam.
c. **Quizzes**: There will be 1-2 quizzes and each will be counted as a HW.
d. **Homework**: Each homework is usually centered around an application and has both written and programming parts.

University Policy S16-9 (http://www.sjsu.edu/senate/docs/S16-9.pdf) states that:
Success in this course is based on the expectation that students will spend, for each unit of credit, a minimum of 45 hours over the length of the course (normally three hours per unit per week) for instruction, preparation/studying, or course related activities, including but not limited to internships, labs, and clinical practica. Other course structures will have equivalent workload expectations as described in the syllabus.

Grading Information

Course weightings will be as follows:

- Homework and Quizzes; 35%
- Project: 25%

- Exams: 20% each.

Exams may be curved (up) to raise their grades if needed. There will be opportunity for extra credit throughout the course

Your course grade will be determined by your final weighted average:
A plus = 97% or higher
A = 93% up to 97%
A minus = 90% to 93%
B plus = 87% to 90%
B = 83% to 87%
B minus = 80% to 83%
C plus = 77% to 80%
C = 73% to 77%
C minus = 70% to 73%
D plus = 67% to 70%
D = 63% to 67%
D minus = 60% to 63%
F = 0% to 60%
Boundary cases count as the higher of the two grades.

All students have the right, within a reasonable time, to know their academic scores, to review their grade-dependent work, and to be provided with explanations for the determination of their course grades. See University Policy S20-2 for more details.

Classroom Protocol

- You are expected to attend classes. If you cannot attend, it is your responsibility to get a copy of the lecture notes and class announcements from a reliable classmate. The instructor reserves the right to ignore frivolous or inappropriate e-mail inquiries. Students are expected to participate actively to provide improvement to presentations by other classmates.

University Policies

Per University Policy S16-9, relevant university policy concerning all courses, such as student responsibilities, academic integrity, accommodations, dropping and adding, consent for recording of class, etc. and available student services (e.g. learning assistance, counseling, and other resources) are listed on Syllabus Information web page (https://www.sjsu.edu/curriculum/courses/syllabus-info.php). Make sure to visit this page to review and be aware of these university policies and resources.

Course Schedule

(Subject to change with fair notice, which will be posted in Canvas)

Week	Topics
2	AI History Ethics and Responsibility
2	Refresher on probability, linear algebra and python
3	Machine learning - Linear regression, Linear Classification
3	Machine learning - Stochastic gradient descent
4	Non Linear Features, Group DRO
4	Feature templates, Generalization and Best Practices
5	Neural Networks, Backpropagation
5	k means, Differentiable programming
6	Modeling search, Tree search problem
6	Uniform cost search, A* and A* relaxation
7	Markov Decision processes - modeling and policy evaluation
7	Reinforcement Learning, Model based Monte Carlo
8	Midterm review
8	Midterm
9	Game modeling, game evaluation
9	Evaluation functions
10	Simultaneous games and non-zero-sum games
10	Constraint Satisfaction problems - modeling
11	Beam search local search
11	Markov Networks, Gibbs sampling
12	Conditional independence, Bayesian networks overview
12	Probabilistic programming
13	HMM - Forward backward algorithms, particle filtering
13	Supervised learning

Week	Topics
14	EM algorithm
14	Special topics in AI research
15	Project evaluations
15	Review for final
Final Exam	12/13/2023