SJSU SAN JOSÉ STATE UNIVERSITY

College of Science · Computer Science

Topics in Artificial Intelligence Section 02

CS 256

Spring 2024 3 Unit(s) 01/24/2024 to 05/13/2024 Modified 01/21/2024

Contact Information

Instructor(s): Vidya Rangasayee

Office Location: TBD

Email: vidya.rangasayee@sjsu.edu

Office Hours: By appointment only

Class Days/Time: TR 4:30PM - 5:45PM

Classroom: MH 224

Prerequisites: CS 156 and Graduate standing. Programming in Python, discrete math, probability.

Allowed Declared Major: Computer Science, Bioinformatics, Data Science. Or instructor consent.

Lecturer: Ms Vidya Rangasayee

Email: vidya.rangasayee@sjsu.edu
Office: TBD

Course Information

Lecture/Discussion

Tuesday, Thursday, 4:30 PM to 5:45 PM, MH 224

Course Description and Requisites

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.

Prerequisite(s): CS 156 and Graduate standing. Allowed Declared Major: Computer Science, Bioinformatics, Data Science. Or instructor consent.

Letter Graded

Classroom Protocols

Attendance

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. As a graduate level course, active learning relies a lot on in-class discussion and participation.

Collaboration

Unless mentioned otherwise, all assignments are to be completed individually. You may discuss general topics/strategies with your classmates but all submitted work must be done individually.

Program Information

Diversity Statement - At SJSU, it is important to create a safe learning environment where we can explore, learn, and grow together. We strive to build a diverse, equitable, inclusive culture that values, encourages, and supports students from all backgrounds and experiences.

Course Goals

Artificial intelligence (AI) aims to tackle complex real-world problems with rigorous mathematical tools. In this course you will learn the foundational principles and practice of implementing various AI systems. Specific topics include machine learning, search, Markov decision processes, game playing, constraint satisfaction, graphical models, and logic

Course Learning Outcomes (CLOs)

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.

Course Materials

Suggested Textbooks

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 homework.

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

Throughout this semester, you will be assigned additional reading material that will be made available to you through Canvas.

Availability: Canvas

Course Requirements and Assignments

General Information: General Information

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.

Final 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.

Exams

There will be one midterm and one final exam. Please see the detailed schedule for the date of the midterm. The final exam will be as per University Finals Schedule

Homework

Each homework is usually centered around an application and has both written and programming parts. There will be a total of 4-5 homework assignments on different topics.

Reading Discussion

There will be 1 Reading discussion at the end of every topic. The format for this may either be a in-class discussion or a summary report on canvas. There will 4-5 such reading discussions throughout the semester

✔ Grading Information

Breakdown

Grade	Range	Notes
A	90-100	A- 90-93 A 93-99 A+ 99 and above

Grade	Range	Notes
B grades	80-89.99	B- 80-83 B 83-87 B+ 87-89.99
C grades	70-79.99	C- 70-73 C 73-77 C+ 77-79.99
D grades	60-69.99	D- 60-63 D 63-67 D+. 67-69.99
F grade	0-59.99	

Criteria

Type	Weight	Topic	Notes
Homework	35%		This includes both homeworks and the reading discussions
Final Project	25%		
Midterm Exam	20%		
Final Exam	20%		

University Policies

Per University Policy S16-9 (PDF) (http://www.sjsu.edu/senate/docs/S16-9.pdf), 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 the Syllabus Information (https://www.sjsu.edu/curriculum/courses/syllabus-info.php) web page. Make sure to visit this page to review and be aware of these university policies and resources.

Course Schedule

When	Topic	Notes

When	Topic	Notes
Week 1	Introduction	First Reading assignment due by Week 2 Wednesday
Week 2	AI History and background	AI History Ethics and Responsibility Refresher on probability, linear algebra and python
Week 3	Machine Learning	Linear Regression and Classification
Week 4	Machine Learning	• Stochastic gradient descent • Non-Linear features • Group DRO • Feature Templates
Week 5	Neural Networks	• Backpropagation • Differentiable Programming
Week 6	Neural Networks	• Differential Programming • Various NN architectures and Applications
Week 7	Search	• Modeling Search • Tree Search • UCS, A*, A* relaxation
Week 8	Decision Under Uncertainty	• Markov Decision Processes • Reinforcement Learning • Model based Monte Carlo
Week 9	Game Modeling	• Midterm • Game Modeling • Game Evaluation
Week 10	Game Modeling	• Evaluation Functions • Simultaneous Games • Non-zero sum games
Week 11	Spring Break	NO CLASS
Week 12	Constraint Satisfaction Problems	• Modeling • Beam Search • Local Search
Week 13	Probabilistic Models	• Markov networks • Conditional independence • Bayesian Networks

When	Topic	Notes
Week 14	Probabilistic Models	• Probabilistic Programming • Hidden Markov Models • Forward Backward Algorithm
Week 15	Special Topics in AI	• TBD
Week 16	Special Topics in AI and Project Evaluations	• Topics TBD • Project Presentations in Class