CS 5325 – Reinforcement Learning
Course Syllabus and Details

Instructor: Dr. Aniruddha Bora
Email: aniruddha_bora@txstate.edu
Web Page: https://aniruddhabora.github.io/
Office: Comal 301D
Office Hours: Tu/Th 9:30 am - 12:30 pm

Lectures: MW 9:30 am - 10:50 am
Classroom: Nueces 00201
Course page: Canvas: TBA
Website: TBA

Course Description

Reinforcement Learning (RL) powers some of today's most advanced AI systems, from game-playing agents like AlphaGo to conversational models like ChatGPT. This graduate-level course introduces the core principles of RL, where intelligent agents learn to make decisions by interacting with their environment to maximize long-term rewards. You'll explore fundamental topics such as Markov Decision Processes, Q-Learning, and Policy Gradients, before advancing to Deep Reinforcement Learning and modern techniques like Proximal Policy Optimization (PPO), which is part of the same family of methods used to align AI models, such as ChatGPT, with human feedback. The course blends theoretical foundations with hands-on projects using Python, PyTorch, and OpenAI Gym.

Course Objectives

- The student will be able to explain the core principles of reinforcement learning and Markov decision processes
- The student will be able to implement key RL algorithms, including Q-Learning, Policy Gradients, and PPO
- The student will be able to apply deep reinforcement learning techniques using PyTorch and OpenAI Gym
- The student will be able to analyze and compare the performance of RL methods in simulated environments
- The student will be able to understand the role of RL in AI alignment and real-world applications, including RLHF as used in systems like ChatGPT
- The student will be able to design and complete hands-on RL projects demonstrating theoretical and practical understanding

Course Prerequisites

- None.

Mode of Teaching

The class meets twice a week (Mondays & Wednesdays 9:30 am – 10:50 am). This course follows a blended, interactive teaching approach combining lectures, hands-on coding demonstrations, and project-based learning. Core theoretical concepts will be introduced through lectures with active discussions to promote conceptual understanding. Practical sessions will focus on implementing RL algorithms using Python, PyTorch, and OpenAI Gym, ensuring students gain real-world experience. Slides, class notes and other lecture material will be available on Canvas. A combination of presentation slides and coding examples will be used in lectures. At the end of each topic, students must attempt to solve the assignment problems. Solutions to exercise and homework problems will be discussed in class. All students are expected to work on these problems and participate in the class discussions. In this course, you will be encouraged to think on your own and to discuss solutions with your peers. Students are strongly encouraged to use reference books and course materials provided by the professor.

Course Topics

- Foundations of Reinforcement Learning
  - Markov Decision Processes (MDP)
  - State-Value and Action-Value Functions
  - The Exploration vs. Exploitation Dilemma

- Tabular RL Methods
  - Dynamic Programming for Policy Evaluation and Improvement
  - Monte Carlo Methods
  - Q-Learning and SARSA

- Policy Gradient Methods
  - Stochastic Policies and the REINFORCE Algorithm
  - Actor-Critic Architectures

- Deep Reinforcement Learning
  - Deep Q-Networks (DQN) and Extensions
  - Proximal Policy Optimization (PPO)

- Modern RL Applications
  - RL for Game-Playing Agents
  - RL in Robotics and Autonomous Systems
  - RLHF (Reinforcement Learning from Human Feedback) and its role in aligning AI models like ChatGPT

Textbook

The course has no required textbook. Recommended course materials come from the following sources:

Recommended Textbooks:

- Sutton, R. S., & Barto, A. G. (2018). *Reinforcement Learning: An Introduction* (2nd Ed.). MIT Press.
- Lapan, M. (2020). *Deep Reinforcement Learning Hands-On* (2nd Ed.). Packt Publishing.

Optional Supplementary Resource:

- Morales, M. (2020). *Grokking Deep Reinforcement Learning*. Manning Publications.
- Laura Graesser and Wah Loon Keng (2020). Foundations of Deep Reinforcement Learning: Theory and Practice in Python. Pearson.
- **David Silver**, *Reinforcement Learning* (online lecture series, UCL).
- **OpenAI Gym Documentation and Tutorials**: https://gymnasium.farama.org/

Assessment

Class participation: 10%

Active participation in lectures, discussions, and coding exercises is essential to reinforce key concepts and promote collaborative learning.

Assignments: 50%

Regular assignments will include mostly coding tasks, and implementation of reinforcement learning algorithms using Python and PyTorch and some theoretical excercise. These are designed to progressively build both conceptual understanding and practical skills.

Final Project: 40%

Students will undertake a comprehensive final project applying reinforcement learning techniques to a real-world or simulated task of their choice. The project will demonstrate mastery of course concepts and practical problem-solving abilities. Details and deadlines will be provided during the course.

Late submissions: Late assignment submissions will incur 10% penalty per day, for up to 3 days. After the 3 days, no submission will be accepted.

Missed Exams and Makeup Work: Talk to the instructor if you miss an exam or quiz due to unavoidable circumstances (e.g., health).

Attendance and Drop Policy: Attendance though not mandatory, is HIGHLY encouraged. Class participation is important to your grade in the 'Homework and Quizzes' component.

Grade Grievance Policy: If a student believes a mistake has been made in grading an assignment, the student has **one week** after an assignment is returned to resubmit an assignment for re-grading if they believe there is an error.

Drop Policy: You must follow the withdrawal and drop policy set up by the University and the College of Science. You are responsible for checking the drop deadlines and making sure that the drop process is complete. http://www.registrar.txstate.edu/registration/drop-a-class.html

Accommodations for students with disabilities: Any student with special needs, requiring special accommodations, should inform me during the first two weeks of classes. The student should also contact the office of disability services at the LBJ student center.

Academic Honesty: You are expected to adhere to both the University's Academic Honor Code as described here: http://www.txstate.edu/effective/upps/upps-07-10-01.html, as well as the Computer Science Department Honor Code, described here: 2013 0426 HonestyPolicy CSPPS.doc.

- Except where explicitly and specially allowed (such as group project), all work submitted in the class is expected to be your individual work. Plagiarism will not be tolerated and if detected will result in automatic "F" grade.

Do not include code (or other materials) obtained from the Internet in your assignments (except what is provided or allowed by the instructor).

- Do not email your program to anyone (except your partner or the instructor).

The penalty for submitting a program that has been derived from the internet or any other non-approved source will be a 0 for that assignment. Violators will be reported to the Texas State Honor Code Council (http://www.txstate.edu/honorcodecouncil/).

Mission: Texas State University is a doctoral-granting, student-centered institution dedicated to excellence and innovation in teaching, research, including creative expression, and service. The university strives to create new knowledge, to embrace a diversity of people and ideas, to foster cultural and economic development, and to prepare its graduates to participate fully and freely as citizens of Texas, the nation, and the world.

Shared Values: In pursuing our mission, we, the faculty, staff, and students of Texas State University, are guided by a shared collection of values:

- Teaching and learning based on research, student involvement, and the free exchange of ideas in a supportive environment;
- Research and creative activities that encompass the full range of academic disciplines—research with relevance, from the sciences to the arts, from the theoretical to the applied;
- The cultivation of character, integrity, honesty, civility, compassion, fairness, respect, and ethical behavior in all members of our university community;
- A diversity of people and ideas, a spirit of inclusiveness, a global perspective, and a sense of community as essential conditions for campus life*;*
- A commitment to service and leadership for the public good;
- Responsible stewardship of our resources and environment; and
- Continued reflection and evaluation to ensure that our strengths as a community always benefit those we serve.

Campus Health, Wellness, and Safety

- Reminder about the 12 Health and Safety Guidelines at Texas State, including those to wear a face covering, practice physical distancing, perform a self-assessment before coming to campus, stay home when sick, get tested for COVID-19, and report any positive COVID-19 test to Bobcat Trace as soon as possible. Please note the university recently added two new guidelines to the original ten guidelines.
- Importance of the Bobcat Pledge, including the shared responsibility to practice healthy behaviors and follow the health and safety guidelines, which shows respect for others and helps prevent the spread of COVID-19 on campus and in the surrounding community.
- Link to the Student Roadmap for more information on students' return to campus.

Statement on Civility and Compliance in the Classroom

Civility in the classroom is very important for the educational process and it is everyone's responsibility. If you have questions about appropriate behavior in a particular class, please address them with your instructor first. Disciplinary procedures may be implemented for refusing to follow an instructor's directive, refusing to leave the classroom, not following the university's requirement to wear a cloth face covering, not complying with social distancing or sneeze and cough etiquette, and refusing to implement other health and safety measures as required by the university. Additionally, the instructor, in consultation with the department chair/school director, may refer the student to the Office of the Dean of Students for further disciplinary review. Such reviews may result in consequences ranging from warnings to sanctions from the university. For more information regarding conduct in the classroom, please review the following policies at AA/PPS 02.03.02, Section 03: Courteous and Civil Learning Environment, and Code of Student Conduct, number II, Responsibilities of Students, Section 02.02: Conduct Prohibited.

Emergency Management

In the event of an emergency, faculty, students, and staff should monitor the Safety and Emergency Communications web page. This page will be updated with the latest information available to the university, in addition to providing links to information concerning safety resources and emergency procedures. Faculty, students, and staff are encouraged to sign up for the TXState Alert system.

Sexual Misconduct Reporting (SB 212)

Effective January 2, 2020, state law (SB 212) requires all university employees, acting in the course and scope of employment, who witness or receive information concerning an incident of sexual misconduct involving an enrolled student or employee to report all relevant information known about the incident to the university's Title IX Coordinator or Deputy Title IX coordinator. According to SB 212, employees who knowingly fail to report or knowingly file a false report shall be terminated in accordance with university policy and The Texas State University System Rules and Regulations.

AI Tools Usage Policy

Students may use AI tools such as ChatGPT, GitHub Copilot, or similar technologies as a resource for brainstorming, concept clarification, or coding assistance. However, the following guidelines apply:

- Treat ChatGPT and similar AI tools like a fellow student in this class: Ask questions — do not copy the answers.
- You may seek help or explanations, but you should not copy code or written responses 1-to-1.
- Copying entire solutions from AI tools, the internet, or other sources is considered academic dishonesty and will be subject to penalties outlined in the university's academic integrity policies.
- Think of using AI tools as similar to asking a peer for help — it's acceptable to get guidance, but copying an entire solution is not.
- If in doubt, paste your code or answer into a search engine or ChatGPT and check for near-identical matches to ensure originality.

Responsible AI usage is encouraged to enhance your learning, but your work must reflect your own understanding and effort.

Acknowledgement

Portions of this syllabus, including structural elements and formatting, are adapted from the course "CS 4347 - Introduction to Machine Learning (Spring 2023)" taught by Dr. Vangelis Metsis in the Department of Computer Science at Texas State University. Modifications have been made to reflect the specific content and requirements of this course.
Additionally, some of the recommended resources included in this course were also suggested by Dr. Vangelis Metsis.