CS 59200: Reinforcement Learning

Lecture: 4:30-5:45 MW, LWSN B134
Semester: Fall 2024
Instructor: Joseph Campbell
Email: joecamp@purdue.edu
Prerequisites:

- Basic proficiency in Python. Experience with Pytorch is recommended but not required.
- Introductory-level knowledge of machine learning and linear algebra.

Course Description:

The ability to autonomously acquire skills is a hallmark of an intelligent agent. This can be achieved through a machine learning paradigm known as *reinforcement learning*, in which an agent learns by repeatedly interacting with an environment. This course will explore topics related to reinforcement learning, including deep reinforcement learning, model-based and model-free learning, intrinsically motivated learning, applications, and open challenges.

This course is structured to introduce methods and research topics in reinforcement learning. It consists of lectures, student presentations, and a course project. You will be expected to work in small groups to formulate and carry out a short-term research project related to reinforcement learning. Additionally, you will individually read, critically analyze, and present research papers throughout the semester.

Learning Outcomes:

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

1. Understand how agents can acquire skills through reinforcement learning.
2. Know the strengths and weaknesses of various classes of reinforcement learning algorithms and understand how they can be applied to real-world problems.
3. Effectively read and critically analyze research papers in the field.
4. Implement a reinforcement algorithm and apply it to a research problem.
5. Clearly articulate and present project findings.

Course Topics (Tentative):

- **Introduction to reinforcement learning:** Markov Decision Processes, multi-armed bandits, exploration-exploitation trade-off
- **Dynamic programming:** policy and value iteration
- **Monte Carlo and temporal difference learning:** SARSA, Q-learning, on-policy and off-policy learning
- **Policy gradient methods:** REINFORCE, actor-critic methods
- **Model-based reinforcement learning**

- **Offline reinforcement learning**
- **Intelligent exploration and reward-free learning**
- **Applications and challenges:** robotics, language, sim-to-real, value alignment

Grading:

- Paper presentations: 30%
    - 2x presentation: 15% each
- Course project: 60%
    - Midterm update: 10%
    - Project report (plus deliverables): 30%
    - Poster presentation: 20%
- Class participation and paper discussion: 10%

Please note that the instructor may revise this syllabus as needed. Changes will be communicated to students in a timely manner.