Course Syllabus for SIADS 644: Reinforcement Learning Algorithms

Course Overview and Prerequisites

This course covers the basic principles of reinforcement learning and popular modern reinforcement learning algorithms. Students will develop familiarity with both model-based and model-free reinforcement learning algorithms, including Q-learning, Actor-Critic algorithms, policy-gradient, and Alpha-zero.

The prerequisites for SIADS 644 include:

- Math Methods for Data Science (502),
- Deep Learning (642)
- Data Mining II (632)

Instructor and Course Assistants

Instructor:

- Lei Ying - leiying@umich.edu Professor, Electrical Engineering and Computer Science, Michigan Engineering

Course Assistants:

- Zixian Yang - zixian@umich.edu
- Jamie White- jqwhite@umich.edu

How to Get Help

If you have questions concerning the degree program, encounter a technical issue with Coursera, or issues using Slack, please submit a report to the ticketing system at umsimadshelp@umich.edu.

If you have an issue specific to the Coursera environment, you can receive support through the Coursera Learner Help Center (you may be asked to log in to your Coursera account). Live chat support is also available through the help center.

For questions regarding course content, refer to the **Communications Expectations** section below.

Course Communication Expectations

Slack is the preferred communication tool for this course. If you have questions about course content (e.g. lecture videos or assignments), please make sure to use Slack. Instructor and course assistant response time to Slack messages will be within 24 hours.

Personal communication that may involve sensitive information may be emailed directly to the instructor or course assistant. If you email the instructor or course assistant, please include SIADS 644 in the email subject. Instructor and course assistant response time to email messages will be within 24 hours.

Instructor grading response time will be 4 days after assignments are submitted.

Learning Outcomes

- Understand the principle of optimality and the Bellman equation
- Understand Q-function and model-free reinforcement learning.
- Program deep Q-learning algorithms for problems with discrete action spaces
- Program deep actor-critic algorithms for problems with continuous action spaces
- Program simple Alpha-zero

Course Schedule

- **This course begins on July 26, 2022, and ends on August 22, 2022.**
- Weekly assignments will be **due on Mondays at 11:59 pm** (Ann Arbor, Michigan time; U.S. Eastern time zone).

Week 1 focuses on the mathematical foundation of reinforcement learning. Students will learn dynamic programming and the principle of optimality. Students will focus on the Bellman equations for different types of reinforcement learning problems.

Week 2 focuses on the algorithmic foundation. Students will learn value iteration and policy iteration algorithms. After that, model-free reinforcement learning algorithms such as Q-learning and SARSA will be introduced. Students will also learn linear function approximation for complex reinforcement learning problems. Students will be asked to use linear function approximation to play simple Tetris.

Week 3 focuses on deep learning algorithms such as Q-learning, actor-critic algorithms, and deep policy gradient algorithms. Students will be asked to train a BipedalWalker in OpenAI Gym.

Week 4 focuses on the exploration-exploitation tradeoff and the ideas behind Alpha-zero. Students will learn upper-confidence bound and Monte Carlo Tree Search. Students will be asked to program a simple Alpha-zero based on self-play.

Weekly Office Hours via Zoom (Ann Arbor, Michigan time):

Your instructor will hold weekly, synchronous office hours using the video-conferencing tool, Zoom. The schedule of office hours can be found by clicking on the **Live Events** link in the left-hand navigation menu. Additionally, all office hours will be recorded and archived so that you can retrieve them at a later date. Office hour recordings will be available in the course. Password for all office hours will be 644 and take place in EST time.

- Lei Ying: Thursdays at 2pm
- Zixian Yang: Thursdays at 8pm
- Jamie White: Saturdays at 2pm

Grading

Course Item	**Percentage of Final Grade**	**Due**
Week 1 Homework - Part 1 - Part 2	25%	Monday, August 1, 2022 - 11:59pm (U.S. Eastern time)
Week 2 Homework - Part 1 - Part 2	25%	Monday, August 8, 2022 - 11:59pm (U.S. Eastern time)
Week 3 Homework - Part 1 - Part 2	25%	Monday, August 15, 2022 - 11:59pm (U.S. Eastern time)
Week 4 Homework	25%	Monday, August 22, 2022 - 11:59pm (U.S. Eastern time)
Total	**100%**	

Note: All assignments are required to earn credit for this course.

Letter Grades, Course Grades, and Late Submission Policy

Late assignment submissions will be reduced by 25% per day for 3 days, then 0 points awarded after 4 days.

The grading scale for this course is as follows:

A+	95%
A	90%
A-	85%
B+	80%
B	75%
B-	70%
C+	65%
C	60%
C-	55%
D+	50%

D	45%
D-	40%
F	0%

Academic Integrity/Code of Conduct

Refer to the Academic and Professional Integrity section of the UMSI Student Handbook.

Accommodations

Refer to the Accommodations for Students with Disabilities section of the UMSI Student Handbook. Use the Student Application Form in Accommodate to begin the process of working with the University's Office of Services for Students with Disabilities.

Accessibility

Refer to the Screen reader configuration for Jupyter Notebook Content document to learn accessibility tips for Jupyter Notebooks.

Library Access

Refer to the U-M Library's information sheet on accessing library resources from off-campus. For more information regarding library support services, please refer to the U-M Library Resources section of the UMSI Student Handbook (access to the Student Orientation course required).

Student Mental Health

Refer to the University's Resources for Stress and Mental Health website for a listing of resources for students.

Student Services

Refer to the Introduction to UMSI Student Life section of the UMSI Student Handbook (access to the Student Orientation course required).

Technology Tips

- Recommended Technology
    - This program requires Jupyter Notebook for completion of problem sets and Adobe or other PDF viewer for reading articles.
- Working Offline

While the Coursera platform has an integrated Jupyter Notebook system, you can work offline on your own computer by installing Python 3.5+ and the Jupyter software packages, including pyspark. For more details, consult the Jupyter Notebook FAQ.