TO: The Faculty of the College of Engineering

FROM: Elmore Family School of Electrical and Computer Engineering

RE: New Graduate Course, ECE 50271 Reinforcement Learning – Theory and Algorithms

The faculty of the School of Electrical and Computer Engineering has approved the following new course. This action is now submitted to the Engineering Faculty with a recommendation for approval.

ECE 50271 Reinforcement Learning – Theory and Algorithms
Sem. 1, Lecture 3, Cr. 3.
Prerequisite: Undergraduate understanding of Linear Algebra, Probability, and Calculus OR graduate student standing.

Description: This course introduces the foundations and the recent advances of reinforcement learning, an area of machine learning closely tied to optimal control that studies sequential decision-making under uncertainty. This course aims to create a deep understanding of the theoretical and algorithmic foundations of reinforcement learning while discussing the practical considerations and various extensions of reinforcement learning.

Reason:
This course has focus on basics and framework related to reinforcement learning. This course provides fundamental knowledge about reinforcement learning, which is an active area of research. It is expected that it will attract many students at both undergraduate and graduate levels within ECE and beyond that.

Course History: Fall 2023 – 84, Fall 2024 – 89

Mithuna Thottethodi,
Associate Head for Teaching and Learning
Elmore Family School of Electrical and Computer Engineering

ECE 50271: Reinforcement Learning - Theory and Algorithms

Fall 2024

Instructor: Prof. Mahsa Ghasemi (mahsa@purdue.edu), Office at MSEE 238

Teaching Assistant: Qasim Elahi (elahi0@purdue.edu)

Course Credits: 3

Teaching Mode: On-campus and online

Lectures: Tuesday and Thursday, 1:30-2:45 pm at Seng-Liang Wang Hall 2599

Office Hours: The current schedule of the office hours is as follows.

- **Instructor's office hours:**
  Tuesday, 3-4 pm in hybrid format, at MSEE 238 and over Zoom
  Thursday, 4-6 pm in hybrid format, at MSEE 238 and over Zoom
  For meeting outside of these time frames, please schedule a time via email.
  Zoom link: https://purdue-edu.zoom.us/j/93301419449?pwd=SmN5LzJSSEJPS0Z4RTdFMjZGd2V2dz09

- **TA's office hours:**
  Monday, 11 am-1 pm in hybrid format, at BHEE 208 (Table 2) and over Zoom
  Friday, 6-7 pm in online format, over Zoom
  Zoom link: https://purdue-edu.zoom.us/j/9924399450

Educational Platforms: The course will utilize Brightspace, Gradescope, and Piazza. The students are expected to ensure they have access to all three platforms and will receive all relevant communication through these platforms.

- Brightspace: https://purdue.brightspace.com/d2l/login
  The course materials, including lecture notes and lecture recordings, assignment announcements, and grades will be posted on Brightspace.

- Gradescope: https://www.gradescope.com/
  All assignments and exams should be submitted through Gradescope. The grading of assignments and exams, as well as the regrade requests, will be handled through Gradescope.

- Piazza: https://piazza.com/purdue/fall2024/ece59500rlt/home
  We will use Piazza for class discussion. If you have any questions about the course content, assignments, or project, we encourage you to post them on Piazza. It is a shared discussion forum, where you can get help fast and efficiently from your classmates, the TA, and the instructor while benefitting other students with similar questions.

Prerequisites: Undergraduate understanding of linear algebra (MA 26500 or equivalent), probability (ECE 30200 or equivalent), calculus (MA 26100 or equivalent)

Course Description: This course introduces the foundations and the recent advances of reinforcement learning, an area of machine learning closely tied to optimal control that studies sequential decision-making under uncertainty. This course aims to create a deep understanding of the theoretical and algorithmic foundations of reinforcement learning while discussing the practical considerations and various extensions of reinforcement learning.

Required Texts: None

Recommended Texts:

- Reinforcement Learning: An introduction, Richard S. Sutton, and Andrew G. Barto, MIT press, 2018.
- Markov Decision Processes: Discrete Stochastic Dynamic Programming, Martin L. Puterman, John Wiley & Sons, 2014.
- Dynamic Programming and Optimal Control, Dimitri P. Bertsekas, Belmont, MA: Athena Scientific, 2011.
- Neuro-Dynamic Programming, Dimitri P. Bertsekas, and John N. Tsitsiklis, Athena Scientific, 1996.
- Bandit Algorithms, Tor Lattimore, and Csaba Szepesvári, Cambridge University Press, 2020.
- Reinforcement Learning: Theory and Algorithms, Alekh Agarwal, Nan Jiang, Sham M. Kakade, and Wen Sun, 2019.
- Foundations of Deep Reinforcement Learning, Laura Graesser, and Wah Loon Keng, Addison-Wesley Professional, 2019.

Learning Outcomes:

A student who successfully fulfills the course requirements will have demonstrated an ability to:

- Understand different problem formulations for reinforcement learning
- Apply various algorithmic solutions to a wide range of sequential decision-making problems
- Analyze the performance capabilities and limitations of different algorithm for sequential decision-making
- Conduct a research project by collaborating with one or more partners and write a scientific report of the research findings

Assessment and Grading: The course assessment will be based on the following components. The cut-offs for the final letter grade will be determined at the end of the semester according to the grade distribution but will not be released to the students.

- Homework --- 40%
- Midterm Exam --- 30%
- Final Project --- 30%
- Bonus --- Up to 4% bonus points may be considered for engagement in class and over Piazza, as well as completing course evaluation surveys

Homework: There will be around 5-7 assignments, containing a mixture of theory-oriented questions and applied questions, throughout the semester. The assignments will be due every 1-2 weeks. Unless otherwise specified, assignments are due at 11:59 PM on the deadline. Late submissions will be generally accepted up to two days late, docking 15% from the score per day late. Submissions more than two days late will be assigned a score of 0. Except for medical and family emergencies (accompanied by verification), there will be no individual extensions granted for assignments. The assignments should be submitted through Gradescope for grading. All regrade requests will be handled through Gradescope as well. All such requests must be made within one week of when the grade was posted. Regrade requests arriving later than this will not be considered. Please note that during the regrading process, we reserve the right to regrade all portions of your submission.

Midterm Exam: There will be a midterm exam in a take-home format to assess the students' understanding of the fundamental concepts covered roughly in the first half of the semester. Tentatively, it will take place during the 9th week of the semester. No form of discussion or collaboration among the students is allowed for the midterm exam. The exam should be submitted through Gradescope for grading. All regrade requests will be handled through Gradescope as well. All such requests must be made within one week of when the grade was posted. Regrade requests arriving later than this will not be considered. Please note that during the regrading process, we reserve the right to regrade all portions of your submission.

Final Project: The final project will be delivered in three stages. In the first stage, the students will submit a short research proposal on their topic of choice related to the course content. In the second stage, the students will submit a progress report to describe the progress of their project and potential changes. In the third stage, the students will submit a final report, code required to replicate the results, and a recorded presentation. The final report should follow the format guidelines of a journal or a conference paper. The code should be accompanied by a text file explaining how it may be run. The recorded presentation is expected to be around 10 minutes. The students may do the final project in groups of two or three; the expected effort and output will be proportional to the size of the group.

Computational Resources: The students may run their code on the following platforms, if desired.

- Scholar cluster at Purdue RCAC: https://www.rcac.purdue.edu/compute/scholar
  You will have a Scholar account created for you at the beginning of the semester.

- Google Colab: https://colab.research.google.com/

Lecture Outline: Below is a tentative outline of the lectures.

- Fundamentals:
  - Introduction, Motivation, Overview of Relevant Background
  - Markov Models and Markov Decision Processes
  - Dynamic Programming and Policy Evaluation
  - Value Iteration and Policy Iteration
  - Finite-Horizon and Infinite-Horizon Problems
  - Monte Carlo and Temporal Difference Methods
  - Off-policy Methods
- Learning Problems and Algorithms:
  - Model-Based and Model-Free Reinforcement Learning
  - Deep Reinforcement Learning Methods
  - Approximate Policy Iteration
  - Conservative Policy Iteration
  - Stochastic Gradient Descent and Policy Gradient Methods
  - Trust Region Methods
  - Multi-Armed Bandits and Exploration in Reinforcement Learning
- Special Topics:
  - Linear Quadratic Regulators and Optimal Control
  - Reinforcement Learning with Partial Observability
  - Risk-Averse Reinforcement Learning
  - Inverse Reinforcement Learning
  - Reinforcement Learning from Human Feedback
  - Meta-Learning and Transfer Learning
  - Multi-agent Reinforcement Learning

Attendance Policy:

The students registered for the in-person section are expected to attend all classes in person unless they are ill or otherwise unable to attend class. The students registered for the online section are expected to watch the recordings of all lectures soon after they are posted.

Academic Honesty:

Unless expressly allowed, you are expected to complete all assignments by yourself and submit your own independent solution. However, you may discuss the problems and high-level solution approaches with other students. Use of AI tools for solving the problems in the assignments and exams or advancing the project is not allowed.

Punishments for academic dishonesty are severe, including receiving an F in the course or being expelled from the University. By departmental rules, all instances of cheating will be reported to the Dean of Students.

Classroom Guidance Regarding Protect Purdue:

The Protect Purdue Plan, which includes the Protect Purdue Pledge, is campus policy and as such all members of the Purdue community must comply with the required health and safety guidelines. Students who do not comply with the required health behaviors are violating the University Code of Conduct and will be reported to the Dean of Students Office with sanctions ranging from educational requirements to dismissal from the university.

Any student who has substantial reason to believe that another person in a campus room is threatening the safety of others by not complying may leave the room without consequence. The student is encouraged to report the behavior to and discuss next steps with their instructor. Students also have the option of reporting the behavior to the Office of the Student Rights and Responsibilities. See also Purdue University Bill of Student Rights.

If you feel ill, have any symptoms associated with COVID-19, or suspect you have been exposed to the virus, you should stay home and contact the Protect Purdue Health Center (496-INFO). The student still needs to inform the instructor of any conflict that can be anticipated and will affect the submission of an assignment or the ability to take an exam. When conflicts can be anticipated, such as for many university-sponsored activities and religious observations, the student should inform the instructor of the situation as far in advance as possible. For unanticipated or emergency conflict, when advance notification to an instructor is not possible, the student should contact the instructor as soon as possible by email. When the student is unable to make direct contact with the instructor and is unable to leave word with the instructor's department because of circumstances beyond the student's control, and in cases of bereavement, quarantine, or isolation, the student or the student's representative should contact the Office of the Dean of Students via email or phone at 765-494-1747. Our course Brightspace includes a link on Attendance and Grief Absence policies under the University Policies menu.

Campus Interruptions:

In the event of a major campus emergency, course requirements, deadlines and grading percentages are subject to changes that may be necessitated by a revised semester calendar or other circumstances beyond the instructor's control.

Nondiscrimination Statement:

Purdue University is committed to maintaining a community which recognizes and values the inherent worth and dignity of every person; fosters tolerance, sensitivity, understanding, and mutual respect among its members; and encourages each individual to strive to reach his or her own potential. In pursuit of its goal of academic excellence, the University seeks to develop and nurture diversity. The University believes that diversity among its many members strengthens the institution, stimulates creativity, promotes the exchange of ideas, and enriches campus life. More details are available on our course Brightspace table of contents, under University Policies.

Accessibility:

Purdue University strives to make learning experiences as accessible as possible. If you anticipate or experience physical or academic barriers based on disability, you are welcome to let me know so that we can discuss options. You are also encouraged to contact the Disability Resource Center at: drc@purdue.edu or by phone: 765-494-1247. More details are available on our course Brightspace under Accessibility Information.

Mental Health Statement:

If you find yourself beginning to feel some stress, anxiety and/or feeling slightly overwhelmed, try WellTrack. Sign in and find information and tools at your fingertips, available to you at any time.

If you need support and information about options and resources, please contact or see the Office of the Dean of Students. Call 765-494-1747. Hours of operation are M-F, 8 am- 5 pm.

If you find yourself struggling to find a healthy balance between academics, social life, stress, etc. sign up for free one-on-one virtual or in-person sessions with a Purdue Wellness Coach at RecWell. Student coaches can help you navigate through barriers and challenges toward your goals throughout the semester. Sign up is completely free and can be done on BoilerConnect. If you have any questions, please contact Purdue Wellness at evans240@purdue.edu.

If you're struggling and need mental health services: Purdue University is committed to advancing the mental health and well-being of its students. If you or someone you know is feeling overwhelmed, depressed, and/or in need of mental health support, services are available. For help, such individuals should contact Counseling and Psychological Services (CAPS) at 765-494-6995 during and after hours, on weekends and holidays, or by going to the CAPS office of the second floor of the Purdue University Student Health Center (PUSH) during business hours.

Disclaimer:

This syllabus is tentative and subject to change. Any substantive change will be communicated to the students.