THE UNIVERSITY OF TEXAS AT AUSTIN
Cockrell School of Engineering

ASE 389: Thereotical Foundations of Reinforcement Learning Algorithms
Spring 2026

SYLLABUS

Unique Number:	15505
Instructor:	Thinh T. Doan
	Office: ASE 3.206 Email: thinhdoan@utexas.edu
Time:	MW 8:30-9:45 CT
Lecture location:	ASE 1.124
Office hour:	W 1-2PM CT in the conference room ASE 3.204
Web Page:	Canvas

Catalog Description:
Reinforcement learning (RL) has emerged as one of the predominant frameworks for real-time decision making and control. It has been the driving force behind several recent high-profile successes of artificial intelligence, enjoying success in areas as diverse as, computer games, robotic control, training large language models, etc. This course is about the mathematical foundation of reinforcement learning algorithms where we will cover three related themes: 1) mathematical models of RL problems; 2) developing and analyzing the complexity of RL algorithms; and 3) advanced topics in RL including multi-task learning, game theory, and distributed learning. Students will also have opportunities to work on several hands-on programming assignments, followed by final projects on the applications of RL.

Course Objectives:
The course is designed for graduate students with strong interests in theory of reinforcement learning, control, and optimization. Technical theorems and proofs of optimization and learning algorithms will be provided. Overall, the ultimate goal is to:

- Introduce mathematical models of reinforcement learning.
- Develop learning algorithms and derive their performance complexity.
- Apply and simulate reinforcement learning to various applications.
- Expose students to recent advances in the literature of reinforcement learning.

Prerequisites:
Sufficient background in linear algebra and real analysis (e.g., M341, M346, M361K, M365C),

Date	Topics	Lectures
1/12 - 1/21	Markov decision processes: Bellman equations, value iterations, and policy iterations	Lecture 1
1/26 - 2/2	Policy evaluation problems: Temporal-difference (TD) learning and gradient TD	Lecture 2
2/4 - 2/9	Q-learning	Lecture 3
2/11 - 2/18	Stochastic Approximation: Single-time-scale and multi-time-scale variants	Lecture 4
2/23 - 3/4	Policy Optimization: Policy gradient and actor-critic methods	Lecture 5
3/9 - 3/18	Function approximation in RL: Linear function approximation vs. neural networks	Lecture 7
3/23 - 3/30	Stochastic optimization: Stochastic gradient and mirror descent methods	Lecture 6
4/1 - 4/15	Advanced topics: Multi-agent RL, POMDP, etc.	Lectures 8 & 9 & 10
4/20 - 4/27	Course project presentations	

Table 1: Tentative topics and schedule

probability (e.g., M362K, M362M, ECE 351K), and control theory (e.g., ASE 381P). Knowledge in graduate classes in stochastic decision processes (e.g., ECE 381J) and optimization (ECE 381K) will be very useful. Students are encouraged to discussed with the instructor if they have concerns about the prerequisites

Attendance:
Regular attendance is expected. If you are sick, please do not attend so that your classmates stay healthy!

Tentative Topics and Class Schedule:
A list of tentative topics and class schedule is presented in Table 1. The content of each lecture and assignment due dates are subject to change. There will be no exams. The final project writeups will be due on 5/1 at 11:59pm CT.

Text:
There is no required textbook for this class. Lecture notes will be provided. However, below are some useful references, where some useful materials will be used for our class. Other advanced theoretical results will be taken from research papers.

References:

1. Andrew, Barto, and Sutton Richard S. "Reinforcement Learning: An Introduction." (2018).
2. Martin L. Puterman, Markov Decision Processes: Discrete Stochastic Dynamic Programming, John Wiley & Sons Inc., 1994.
3. Bertsekas, D. P. "Neuro-dynamic programming." Athena Scientific (1996).
4. Meyn, Sean. Control systems and reinforcement learning. Cambridge University Press, 2022.
5. Tamer Basar, Geert Jan Olsder. Dynamic Noncooperative Game Theory.

6. Vivek S. Borkar. Stochastic Approximation A Dynamical Systems Viewpoint.
7. Lan, Guanghui. First-order and stochastic optimization methods for machine learning. Vol. 1. Cham: Springer, 2020.
8. Nesterov, Yurii. Lectures on convex optimization. Vol. 137. Berlin: Springer, 2018.
9. Beck, Amir. First-order methods in optimization. Society for Industrial and Applied Mathematics, 2017.

Assignments:
There will be four or five assignments, which include both theoretical questions and programming tasks. Each student is expected to write their own solution. In lieu of exams, we will have a final project which will include both a written and oral component. The written report will be similar to a formal conference paper submission.

Grading:
Grades will be computed according to the following proportions:

- Final project: 40%.
- Assignments: 50%.
- Class presentations: 10%.

Grades will be assigned as letter grades with plus/minus modifiers as appropriate.

Homework Policy:
Assignments should be completed and submitted independently by each student. Late submissions will not be accepted. Exceptions may be granted at the instructors discretion.

Examinations:
There are no exams in this class.

Important Dates:
Please refer to the university academic calendar for important administrative dates.

Special Notes:
The University of Texas at Austin provides upon request appropriate academic adjustments for qualified students with disabilities. For more information, contact the Office of the Dean of Students at 471-6259, 471-4641 TDD or the Cockrell School of Engineering Director of Students with Disabilities at 471-4321.

Evaluation:
The Measurement and Evaluation Center forms for the Cockrell School of Engineering will be used during the last week of class to evaluate the course and the instructor. They will be conducted in an electronic format.

Class Recordings:
Any recordings which are made are reserved only for students in this class for educational purposes and are protected under FERPA. The recordings should not be shared outside the class in any form. Violation of this restriction by a student could lead to Student Misconduct proceedings. See below for more details.

"Audio or visual recording devices may not be used in University classrooms or laboratories unless specifically approved by the instructor or Disability and Access. Instructors may create electronic recordings of their classes and have the discretion to allow or disallow students or visitors to make electronic recordings of their classes. Any recording(s) must be limited to legitimate educational purposes and may not be shared with or distributed to others in any form unless otherwise permitted by HOP 9-1610, Student Rights Under the Family Educational Rights and Privacy Act (FERPA). If an instructor permits a student to electronically record their class, the recording(s) must be used solely for the educational use of the student or other students in that section and year of the class. Individuals who violate these limitations may be subject to disciplinary proceedings. An instructor must allow a student to electronically record their class if the student has received an accommodation through Disability and Access that permits the student to electronically record class.

DISABILITY & ACCESS (D&A):

The University of Texas at Austin is committed to creating an accessible and inclusive learning environment consistent with university policy and federal and state law. Please let me know if you experience any barriers to learning so I can work with you to ensure you have equal opportunity to participate fully in this course. If you are a student with a disability, or think you may have a disability, and need accommodations please contact Disability & Access (D&A). Please refer to the D&A website for more information: http://diversity.utexas.edu/disability/. If you are already registered with D&A, please deliver your Accommodation Letter to me as early as possible in the semester so we can discuss your approved accommodations and needs in this course.. If you are already registered with D&A, please deliver your Accommodation Letter to me as early as possible in the semester so we can discuss your approved accommodations and needs in this course.

University Resources:

For a list of university resources that may be helpful to you as you engage with and navigate your courses and the university, see the University Resources Students Canvas page.