ECE 414/517: Reinforcement Learning

Fall 2026 Syllabus

University of Tennessee, Knoxville

Course meetings:	Tuesdays and Thursdays, 12:55-2:10 PM
Location:	Min Kao Building, Room 405 (MKB 405)
Instructor:	Fei Liu, fliu33@utk.edu
Office hours:	Thursdays, 2:15-3:00 PM, Min Kao Building, Room 612; or by appointment
TA:	Farong Wang, fwang31@vols.utk.edu
Website:	https://lnnx2006.github.io/teaching/ece414.517_fall26/index.html

Course Description

This course develops the foundations of reinforcement learning (RL): how an agent can learn to make sequential decisions through interaction with an environment. The course begins with multi-armed bandits and Markov decision processes, then develops dynamic programming, Monte Carlo methods, temporal-difference learning, tabular control, planning, function approximation, and policy-gradient methods. The final portion introduces deep reinforcement learning, reward design and learning from preferences, and reinforcement learning beyond the Markov property.

The emphasis is on understanding the assumptions, derivations, and behavior of core algorithms rather than simply applying software libraries. Homework, coding assignments, and a semester-long project connect the theory to complete RL learning systems and evaluation.

Learning Objectives

By the end of the course, students should be able to:

- formulate sequential decision problems as multi-armed bandits or Markov decision processes;
- derive and apply Bellman equations and dynamic-programming methods;
- compare Monte Carlo and temporal-difference prediction and control;
- explain on-policy, off-policy, model-free, and model-based learning;
- analyze value-function approximation, policy-gradient, actor-critic, and deep-RL methods;
- implement representative RL algorithms and evaluate them using reproducible experiments; and
- diagnose learning behavior and communicate experimental findings clearly.

Prerequisites

Familiarity with probability, linear algebra, basic optimization, and programming is expected. The course does not include a separate mathematics or coding foundations unit; students should refresh these topics independently as needed.

Required and Recommended Materials

Primary text: Richard S. Sutton and Andrew G. Barto, *Reinforcement Learning: An Introduction*, second edition. The book is available from the authors at the Sutton-Barto book webpage.

Additional reference: David Silver, *Reinforcement Learning Course*, University College London.

Course-specific materials, announcements, assignment specifications, and grades will be posted on Canvas. Readings listed in the schedule refer to Sutton and Barto (S&B), second edition.

Assessment and Grading

Component	Weight
Class participation and in-class quizzes	5%
Three written homework assignments	30%
Three coding assignments	45%
Combined midterm/final project	20%
Total	**100%**

Each written homework assignment and its corresponding coding assignment will be released together. Unless an assignment specification states otherwise, submissions will be made through Canvas.

Late policy: Work submitted up to 24 hours after the deadline will receive a 25% penalty. Work submitted more than 24 hours late will not be accepted unless an exception has been approved by the instructor in advance or is required by university policy.

Letter-grade scale: 90-100 = A, 80-90=B, 70-80=C, 60-70=D, 0-60=F

Combined Course Project

The midterm and final are two milestones of one combined course project.

1. **Midterm milestone-build a complete RL system.** Students will define the environment interface, state and action spaces, reward signal, agent interface, experiment loop, logging, and evaluation pipeline. Project integration takes place on October 13, followed by system demonstrations, testing, and milestone review on October 15.
2. **Final milestone-train and evaluate an RL agent.** Students will use the theory and algorithms learned in the course to train an RL agent in simulation, establish appropriate baselines and evaluation metrics, analyze experimental results, and communicate lessons learned. A short progress report is due November 10. Final posters and live demonstrations will be presented on December 1.

Detailed project requirements, group rules, deliverables, rubrics, and final-report deadlines will be announced in class and posted on Canvas.

Course Expectations and Policies

Participation and attendance. Students are expected to attend class, participate in discussions and activities, and complete in-class quizzes. Students who must miss class for an approved reason should notify the instructor as early as possible and remain responsible for the material and announcements.

Collaboration. Unless explicitly authorized, submitted solutions, code, analyses, plots, and reports must represent the student's own work. Students must identify permitted collaborators and external resources as directed by the assignment.

Generative AI and coding assistants. No AI-generated deliverables allowed.

Communication. Course announcements and assignment updates will be posted on Canvas. Students are responsible for checking Canvas and their university email regularly.

Academic integrity. Students are expected to follow the University of Tennessee's Honor Statement and current academic-integrity policies. Suspected violations will be handled according to university procedures.

Accessibility and accommodations. Students who need accommodations should contact the appropriate UTK accessibility office and provide the instructor with the official accommodation notice as early as possible.

Changes to the syllabus. The instructor may adjust the schedule or course policies when necessary. Material changes will be communicated in class and through Canvas.

Last updated: August 18, 2026