Topics in Multiagent Learning

MIT 6.7980 — Fall 2026

Lecture: Tuesdays and Thursdays, 11:00 am–12:30 pm, in room `E25-111`.

Instructors:

- Prof. Gabriele Farina, office `45-501F` (in the College of Computing building)
  ✉ `gfarina@mit.edu` URL: `www.mit.edu/~gfarina`
- Prof. Constantinos Daskalakis, office `32-G694` (in the Stata building)
  ✉ `costis@csail.mit.edu` URL: `people.csail.mit.edu/costis`

We are happy to meet with students by appointment.

Teaching assistants:

- Kat Federova (✉ `fedorova@mit.edu`). Office hours: TBD.
- Mingyang Liu (✉ `liumy19@mit.edu`). Office hours: TBD.
- Daniel Xia (✉ `dxia03@mit.edu`). Office hours: TBD.
- Rui Yao (✉ `rayyao@mit.edu`). Office hours: TBD.

Grading:

- Attendance and participation (see below) — 20%.
- Improving material (see below) — 30%.
- Project (see below) — 50%.

Attendance: We expect everyone to attend at least 50% of the lectures. The "Attendance and participation" component of grading above reflects (in a binary manner) whether that was met. We will measure attendance using random quizzes during classes.

Coursework: There are no assigned homework sets. Students will contribute to the shared course materials and complete a project, as described below.

GitHub: Main course repository

Prerequisites: Discrete Mathematics and Algorithms at the advanced undergraduate level; mathematical maturity.

Lecture notes: Lecture notes are available as HTML and PDF on the course website. Announcements and administrative materials will be posted on Canvas.

Collaboration policy: We encourage working together on the course materials, projects, and discussion of the ideas. Contributions and project work should reflect *your own* understanding. Acknowledge collaborators and sources, and do not present somebody else's work as your own.

Acceptable AI use: [1] Students may use generative AI tools (e.g., LLMs) to support learning, brainstorming, editing, debugging, or generating explanations. We believe that *your human understanding* is still the ultimate goal of education more broadly. We should aim to expand our knowledge, deepen our understanding, and sharpen our thinking. Generative AI might be a great tool for aiding this process, but learning and growth require productive struggle. Don't let generative AI remove or reduce your productive struggle. We also expect (and require) *full transparency* regarding your use of AI. Always indicate if and how you used generative AI, i.e., which tools you used and during which phases of your work.

[1] This policy is inspired *in part* from the template policy provided by the MIT Teaching + Learning Lab.

Description

This course studies multiagent systems through game theory, optimization, and learning theory. We cover foundational topics such as Nash equilibria, regret minimization, learning dynamics, and extensive-form games.

We also explore modern topics: multiagent deep reinforcement learning; information and mechanism design; team games and hidden-role games; alignment; high-dimensional and kernelized learning; nonconvex games; calibration; and the complexity of finding equilibria. Applications and open research questions connect the theory to multiagent AI.

Improving Material

We would like to make the lecture notes available to as many people as possible. You can now read them in a browser, follow links between sections and references, and move between the notes and their source. We would like everyone's help to make this a useful resource for learners around the world.

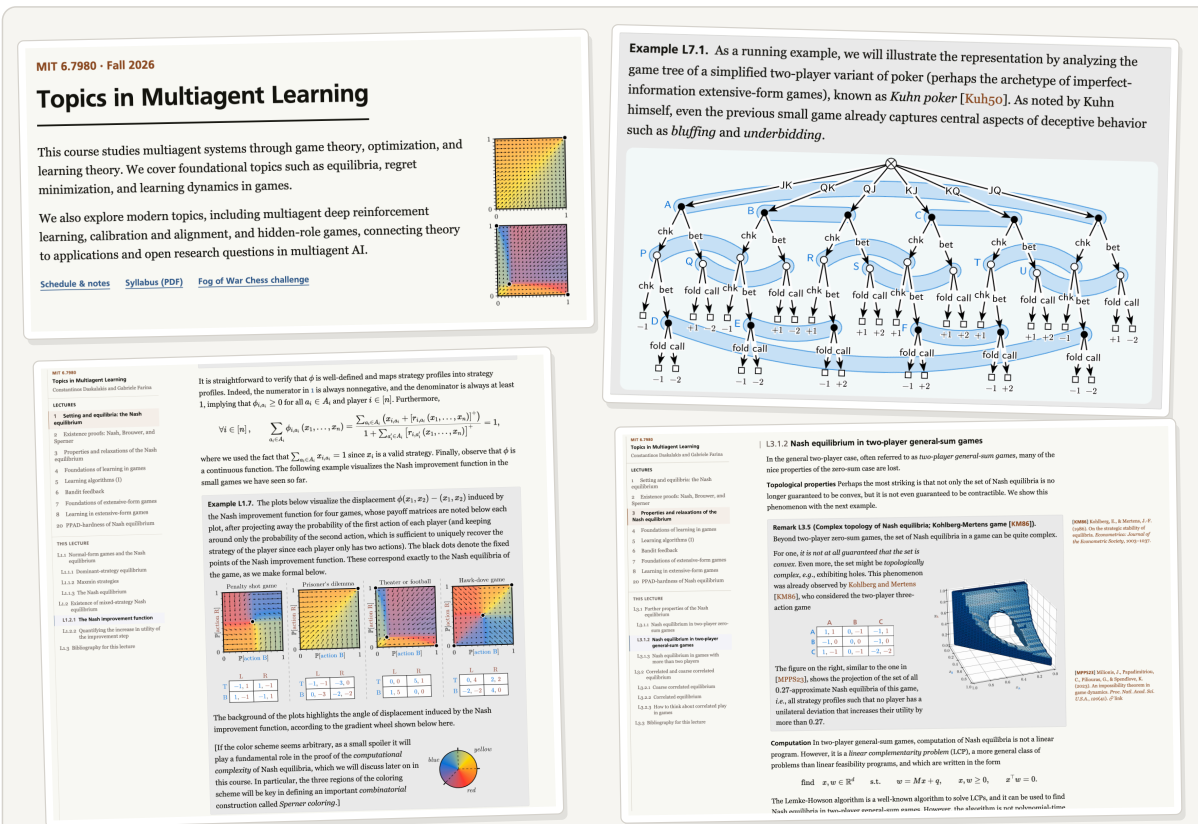

We will divide the class into groups, each focusing on a different part of the material. Using the class GitHub repository, each group can open issues to identify improvements and submit pull requests to implement them. We will improve the material together, reviewing and building on one another's contributions.

Contributions can include clarifying explanations and proofs, fixing errors, adding examples and homework-style exercises for future readers, and polishing figures, organization, and presentation. If anyone is brave enough, we would also love interactive components that let readers experiment with the ideas.

On the bright side, there is no homework! :-) Improving the shared material accounts for 30% of the course grade.

Project

Projects may be completed individually or in groups of 2-5 students and will include a presentation. We will offer three project directions:

Fog of War Chess Challenge. Build and evaluate an agent that plays with partial information. Explore how it uses observations, reasons about uncertainty, and chooses strategic actions. Each bot sandbox is allocated two CPU cores and 4 GiB of memory. A dedicated document will describe the challenge, including the rules, starter code, and how to access the arena.

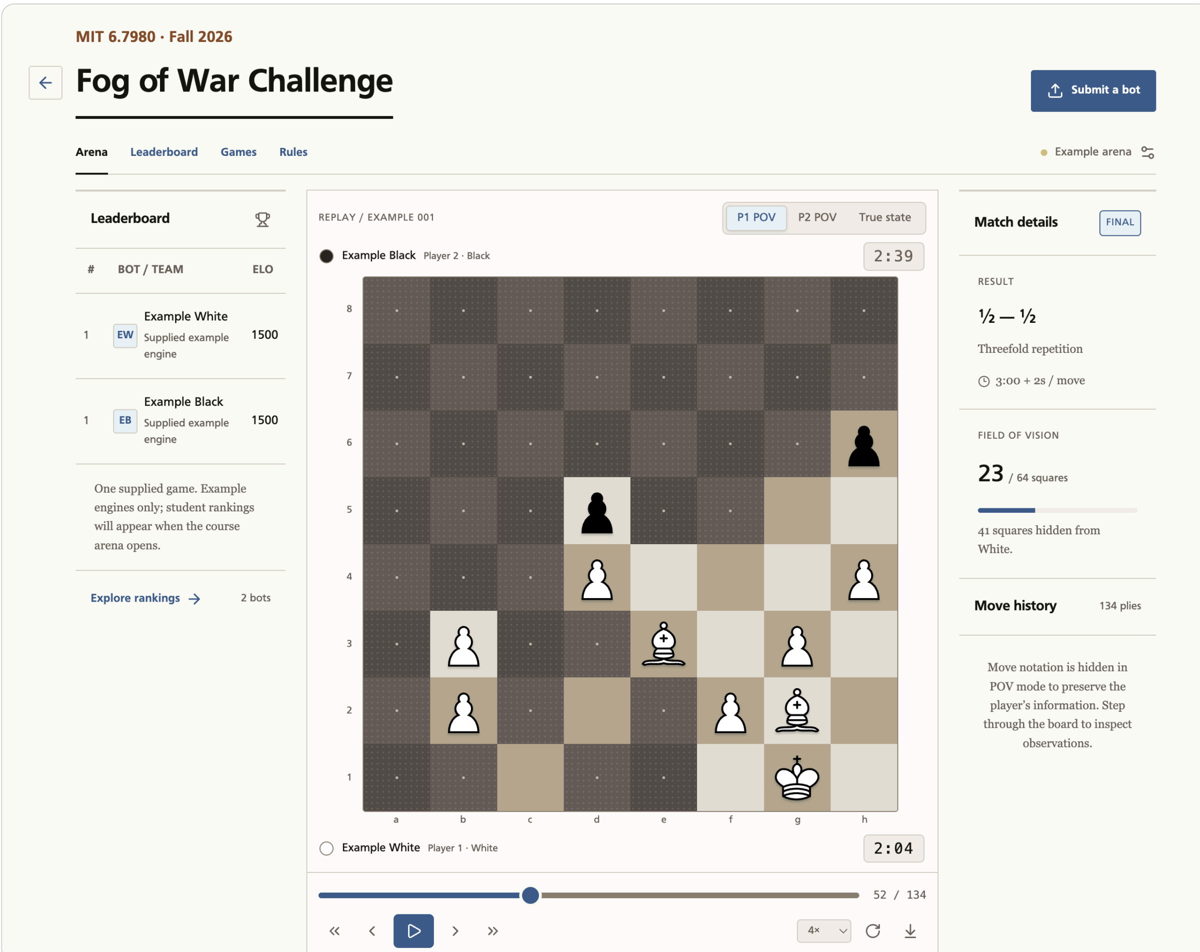

Modeling questions. Formulate a multiagent problem by specifying the players, objectives, information, and available actions. Study how modeling choices affect the resulting strategic behavior. We will provide a separate document with possible modeling questions and leads to explore.

Theory questions. Investigate a mathematical question about equilibria, learning dynamics, or computational complexity. Develop rigorous proofs, bounds, or counterexamples that clarify the behavior of multiagent systems. We will provide a separate document with possible theory questions and leads to explore.

The project is the central component of the course and accounts for 50% of the final grade. We will therefore be "robust" in our grading: we will look carefully at the depth of your understanding, the quality and substance of your work, and how clearly you explain your results.

Tentative Schedule

0	Sep 10th	**Course Overview** We will discuss the syllabus, projects, administrative details, and an overview of the topics covered.

PART I: FOUNDATIONS

1	Sep 15th	**Setting and equilibria: the Nash equilibrium** Nash's existence theorem and its connection to fixed-point theorems.
2	Sep 17th	**Brouwer and Sperner** Sperner's lemma, Brouwer's theorem, and combinatorial proofs of equilibrium existence.
3	Sep 22nd	**Properties of Nash equilibrium** Topological and computational properties. Zero-sum games and linear programming. Correlated and coarse correlated equilibria.
4	Sep 24th	**Learning in games: Foundations** Regret and hindsight rationality. Regret minimization and its relationships with equilibrium concepts.
5	Sep 29th	**Learning in games: Algorithms** General principles for learning algorithms. Follow-the-leader, regret matching, multiplicative weights, and online mirror descent.
6	Oct 1st	**Learning with bandit feedback** Partial feedback and exploration. From multiplicative weights to `Exp3`; regret guarantees.
7	Oct 6th	**Modeling extensive-form games** Perfect and imperfect information. Kuhn's theorem. Normal-form and sequence-form strategies.
8	Oct 8th	**Learning in extensive-form games** No-regret learning, counterfactual utilities, and counterfactual regret minimization (CFR).
	Oct 13th NO CLASS	**No class** MIT follows a Monday schedule.
9	Oct 15th	**Multiagent deep RL** Reinforcement learning in games. Self-play and deep learning methods for perfect-information games.
10	Oct 20th PROJECT	**Taking stock** We will discuss big open questions in the field and possible project ideas.

PART II: INFORMATION, COMMUNICATION, ALIGNMENT

11	Oct 22nd	**Information and mechanism design** Designing information and incentives in strategic interactions.

12	Oct 27th	**Team games and hidden-role games** Coordination in teams and games with hidden roles.
13	Oct 29th	**Alignment (part I)** Reinforcement learning from human feedback (RLHF) and alignment.
14	Nov 3rd	**Alignment (part II)** Regularized RLHF and direct preference optimization (DPO).
	PART III: ADVANCED LEARNING	
15	Nov 5th	**Forecasting and calibration** Calibrated prediction and its connections to learning in games.
16	Nov 10th	**High-dimensional games** Learning with large strategy spaces. Kernelized methods and multiplicative weights.
17	Nov 12th	**Nonconvex games** Nonconvexity, learning dynamics, and local equilibrium concepts.
	PART IV: COMPUTATIONAL COMPLEXITY	
18	Nov 17th	**Total search and TFNP** Total search problems, the TFNP framework, and the PPAD complexity class.
19	Nov 19th	**PPAD-hardness of Nash equilibrium** Reductions and the computational hardness of finding Nash equilibria.
	PROJECT WORK AND PRESENTATIONS	
	Nov 24th BREAK	**No class** Project break
	Nov 26th NO CLASS	**No class** Thanksgiving holiday.
20	Dec 1st PROJECT	**Project presentations** Show us your cool work!
21	Dec 3rd PROJECT	**Project presentations** Show us your cool work!
22	Dec 8th PROJECT	**Project presentations** Show us your cool work!
23	Dec 10th PROJECT	**Project presentations** Show us your cool work!