# Topics in Multiagent Learning

## MIT 6.S890 — Fall 24

**Lecture:** Tuesdays and Thursdays, 11:00am-12:30pm, 3-333.

**Instructor:** Prof. Gabriele Farina, office 45-501F (in the new College of Computing building)
✉ gfarina@mit.edu URL: www.mit.edu/~gfarina
I am happy to meet with students by appointment.

**Teaching assistants:**

- Zhiyuan Fan (✉ fanzy@mit.edu). Office hours: Fridays 2:00–3:00pm, Room 45-509.

**Grading:** Homework 50% — Project 50%.

**Homework:** This course will include 2-3 homework sets and a project presentation. Projects may be done individually or in groups of 2-3 students and can be theoretical or experimental. For the problem sets, we will use Gradescope as the grading platform.

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

**Lecture notes:** All handouts, including homework, will be posted in the course's Canvas page.

**Collaboration policy:** We encourage working together whenever possible: in the coding exercises, problem sets, and general discussion of the material and assignments. Keep in mind, however, that for the problem sets the solutions you hand in should reflect *your own* understanding of the class material, and *should be written solely by you.* It is not acceptable to copy a solution that somebody else has written.

## Description

While machine learning techniques have had significant success in single-agent settings, an increasingly large body of literature has been studying settings involving several learning agents with different objectives. In these settings, standard training methods, such as gradient descent, are less successful and the simultaneous learning of the agents commonly leads to nonstationary and even chaotic system dynamics.

Motivated by these challenges, this course presents the foundataions of multi-agent systems from a combined game-theoretic, optimization and learning-theoretic perspective. We start with basic matrix games, such as rock-paper-scissors, and advance to more complex forms including imperfect information games and structured games like combinatorial games, polymatrix games, and stochastic games. Topics will cover various equilibrium concepts, aspects of equilibrium computation and learning, as well as the computational complexity of finding equilibria. Additionally, we will examine how these models and methods have driven recent breakthroughs in AI, producing human- and superhuman-level agents for well-known games such as Go, Poker, Diplomacy, and Stratego. A tentative course syllabus can be found below.

## Tentative schedule

| | | |
|---|---|---|
| 1 | Sep 5th | **Introduction** |
| | Part I: Normal-form games | |
| 2 | Sep 10th | **Setting and equilibria: the Nash equilibrium**<br>Definition of normal-form games. Solution concepts and Nash equilibrium. Nash equilibrium existence theorem. |
| 3 | Sep 12th | **Setting and equilibria: the Correlated equilibrium**<br>Topological and computational properties of the set of Nash equilibria in normal-form games. Connections with linear programming. Definition of correlated and coarse correlated equilibria; relationships with Nash equilibria. |
| 4 | Sep 17th<br>HW1 OUT | **Learning in games: Foundations**<br>Regret and hindsight rationality. Definition of regret minimization and relationships with equilibrium concepts. |
| 5 | Sep 19th | **Learning in games: Algorithms (part I)**<br>General principles in the design of learning algorithms. Follow-the-leader, regret matching, multiplicative weights update, online mirror descent. |
| 6 | Sep 24th | **Learning in games: Algorithms (part II)**<br>Optimistic mirror descent and optimistic follow-the-regularized-leader. Accelerated computation of approximate equilibria. |
| 7 | Sep 26th | **Learning in games: Bandit feedback**<br>From multiplicative weights to Exp3. General principles. Obtaining high-probability bounds. |
| 8 | Oct 1th | **Learning in games: Φ-regret minimization**<br>Gordon, Greenwald, and Marks (2008); Blum and Mansour; Stolz-Lugosi. |
| | Part II: Extensive-form games | |
| 9 | Oct 3rd | **Foundations of extensive-form games**<br>Complete versus imperfect information. Kuhn's theorem. Normal-form and sequence-form strategies. Similarities and differences with normal-form games. |
| 10 | Oct 8th<br>HW2 OUT | **Learning in extensive-form games**<br>No-regret algorithms for extensive-form games. Counterfactual utilities and counterfactual regret minimization (CFR). |
| 11 | Oct 10th | **Equilibrium refinements**<br>Sequential irrationality. Extensive-form perfect equilibria and quasi-perfect equilibrium. |
| 12 | Oct 15th | HOLIDAY **No class** (Student holiday) |
| 13 | Oct 17th | PROJECT **Project ideas and brainstorming** |

| | | |
|---|---|---|
| 14 | Oct 22nd | PROJECT **Project break**<br>Coincides with INFORMS 2024 Annual Meeting |
| 15 | Oct 24th | PROJECT **Project break**<br>Coincides with INFORMS 2024 Annual Meeting |
| 16 | Oct 29th | **Deep reinforcement learning for large-scale games (part I)**<br>Rough taxonomy of deep RL methods for games. Decision-time planning in imperfect-information games, construction of superhuman agents for no-limit Hold'em poker. Public belief states techniques (ReBeL). |
| 17 | Oct 31st | **Deep reinforcement learning for large-scale games (part II)**<br>PPO and magnetic mirror descent. |
| | Part III: Other structured games | |
| 18 | Nov 5th | **Combinatorial games and Kernelized MWU (Part I)**<br>Example of combinatorial games. Kernelized multiplicative weights update algorithm. |
| 19 | Nov 7th | **Combinatorial games and Kernelized MWU (Part II)**<br>Example of combinatorial games. Kernelized multiplicative weights update algorithm. |
| 20 | Nov 12th<br>HW3 OUT | **Computation of exact equilibria**<br>A second look at the minimax theorem. Hart and Schmeidler's proof of existence of correlated equilibria. Ellipsoid against hope algorithm. |
| 21 | Nov 14th | **Markov (aka stochastic) games**<br>Setting and taxonomy of equilibria. Stationary Markov Nash equilibria in inifinite-horizon games. Shapley's minimax theorem. |
| | Part IV: Complexity of equilibrium computation | |
| 22 | Nov 19th | **PPAD-completeness of Nash equilibria (part I)**<br>Sperner's lemma. The PPAD complexity class. Nash $\in$ PPAD. |
| 23 | Nov 21st | **PPAD-completeness of Nash equilibria (part II)**<br>Arithmetic circuit SAT. PPAD-hardness of Nash equilibria. |
| | Project break & presentations | |
| 24 | Nov 26th | PROJECT **Project break** |
| 25 | Nov 28th | HOLIDAY **No class** (Thanksgiving) |
| 26 | Dec 3rd | PROJECT **Project presentations** |
| 27 | Dec 5th | PROJECT **Project presentations** |