

EE-735

Online learning in games

Cevher Volkan

Cursus	Sem.	Type
Computer and Communication Sciences		Opt.
Electrical Engineering		Opt.

Language of teaching	English
Credits	4
Session	
Exam	Oral presentation
Workload	120h
Hours	70
Lecture	28
Practical work	42
Number of positions	30

Frequency

Every 2 years

Remark

Next time: Spring 2024

Summary

This course provides an overview of recent developments in online learning, game theory, and variational inequalities and their point of intersection with a focus on algorithmic development. The primary approach is to lay out the different problem classes and their associated optimal rates.

Content

The dynamics of games have been studied by many different communities, such as online learning, optimization, and game theory. This course seeks to uncover the similarities and subtle differences of these perspectives towards a fundamental understanding of the underlying algorithms.

The course will broadly cover the following themes:

- Equilibrium concepts: source problems, solution concepts and their relationships.
 - Online learning: the concept of regret, regret lower bounds, online learning algorithms, and online to batch conversion.
 - Variational inequalities: monotone operators, algorithms, and their relationships to online learning
 - Limited feedback in online learning: (semi-)bandit feedback, best of both worlds, stochastic approximation
 - Extensions: linear bandits, combinatorial online learning, Markov decision processes, dynamics regret, adaptive regret
- The course proceeds through a combination of books and recent papers and re-derive them under a common language. Concretely we will use the following books and papers:

Lecture 1 Introduction to online learning

- Online learning setting
- Follow the regularized Leader (FTRL) & Online Mirror Descent (OMD)
- A special cases: Hedge Algorithm, online gradient descent
- Lower bounds for online learning problems

Lecture 2 Online learning with bandit feedback

-

Introduction to bandit feedback

-
- Bandit online learning algorithms
-
- Lower bounds for bandit online learning algorithms

Lecture 3 Variational inequalities

-
- An operator view of minimax problems and the associated solution concepts
-
- Gradient descent-ascent based on FTRL
-
- Gradient descent-ascent under co-coercivity
-
- Extragradient (EG)

Lecture 4 Variational inequalities

-
- EG friends: Single-call variants
-
- Error bounds and faster rates
-
- Last iterate analysis of EG

Lecture 5 Learning in games

-
- Introduction to game theory
-
- Coarse correlated equilibrium
-
- Optimistic methods in games
-
- Special cases: optimistic hedge and friends

Lecture 6 Online learning in games with adaptivity and stochastic feedback

- Recent advances, including

1. Hsieh, Yu-Guan, Kimon Antonakopoulos, and Panayotis Mertikopoulos. "Adaptive learning in continuous games: Optimal regret bounds and convergence to Nash equilibrium." Conference on Learning Theory. PMLR, 2021.
2. Mertikopoulos, Panayotis, Ya-Ping Hsieh, and Volkan Cevher. "Learning in games from a stochastic approximation viewpoint." arXiv preprint arXiv:2206.03922 (2022).
3. Hsieh, Yu-Guan, et al. "No-Regret Learning in Games with Noisy Feedback: Faster Rates and Adaptivity via Learning Rate Separation." arXiv preprint arXiv:2206.06015 (2022).

Lecture 7 Stochastic-adversarial bandits (best of both worlds)

- Key advances, including

1. Bubeck, Sébastien, and Aleksandrs Slivkins. "The best of both worlds: Stochastic and adversarial bandits." Conference on Learning Theory. JMLR Workshop and Conference Proceedings, 2012.
2. Lykouris, Thodoris, Vahab Mirrokni, and Renato Paes Leme. "Stochastic bandits robust to adversarial corruptions." Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, 2018.

Lecture 8 Computationally efficient online learning

- Efficient online learning in combinatorial domains, including

1. Kalai, Adam, and Santosh Vempala. "Efficient algorithms for online decision problems." Journal of Computer and

System Sciences 71.3 (2005): 291-307.

2. Awerbuch, Baruch, and Robert D. Kleinberg. "Adaptive routing with end-to-end feedback: Distributed learning and geometric approaches." Proceedings of the thirty-sixth annual ACM symposium on Theory of computing. 2004.

3. Hassani, Hamed, Mahdi Soltanolkotabi, and Amin Karbasi. "Gradient methods for submodular maximization." Advances in Neural Information Processing Systems 30 (2017).

Lecture 9 Online learning and reinforcement learning

- Recent advances, including

1. Jin, Chi, et al. "Is Q-learning provably efficient?." Advances in neural information processing systems 31 (2018).

2. Jin, Chi, et al. "Provably efficient reinforcement learning with linear function approximation." Conference on Learning Theory. PMLR, 2020.

Lecture 10 Stronger notions of regret

-

Dynamic regret

-

Adaptive regret

Lecture 11-12 Reading group

Note

The students are expected to build/present Lectures 6 and onwards at the end of the semester for grade.

Keywords

Online learning, bandits, game theory, variational inequalities, adaptivity, monotone operators, regret, lower-bounds

Learning Prerequisites

Recommended courses

EE-556 Mathematics of Data is recommended.

Important concepts to start the course

Basic probability and linear algebra.

Learning Outcomes

By the end of the course, the student must be able to:

- Choose
- Analyze algorithms
- Explain regret
- Produce a presentation
- Theorize appropriate structures in optimization
- Present concepts in game theory

Teaching methods

Lecture + active learning

Expected student activities

Build and present part of the material with the teaching team in form of a lecture.

Resources

Moodle Link

- <https://go.epfl.ch/EE-735>