

CREST - GENES

Cours doctoraux 2024 – 2025

ON THE FOUNDATIONS OF INTERACTIVE DECISION-MAKING

Alexander Rakhlin

MIT

SCHEDULE	Monday	3 March 2025 5 March 2025	10:15 à 13:15	ROOM 3001
	Wednesday	10 March 2025 12 March 2025	10:15 à 13:15	ROOM 3001

Topics

In this course, we will study general principles for interactive decision making. We will aim to cover the following topics:

1. Basic results from Statistical Learning and Online Learning
2. Multi-Armed Bandits: epsilon-Greedy and Upper Confidence Bound approaches
3. Contextual Bandits: optimism and its failures; epsilon-Greedy; Inverse Gap Weighting as a minimax optimal strategy.
4. Structured Bandits and the Decision Estimation Coefficient.
5. Characterization of bandit learnability.
6. General decision making with a model class. Beyond reward observability. Connection to the Donoho-Liu modulus of continuity for estimation of linear functionals.
7. Lower bounds: Fano, Le Cam, and Assouad with interaction.
8. Extensions to non-stationary and adversarial environments. The Exploration-by-Optimization algorithm.
9. Decision making with Local Differential Privacy and Statistical Queries.

References

D. Foster and A. Rakhlin, “Foundations of Reinforcement Learning and Interactive Decision Making,” <https://arxiv.org/abs/2312.16730>