復旦大學
FUDAN UNIVERSITY

Applied Math Ph.D. Seminar

Batched Stochastic Bandit for Nondegenerate Functions

Speaker: Yu Liu (Fudan University)

Time: 2024-11-14, 16:10 to 17:00

Location: Rm 1801, Guanghua East Tower

Advisor: Tianyu Wang (Fudan University)

Abstract: In batched stochastic bandits, an agent collects noisy rewards or losses in batches, aiming to identify the optimal option while efficiently exploring the decision space. Beyond merely minimizing regret—which measures policy effectiveness—the agent also seeks to minimize the number of batches used. In this talk, we examine batched stochastic bandits for a significant class of functions over a compact doubling metric space, called "nondegenerate functions", which may exhibit nonconvexity, nonsmoothness, or discontinuities. To address this problem, we introduce an algorithm called Geometric Narrowing (GN). The regret bound for GN is of order $\widetilde{\mathcal{O}}(A_+^d\sqrt{T})$, where d is the doubling dimension of the metric space and A_+ is a constant independent of both d and the time horizon T. Notably, GN requires only $\mathcal{O}(\log\log T)$ batches to achieve this regret bound. We further present a lower bound analysis, demonstrating that the GN algorithm attains near-optimal regret with minimal number of batches.