

NUMERICAL ANALYSIS AND SCIENTIFIC COMPUTING
SEMINAR

Top-k Extreme Contextual Bandits with Arm Hierarchy

Lexing Ying
Stanford University

Abstract: Contextual bandit is an online decision making framework that has found many applications in recommendation systems and search tasks. In this talk, we consider the extreme contextual bandit problem where the enormous number of arms poses the main theoretical and algorithmic challenges. This setting is particularly relevant to the mission of Amazon Search but yet rather under-explored in the literature. To address the large arm space, we introduce two new techniques. The first is an extension of the inverse gap weighting to the multiple arm feedback case, where the optimal regret bound is shown under realizability condition. The second is an arm space hierarchy that exploits the potential similarity between the arms. By combining these two techniques, we develop an extreme top-k contextual bandit algorithm that scales logarithmically in terms of the number of arms. Joint work with Rajat Sen, Alexander Rakhlin, Rahul Kidambi, Dean Foster, Daniel Hill, and Inderjit Dhillon.

Friday, March 26, 2021, 1:30 pm
<https://emory.zoom.us/j/95900585494>

MATHEMATICS
EMORY UNIVERSITY