Abstract:
In this talk we discuss two important problems related to the practice of data-driven decision making: learning with limited data and interpretable machine learning models.

First, we consider the problem of sequential product recommendations when customer preferences are unknown; i.e. very little prior data is available to estimate the customer's preference. First, we present empirical evidence of customer disengagement using a sequence of ad campaigns from a major airline carrier. In particular, customers decide to stay on the platform based on the relevance of the recommendations. We then formulate this problem as a linear bandit problem, with the notable difference that the customer's horizon length is a function of past recommendations. We prove that any algorithm in this setting achieves linear regret. As a result, no algorithm can keep all customers engaged; regardless, we can hope to keep a subset of customers engaged. Unfortunately, we find that classical bandit learning as well as greedy algorithms provably over-explore, thereby incurring linear regret for every customer. As a result, we modify bandit learning strategies by constraining the action space upfront using integer optimization. We first show that the algorithm is optimal by mapping the problem to a classical scheduling problem. Then, we prove that this modification allows our algorithm to achieve sublinear regret for a significant fraction of customers. Furthermore, numerical experiments on real movie recommendations data demonstrate that our algorithm can significantly improve customer engagement.

Then in the second part of the talk we briefly discuss XSTrees (Extended Sampled Trees), a tree sampling framework for classification and regression that is interpretable and scalable. We present extensive numerical experiments on synthetic data and public available datasets to show that the algorithm achieves comparable performance in terms prediction accuracy, while providing interpretability for ensemble methods.