Stochastic Bandits for Sticky Recommendations

Theja Tulabandhula, Assistant Professor, University of Illinois Chicago

Area: Production & Operations Management, Decision Sciences

Date: 10.06.2019, Venue: P22 @ 11.30AM

Abstract

We consider sequential decision problems related to making recommendations. A platform needs to show good and timely recommendations to its users to engage them and increase revenues, while not knowing their behavioral patterns a priori. We consider two behavioral effects that modulate relevance: (a) the users have a latent propensity to act on a recommendation based on its position in a sequence of recommended items, and (b) the users have a latent propensity to act on a recommendation only if it has been shown to them repeatedly. In both settings, the platform has to simultaneously learn the quality of its recommendations and the corresponding user behavior, while exploiting the information it knows so far. We develop new bandit algorithms with regret guarantees when considering both these effects, and validate their performance with experiments. (Papers: https://arxiv.org/abs/1901.07734 and https://arxiv.org/abs/1811.09026)

Speaker Profile:

Theja is an Assistant Professor at the University of Illinois Chicago. Previously he was a Research Scientist at Xerox Research (Machine Learning and Statistics group) from 2014-2016. He received a PhD in Electrical Engineering and Computer Science from MIT, and a Dual degree in Electrical Engineering from IIT Kharagpur where he was awarded the Prime Minister's Gold Medal. His research interests include designing new machine learning and optimization methods for sequential decision making, with a focus on applications in retail and transportation.