Title: Sequential learning in a stochastic multi-armed bandit framework

Speaker: Sandeep Juneja, TIFR and Google India

Area: DS

Date: 15.06.2023, Venue: K21 @ 2.30PM

Abstract:

The classic stochastic multi armed bandit framework involves finitely many unknown probability distributions that can be sequentially sampled to generate independent rewards. In this talk we consider two foundational problems: First one corresponds to sampling to minimize the expected regret, or equivalently, to maximize the expected total reward. The second one corresponds to the best arm identification, i.e., identifying the arm with the largest mean, or any other performance measure, using as few samples as possible while providing explicit probabilistically correct selection guarantees. These problems form the bedrock of algorithms used in web design and advertising, recommendation systems, clinical trials and many other exciting applications.

In this talk we review some of the popular algorithms used for these problems emphasizing the intuition underlying the elegant ideas. Technically speaking, these problems have been well studied under the restrictive assumption that arm distributions belong to a single parameter exponential family, that includes distributions such as Bernoulli and Gaussian with known variance. Under these settings, lower bounds on samples needed are developed using ideas from hypothesis testing, and algorithms are proposed that match the lower bound. We propose optimal algorithms that match the lower bounds even to a constant for general probability distributions under minimal restrictions. We discuss further enhancements in the presence of offline data that needs to be combined with online data. We also propose a new algorithm in the best arm identification setting that along with minimizing sample complexity, is also computationally efficient. We conduct mean field analysis that insightfully explains the algorithm behavior.

Speaker Profile:

Sandeep is a senior professor at the School of Technology and Computer Science in Tata Institute of Fundamental Research in Mumbai. He received his B. Tech. in Mechanical Engineering from IIT Delhi (1989) and his M. S. in Statistics and Ph.D. in Operations Research from Stanford University (1993). He then worked for a financial credit insurance company (American Credit Indemnity) in Baltimore, US (2 years) followed by one year stint in Andersen Consulting in India. From December 1996 to December 2002 he was a faculty in the Operations Research Group in the Department of Mechanical Engineering at IIT Delhi. Thereafter he has been at TIFR. He has held visiting positions at many places including at Columbia University, Stanford University and Indian School of Business. In the year 2008, he was on leave from TIFR and headed the quantitative activity in Bank of America's Indian operations. He was then also a member of the Bank of America's executive quantitative council. He is currently on the editorial board of Stochastic Systems. Earlier he has been on editorial boards of Mathematics of Operations Research, Management Science and ACM TOMACS. He is a recipient of IBM faculty partnership award in the year 2001-02 and he co-authored papers that received best paper awards at 4 th as well as 6th

International ICST Conference on Performance Evaluation Methodologies and Tools (in 2009 and 2012). He spent a sabbatical and was an adjunct at Centre for Advanced Financial Research and Learning (CAFRAL), a research wing of Reserve Bank of India (2015-16). Currently he is visiting the Machine Learning and Optimization Group at Google Research in Bangalore. He has consulted with research labs as well as financial firms in India. He has taught financial mathematics to quantitative associates of many leading multinational investment banks based in India. His research interests lie in applied probability including in sequential learning, mathematical finance, Monte Carlo methods, and game theoretic analysis of queues. Lately, he has been involved in modelling Covid-19 spread in Mumbai, and in mathematics of certain epidemiological models.

Webpage Link: https://www.tcs.tifr.res.in/~sandeepj/