

Improving the Knowledge Gradient Algorithm

Ms Le Yang

PhD Student
Department of Systems Engineering
City University of Hong Kong

8 January 2024 (Mon)

10:30 am

YEUNG-P7303

Abstract

The knowledge gradient (KG) algorithm is a popular policy for the best arm identification (BAI) problem. It is built on the simple idea of always choosing the measurement that yields the greatest expected one-step improvement in the estimate of the best mean of the arms. In this research, we show that this policy has limitations, causing the algorithm not asymptotically optimal. We next provide a remedy for it, by following the manner of one-step look ahead of KG, but instead choosing the measurement that yields the greatest one-step improvement in the probability of selecting the best arm. The new policy is called improved knowledge gradient (iKG). iKG can be shown to be asymptotically optimal. In addition, we show that compared to KG, it is easier to extend iKG to variant problems of BAI, with the ϵ -good arm identification and feasible arm identification as two examples. The superior performances of iKG on these problems are further demonstrated using numerical examples.

About the Speaker

Le Yang is a PhD student in the SYE department at City University of Hong Kong. She received her bachelor's and master's degrees in statistics from Shandong University in 2017 and 2020 respectively. Her research interest is centered around theory of machine learning and simulation optimization.