



THE HONG KONG UNIVERSITY OF SCIENCE & TECHNOLOGY

Department of Mathematics

**SEMINAR ON STATISTICS AND  
DATA SCIENCE**

**Exponential Lower Bounds and Fast Convergence  
for Policy Optimization**

By

**Prof. Yuting WEI**  
The Wharton School  
University of Pennsylvania

Abstract

Policy gradient (PG) methods and their variants lie at the heart of modern reinforcement learning. Due to the intrinsic non-concavity of value maximization, however, the theoretical underpinnings of PG-type methods have been limited even until recently. In this talk, we discuss both the ineffectiveness and effectiveness of nonconvex policy optimization. On the one hand, we demonstrate that the popular softmax policy gradient method can take exponential time to converge. On the other hand, we show that employing natural policy gradients and enforcing entropy regularization allows for fast global convergence.

Biography

*Yuting Wei is currently an assistant professor in the Statistics and Data Science Department at the Wharton School, University of Pennsylvania. Prior to that, Yuting spent two years at Carnegie Mellon University as an assistant professor of statistics, and one year at Stanford University as a Stein Fellow. She received her Ph.D. in statistics at the University of California, Berkeley. She was the recipient of the 2022 NSF Career award, the Erich L. Lehmann Citation from the Berkeley statistics department in 2018. Her research interests include high-dimensional and non-parametric statistics, statistical machine learning, and reinforcement learning.*

**Date : 17 October 2022 (Monday)**

**Time : 11:00am – 12:00nn**

**Zoom Meeting : <https://hkust.zoom.us/j/94883840530> (Passcode: hkust)**

*All are Welcome!*