

# SCHOOL OF DATA SCIENCE

## SEMINAR SERIES

### Markov Decision Processes: Robustness and Modern Applications

Date: 25 January 2019 (Friday)

Time: 10:00am to 11:00am

Venue: P7510, 7/F, Yeung Kin Man Academic Building (YEUNG),  
City University of Hong Kong

**Dr. Ho, Clint Chin Pang**  
Junior Research Fellow  
Imperial College Business School  
Imperial College London



#### Guest Speaker's profile

Clint Chin Pang Ho is a Junior Research Fellow (now known as Imperial College Research Fellow) in the Imperial College Business School. Before that, he received a BS in Applied Mathematics (Summa Cum Laude) from the University of California, Los Angeles (UCLA) in 2011, an MSc in Mathematical Modeling and Scientific Computing from the University of Oxford in 2012, and a PhD in computational optimization from Imperial College London in 2016. Clint studies optimization algorithms and computational methods for structured problems, as well as their applications in machine learning, operations research, and business analytics.

#### Abstract

Since the 1950s, Markov decision processes (MDPs) have been a mainstream approach to model dynamic decision problems under uncertainty. Despite their long history and many successful practical implementations, the application of MDPs is challenged by the requirement to specify the underlying stochastic process exactly.

In the first part of my talk, I propose robust MDPs to alleviate the burden of precisely specifying this stochastic process. Instead of requiring the transition probabilities to be known precisely, robust MDPs assume that these probabilities are merely known to be contained in an uncertainty set (that can be estimated from domain knowledge or past data). By optimizing in view of the worst transition probabilities within this uncertainty set, robust MDPs yield decisions that hedge against estimation errors. Unfortunately, robust MDPs pose significant computational challenges if solved with traditional techniques. To this end, I will propose several efficient algorithms that solve robust MDPs in nearly at the same time complexity as standard (non-robust) MDPs.

In the second part of this talk, I will discuss a real-world application in which a personalized marketing strategy has to be derived from customer data. I model this problem via customer-specific MDPs, and I will discuss how efficient solutions can be determined even for instances with large numbers of customers.