MICHIGAN STATE UNIVERSITY
Department of Statistics and Probability

COLLOQUIUM

Xun Huan
University of Michigan

Simulation-based Bayesian Sequential Design Using Reinforcement Learning

Tuesday, February 4, 2020
10:20 AM - 11:10 AM
Refreshments 10:00 AM
C405 Wells Hall

Abstract

Experiments are indispensable for learning and developing models in science and engineering, and a careful design of these often limited and expensive data-acquisition opportunities is crucial. Where simulation models are available, optimal experimental design seeks to maximize a statistical criterion—reflecting the value gained from an experiment—leveraging these model predictions. We focus on the design of a finite sequence of experiments, seeking design policies (strategies) that can (a) adapt to newly collected data during the sequence (i.e. feedback) and (b) anticipate future changes (i.e. lookahead). We cast this sequential learning problem in a Bayesian setting with information-based utilities, and solve it numerically via policy gradient methods from reinforcement learning. In particular, we directly parameterize the policies and value functions, thus adopting an actor-critic approach, and improve them using gradient estimates produced from simulated design sequences. The overall method is demonstrated on an algebraic benchmark and a sensor placement application for source inversion. The results provide intuitive insights on the benefits of feedback and lookahead, and indicate substantial computational advantages compared to previous numerical methods based on approximate dynamical programming.

To request an interpreter or other accommodations for people with disabilities, please call the Department of Statistics and Probability at 517-355-9589.