Title: Optimizing Decision Making via Submodular Surrogates
Speaker: Yuxin Chen (University of Chicago)

Abstract: How can we intelligently acquire information for decision making, when facing a large volume of data? In this talk, I will introduce the decision-theoretic value of information problem, and show that for a large class of combinatorial optimization problems that are known to be NP-hard, one could devise efficient surrogate objectives that are amenable to greedy optimization, while still achieving strong approximation guarantees. A key property we seek in the design of the greedy heuristics is submodularity, a natural diminishing returns condition common to a broad class of decision-making problems. When it is challenging to construct such surrogate functions, we further introduce a data-driven optimization framework based on a novel loss function---namely the "submodular-norm" loss---which encourages the resulting objective to exhibit diminishing returns. We show that our model can be easily integrated with modern deep imitation learning pipelines for sequential prediction tasks. We demonstrate our algorithms on a variety of batched and sequential optimization tasks, including active learning, robotic manipulation, and sequential experimental design for protein engineering.

The talk will mainly be based on the following work:

- Submodular Surrogates for Value of Information
  Yuxin Chen, Shervin Javdani, Amin Karbasi, Drew Bagnell, Siddhartha Srinivasa, Andreas Krause. In the 29th AAAI Conference on Artificial Intelligence (AAAI) , Austin, TX, January 2015.
  [bibtex] [pdf] [long version] [poster]

- Learning to Make Decisions via Submodular Regularization
  Ayya Alieva, Aiden Aceves, Jialin Song, Stephen Mayo, Yisong Yue, Yuxin Chen.
  In the International Conference on Learning Representations (ICLR), May 2021
  [pdf] [poster]