

Learning Submodular Functions with Applications to Multi-Agent Systems

Maria Florina Balcan
School of Computer Science
Carnegie Mellon University
ninamf@cs.cmu.edu

Abstract

Machine learning can provide powerful tools for the design and analysis of multi-agent systems, as well as a novel lens on important questions in the area. In this talk I will focus on the use of machine learning for understanding submodular functions, an important class of discrete functions that model laws of diminishing returns and enjoy many important applications in multi-agent settings. For example, submodular functions are commonly used to model valuation functions for bidders in auctions, the influence of various subsets of agents in social networks, and the benefits of performing different actions in a variety of situations. Traditionally it is assumed that these functions are known to the decision maker; however, for large scale systems, in the age of big data, it is often the case they must be learned from observations.

In this talk, I will discuss a recent line of work on studying the learnability of submodular functions, and highlight its applications to the analysis of multi-agent systems. I will discuss both general algorithms for learning such functions, as well as even better guarantees that can be achieved for important classes appearing in multi-agent scenarios that exhibit additional structure. These classes include: probabilistic coverage functions that can be used to model the influence function in classic models of information diffusion in networks; functions with bounded complexity used in modeling bidder valuation functions in auctions, including XOS and gross-substitutes; and classes of functions appearing in cooperative game theory for expressing the values of various types of coalitions.

I will additionally discuss a large scale application of our algorithms for learning the influence functions in social networks, that significantly outperforms existing approaches empirically in both synthetic and real world data.

ACM Classification: I.2.6 [Learning], J.4 [Social and Behavioral Sciences]: Economics, F.2 [Analysis of Algorithms and Problem Complexity]

Keywords: Diminishing returns; learning submodular functions; influence of agents in social networks; valuation functions in auctions; values of coalitions



Short Bio

Maria Florina Balcan is an Associate Professor in the School of Computer Science at Carnegie Mellon University. Her main research interests are machine learning, computational aspects in economics and game theory, and algorithms. Her honors include the CMU SCS Distinguished Dissertation Award, an NSF CAREER Award, a Microsoft Faculty Research Fellowship, a Sloan Research Fellowship, and several paper awards at COLT. She is currently a board member of the International Machine Learning Society and was recently Program Committee chair for COLT 2014.