

The University of Chicago Computer Science Department

PRESENTS:

“First-Order Methods Unleashed: Scalable Optimization in the Age of Big Data”



Lorenzo Orecchia
Boston University

Abstract:

First-order methods are a fundamental tool in the design of efficient algorithms for large-scale computational problems. Besides being the optimization workhorse of machine learning, first-order methods have recently served as a springboard for a number of algorithmic advances in discrete optimization, including submodular optimization and maximum flow problems. In this talk, I will showcase a number of results from my research that demonstrate the power of first-order methods as a generic framework for algorithm design.

In the first part, I will describe my view of first-order methods as discretizations of continuous dynamical systems over curved spaces. For convex optimization, such dynamics conserve a specific quantity -- the product of time and a notion of duality gap -- which immediately guarantees convergence to optimum. This primal-dual view helps us to both design novel algorithms and simplify the analyses of existing ones. In particular, I will discuss how it yields a simple, intuitive analysis of accelerated algorithms and how it allows us to port such algorithms to contexts that do not squarely match standard smoothness assumptions.

In the second part, we will see how to exploit problem-specific structure by preconditioning, i.e., by endowing the space with a curved geometry that facilitates the convergence of the dynamics above. In particular, I will describe how different random-walk-based algorithms for graph partitioning arise from different preconditionings of the same optimization problem, and how combinatorial preconditioners yield nearly-linear-time algorithms for flow problems over undirected graphs.

Bio:

Lorenzo Orecchia is an assistant professor in the Department of Computer Science at Boston University. Lorenzo's research focuses on the design of efficient algorithms for fundamental computational challenges in machine learning and combinatorial optimization. His approach is based on combining ideas from continuous and discrete optimization into a single framework for algorithm design. Lorenzo obtained his PhD in computer science at UC Berkeley under the supervision of Satish Rao in 2011, and was an applied mathematics instructor at MIT under the supervision of Jon Kelner until 2014. He was a recipient of the 2014 SODA Best Paper award and a co-organizer of the Simons semester "Bridging Continuous and Discrete Optimization" in Fall 2017.

Wednesday, April 3, 2019

1:00 pm

JCL 390

Host: Rebecca Willett