Shifted divergences for sampling, privacy, and beyond

Jason Altschuler
NYU / UPenn, USA
ja4775@nyu.edu

Shifted divergences provide a principled way of making information theoretic divergences (e.g. KL) geometrically aware via optimal transport smoothing. In this talk, I will argue that shifted divergences provide a powerful approach towards unifying optimization, sampling, differential privacy, and beyond. For concreteness, I will demonstrate these connections via three recent highlights. (1) The fastest high-accuracy algorithm for sampling from log-concave distributions. (2) Resolving the mixing time of the Langevin Algorithm to its stationary distribution for log-concave sampling. (3) Resolving the differential privacy of Noisy-SGD, the standard algorithm for private optimization in both theory and practice. A recurring theme is a certain notion of algorithmic stability, and the central technique for establishing this is shifted divergences.

Joint work with Kunal Talwar (Apple) and Sinho Chewi (MIT).