Reconstructing Latent Similarities in a Multiplex Social Network

Alex Slivkins, Microsoft Research

It is commonly assumed that individuals tend to be more similar to their friends than to strangers. Thus, we can view an observed social network as a noisy signal about the latent underlying "social space": the way in which individuals are (dis)similar. This naturally raises the inverse question: given a social network, how accurately can we reconstruct the social space?

We begin to address this problem formally. We assume that each category (e.g., geography, profession, hobbies) is characterized by a latent metric capturing (dis)similarities in this category, and gives rise to a separate social network: a random graph parametrized by this metric. The algorithm only observes the unlabeled union of these graphs, and reconstructs each metric with provably low distortion.

Joint work with Ittai Abraham, Shiri Chechik and David Kempe.
http://arxiv.org/abs/1202.0922