Matrix Learning: A Tale of Two Norms

Nathan Srebro, Technion

There has been much interest in recent years in various ways of constraining the complexity of matrices based on factorizations into a product of two simpler matrices. Such measures of matrix complexity can then be used as regularizers for such tasks as matrix completion, collaborative filtering, multi-task learning and multi-class learning. In this talk I will discuss two forms of matrix regularization which constrain the norm of the factorization, namely the trace-norm (aka nuclear-norm) and the so-called max-norm (aka $\gamma_2:\ell_1\rightarrow\ell_\infty$ norm). I will both argue that they are independently motivated and often better model data than rank constraints, as well as explore their relationships to the rank. In particular, I will discuss how simple low-rank matrix completion guarantees can be obtained using these measures, and without various "incoherence" assumptions. I will present both theoretical and empirical arguments for why the max-norm might actually be a better regularizer, as well as a better convex surrogate for the rank.

Based on joint work with Rina Foygel, Jason Lee, Ben Recht, Russ Salakhutdinov, Ohad Shamir, Adi Shraibman and Joel Tropp and others.