

Computer Science Seminar

Convex Risk Minimization and Conditional Probability Estimation

Matus Telgarsky (University of Michigan)

Abstract: This talk will investigate and strengthen the link between convex risk minimization and conditional probability estimation, a connection already notable for establishing consistency results (Friedman et al., 2000; Zhang, 2004; Bartlett et al., 2006). Specifically, this talk will first show that a differentiable loss function, linear space of predictors, and probability measure together define a unique optimal conditional probability model, moreover one which may be attained by the usual convex risk minimization. This result is proved in infinite dimensions, and thus gives a concrete convergence target for unregularized methods like boosting which can fail to have minimizers. Second, this convergence result is refined in finitely many dimensions to hold for empirical risk minimization. This uniform convergence result exhibits no dependence on the norms of its predictors, and thus can justify the practical effectiveness of minimally-regularized optimization schemes.

This is joint work with Miroslav Dudik and Robert Schapire

Wednesday, May 6 at 3:00 PM in SEO 427
--