Title: Reliable Learning via Distributionally Robustness

Abstract: Spurred by recent advances, machine learning methods are beginning to prescribe decisions in high-stakes domains, including autonomous driving and medical diagnoses. The standard machine learning paradigm that optimizes average-case performance, however, often yields models that perform poorly on tail subpopulations (e.g. underrepresented demographic groups). We present convex procedures that optimize the worst-case subpopulation performance, thereby guaranteeing a uniform level of performance over all subpopulations. We prove finite-sample minimax upper and lower bounds, showing that uniform subpopulation performance comes at a cost in convergence rates. Empirically, our procedure improves performance on tail subpopulations, and provides a uniform level of performance through time by controlling latent minority proportions.