Proximity Without Consensus in Online Multiagent Optimization

Abstract:

We consider stochastic optimization problems in multiagent settings, where a network of agents aims to learn parameters that are optimal in terms of a global convex objective, while giving preference to locally observed streaming information. To do so, we depart from the canonical decentralized optimization framework where agreement constraints are enforced, and instead formulate a problem where each agent minimizes a global convex objective while enforcing network proximity constraints. This formulation includes online consensus optimization as a special case, but allows for themore general hypothesis that there is data heterogeneity across the network. To solve this problem, we propose using a stochastic saddle point algorithm inspired by Arrow and Hurwicz. This method yields a decentralized algorithm for processing observations sequentially received at each node of the network. UsingLagrange multipliers to penalize the discrepancy between them, only neighboring nodes exchange model information.We establish that under a constant step-size regime the time-average suboptimality and constraint violation are contained in a neighborhood whose radius vanishes with increasing number of iterations. As a consequence, we prove that the time-average primal vectors converge to the optimal objective while satisfying the network proximity constraints.We apply this method to the problem of sequentially estimating a correlated random field in a sensor network, as well as an online source localization problem, both

of which demonstrate the empirical validity of the aforementioned convergence results.