Learning to predict Nash equilibria from data using fixed point networks.

Daniel McKenzie

University of California, Los Angeles, Department of Mathematics, USA

We study the problem of predicting the outcome of a contextual game, given only the context, and assuming that the player's cost functions are unknown. We use the recently introduced Fixed Point Network (FPN) framework to phrase this as a learning problem using historical data consisting of pairs of context and game outcomes. Using several "tricks" (e.g. Davis-Yin operator splitting, constraint decoupling) we improve the efficiency of this scheme to the extent that it can be readily applied to large games with complicated constraint sets. Finally, we demonstrate the efficacy of our proposed scheme (dubbed Nash-FPNs) on a collection of real-world traffic routing problems.