

Constantinos Daskalakis, EECS and CSAIL, MIT

Title: Games with Non-Concave Utilities and Machine Learning

Abstract: A pervasive assumption in Game Theory is that players' utilities are concave, or at least quasi-concave, with respect to their own strategies. While mathematically instrumental, enabling the existence of many kinds of equilibria in many kinds of settings, (quasi-)concavity of payoffs is too restrictive an assumption. For the same reasons that (quasi-)concave utilities can only go so far in capturing single-agent optimization problems, they can only go so far in modeling the considerations of an agent in a strategic interaction. Besides, the study of games with non-concave utilities is increasingly coming to the fore as Deep Learning ventures into multi-agent learning applications. In this article, we study what types of equilibria exist in such games, and whether they are computationally tractable, proposing new frontiers at the interface of Game Theory and multi-agent learning.