Equilibrium Computation, Deep Learning, and Multi-Agent Reinforcement Learning

Constantinos Daskalakis
EECS and CSAIL, MIT, Cambridge, MA, USA

Abstract

Machine Learning has recently made significant advances in challenges such as speech and image recognition, automatic translation, and text generation, much of that progress being fueled by the success of gradient descent-based optimization methods in computing local optima of non-convex objectives. From robustifying machine learning models against adversarial attacks to causal inference, training generative models, multi-robot interactions, and learning in strategic environments, many outstanding challenges in Machine Learning lie at its interface with Game Theory. On this front, however, gradient-descent based optimization methods have been less successful. Here, the role of single-objective optimization is played by equilibrium computation, but gradient-descent based methods commonly fail to find equilibria, and even computing local approximate equilibria has remained daunting. We shed light on these challenges through a combination of learning-theoretic, complexity-theoretic, game-theoretic and topological techniques, presenting obstacles and opportunities for Machine Learning and Game Theory going forward. I will assume no Deep Learning background for this talk and present results from joint works with S. Skoulakis and M. Zampetakis [2] as well as with N. Golowich and K. Zhang [1].

2012 ACM Subject Classification Theory of computation → Multi-agent learning; Theory of computation → Multi-agent reinforcement learning; Theory of computation → Solution concepts in game theory; Theory of computation → Exact and approximate computation of equilibria

Keywords and phrases Deep Learning, Multi-Agent (Reinforcement) Learning, Game Theory, Nonconvex Optimization, PPAD

Digital Object Identifier 10.4230/LIPIcs.ICALP.2022.2

Category Invited Talk

Funding This work is supported by NSF Awards CCF-1901292, DMS-2022448 and DMS2134108, a Simons Investigator Award, the Simons Collaboration on the Theory of Algorithmic Fairness, a DSTA grant, and the DOE PhILMs project (DE-AC05-76RL01830).

References

1 Constantinos Daskalakis, Noah Golowich, and Kaiqing Zhang. The complexity of Markov equilibrium in stochastic games. *arXiv preprint*, 2022. `arXiv:2204.03991`.

2 Constantinos Daskalakis, Stratis Skoulakis, and Manolis Zampetakis. The complexity of constrained min-max optimization. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*, pages 1466–1478, 2021.