



NOAM BROWN

Equilibrium Finding for Large Adversarial Imperfect-Information Games

Wednesday, September 16, 2020 – 3:00 p.m. – REMOTE

Imperfect-information games model interactions between multiple agents with private information. A typical goal in this setting is to approximate an equilibrium in which all agents' strategies are optimal. This thesis describes a number of advances in the computation of equilibria in large adversarial imperfect-information games.

We begin by introducing improvements to counterfactual regret minimization (CFR), an iterative algorithm that converges to a Nash equilibrium in two-player zero-sum games. We describe new variants of CFR that use discounting to dramatically speed up convergence. These new CFR variants are now the state-of-the-art equilibrium-finding algorithms for large adversarial imperfect-information games. We also introduce the first general technique for warm starting CFR. Finally, we introduce theoretically sound pruning techniques that can speed up convergence by orders of magnitude in large games.

Next, we describe new ways to scale CFR to extremely large games via automated abstraction and function approximation. In particular, we introduce the first algorithm for discretizing continuous action spaces in imperfect-information games that is provably locally optimal. We extend this into an algorithm that solves games with continuous action spaces. Finally, we introduce Deep CFR, a form of CFR that uses neural network function approximation rather than bucketing-based abstractions. Deep CFR was the first non-tabular form of CFR to scale to large games and enables CFR to be deployed in settings with little domain knowledge.

Finally, we present new search techniques for imperfect-information games that ensure the agent's search strategy cannot be exploited by an adversary. These new forms of search outperform past approaches both in theory and in practice. Next, we introduce a method for depth-limited search that is drastically computationally cheaper than prior approaches. Finally, we present an algorithm that combines reinforcement learning with search at both training and test time, which takes a major step toward bridging the gap between research on perfect-information games and imperfect-information games.

Together, these new techniques made it possible for the first time for an AI agent to defeat top human professionals in full-scale poker, which had for decades served as a grand challenge problem in the fields of AI and game theory.

Thesis Committee:
Tuomas Sandholm, Chair
Geoff Gordon
Ariel Procaccia
Satinder Singh, University of Michigan
Michael Wellman, University of Michigan