Abstract

From the onset of AI research, games have played an important part, serving as a benchmark for progress in artificial intelligence. Recent approaches using search in combination with learning from self-play have shown strong performance and the ability to generalize across a wide range of perfect information games. In contrast, the leading algorithms for imperfect information traditionally used a small, abstract version of a game and solved this abstraction in one go. This thesis introduces a chain of improvements for imperfect information algorithms that culminates in two significant milestones that helped bridge the gap between perfect and imperfect information games. The first milestone is DeepStack — the first agent that successfully used a combination of sound search and a learned value function in imperfect information games. This led to the first AI to achieve victory over human professional players in no-limit poker. The second milestone is Player of Games — a universal algorithm that can master both perfect and imperfect information games starting from scratch.