AAAI-15 Workshop on Computer Poker and Imperfect Information

Sam Ganzfried (sam.ganzfried@gmail.com)

Abstract

The AAAI Workshop on Computer Poker and Imperfect Information took place on January 26, 2015 in Austin, Texas. The workshop brought together researchers studying new theoretical and practical approaches for imperfect-information games, with application to a variety of domains.

Poker has emerged as a major AI challenge problem. Since 2006, there has been a competition between the strongest computer poker agents held annually at AAAI. Building strong poker agents involves solving many of the most challenging issues faced by all problems in artificial intelligence, including dealing with enormous state spaces, the presence of multiple self-interested agents, imperfect information, stochastic events, balancing between exploitation of opponents' mistakes and minimizing our own "exploitability" (ensuring strong agents cannot take advantage of our own mistakes), and performing robust, large-scale optimization. Poker is not simply a toy game; it is tremendously popular for humans, and online poker is a multi-billion dollar industry. The version of two-player no-limit Texas Hold'em played in the AAAI competition has approximately 10^{165} states in its game tree. Other domains where algorithms for imperfect-information games are pivotal include business (e.g., auctions and negotiations), medicine (e.g., developing robust treatment policies to combat diseases), and (cyber)security.

The AAAI-15 Workshop on Computer Poker and Imperfect Information brought together researchers studying a variety of topics pertaining to imperfect-information games, ranging from theoretical analysis and general-purpose algorithms to the design of agents for poker and other domains. The workshop was conceived at AAAI in 2012, and has been running for four consecutive years now.

The workshop was comprised of nine oral presentations followed by a poster session and round table discussion. The standard paradigm used by most strong agents for large imperfect-information games is to first apply an abstraction algorithm to compute a smaller approximation of the game, and then to compute an approximate equilibrium in the abstraction; these are done offline, and then the strategies are implemented in real time by performing a table lookup. Several of the talks proposed novel approaches constituting significant departures from this paradigm. Kevin Waugh from Carnegie Mellon University described a new game-solving approach called "functional regret estimation" that combines the abstraction and equilibrium-finding components by employing an online regressor from domain features to approximate regrets used by a regret-minimizing equilibrium-finding algorithm. This, in effect, generalizes the standard abstraction approaches by allowing us to learn good abstractions from features, as opposed to computing a single fixed abstraction in advance of equilibrium finding. For the application of the approach to poker, these features included expected hand strength, pot size, and number of actions in the current hand.

Sam Ganzfried from Carnegie Mellon University described another approach that marks a significant departure from the traditional paradigm, involving solving portions of the game that are actually reached in real time to a greater degree of accuracy than is done in offline computations. A new efficient algorithm for performing such "endgame solving" in large imperfect-information games was presented, and the approach was shown to lead to significant performance improvements in no-limit Texas Hold'em against the strongest competition agents. Theoretical analysis showed that this approach may produce highly exploitable strategies in certain games; though it guarantees low exploitability in others, and a new framework was presented for assessing its broader applicability in different game classes.

As is typical for this workshop, most of the talks presented novel theoretical analysis and domain-independent approaches, and several domains were considered in addition to poker. Hsuan-Tien Lin from National Taiwan University presented an approach for learning new bidding conventions in the domain of bridge. Training on random deals, their algorithm was able to learn a bidding system that performed competitively with a champion bridge agent that uses the Standard American human bidding system. The conventions learned by the algorithm differed significantly from the standard system; for example, for its initial action the new approach bids one spade with a hand with nearly-balanced suits and never bids one notrump, while the standard approach bids one spade with 5 or more spades, and frequently bids one notrump.

Kevin Waugh from Carnegie Mellon University presented the results from the poker competition. Due to the close proximity to the most recent competition which took place in July 2014, this competition was a significantly scaled-down version, featuring only 3-player Kuhn poker. A full-scale competition is planned for AAAI 2016. Four agents were submitted to the 2015 3-player Kuhn poker competition; the winning agent was Umpa, from Ariel University. The agent used an opponent modeling approach where each opponent was represented by a vector of 48 real numbers in [0,1] corresponding to each of the possible 48 states of the game (where each state consists of a private card and betting history for the current round). The models were updated based on new information revealed after each hand, utilizing a genetic algorithm to tune the learning rate.

The round table discussion focused on possible rule changes and integration of new events for the 2016 poker competition. There was a strong contingency in favor of adding a six-player no-limit Texas Hold'em competition. Some of the arguments in favor were that it is an extremely popular variant for humans, that it would involve drastically new algorithmic challenges to deal with the massively larger state space, and that it would involve conceptual challenges, as Nash equilibrium is not as well motivated as in two-player games. Some arguments against this proposal were the increased computational resources required, as well as the challenges with obtaining statistical significance. However, such challenges also open up interesting research questions on tournament design and variance reduction.

There was also interest in adding a "lightweight" two-player no-limit Texas Hold'em competition, with a significantly smaller limit on the size of agents. This would potentially make the competition more easily accessible to newer entrants and approaches that may have more limited access to computational resources than some existing teams.

Proposals for these and other events should be submitted by early March through the competition Google group, maintained by the 2015 poker competition organizers—Neil Burch from the University of Alberta and Kevin Waugh from Carnegie Mellon University. Kevin Waugh will be one of the organizers for the 2016 competition, and the other organizer is still to be determined.

Sam Ganzfried served as the chair of this workshop. The papers of the workshop will be published as AAAI Technical Report WS-15-08.

- Sam Ganzfried is a PhD student in the Computer Science Department at Carnegie Mellon University.