

Generating Trading Agent Strategies

Daniel Reeves

Advisor: Michael Wellman

University of Michigan Artificial Intelligence Lab.

1101 Beal Av, Ann Arbor, MI 48109-2110 USA

<http://ai.eecs.umich.edu/people/dreeves/dreeves@umich.edu>

My thesis work concerns the generation of trading agent strategies—automatically, semi-automatically, and manually. Automatic generation of an agent strategy means creating a system that can read the description of some mechanism (i.e., a game) and output a strategy for a participating agent—i.e., a mapping from percepts to actions in the environment defined by the mechanism. To make this more concrete, consider an extremely simple auction mechanism: a two-player first-price sealed-bid auction. This is a game in which two players each have one piece of private information—their valuations for the good being auctioned. Each agent also has a continuum of possible actions—its bid amount. The payoff to an agent is its valuation minus its bid, if its bid is highest, and zero otherwise. My current system can take such a game description and output the optimal strategy, i.e., the Nash equilibrium. (In this case, that strategy is to bid half of your valuation.) Existing game solvers (Gambit and Gala) can only solve games with an enumerated (finite) set of actions, and this limitation makes it impossible to even approximate (i.e., by discretizing) the solution to games with a continuum of actions because the size of the game tree quickly explodes. Of course, the optimal strategy for the first-price sealed-bid auction was computed before game solvers existed; however, my algorithm can automatically solve any of a class of games (with certain caveats) that current solvers can't. In addition to this algorithm for exact solutions, I have an approximation algorithm using monte carlo simulation that can handle a more general class of games (e.g., arbitrary payoff functions and any number of players) albeit at high computational cost.

Both of the above methods are only tractable for quite simple games. For example, almost any mechanism that involves iterated bidding and multiple auctions is likely not to be tractable for strictly game-theoretic analysis, regardless of whether exact or approximate solutions are sought. An example of such a mechanism that we are analyzing is a simultaneous ascending auction for scheduling resources among a group of agents. In this domain, every agent has certain preferences for acquiring time slots (say, for use of a resource in a factory) and simultaneous English auctions are held for every slot until bidding stops and the slots are allocated. Since this mechanism is too complicated for the fully

automatic techniques described above, we are attempting to generate strategies semi-automatically. We started with a baseline strategy and generalized it via a set of parameters, defining a restricted space of agent strategies. Our current work involves automating the search over that space by co-evolving strategies until an evolutionarily stable profile of strategies emerges. Such a profile will be a Nash equilibrium for the restricted game in which agents can only choose from a limited number of strategies.

Finally, there are domains in which even the semi-automatic method for generating strategies is not yet feasible. These include many real-world market scenarios, but also a particular structured domain—the Trading Agent Competition (TAC). This competition was created by our research group and first held in 2000. It will be held this year for the third time, in July in Edmonton. The domain involves shopping for travel packages for a group of hypothetical clients with varying preferences over length of trip, hotel quality, and entertainment options. The shopping involves participating in a myriad of auctions of different types. For example, hotels are sold in multi-unit English auctions while entertainment tickets are bought and sold in continuous double auctions (like the stock market). An agent's payoff is the total utility it achieves for its clients, minus its net expenditure. After hosting this competition for two years, this year we will be submitting our own agent to compete in TAC. Past competitors have made many worthwhile research advances in automated trading, and we hope our efforts will contribute as well, especially as another test for the methodology I'm developing in my thesis work. We believe this work in designing trading agent strategies for very rich domains will shed light on how to extend our results on the automatic generation of trading agent strategies.