CS 396 Special Topics in Artificial Intelligence
Project Deliverable 1

Ari Wilson

February 6th, 2007

Overall description

I intend to study the effectiveness of different methods of decision making and task learning within the context of a particular multiagent system. Electronic markets is the chosen domain because a properly developed market system can be transformed into any number of useful negotiation and coordination protocols in other domains. Electronic market systems are now the dominant metaphor for the exchange of goods in many large human-agent systems, including Amazon.com and eBay.

Algorithms

Automated systems that make decisions have many of the same problems as humans when making decisions with limited resources, risk, and uncertainty. Game theory and Bayesian decision-making are two models that can be used to analyze such systems. Some of the algorithms that I wish to implement, therefore, include methods of analyzing stochastic utility functions and computing equilibrium states of general games.

The classic theories of game theory and statistical behavior, however, do not adapt for the behavior of adversaries and collaborators. Thus, I would like to use machine learning to have agents make better predictions about the complex world they (and we) find ourselves in. However, such lofty goals may be limited by the amount of material the course is able to cover. Some material that looks interesting from later chapters in our textbook includes reinforcement learning in stochastic games, no-regret learning (based on the material on Amy Greenwald's website), and evolutionary learning.

Development method

The TeamBots package (available from `http://www.teambots.org`) allows for the simulation of multiple hardware or software agents communicating and operating in a common environment and is written in Java. Within this software environment, agents will need to collect resources that are desirable to other agents that have the resources that they want. Once all items have been collected, the bidding phase will begin. Agents will bid on these items using their items as currency and the highest bidder, according to the internal benefit formula of the seller, will win. Bidding will be done over multiple rounds until no agent believes that it can further increase the relative value of its possessions (the system has reached equilibrium). Agents will be ranked by how well the items they end up with match the goals set for them at the beginning of collection and bidding.

Agents will be written in Java, and the mechanisms of inheritance and interfaces will be used to extend agents based on performance measurements to create new ones. Necessarily, internal data

storage will consist of internal representations of the state and reasoning of other agents, which will be done using standard first order logic and associated probabilities. In some sense then, agents will have their own "personality" consisting of their worldview and associated reasoning processes. Some agents will attempt to simulate these complexities of other agents in order to predict their behavior and thus maximize their own utility.

Evaluation

The agents created in this project will be evaluated over a large number of trials with differing agent goals and randomly located resources. A large number of trials are done to ensure that no bias due to the particular environment is present. These results will be averaged to provide a relative ranking of agents based on their overall success in bidding for items. For an agent a desiring a set of items I with respective desirabilities measured by the function $D(i)$ $(\forall i(D(i) \geq 0))$ after a trial, a's success can be measured after a single trial ending with the possession of a set of items P by the formula:

$$u(a) = \frac{\sum_{x \in I \cap P} D(x) - \sum_{x \in I - P} D(x)}{\sum_{x \in I} D(x)}$$

Schedule

- *Project deliverable 2*: This deliverable will consist of a working TeamBots implementation of the simulation environment with sample simple agents as well as accompanying documentation.
- *Project deliverable 3*: This deliverable will consist of an implementation of my relative performance evaluator and a demonstration that it is unbiased. As well, the simple agents from deliverable 2 will be extended to make use of different agent architectures derived from topics covered in chapters 2, 3, and 4 of our textbook (Bayesian and game-theoretic agent design) and will be presented with accompanying documentation and comparative evaluations.
- *Project deliverable 4*: This deliverable will consist of first renditions of agents implemented with learning algorithms (chapter 13 of our textbook), comparative evaluation results with the agents of deliverable 3, and associated documentation.
- *Project deliverable 5*: This deliverable will consist of final renditions of a number of agents as refined from deliverables 3 and 4 with updated documentation.
- *Post-deliverable 5*: This work will consist of performance evaluations and analysis, a presentation detailing exactly how the project went, and a final write-up of lessons learned and achievements over the course of the project.

Conference suitability

At this time, I do not know whether or not I will attempt to write up this project as a conference paper. Ultimately, this depends on whether or not I come up with provably useful modifications to the existing algorithms implemented for the electronic market domain in this project. As the course has not yet proceeded to a point where we have discussed many of the ideas I wish to explore, I cannot say that such a thing will happen.

Web link

http://www.ariwilson.com/daiproj.aspx.