Designing Adaptive Trading Agents

DAVID PARDOE
Yahoo! Labs
and
PETER STONE
The University of Texas at Austin

This extended abstract summarizes the research presented in Dr. Pardoe's recently-completed Ph.D. thesis [Pardoe 2011]. The thesis considers how adaptive trading agents can take advantage of previous experience (real or simulated) in other markets while remaining robust in the face of novel situations in a new market. Its contributions are at the intersection of machine learning and electronic commerce, with particular focus on transfer learning and fully autonomous trading agents.

Categories and Subject Descriptors: I.2.6 [**Computing Methods**]: Artificial Intelligence

General Terms: Algorithms, Design, Economics, Experimentation

Additional Key Words and Phrases: agents, auctions, machine learning, Trading Agent Competition

Along with the growth of electronic commerce has come an interest in developing autonomous trading agents. In many situations, such agents must interact directly with other (human or automated) market participants, and so the behavior of these participants must be taken into account when designing agent strategies. There are two commonly used approaches to doing so: the game-theoretic approach of finding an equilibrium, and the empirical approach of using historical market data to create a complete model of the market. These approaches may not be useful in certain cases, however, such as when either historical market data or information about the preferences of other market participants is unavailable. The focus of our work is such a case: that of an agent entering a new market in which it has no previous experience. While the agent could adapt by learning about the behavior of other market participants, it would need to do so in an online fashion while interacting with the market. The agent would not necessarily have to learn from scratch, however. If the agent had previous experience in similar markets, or could simulate a plausible range of experience in the new market, the agent could use this experience to tailor its learning approach to its particular situation.

Our work focuses on answering the question of exactly *how* an agent should make use of this experience. Within the context of a specific market scenario, the answer depends on the following four questions:

(1) What does the agent need to learn about to adapt?

(2) What sources of previous market experience are available? More specifically:
—If using existing data from other markets, which markets are relevant?

Authors' email: dpardoe@yahoo-inc.com, pstone@cs.utexas.edu

—If using simulation, what behaviors will be used by the simulated market participants?

(3) What form of learning will be used to adapt?

(4) How will this previous market experience be used to improve the adaptive approach?

The answers to the first three questions depend on the particular market scenario being considered, and our work spans distinct scenarios for which the answers to these questions differ. Our ultimate goal is to provide answers to the fourth question by presenting methods that can be used to improve the performance of specific learning algorithms in a domain-independent way. These methods can be divided into four general categories (which are not necessarily mutually exclusive and may be combined in some cases):

A. Learn from the new market only.

B. Attempt to identify the new market from a space of possible markets, or maintain a distribution over possible markets. In this case, information about previous or simulated markets helps us choose this space, while information from the new market is used to identify the market.

C. Choose a learning method that considers only information from the new market, but is tuned through experience in previous markets. This tuning could take a variety of forms, such as choosing initial parameters or the structure of a model.

D. Learn from both the new market and the previous markets by using a combination of either data or models from all sources. In either case, the challenge is to identify which data/models are relevant, and then determine how to perform the combination.

The two main contributions of our work are thus:

—the introduction and experimental validation of a number of novel algorithms for market adaptation fitting these categories; and

—an exploration of the degree to which the nature of market experience (quantity of experience, similarity of previous and current markets, non-stationarity, etc.) impacts the relative performance of methods from these categories.

We consider a variety of different trading agent settings. In one, an agent acting as an auctioneer wishes to maximize its revenue by adapting the parameters of an auction mechanism over the course of many auctions in response to bidder behavior, and the problem of adaptation can be framed as a reinforcement learning problem. We show that for a particular choice of reinforcement learning algorithm, a metalearning process can be used to select the learning parameters of the algorithm in such a way that the agent can quickly identify optimal auction parameters when bidder behavior is similar to previous experience, yet the agent remains robust to unexpected situations.

The remainder of our work takes place in two scenarios from the Trading Agent Competition[1]: supply chain management and ad auctions. In the supply chain

[1]`http://tradingagents.org`

management scenario, six agents compete as computer manufacturers and must purchase computer components, assemble computers, and sell the computers to customers. In the ad auction scenario, eight agents compete as advertisers to see who can make the most profit from selling a limited range of home entertainment products. Agents negotiate with a search engine to have their ads shown alongside search results by bidding in keyword auctions. Our agent TacTex is the reigning champion in both scenarios (winner of the 2010 competitions). The high-level design of both agents is similar: TacTex makes predictions about the prices it will encounter in the market and then optimizes its actions with respect to these predictions. In both scenarios, price prediction is accomplished by training regression models using supervised learning algorithms. The accuracy of these predictions has a significant impact on agent performance, and so developing accurate price models has been a key focus of our work.

One problem in both TAC and real world markets is that data for training models for the current market (i.e., games against the current set of competing agents in TAC) may be limited, even though data from experience in previous markets (i.e., games against different sets of TAC agents) may be plentiful. This problem can be framed as an instance of *transfer learning*: learning a concept of interest given some data representing that concept and some data representing different, possibly related concepts. While transfer learning has attracted significant recent interest in the machine learning community, the development and evaluation of algorithms for regression transfer that would be suitable for learning price models in TAC has been limited. Our research agenda has thus included the development and comparison of several regression transfer algorithms (some novel and some based on existing classification transfer algorithms) fitting the four categories of adaptation described previously. In particular, we identified a boosting-based regression transfer algorithm that was consistently the top performer across many diverse TAC and non-TAC data sets.

Overall, we expect the thesis to be interesting and useful to anyone who is interested in applying machine learning to agent-based automated trading in electronic markets. It will be particularly useful to people interested in participating in the ongoing series of Trading Agent Competitions, especially the supply chain management and ad auction scenarios, but also the upcoming electric power market scenario. In addition to the full thesis, focused publications regarding the separate contributions of the thesis are available online at `http://www.cs.utexas.edu/~TacTex/dpardoe`.

REFERENCES

Pardoe, D. 2011. Adaptive trading agent strategies using market experience. Ph.D. thesis, The University of Texas at Austin. UTCS Tech Report TR-11-26. Available from `http://www.cs.utexas.edu/~TacTex/dpardoe`.