Decision Validation for Agents in Trading Competitions

Brett Borghetti
Department of Computer Science and Engineering
University of Minnesota
borg@cs.umn.edu

ABSTRACT

This research focuses on answering the question: How do we validate the decisions of autonomous trading agents? We start with a holistic approach, exploring evaluation techniques for measuring the performance of competitive autonomous agents during controlled competitive simulations. We focus on cases where the number of opponents is small such as in the Trading Agent Competition for Supply Chain Management (TAC-SCM). In these situations each agent has a large influence on the outcome of the game. We create a methodology, provide benchmarking tools and metrics for developers to reveal and resolve undesirable decision-making behavior in their agents. This framework provides a basis from which to scale the research to the larger context of real world trading agents.

1. INTRODUCTION

This paper describes the first step of a larger area of research in which we address the question: How do we validate the decisions made by an autonomous trading agent? During this first phase, we focus on exploring the concept of error in agent decisions and the relationship between error and the meaning of a valid decision in a small, well defined competitive environment. This first phase is composed by the following steps:

1. Building a stable isolated environment for exploring agent performance.
2. Identifying internal measurements that are indicative of active decision-making.
3. Defining the concept of decision error and how error relates to decision validation in the context of both game rules and general economic principles.
4. Developing methods of measuring the error.
5. Determining whether there is a strong connection between validated decisions in the isolated environment and performance in a real competition.

To begin phase one, we chose a well known framework in which to perform this research: the Trading Agent Competition for Supply Chain Management (TAC-SCM). TAC-SCM is one of the most prominent proving grounds for current research in autonomous trading agents. In this yearly international event, autonomous agents battle for supremacy in a simulation where the highest profit-earning agent wins. In TAC-SCM, agents make all the decisions to run a virtual computer-manufacturing operation. They negotiate to purchase parts from suppliers, optimize their assembly lines, and sell by auction their products to customers. They manage parts and product inventories, minimize costs, optimize profit, and try to beat their competitors.

Agents compete against five other adversaries in each game. Different combinations of competitors, in addition to the randomness in the game, cause different market conditions to arise. An agent must perform well under many different market conditions to succeed. While strong competitors are assured in the yearly competition, developers do not have access to adversary agents prior to the event. This poses a problem which has led to the first step in phase one of our research: How does one simulate tough competition and measure expected performance of a potential agent outside of the actual event? The answer lies in extracting and modeling the behavior of strong competitors outside of a competition.

To explore an agent's decision-making behavior under various levels of competition, we developed a pair of agents that can simulate competition in our isolated environment by modeling the behavior of competitors. These stand-alone market manipulator agents to create a configurable level of pressure so that the entire spectrum of competition can be simulated in our isolated environment.

2. RELATED WORK

TAC-SCM allows research teams to develop trading agents and compare their relative performance in a complex, standardized environment [1]. Several teams developed methods of analyzing agent performance. The University of Southampton team examined running variations of their competitive agent with different risk strategies for pricing [4] to show that their competitive agent made the highest profit among the variations. The University of Michigan team analyzed post-competition performance of the TAC-SCM winning agents and explored relationships between total profit and other measurements of performance [3]. The team at the University of Texas at Austin focused on comparing multiple types of learning agents competing against each other [5]. While each of these teams characterized their agent's performance against other agents, no one developed a isolated stable, controlled competition environment in which to measure agent performance. We felt we must explore this region of agent performance analysis.

3. APPROACH

There are two key challenges in designing a useful isolated environment for TAC. The first is how to build controlled competitive

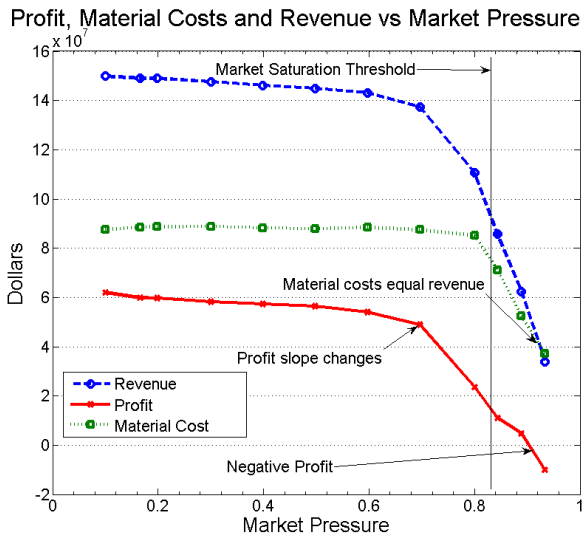

Figure 1: Market Pressure Effects on MinneTAC

environments for observing agent behavior. The second is how to manage the randomness and repeatability of the games. While we addresses both of these challenges in other work, in this paper we focus on the first.

To simulate the market conditions existing in a TAC-SCM competition, we developed two new agents, the *Market Relief* or "do nothing" agent, and the *Market Pressure* agent. The Market Relief Agent occupies one or more of the six slots in a TAC simulation without making any financial transactions. This agent provides relief to the market from the perspective of the other agents: it reduces demand on the suppliers which leads to a lowering of supplier prices; and it reduces available supply for the customers which leads to an increase in what customers were willing to pay for computers from other agents.

Conversely, the configurable Market Pressure Agent does the opposite: it increases available supply to customers, allowing them to pay less for computers while simultaneously putting more demand on suppliers, encouraging them to increase their prices. This Market Pressure Agent operates by continually adjusting its customer offer prices to achieve a desired market share on the interval 0% to 100% . Using combinations of Market Relief and Market Pressure Agents, we can simulate the entire spectrum of competition in an isolated environment

4. EXPERIMENTS

We examine what happens to an agent when altering the level of competition in the marketplace. As shown in Figure 1, our MinneTAC [2] agent earns a relatively constant profit (shallow slope) until the Market Pressure Agent absorbs significant market share near the *Market Saturation Threshold*: 83% $(5/6)$ of the available market. Market Saturation represents the point at which there is little unmet customer demand remaining: if one agent wishes to sell more computers and gain market share from another agent in this highly competitive environment, price wars occur and profit margins suffer.

Notice the profit slope changes abruptly around 70% market pressure. The point at which the slope of agent profit changes in this region is relevant: a better performing agent has a shallower slope through a higher market pressure, so if we make adjustments to our agent, we should chose adjustments that slide this point right to a higher market pressure and make the slope more shallow. Ideally we would like to make our agent attain the shallowest profit slope as far right as possible.

We examined two other agents: TacTex and MerTACor, in separate controlled experiments using the market pressure agents to simulate competition. The results vary in profit slopes and the point at which each agent crosses the $0 profit threshold. Interestingly, under identical conditions, the TacTex agent earns more profit than either of the other agents by a wide margin, yet crosses the $0 profit threshold at a much lower market pressure than the other agents. Since TacTex won the 2005 TAC-SCM competition, our intuition is that the pressure exerted on TacTex by the other agents was below this threshold, and they were able to operate in an area where TacTex profit exceeded that of the other agents.

5. CONCLUSION AND FUTURE WORK

We developed a framework of tools to control the simulation environment and measure the performance of trading agents under the full spectrum of market conditions. These tools allow research teams to use controlled experiments to examine how their agents will perform in various market conditions.

The importance of this style of analysis is that it reveals that while the aggregate outcomes of an agents' decisions vary widely over the spectrum of competition, certain trends are evident. The limitation of this first step is that it reveals these traits but provides no explanations why they occur. As we continue with phase one research, in the next step we will measure the internal decisions of the agents so that we can dissect these aggregate outcomes and determine why the agent behaves as it does. We will continue to explore and define the concept of decision error and validity of decisions in the context of the environment

In the second phase of the research we will explore the scalability of this framework to real world problems in which the number of competitors is unbounded and the effects of individual actions on the marketplace is more subtle. Despite the differences, we hope to find generalities that describe and codify the relationships between decision error and validity.

6. REFERENCES

[1] J. Collins, R. Arunachalam, N. Sadeh, J. Ericsson, N. Finne, and S. Janson. The supply chain management game for the 2006 trading agent competition. Technical Report CMU-ISRI-05-132, Carnegie Mellon University, Pittsburgh, PA, November 2005.

[2] W. Ketter, E. Kryzhnyaya, S. Damer, C. McMillen, A. Agovic, J. Collins, and M. Gini. Analysis and design of supply-driven strategies in TAC-SCM. In *AAMAS-04 Workshop on Trading Agent Design and Analysis*, pages 44–51, New York, July 2004.

[3] C. Kiekintveld, Y. Vorobeychik, and M. P. Wellman. An analysis of the 2004 supply chain management trading agent competition. In *IJCAI-05 Workshop on Trading Agent Design and Analysis TADA-05*, Edinburgh, Scotland, August 2005.

[4] D. E. Minghua He, Alex Rogers and N. R. Jennings. Designing and evaluating an adaptive trading agent for supply chain management applications. In *IJCAI-05 Workshop on Trading Agent Design and Analysis TADA-05*, Edinburgh, Scotland, August 2005.

[5] D. Pardoe and P. Stone. Bidding for customer orders in TAC SCM. In *AAMAS 2004 Workshop on Agent Mediated Electronic Commerce VI*, July 2004.