An Empirical Study of Trading Agent Robustness (Extended Abstract)

Shai Hertz
Tel Aviv University
shaihertz@tau.ac.il

Mariano Schain
Tel Aviv University
marianos@tau.ac.il

Yishay Mansour
Tel Aviv University
mansour@tau.ac.il

ABSTRACT

We study the empirical behavior of trading agents participating in the Ad-Auction game of the Trading Agent Competition (TAC-AA). Aiming to understand the applicability of optimal trading strategies in synthesized environments to real-life settings, we investigate the robustness of the agents to deviations from the game's specified environment. Our results indicate that most agents, especially the top-scoring ones, are surprisingly robust. In addition, using the game logs, we derive for each agent a *strategic fingerprint* and show that it almost uniquely identifies it. Finally, we show that although the Machine Learning modeling in TAC-AA is inherently inaccurate, further improvement in modeling accuracy is likely to have only a limited contribution to the overall performance of TAC-AA agents.

Categories and Subject Descriptors

I.2.6 [**ARTIFICIAL INTELLIGENCE**]: Learning; I.2.11 [**ARTIFICIAL INTELLIGENCE**]: Distributed Artificial Intelligence—*Intelligent agents, Multiagent systems*

Keywords

Agents; Robustness; Learning

1. AGENT'S ROBUSTNESS

The Ad-Auctions (AA) game in the yearly Trading Agent Competition (TAC) presents a sponsored search scenario that employs an ad auction mechanism and a structured model of users [2]. Competitors in this game implement retailers that aim to maximize their profit through the use of sponsored search advertising. The competing agents are developed to take advantage of various features of this synthetic environment. Ad-Auctions games have been held since 2009, and in the course of time the agents improved their performance by employing complex techniques and strategies [4, 1, 5].

The TAC setting facilitates research of agent strategies in a multi-agent competitive environment. Furthermore, it can be used to draw more general conclusions about ad-auction mechanisms and sponsored search [3]. A key motivation of this paper is therefore to show that the TAC-AA game achieves this purpose, and we do so by establishing the robustness of the top performing TAC-AA agents: By testing whether the agents can adapt to a different environment and still perform well, as expected from agents in a real world. To that end, we modified the game parameters and tested the effect of this modification on the performance of some recent TAC-AA agents. We show that although (as expected) most of the agents are tailored to the specific game parameters, the top performing agents exhibit robustness and perform well even when the parameters are changed.

In order to assess the robustness of TAC-AA agents, we ran several experiments in which we varied some of the game parameters (without revealing this fact to the agents) and ran a standard 48-game competition in each new setting: We modified the user model (i.e. user state transition probabilities), the click-through rate range, the conversion rate, the number of ads (specifically, we simulated a display-ad setting by reducing the number of advertising slots from 5 to 1), and the population sizes.

For each experiment, we compared the score of each agent to its score in a benchmark competition, and noted the difference in the agent's position. We ran t-tests with 0.95 confidence level to find the statistical significance of this difference. We also compared the median score in each experiment to the median score of the benchmark competition, in order to understand the general effect of the changes we made.

Our experiments show that most of the TAC-AA agents adapt well to different settings, despite being optimized to the exact game parameters. The top performing agents of TAC-AA - TacTex, Schlemazl and tau11 - are rather robust to most changes, but when the setting is changed drastically, as in the last two experiments, their performance deteriorates. This robustness result is somewhat surprising, since one could expect that the top performing agents in the TAC-AA would be more optimized to the exact game parameters and thus will be more affected by changes in these parameters. However, the experiments show that most agents are less overfit to the game parameters than expected.

2. AGENTS BEHAVIORAL IDENTIFICATION

Another objective of our research is to define a *strategic fingerprint* of a TAC-AA agent and characterize its behavior. To achieve this objective, we define several observable attributes (such as average ad position, standard deviation of daily profits, quary distribution of impressions, and more) that are calculated from the game logs for each agent in each

Appears in: *Proceedings of the 12th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2013)*, Ito, Jonker, Gini, and Shehory (eds.), May, 6–10, 2013, Saint Paul, Minnesota, USA.

game, and we incorporate them into an attribute vector we call a strategic fingerprint.

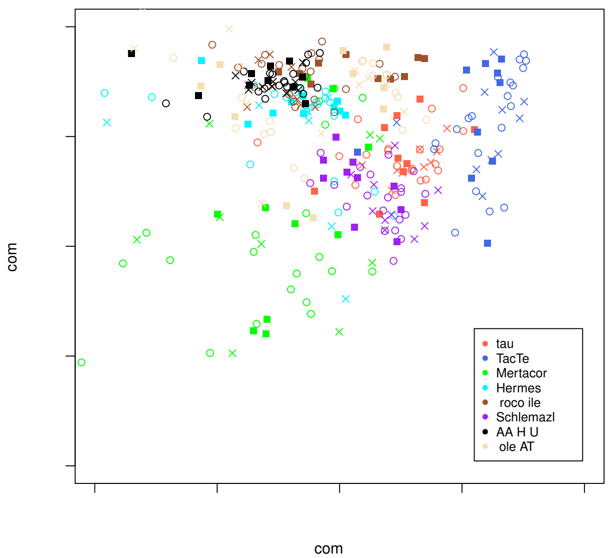

Figure 1: The Strategic Fingerprints of TAC-AA 11 finalists, projected on the 2 most principal components

To visually illustrate the strategic ranges of different agents, we used Principal Components Analysis (PCA) and projected the agents' vectors on the first 2 components. The result is shown in Figure 1, and we can see that each agent maintains a slightly different zone. To demonstrate agent identification using its strategic fingerprint, we used a simple 3-Nearest Neighbor model that classifies agents based on their strategic fingerprints. The error rate of this model was 5.9%. Most of the errors of the model are due to a specific agent, Mertacor, whose strategic range is rather wide, while the other agents are more accurately classified.

Since each agent has a typical range of strategic fingerprint vectors, we conclude that our strategic fingerprints model well the behavior of TAC-AA agents. We show that this strategic fingerprint not only identifies agents with high precision, but is also well correlated with their profit. In addition, it reflects the vulnerability of simple log anonymization, and demonstrates that it can be overcome using simple Machine Learning (ML) tools.

3. MACHINE LEARNING

TAC-AA agents employ different Machine Learning techniques trying to estimate different game parameters and states, and perform well in the game despite the fact that the ML models used are inaccurate. Our third goal was to estimate the benefit that agents can obtain from improving their Machine Learning components, that is, the potential score improvement that can be achieved by using more accurate models. To that end, we modified the game server so it will send to one of the agents some of the unobservable parameters the agents try to learn, so this agent has a perfect knowledge of them. Specifically, we revealed in one experiment the Cost Per Click (CPC), click-through rates (CTR), and effective number of impressions. In a second experiment we revealed user distributions over states. This information obviates the need for certain ML models, and in fact simulates the use of perfect models (For example, to assess the user distribution over states our original agent uses a Nearest-Neighbor classifier that has a 25% error rate), Hence, these tests enable us to assess the additional profit one could hope to gain by improving these ML models.

Our results indicate that even though the Machine Learning models are inherently inaccurate, eliminating their error rates completely has only a minor effect on the performance of the agent: In each of the settings above, the informed agent's average score improvement was insignificant. However, when both information sets were revealed the agent's performance was improved by 5%, an improvement which is statistically significant ($p < 0.05$). This result is somewhat surprising when considering the minor effect that each set of information had by itself. It seems that improving only one model has little contribution, but having perfect models in both domains is more significant. The results of these three experiments suggest that improving the accuracy of the ML models employed in TAC-AA has a limited contribution to the final result. Apparently, the high error rate of the ML models does not lead to impaired performance, and this error rate is somehow overcome during the optimization process of the agent. Presumably, the estimation process in the real world is even harder than in the limited setting of TAC-AA. However, our findings suggest that top performance may be achieved by trading agents even in the absence of accurate estimation.

4. ACKNOWLEDGMENTS

This research was supported in part by the Google Inter-university center for Electronic Markets and Auctions, by The Israeli Centers of Research Excellence (I-CORE) program, (Center No. 4/11), by a grant from the Israel Science Foundation, by a grant from United States-Israel Binational Science Foundation (BSF), and by a grant from the Israeli Ministry of Science (MoS).

5. REFERENCES

[1] J. Berg, A. Greenwald, V. Naroditskiy, and E. Sodomka. A knapsack-based approach to bidding in ad auctions. In *ECAI*, pages 1013–1014, 2010.

[2] P. R. Jordan and M. P. Wellman. Designing an ad auctions game for the trading agent competition. In *AMEC*. 2010.

[3] P. R. Jordan, M. P. Wellman, and G. Balakrishnan. Strategy and mechanism lessons from the first ad auctions trading agent competition. In *ACM Conference on Electronic Commerce*, pages 287–296, 2010.

[4] D. Pardoe, D. Chakraborty, and P. Stone. Tactex09: a champion bidding agent for ad auctions. In *AAMAS*, pages 1273–1280, 2010.

[5] D. Pardoe and P. Stone. A particle filter for bid estimation in ad auctions with periodic ranking observations. In *AAMAS*, pages 887–894, 2011.