ECAI 2008
M. Ghallab et al. (Eds.)
IOS Press, 2008

doi:10.3233/978-1-58603-891-5-861

An Empirical Investigation of the Adversarial Activity Model

Inon Zuckerman and **Sarit Kraus** [1] and **Jeffrey S. Rosenschein** [2]

Abstract.

Multiagent research provides an extensive literature on formal Belief-Desire-Intention (BDI) based models describing the notions of teamwork and cooperation, but adversarial and competitive relationships have received very little formal BDI treatment. Moreover, one of the main roles of such models is to serve as design guidelines for the creation of agents, and while there is work illustrating that role in cooperative interaction, there has been no empirical work done to validate competitive BDI models.

In this work we use the *Adversarial Activity* model, a BDI-based model for bounded rational agents that are operating in a general zero-sum environment, as an architectural guideline for building bounded rational agents in two adversarial environments: the *Connect-four* game (a bilateral environment) and the *Risk* strategic board game (a multilateral environment). We carry out extensive simulations that illustrate the advantages and limitations of using this model as a design specification.

1 Introduction

Formal Belief-Desire-Intention (BDI) [1] based models of cooperation and teamwork have been extensively explored in multiagent worlds. They provide firm theoretical foundations and guidelines for the design of cooperative automated agents [4, 2]. However, as cooperation and teamwork led the research agenda, little work was done on providing BDI-based models for adversarial or competitive interactions that naturally occur in multiagent environments. The desire to adapt BDI-based models for competitive interactions comes from their successful implementation in teamwork domains [5] and the limitations of classical solutions in complex adversarial interactions.

Recently, the *Adversarial Activity* (AA) model [6] was presented: a formal BDI-based model for bounded rational agents in zero-sum adversarial environments. Alongside the model were also presented several behavioral axioms that should be used when an agent finds itself in an *Adversarial Activity*. However, the discussion in [6] lacked empirical work to validate the advantages as well as the limitations of those behavioral axioms in adversarial domains. Our aim here is to fill that gap, demonstrate how the AA model can be used as a design specification, and investigate its usefulness in bounded rational agents. We will explore whether AA-based agents can outperform state of the art solutions in various adversarial environments.

2 Overview of the *Adversarial Activity* Model

The AA model provides the specification of capabilities and mental attitudes of an agent in an adversarial environment from a single adversarial agent's perspective. The model describes both bilateral and multilateral instantiations of zero-sum environments, in which all agents are adversarial (i.e., there are no cooperative or neutral agents). Alongside the model, there exist several behavioral axioms that the agent can follow:

A1. Goal Achieving Axiom. This axiom is a simple and intuitive one, stating that if the agent can take an action that will achieve its main goal (or one of its subgoals), it should take it.

A2. Preventive Act Axiom. This axiom relies on the fact that the interaction is zero-sum. It says that the agent might take actions that will prevent its adversary from taking future high beneficial actions, even if they do not explicitly advance the agent towards its goal.

A3. Suboptimal Tactical Move Axiom. This axiom relies on the fact that the agent's reasoning resources are bounded, as is the knowledge it has about its adversaries. In such cases the agent might decide to take actions that are suboptimal with respect to its limited search boundary, but they might prove to be highly beneficial actions in the future, depending on its adversaries reactions.

A4. Profile Manipulation Axiom. This provides the ability to manipulate agents' profiles (the knowledge one agent holds about the other), by taking actions such that the adversary's reactions to them would reveal some of its profile information.

A5. Alliance Formation Axiom This axiom allows the creation of temporary task groups when, during the interaction, several agents have some common interests that they wish to pursuit together.

A6. Evaluation Maximization Axiom. In a case when all other axioms are inapplicable, the agent will proceed with the action that maximizes the heuristic value as computed in its evaluation function.

3 Empirical Evaluation

We will use two different experimental domains. The first one is the *Connect-Four* board game, which will allow us to evaluate the model in a bilateral interaction. The second domain is the well-known *Risk* strategic board game of world domination.

The embedding of behavioral axioms into the agent design, in both domains, was done by providing new functions, one for each of the implemented axioms (denoted as $AxiomNValue()$, where N is the number of the axiom in the model). These functions return a possible action if its relevant precondition holds. The preconditions are the required beliefs, as stated in the axiom formalizations, formulated according to the relevant domain. The resulting architecture resembles a rule-based system, where each function returns its value and the final selection among the potential actions is computed in a "*Decide*" function, whose role is to select among the actions (if there is more than a single possible action) and return its final decision.

3.1 A Bilateral Domain—Connect4

We built an experimental environment where computer agents play the *connect-four* game against one another, and we have control over

[1] Bar-Ilan University, Israel, email: {zukermi,sarit}@cs.biu.ac.il
[2] The Hebrew University, Israel, email: jeff@cs.huji.ac.il

the search depth, reasoning time, and other variables. We built six different agents, each with a different evaluation function (H1-H6), ranging from a naive function to a reasonable function that can win when playing against an average human adversary.

We had 12 different agents: 6 alpha-beta and 6 axiom-augmented agents, each using one of the evaluation functions. We staged a round-robin tournament among all agents, where each agent played with 3 different depth searches (3, 5, and 7) against all other agents and possible search depths. The tournament was played twice: once for the agents playing as the first player (yellow), and the other time for them playing as the second (red) player (i.e., 11 opponents * 3 own depth * 3 opponent depth * 2 disc colors = 198 games).

The results of the tournament are summarized in Figure 1. The figure shows the percentage of winning games for each of the 12 agents, where the agent names are written as R_1 for regular agent using H1, and A_3 shows the results for axiom-embedded agents using H3. The results clearly indicate that all agents improved their performance following the integration of axioms. The agents with naive heuristics (A_1 and A_2) showed only a small improvement, which usually reflected additional wins over their "regular" versions (R_1 and R_2), while the mid-ranged functions ($H4$ and $H5$) showed the largest improvement, with additional wins over different agents that were not possible prior to the embedding of axioms. Overall, we see that the best two agents were A_4 and A_6, with a single win advantage for the A_6 player, which in turn led A_5 by 7 wins.

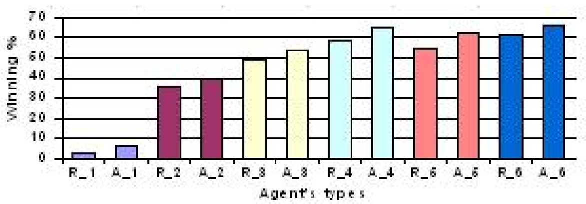

Figure 1. Connect-Four experiment results

3.2 A Multilateral Domain—Risk

Our next domain is a multilateral interaction in the form of the *Risk* board game. The game is a strategy board game that incorporates probabilistic elements and strategic reasoning in various forms. *Risk* is too complicated to solve using classical search methods.We used the *Lux Delux*[3] environment which provides a large number of computer opponents implemented by different programmers and employing varying strategies. We chose to work with exactly the same subset of adversaries that was used in [3], which contains 12 adversaries of different difficulty levels (Easy, Medium, and Hard): (1) *Angry* (2) *Stinky* (3) *Communist* (4) *Shaft* (5) *Yakool* (6) *Pixie* (7) *Cluster* (8) *Bosco* (9) *EvilPixie* (10) *KillBot* (11) *Que* (12) *Nefarious*.

The basic agent implementation and evaluation function were based on the one described in [3], as it proved to be a very successful evaluation function-based agent, which does not use expert knowledge about the strategic domain. The next step was to augment the original agent with the implementation of the adversarial axioms (we used continent ownership as a subgoal).

Experiment 1: The game map was "Risk classic", card values were set to "5, 5, 5, . . . ", the continent bonus was constant, and starting position and initial army placement were randomized. Each game had 6 players, randomized from the set of 14 agents described above. Figure 2 shows results of running 1741 such games, with the winning percentage of each of the agents (we use the agent number from the above list instead of their names). The worst agent was *Angry* (#1) with a 0.44% win percentage, while the best was *KillBot* (#10) with 32.54%. Looking at our agents, we can see that the basic heuristic agent (denoted as "He" and whose bar is colored in blue) managed to achieve only 11.79%, whereas its axiom-augmented version *Ax* (colored red on the graph) climbed all the way up to 26.84%, more than doubling the winning percentage of its regular version.

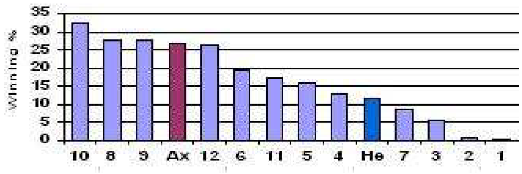

Figure 2. Winning percentage on "Risk classic" map

Experiment 2: In the second experiment we compared the performance of both kinds of agents on randomly-generated world maps. The results show approximately the same improvement, from 9.16% with the regular heuristic agent, to a total of 21.36% with its axiom-augmented version.

Experiment 3: We fixed a five-agent opponent set (agent 1 through 5), and ran a total of 2000 games on the classic map setting: 1000 games with agent *He* and the opponent set, and 1000 games with agent *Ax* and the opponent set. The results show that even when playing against very easy opponents, in which the regular heuristic agent led the group with a winning percentage of 31.8%, the integration of the axioms managed to lift the agent to an impressive winning percentage of 57.1%.

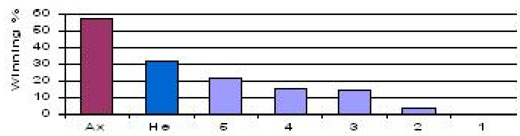

Figure 3. Winning percentage with fixed opponents

4 Conclusions

We have presented an empirical evaluation of the *Adversarial Activity* model for bounded rational agents in a zero-sum environment. Our results show that bounded-rational agents can improve their performance when their original architectures are augmented with the model's behavioral axioms, even as their evaluation functions remained unchanged.

5 Acknowledgments

This work was supported in part by NSF under grant #IS0705587 and ISF under grant #1357/07. Sarit Kraus is also affiliated with UMIACS.

REFERENCES

[1] Michael E. Bratman, *Intention, Plans and Practical Reason*, Harvard University Press, Cambridge, MA, 1987.

[2] Barbara J. Grosz and Sarit Kraus, 'Collaborative plans for complex group action', *AIJ*, **86**(2), 269–357, (1996).

[3] Stefan J. Johansson and Fredrik Olsson, 'Using multi-agent system technology in risk bots.', in *AIIDE*, pp. 42–47, (2006).

[4] H. J. Levesque, P. R. Cohen, and J. H. T. Nunes, 'On acting together', in *Proc. of AAAI-90*, pp. 94–99, Boston, MA, (1990).

[5] M. Tambe, 'Agent architectures for flexible, practical teamwork', in *National Conference on Artificial Intelligence (AAAI)*, (1997).

[6] Inon Zuckerman, Sarit Kraus, Jeffrey S. Rosenschein, and Gal A. Kaminka, 'An adversarial environment model for bounded rational agents in zero-sum interactions', in *AAMAS 2007*, pp. 538–546, (2007).

[3] Downloadable from http://sillysoft.net/lux/.