

Modeling Human Adversary Decision Making in Security Games: An Initial Report

(Extended Abstract)

Thanh H. Nguyen, James Pita,
Rajiv Maheswaran, Milind Tambe
University of Southern California, Los Angeles,
CA 90089, USA
{thanhng,jpita,maheswar,tambe}@usc.edu

Amos Azaria⁺, Sarit Kraus⁺⁺
⁺Bar-Ilan University, Ramat Gan 52900, Israel
⁺⁺Institute for Advanced Computer Studies,
University of Maryland, College Park, MD 20742
{azariaa1,sarit}@cs.biu.ac.il

ABSTRACT

Motivated by recent deployments of Stackelberg security games (SSGs), two competing approaches have emerged which either integrate models of human decision making into game-theoretic algorithms or apply robust optimization techniques that avoid adversary modeling. Recently, a robust technique (MATCH) has been shown to significantly outperform the leading modeling-based algorithms (e.g., Quantal Response (QR)) even in the presence of significant amounts of subject data. As a result, the effectiveness of using human behaviors in solving SSGs remains in question. We study this question in this paper.

Categories and Subject Descriptors

I.2.11 [Artificial Intelligence]: Distributed Artificial Intelligence

General Terms

Algorithms, Experimentation, Security

Keywords

Game Theory, Human Behavior, Quantal Response, Robust Optimization, Bounded Rationality


1. INTRODUCTION

Stackelberg Security Games (SSGs) have drawn great attention in solving real-world security problems in which security agencies (“defender”) have to allocate their limited resources to protect important settings against human adversaries [13, 1, 15, 6]. Multiple recent SSG-based deployments attempt to compute the optimal defender strategy with a key assumption that the adversary will respond optimally, i.e., he (the adversary is “he” by convention) tries to maximize his expected value given the defender’s strategy. Nevertheless, in real-world problems, the adversary’s decision may be governed by his bounded rationality [5, 7]; he may deviate from the optimal action due to the effects of the complexity of the problem or the emotion, etc. Thus, the assumption of perfect rationality is not robust for addressing bounded rationality of human adversaries [3]. As a result, alternative approaches to overcome this limitation in solving SSGs have been proposed [10, 14, 11].

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.

Copyright © 2013, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

Two leading approaches have emerged to handle human bounded rationality in SSGs. One suggests that we integrate human behavior models into algorithms for solving SSGs and is exemplified by the BRQR algorithm which applies Quantal Response (QR) [8] for representing human decision making of the adversary [14]. Another approach computes the optimal strategy for the defender using robust optimization techniques and intentionally avoids creating human behavior models. This approach is exemplified by the MATCH algorithm [11]. In particular, BRQR assumes that the adversary responds stochastically in SSGs following the QR model; where the lower the cost of the deviation in terms of expected value, the more likely that the deviation occurs. In the BRQR algorithm, the key parameter λ is used for measuring the bounded rationality of the adversary. In contrast, MATCH attempts to bound the loss of the defender if the adversary deviates from the optimal action without creating a human behavior model. It maximizes the defender’s expected value with the constraint that the loss of the defender is less than a factor of β times the loss of the adversary in terms of expected value with regard to his deviation. The key parameter β is used for controlling the loss of the defender. Pita et al. have shown that MATCH significantly outperformed BRQR even when a significant amount of data is used for tuning the parameter λ of BRQR, even while no tuning is done for MATCH’s β parameter [11]. Thus, MATCH, with its robust optimization, is now suggested to be the sole dominant algorithm to handle human adversaries in SSGs.



(a) ARMOR

(b) PROTECT

Figure 1: Game-theoretic applications

This result has raised an important open question of whether there is any value in applying human behavior models for solving SSGs. We attempt to answer this question in this paper.

2. SSG-BASED SECURITY SIMULATION

Our work is motivated by multiple game-theoretic applications deployed in security domains such as ARMOR [9] being used by

the LAWA police for protecting LAX; the largest destination airport in the United State (Figure 1a) and PROTECT [12] being used by United States Coast Guard (USCG) in Boston (Figure 1b).

Recall that SSGs are a class of Stackelberg games where a defender acts as a leader and an adversary acts as a follower [4, 2, 6]. While the defender attempts to allocate her (the defender is “she” by convention) limited resources to protect a set of targets, the adversary plans to attack one such target. SSGs are commonly used in real-world security domains because it captures the fact that the defender first commits to a mixed strategy assuming that the adversary can observe that strategy; then, the adversary takes his action.

In SSGs, the information presented to a human subject, who acts as an adversary, for each choice includes: the marginal coverage on target t , the reward and penalty of the adversary, and finally, the reward and penalty of the defender at the target. Let T be the number of targets and K be the number of resources of the defender. The payoffs of both the defender and adversary depend on the attacked target and whether the defender covered that attacked target or not. When the adversary attacks a target t , he will receive a reward R_t^a if the target is not covered by the defender; otherwise, he will receive a penalty P_t^a . On the contrary, the defender will get a penalty P_t^d in the former case and a reward R_t^d in the latter case, respectively. We have $R_t^a, R_t^d > 0$ and $P_t^a, P_t^d < 0$. Let x_t be the coverage probability of the defender at target t . The expected value of the defender and the attacker at target t are given by $U_t^d = x_t R_t^d + (1 - x_t) P_t^d$ and $U_t^a = x_t P_t^a + (1 - x_t) R_t^a$, respectively.

3. EVALUATION

As mentioned earlier, two competing approaches have emerged to handle adversary bounded rationality in SSGs. One attempts to integrate models of human decision making into reasoning about defender strategies in SSGs. In particular, the BRQR algorithm which integrates the QR model is the leading algorithm within this approach [14]. QR models a stochastic adversary response—the greater the expected value of a target the more likely the adversary will decide to attack that target. BRQR computes an optimal strategy by assuming that the adversary will respond against the defender strategy with such a stochastic response. In BRQR, the parameter λ in the QR model represents the amount of noise in the adversary’s response.

MATCH is an alternative robust approach introduced by Pita et al. which does not integrate the human behavior models [11]. MATCH computes the optimal defender strategy with the constraint that the defender’s loss is no worse than a proportion of (β) the loss the adversary receives in terms of his expected value if he deviates from the optimal choice. In MATCH, β is the key parameter deciding how much loss the defender is willing to endure if the adversary responds non-optimally. A comparison of these two algorithms conducted by Pita et al., using over 100 payoff structures, shows that MATCH significantly outperforms BRQR even a significant amount of data is used for re-estimating the parameter λ of BRQR. In particular, they systematically selected 15 payoff structures where MATCH outperformed BRQR significantly and re-estimated the parameter of the Quantal Response model in each of these payoff structures. They re-evaluated the performance of MATCH against BRQR in the 15 selected payoff structures with the re-estimated. Their experimental results showed that MATCH still outperformed BRQR.

We suggest that the Quantal Response model integrated with the expected value function does not capture the decision-making of human adversaries in SSGs. Specifically, we hypothesize that the Quantal Response model models its stochastic response based on the expected value function; and that humans may not be driven

by such expected value. Therefore, BRQR which is based on that model performs poorly in comparison with the robust technique, MATCH.

4. ACKNOWLEDGEMENT

This research was supported by Army Research Office under the grant # W911NF-10-1-0185 and by the Google Inter-university center for Electronic Markets and Auctions, ARO under the grants # W911NF0910206 and # W911NF1110344.

5. REFERENCES

- [1] N. Basilico, N. Gatti, and F. Amigoni. Leader-follower strategies for robotic patrolling in environments with arbitrary topologies. In *AAMAS*, pages 57–64, 2009.
- [2] N. Basilico, N. Gatti, and F. Amigoni. Leader-follower strategies for robotic patrolling in environments with arbitrary topologies. In *AAMAS*, pages 57–64, 2009.
- [3] C. Camerer, T. Ho, and J. Chong. A cognitive hierarchy model of games. *The Quarterly Journal of Economics*, 119(3):861–898, 2004.
- [4] V. Conitzer and T. Sandholm. Computing the optimal strategy to commit to. In *EC*, pages 82–90. ACM, 2006.
- [5] J. Conlisk. Why bounded rationality? *Journal of economic literature*, pages 669–700, 1996.
- [6] J. Letchford and Y. Vorobeychik. Computing randomized security strategies in networked domains. In *AARM Workshop In AAI*, 2011.
- [7] J. March. Bounded rationality, ambiguity, and the engineering of choice. *The Bell Journal of Economics*, pages 587–608, 1978.
- [8] R. McKelvey and T. Palfrey. Quantal response equilibria for normal form games. *Games and economic behavior*, 10(1):6–38, 1995.
- [9] J. Pita, M. Jain, J. Marecki, F. Ordóñez, C. Portway, M. Tambe, C. Western, P. Paruchuri, and S. Kraus. Deployed armor protection: the application of a game theoretic model for security at the los angeles international airport. In *AAMAS*, pages 125–132, 2008.
- [10] J. Pita, M. Jain, M. Tambe, F. Ordóñez, and S. Kraus. Robust solutions to stackelberg games: Addressing bounded rationality and limited observations in human cognition. *Artificial Intelligence*, 174(15):1142–1171, 2010.
- [11] J. Pita, R. John, R. Maheswaran, M. Tambe, R. Yang, and S. Kraus. A robust approach to addressing human adversaries in security games. In *AAMAS*, pages 1297–1298, 2012.
- [12] E. Shieh, B. An, R. Yang, M. Tambe, C. Baldwin, J. DiRenzo, B. Maule, and G. Meyer. Protect: A deployed game theoretic system to protect the ports of the united states. In *AAMAS*, pages 13–20, 2012.
- [13] J. Tsai, C. Kiekintveld, F. Ordonez, M. Tambe, and S. Rathi. Iris—a tool for strategic security allocation in transportation networks. In *AAMAS*, 2009.
- [14] R. Yang, C. Kiekintveld, F. Ordonez, M. Tambe, and R. John. Improving resource allocation strategy against human adversaries in security games. In *IJCAI*, pages 458–464. AAAI Press, 2011.
- [15] Z. Yin, A. Jiang, M. Johnson, M. Tambe, C. Kiekintveld, K. Leyton-Brown, T. Sandholm, and J. Sullivan. Trusts: Scheduling randomized patrols for fare inspection in transit systems. In *Conference on Innovative Applications of Artificial Intelligence (IAAI)*, 2012.