A Robust Approach to Addressing Human Adversaries in Security Games

(Extended Abstract)

James Pita, Richard John, Rajiv Maheswaran, Milind Tambe, Rong Yang
University of Southern California, Los Angeles, CA 90089

Sarit Kraus
Bar-Ilan University, Ramat-Gan 52900, Israel and
Institute for Advanced Computer Studies,
University of Maryland, College Park, MD 20742

ABSTRACT

While game-theoretic approaches have been proposed for addressing complex security resource allocation problems, many of the standard game-theoretic assumptions fail to address human adversaries who security forces will likely face. To that end, approaches have been proposed that attempt to incorporate better models of human decision-making in these security settings. We take a new approach where instead of trying to create a model of human decision-making, we leverage ideas from robust optimization techniques. In addition, we extend our approach and the previous best performing approach to also address human anchoring biases under limited observation conditions. To evaluate our approach, we perform a comprehensive examination comparing the performance of our new approach against the current leading approaches to addressing human adversaries. Finally, in our experiments we take the first ever analysis of some demographic information and personality measures that may influence decision making in security games.

Categories and Subject Descriptors

H.4 [**Computing Methodology**]: Game Theory

General Terms

Security, Algorithms, Performance

Keywords

Human Behavior, Stackelberg Games, Decision-making, Security

1. INTRODUCTION

Game-theory has gained traction in security resource allocation decisions in important settings [4]. Security games refer to a special class of Stackelberg games where there are two agents - the defender (security force) and an attacker - who act as the leader and the follower respectively [9]. Traditionally, Stackelberg games have been used to model these problems because they encapsulate the commitment a defender must make in allocating her security resources before an attacker chooses an attack method.

There exists a number of game-theoretic optimal algorithms for solving security games such as DOBSS [5]. However, one of the key assumptions underlying these approaches is that the attacker is a perfectly-rational player and that the attacker breaks ties in the defender's favor. Thus, these systems optimize their strategy against an expected-value-maximizing opponent and are not robust to deviations from this strategy. It is well known that standard game-theoretic assumptions of expected-value-maximizing rationality are not ideal for addressing human behavior in game-theoretic settings [2]. To that end, a number of approaches have attempted to address these potential deviations by incorporating more realistic models of human decision-making.

COBRA is one such approach that assumes a boundedly-rational opponent and attempts to maximize the defender's utility for the worst-case outcome of any ϵ-optimal response strategy, avoiding the issue of tie breaking by the attacker [6]. One critical issue with COBRA is that if the attacker deviates to any strategy beyond the ϵ-optimal response set then the result can once again be arbitrarily bad for the defender. To address this dilemma, Yang et al. [7] introduced BRQR, which assumes that instead of strictly maximizing expected value, the attacker responds stochastically: the chance of selecting non-optimal strategies increases as the cost of such an error decreases. BRQR thus allows for a more gradual approach to defending against deviations as opposed to the hard-cutoff point. Two issues with BRQR are that it critically depends on the appropriate estimation of λ, which represents the amount of error in the attacker's response function; and that its runtime is slow.

To attempt to address the issues of BRQR and COBRA, we introduce a new approach, MATCH, based on robust optimization [1] where the defender strategy is robust to certain worst-case deviations from the attacker, but modify the traditional worst-case assumption to a new type of graduated optimization. Furthermore, we extend both MATCH and BRQR to address human anchoring biases as it has been shown that this extension is advantageous under limited observation [6]. In order to evaluate our new approach and these extensions we performed a comprehensive experimental study involving 253 human subjects playing 5956 games under three observation conditions (perfect, limited, and no observation). Since we alter the standard assumptions of robust optimization we also include an alternative algorithm, RECON [8], which employs the traditional worst-case robust optimization. In addition, we examine the influence of two personality measures, psychopathy and numeracy, and demographic information, age and gender, on decision-making in security settings. Psychopathy is especially of interest because research has shown that psychopathy is a strong predictor of both criminal behavior and in particular violent crimes [3]. Gaining insight into the influence of such personality measures and demographic information could potentially motivate future algorithmic developments.

Appears in: *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2012)*, Conitzer, Winikoff, Padgham, and van der Hoek (eds.), 4-8 June 2012, Valencia, Spain.

2. METHODOLOGY

Methods for computing MATCH: MATCH is a mixed integer linear program (MILP) that utilizes a new idea of graduated robust optimization. Whereas standard robust optimization robustly guards against a worst-case outcome within some error bound, MATCH assumes a utility maximizing outcome on behalf of the attacker, but constrains the impact of deviations depending on the magnitude of the deviation. That is, MATCH has an adjustable parameter, β, which constrains the defender's loss for a deviation by the attacker to be no worse than a proportion (β) of the loss the attacker incurs for that deviation. For example, if the attacker deviates from the expected-value-maximizing target and loses 2 utility, then the defender should not lose more than $\beta * 2$ for this deviation.

Extending BRQR: In order to extend BRQR to handle an anchoring bias we need to alter the way the adversary perceives his reward. Specifically, if the defender has chosen a strategy x for defending her targets, the attacker will now base his decisions on a strategy x' that accounts for his anchoring biases. Thus, in the new model for BRQR, the adversary will respond stochastically according to x' where the chance of selecting non-optimal strategies increases as the *perceived* cost of such an error decreases. We refer to this new strategy as BRQRA.

Extending MATCH: MATCH originally assumes a perfectly-rational adversary so we chose to extend MATCH to address both an anchoring bias and a boundedly-rational attacker as in COBRA. We refer to this new formulation as COBRA-MATCH. Since MATCH and COBRA are both MILPs we are able to extend MATCH utilizing the same types of constraints originally presented in COBRA [6]. Specifically, as in BRQR, the attacker now makes his decision based on x' rather than x. Furthermore, given his perception of the defender strategy (i.e., x') he is willing to choose any strategy within ϵ of what he perceives to be the expected-utility-maximizing strategy. One important consideration in the COBRA-MATCH formulation is that now we must model the attacker's losses for a deviation according to his perception of his loss (i.e., according to x'), while the defender's loss is still based on the real defender strategy (i.e., according to x). It follows that the defender should only lose a proportion (β) of what the attacker *perceives* he has lost.

3. EVALUATION

We conducted empirical tests with human subjects playing a web-based game to evaluate the performance of defender strategies generated using six candidate algorithms: DOBSS, MAXIMIN, COBRA, BRQR/BRQRA, MATCH/COBRA-MATCH, and RECON. In our experiments, we utilize the same eight-target scenario used by Yang et al. [7].Before beginning, subjects were given a tutorial and a test to ensure that they understood the general game play.

Our experiments were run in Amazon Mechanical Turk and participants were paid a base amount of US $1.50 for participating. In order to motivate the subjects, they were informed that a small sample of their games would be chosen at random and they would be paid an additional US $0.15 for the total points earned in that sample. Also, two obvious games were introduced to ensure subjects were paying attention.If subjects failed to respond correctly in the obvious games then their data was removed from the set.

We tested nine different payoff structures (five new, four from yang et al. [7]) in the unlimited observation condition and four in the limited and unobserved conditions (from yang et al. [7]). For each payoff structure, we generated the mixed strategies for the defender using the six algorithms with a variety of parameter settings. We ran experiments for the unlimited observation condition separately from experiments in the limited and unobserved observation conditions. This was to avoid confusion in the subjects and to keep the experimental conditions controlled. Additionally, the order of game instances played by each subject was randomized to mitigate ordering effects on their response. We also examined runtime performance for MATCH versus BRQR.

4. CONCLUSIONS

To address human adversaries, a number of approaches, including COBRA and BRQR, have been introduced which attempt to include more realistic models of human-decision making. Our work provides five fundamental contributions to this line of research: (i) we develop an approach to addressing human adversaries based on robust optimization rather than relying on finding more appropriate models of human decision-making; (ii) we extend both BRQR and MATCH to address human anchoring biases under limited observation; (iii) we do a comprehensive experimental analysis of the performance of MATCH against previous approaches and runtime analysis showing the efficiency of MATCH; and (iv) we make the first ever evaluation of the influence of some demographic and personality measures on decision-making in security games.

5. ACKNOWLEDGEMENT

This research was supported by the United States Department of Homeland Security through the National Center for Risk and Economic Analysis of Terrorism Events (CREATE) under award number 2010-ST-061-RE0001. However, any opinions, findings, and conclusions or recommendations in this document are those of the authors and do not necessarily reflect views of the United States Department of Homeland Security, or the University of Southern California, or CREATE.

6. REFERENCES

[1] M. Aghassi and D. Bertsimas. Robust game theory. *Math. Program.*, 107(1-2):231–273, 2006.

[2] C. Camerer. In *Behavioral Game Theory: Experiments in Strategic Interaction*. Princeton University Press, 2003.

[3] R. D. Hare and C. S. Neumann. Psychopathy as a clinical and empirical construct. *Annual Review of Clinical Psychology*, 4:217–246, 2008.

[4] M. Jain, J. Tsai, J. Pita, C. Kiekintveld, S. Rathi, F. Ordóñez, and M. Tambe. Software assistants for randomized patrol planning for the LAX airport police and the Federal Air Marshals Service. *Interfaces*, 40(4):267–290, 2010.

[5] P. Paruchuri, J. Marecki, J. Pearce, M. Tambe, F. Ordóñez, and S. Kraus. Playing games for security: An efficient exact algorithm for solving bayesian stackelberg games. In *AAMAS*, 2008.

[6] 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 Journal*, 174(15):1142–1171, 2010.

[7] R. Yang, C. Kiekintveld, F. Ordóñez, M. Tambe, and R. John. Improving resource allocation strategy against human adversaries in security games. In *IJCAI*, 2011.

[8] Z. Yin, M. Jain, M. Tambe, and F. Ordóñez. Risk-averse strategies for security games with execution and observational uncertainty. In *AAAI*, 2011.

[9] Z. Yin, D. Korzhyk, C. Kiekintveld, V. Conitzer, and M. Tambe. Stackelberg vs. nash in security games: Interchangeability, equivalence, and uniqueness. In *AAMAS*, 2010.