Designing Better Resource Allocation Strategy against Human Adversaries in Security Games

(Extended Abstract)

Rong Yang
University of Southern California
Los Angeles, CA, US 90089
yangrong@usc.edu

ABSTRACT

Stackelberg games have garnered significant attention in recent years given their deployment for real world security, such as ARMOR, IRIS and GUARDS. Most of these systems have adopted the standard game-theoretical assumption that adversaries are perfectly rational, which may not hold in real-world security problems due to the bounded rationality of human adversaries and could potentially reduce the effectiveness of these systems.

My thesis focuses on relaxing the assumption of perfectly rational adversaries in Stackelberg security games. In particular, I aim at developing new adversary models incorporating their bounded rationality and building new algorithms for efficiently computing a defender's best response against these new models. To that end, I have developed a new adversary model using quantal response (QR) and a new efficient algorithm (Pasaq) to compute a defender's strategy against such a model in massive real-world security games. Experimental results with human subjects show that this new model gives significantly better defender strategies than the previous leading contender. Furthermore, Pasaq has been deployed in a real-world security application, PROTECT, by the U.S. Coast Guards at the port of Boston. Recently, I started extending the model to incorporate features from more complicated games, including Network Security games and Bayesian Stackelberg games.

Categories and Subject Descriptors

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

General Terms

Algorithms Experimentation Security

Keywords

Bounded Rationality, Stackelberg Games, Decision-making

1. INTRODUCTION

Stackelberg game have recently become important tools for analyzing real-world security resource allocation problems, such as critical infrastructure protection and robot patrolling strategies. In a Stackelberg Security game (SSG), one player (the defender) plays as the leader and commits to a mixed-strategy first; the other player (the attacker) plays as the follower and response to the defender's strategy. SSGs provide a sophisticated approach for generating unpredictable, randomized strategies that mitigate the ability of attackers to find weaknesses using surveillance. Some notable examples of real-world deployments of SSGs include ARMOR, IRIS, GUARDS and PROTECT [3].

One of the key sets of assumptions that existing real-world security systems make is about how attackers choose strategies based on their knowledge of the security strategy. Typically, such systems apply the standard game-theoretic assumption that attackers are perfectly rational and strictly maximize their expected utility. This is a reasonable starting point for the first generation of deployed systems. Unfortunately, this assumption may lead to a defense strategy that is not robust against attackers using different decision procedures, and it fails to exploit known weaknesses in the decision-making of human attackers. Indeed, it is widely accepted that standard game-theoretic assumptions of perfect rationality are not ideal for predicting the behavior of humans in multi-agent decision problems. Thus, it is critical that we integrate more realistic models of human decision-making in solving real-world security problems.

There are several open questions we need to address in moving beyond perfect rationality assumptions. First, a large variety of alternative models have been studied in behavioral game theory and cognitive psychology that capture some of the deviations of human decisions from perfect rationality. However, there is an important empirical question of which model best represents the salient features of human behavior in applied security contexts. Second, integrating any of the proposed models into a decision-support system requires developing new computational methods, since the existing algorithms for security games are based on perfectly rational attackers. Third, many of these models imply mathematically complex representations of the adversary's decision-making procedure (e.g. nonlinear and non-convex function forms), which in general leads to a NP-hard problem of computing the defender's optimal strategy. Therefore, developing efficient algorithms to solve such a computationally complex problem is critical for real-world security problems due to their massiveness. My work provides methods to address these open questions.

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. CONTRIBUTIONS

New models of human adversary: My first step towards addressing these open questions is to develop new mathematical models of the adversary's decision-making by using two fundamental theories of human behavior to predict an attacker's decisions: Prospect Theory (PT)and Quantal Response Equilibrium (QRE) [7]. I then develop new algorithms to compute the defender optimal strategies against these new models. In particular, I develop a Mixed Integer Linear Program (Brpt) to compute the defender optimal strategy against the PT based models of an adversary by representing the non-linear functions from Prospect Theory with piecewise approximations. Furthermore, I present a local search method (Brqr) [7] with random restarts to compute the defender optimal strategy against the QRE based models of the adversary. I then evaluate different defender strategies through experiment with human subjects using an online game designed to simulate a security scenario, which is similar to the one analyzed by ARMOR for the LAX airport. The results show that defender strategies computed using the new method Brqr significantly outperform its competitors, including Brpt and Cobra which was presented as the leading contender by previous work that accounts for human behavior in security games.

New algorithms for real-world security: Quantal Response (QR) is a very important solution concept in addressing human bounded rationality in game-theoretic settings. It assumes errors in human decision making and suggests that instead of strictly maximizing utility, individuals respond stochastically in games: the chance of selecting a non-optimal strategy increases as the associated cost decreases. In order to apply QR model of adversary to solve real-world security problems, two major challenges need to be addressed in computing defender optimal strategy: (1) solving a nonlinear non-convex optimization problem efficiently for massive real-world security games; and (2) addressing resource assignment constraints in security games, which adds to the complexity of computing the optimal defender strategy. To that end, I provide two new algorithms, Gosaq and Pasaq [8]. Gosaq guarantees the global optimal solution in computing the defender strategy against a QR model of the adversary and Pasaq provides an efficient approximated computation of the defender strategy with a nearly-optimal solution quality. I provide both the correctness/approximation-bound proof of the algorithms and experimental analysis on their solution quality and computational efficiency. The results shown that both Gosaq and Pasaq achieve better solution quality than Brqr and that Pasaq achieves much better computational efficiency than both Gosaq and Brqr. Given these results, Pasaq is at the heart of the PROTECT system [1] which is currently being used for the US Coast Guard in the port of Boston and deployed in the port of New York.

Network Security Game: Many real-world security domains have structures that are naturally modeled as graphs. We term such domains *network security games* [6], which are Stackelberg games on graphs with intersections as nodes and roads as edges, where certain nodes are targets for attacks. In a network security game, security agencies must allocate limited resources to protect targets embedded in a network, such as important buildings in a city road network. Previous work [4] has formulated this as an attacker-defender Stackelberg game with a perfectly rational attacker. Given that real-world network security games are often extremely complex and hence very difficult for humans to solve, it becomes even more important to incorporate human behavior modeling when designing defender strategies. In my latest work [6], I focused on addressing human bounded rationality in network security game on small graphs. More specifically, I have applied the QR models to network security games, including models that take into account possible heuristics (QRH) used by humans. I learned the model parameters based on human subject data collected from experiments using a web-based game designed to simulate network security games. I evaluated the learned model in a new set of games with the same experimental platform. The results showed that defender strategies using our models learned from human behavior data achieved significantly better expected utilities compared to the rational model.

3. FUTURE WORK

In the future, I plan on extending the current algorithms for computing defender strategies to handle more complex real-world security problems. First, Pasaq currently assumes: 1) a single type adversary; 2) a set of defender pure strategies that can be completely enumerated. Both assumptions need to be relaxed for more complex real-world security problems. Additionally, the current algorithm for computing defender strategies in network security games relies on local search method and does not scale-up efficiently. New efficient algorithms need to be built to address large graphs. Finally, we have observed from the experimental data that humans dynamically adapt their strategies to different game instances. I plan to develop a hybrid model of adversary decision-making to capture such feature and design better defender strategies based on this new model.

4. REFERENCES

[1] 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*, 2012.

[2] D. O. Stahl and P. W. Wilson. Experimental evidence on players' models of other players. *JEBO*, 25(3):309–327, 1994.

[3] M. Tambe. *Security and Game Theory: Algorithms, Deployed Systems, Lessons Learned*. Cambridge University Press, New York, NY, 2011.

[4] J. Tsai, Z. Yin, J. Kwak, D. Kempe, C. Kiekintveld, and M. Tambe. Urban security: Game-theoretic resource allocation in networked physical domains. *In AAAI*, 2010.

[5] J. R. Wright and K. Leyton-Brown. Beyond equilibrium: Predicting human behavior in normal-form games. *In AAAI*, 2010.

[6] R. Yang, A. X. Jiang, F. Fang, M. Tambe, R. Maheswaran, and K. Rajagopal. Designing better strategies against human adversaries in network security games. *In AAMAS*, 2012.

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

[8] R. Yang, F. Ordonez, and M. Tambe. Computing optimal strategy against quantal response in security games. *In AAMAS*, 2012.