A Unifying Methodology for Confronting Uncertainties in Security Games: Advances and Algorithms

(Doctoral Consortium)

Thanh H. Nguyen
Supervisor: Prof. Milind Tambe
University of Southern California
941 Bloom Walk, SAL 300
Los Angeles, California, 90089
thanhhng@usc.edu

ABSTRACT

Given the real-world applications of Stackelberg security games (SSGs), addressing uncertainties in these games is a major challenge. Two competitive approaches have been pursued by previous work on addressing uncertainties in SSGs, namely: (1) applying robust optimization techniques without requiring a prior distribution; and (2) using probabilistic models to capture uncertainties. In general, the decision of which approach to use is based on the availability of data. While the first approach suits data-sparse domains, the second approach works better for data-rich domains. My thesis will focus on addressing uncertainties in SSGs following these two leading approaches. In particular, with regards to robust methods, I attempt to develop new maximin/minimax regret-based robust algorithms for computing a defender's optimal strategy given uncertainties. I also aim to contribute to probabilistic modeling techniques by developing a new computational model of human decision making to capture the adversary's bounded rationality.

Categories and Subject Descriptors

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

General Terms

Security, Algorithms

Keywords

Computational Game Theory; Security Games; Behavioral Model; Robust Optimization; Machine Learning; Uncertainty

1. INTRODUCTION

Real-world deployed applications of defender-attacker Stackelberg Security Games [8, 1] have led to significant research emphasis on handling uncertainties. Two different approaches have been pursued. The first approach, focused on domains with sparsity of data, uses robust optimization with the maximin method of maximizing defender expected utility under the worst case resulting from such uncertainty [2]. The second approach, focused on data-rich domains, on the other hand, applies probabilistic-modeling methods for capturing uncertainties in SSGs with two key ideas. The first idea assumes a known distribution of payoff uncertainty and solves the resulting Bayesian Stackelberg game models [10]. The another idea predicts the attacker's decision making using behavioral models and computes the defender's optimal strategy assuming the attacker's response follows that model [9].

There are several open issues remain in addressing uncertainties with respect to these leading approaches. First, existing maximin-based robust algorithms compartmentalize uncertainties; the lack of a unified framework implies that existing algorithms suffer losses in solution quality in handling uncertainties in real-world security situations — where multiple types of uncertainties may exist simultaneously. Second, with regard to robust techniques, previous works in security games only focus on the maximin method. We lack an alternative less-conservative robust criterion for addressing uncertainties in SSGs. Finally, in terms of modeling human adversaries's behaviors, the model-free algorithm MATCH [7] has been shown to outperform BRQR, the algorithm based on the Quantal Response (QR) model [9], leading to important open question of whether there is any value in adversary modeling in SSGs.

I present the first unified maximin-based framework for handling the different types of uncertainties explored in SSGs. Based on that framework, I propose a set of new "unified" robust algorithms to address combinations of these uncertainties [4]. I then introduce approximate scalable robust algorithms for handling these uncertainties that leverage insights across compartments. Furthermore, I propose the use of the less conservative minimax regret decision criterion for handling uncertainties in SSGs [5]. More specifically, I present new novel algorithms for computing minimax regret for addressing payoff uncertainty. I also address the challenge of payoff elicitation, using minimax regret to develop the first elicitation strategies for reducing payoff uncertainty. Finally, I explore a new model of human adversary's behavior, SUQR, a novel integration of human behavior model with the subjective utility function [6]. I show that my algorithm, SU-BRQR, based on SUQR, significantly outperforms the model-free algorithm, MATCH.

2. UNIFIED MAXIMIN-BASED METHOD

My first contribution is to remedy weaknesses of state-of-the-art maximin-based algorithms when addressing uncertainties in SSGs, due to uncertainty compartmentalization. I present a unified computational framework – a single core problem representation for handling the different types of uncertainties and their combinations,

Appears in: *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2015), Bordini, Elkind, Weiss, Yolum (eds.), May 4–8, 2015, Istanbul, Turkey.*

including uncertainty in the attacker's payoffs, uncertainty related to the defender's strategy (due to the defender's execution and the attacker's observation), and uncertainty in the attacker's bounded rationality. The first key component of my unified framework is a unified formulation of uncertainty sets for SSGs that captures all major existing approaches. Furthermore, based on this unified framework, I present a unified algorithmic framework from which I can derive different "unified" robust algorithms which address any combination of uncertainties presented in the framework. Also, exploiting new insights from our unified framework, I present fast approximate algorithms for handling different subsets of uncertainties in large-scale security games. Finally, my experiments show the solution quality and runtime advantages of our algorithms. The key insights in my unified robust algorithms include: 1) under any combinations of uncertainties, the corresponding robust optimization problem can be represented as a single maximin problem in which all these uncertainties are encapsulated into a set of adversary's strategies; and 2) the space of the defender's strategies can be partitioned into different sets; for any defender strategy within a set, the adversary's feasible strategies are identical. Thus, I can solve any robust optimization problem for addressing uncertainties – modeled within my framework – as the maxima of all corresponding simpler sub-optimization problems created by this partition.

3. MINIMAX REGRET-BASED METHOD

My second contribution is introducing the use of the leading minimax regret decision criterion for addressing uncertainties in SSGs. I start with developing minimax regret-based methods for *Strict Uncertainty Payoff games (SPAC)*. SPAC refers to where defender payoffs are known but attacker payoffs are assumed only to lie within some interval with no distributional information. Minimax regret focuses on the loss with respect to decision quality over possible payoff realizations, making decisions with the tightest possible optimality guarantees. Indeed, minimax regret has not yet been available to policy makers – my work makes it a viable criterion for generating new, less conservative, candidate defensive strategies. Unfortunately, operationalizing minimax regret involves complex, non-convex optimization for which efficient algorithms do not exist. Thus, I develop novel, efficient algorithms for approximating minimax regret, and provided experimental results demonstrating high solution quality along with fast runtime. Furthermore, I develop a payoff elicitation procedure which is based on minimax regret solutions that can be used to optimize the defender's efforts in assessing payoffs, allowing reduction in the uncertainty of those parameters that most improve decision quality. This is another reason to use minimax regret as my robustness criterion – it has been proven to be an effective driver of elicitation in several domains.

4. ATTACKER BEHAVIORAL MODELING

My third contribution on behavioral modeling builds on the significant support for QR: I hypothesize that QR's stochastic response is crucial in building a human decision making model. Where I part company with the original QR model is in its assumption that human stochastic response is based on expected value. Instead, I propose a new model based on integration of a novel subjective utility function (SU) into QR, called the SUQR model. The novelty of my subjective utility function is the linear combination of the game features such as the defender's coverages and the attacker's payoffs. I show that the SUQR model has superior predictive power compared to the QR model. Then, the defender's optimal strategy is computed by maximizing the defender's expected value given that the adversary chooses to attack each target with a probability according to SUQR. The resulting SUQR-based SU-BRQR algorithm, similar to BRQR, to compute the defender strategy is evaluated through two sets of experiments using an online game with Amazon Mechanical Turk (AMT) workers. The experimental results show that SU-BRQR significantly outperforms the model-free MATCH algorithm as well as improved versions of MATCH. Furthermore, by conducting experiments with new game scenarios, where no human behavior data exists, I show SU-BRQR with its earlier learned parameters still outperforms MATCH; and learning from more data, SU-BRQR performance can be further improved.

5. FUTURE WORK

I focus on solving *green* security games wherein security resources are allocated to patrol a vast geographical area against environmental criminals such as poachers in wildlife security domains. The main objective of park rangers is to prevent poaching by conducting patrols throughout the parks. While game-theoretic approaches have been advocated to generate rangers' patrols, there is significant uncertainty in domain features such as animal density and vegetation, etc. My first goal is to apply my previous research to solve this real-world security problem, in particular, modeling poachers' behaviors and addressing domain uncertainty using robust algorithms. I also aim to explore new challenges in the domain that my previous work cannot handle. Currently, I'm working on a paper which is under submission based on wildlife protection [3].

6. ACKNOWLEDGEMENT

This research was supported by MURI Grant W911NF-11-1-0332 and by CREATE under grant number 2010-ST-061-RE0001.

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] C. Kiekintveld, T. Islam, and V. Kreinovich. Security games with interval uncertainty. In *AAMAS*, 2013.

[3] T. H. Nguyen, F. M. Fave, M. Jain, N. Agmon, R. V. Deventer, and M. Tambe. Behavioral minimax regret for security games and its application for uav planning. In *In submission*, 2015.

[4] T. H. Nguyen, A. Jiang, and M. Tambe. Stop the compartmentalization: Unified robust algorithms for handling uncertainties in security games. In *AAMAS*, 2014.

[5] T. H. Nguyen, A. Yadav, B. An, M. Tambe, and C. Boutilier. Regret-based optimization and preference elicitation for stackelberg security games with uncertainty. In *AAAI*, 2014.

[6] T. H. Nguyen, R. Yang, A. Azaria, S. Kraus, and M. Tambe. Analyzing the effectiveness of adversary modeling in security games. In *AAAI*, 2013.

[7] J. Pita, R. John, R. Maheswaran, M. Tambe, and S. Kraus. A robust approach to addressing human adversaries in security games. In *ECAI*, pages 660–665, 2012.

[8] 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.

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

[10] Z. Yin and M. Tambe. A unified method for handling discrete and continuous uncertainty in bayesian stackelberg games. In *AAMAS*, 2012.