Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence (AAAI-16)

Conquering Adversary Behavioral Uncertainty in Security Games: An Efficient Modeling Robust Based Algorithm

Thanh H. Nguyen, Arunesh Sinha, Milind Tambe
University of Southern California
SAL 300, 941 Bloom Walk, Los Angeles, California, 90089
{thanhhng, aruneshs, tambe}@usc.edu

Introduction

Real-world deployed applications of Stackelberg Security Games (Shieh et al. 2012; Basilico, Gatti, and Amigoni 2009; Letchford and Vorobeychik 2011) have led to significant research emphasis on modeling the attacker's bounded rationality (Yang et al. 2011; Nguyen et al. 2013). One key assumption in behavioral modeling is the availability of a significant amount of data to obtain an accurate prediction. However, in real-world security domains such as the wildlife protection, this assumption may be inapplicable due to the limited access to real-world data (Lemieux 2014), leading to uncertainty in the attacker's behaviors — a key research challenge of security problems.

Recent research has focused on addressing uncertainty in behavioral modeling, following two different approaches: 1) one approach assumes a known distribution of multiple attacker types, each follows a certain behavioral model, and attempts to solve the resulting Bayesian games (Yang et al. 2014); and 2) another considers the existence of multiple attacker types of which behavioral models are perfectly known, but without a known distribution over the types. It then only considers the worst attacker type for the defender (Brown, Haskell, and Tambe 2014). These two approaches have several limitations. First, both still require a sufficient amount of data to precisely estimate either the distribution over attacker types (the former approach) or the model parameters for each individual type (the latter approach). Second, solving the resulting Bayesian games in the former case is computationally expensive. Third, the latter approach tends to be overly conservative as it only focuses on the worst-case attacker type.

This paper remedies these shortcomings of state-of-the-art approaches when addressing behavioral uncertainty in SSG by providing three key contributions. First, we present a new game model with uncertainty in which we consider a single behavioral model to capture decision making of the whole attacker population (instead of multiple behavioral models); uncertainty intervals are integrated with the chosen model to capture behavioral uncertainty. The idea of uncertainty interval is commonly used in literature (Aghassi and Bertsimas 2006) and has been shown to effectively represent uncertainty in SSG (Kiekintveld, Islam, and Kreinovich 2013). Second, based on this game model, we propose a new efficient robust algorithm that computes the defender's optimal strategy which is robust to the uncertainty.

Overall, the resulting robust optimization problem for computing the defender's optimal strategy against the worst case of behavioral uncertainty is a non-linear non-convex fractional maximin problem. Our algorithm efficiently solves this problem based on the following key insights: 1) it converts the problem into a single maximization problem via a non-linear conversion for fractional terms and the dual of the inner minimization in maximin; 2) a binary search is then applied to remove the fractional terms; and 3) the algorithm explores extreme points of the feasible solution region and uses a piece-wise linear approximation to convert the problem into a Mixed Integer Linear Program (MILP). Our new algorithm provides an $O(\epsilon+\frac{1}{K})$-optimal solution where ϵ is the convergence threshold for the binary search and K is the number of segments in the piecewise approximation.

Background

Stackelberg security games (SSG). SSG refer to a class of defender-attacker games in which the defender attempts to optimally allocate her limited security resources to protect a set of T targets from being attacked by the attacker. The key assumption of SSG is that the defender has to commit to a (*mixed*) strategy first and the attacker can observe that strategy and then attack one of the targets (Korzhyk, Conitzer, and Parr 2010; Tambe 2011). Suppose that the defender has R resources ($R \ll T$) and $\mathbf{x} = \{x_i\}$ is a defender's mixed strategy where x_i is the defender's coverage probability at target i, the defender's feasible strategy set is defined as follows: $\mathbf{X} = \{\mathbf{x} : 0 \leq x_i \leq 1, \sum_i x_i = R\}$. Suppose that the attacker attacks target i, he will obtain a reward R_i^a if the defender is not protecting that target, otherwise he will get a penalty P_i^a. Conversely, the defender receives a penalty P_i^d and a reward R_i^d respectively. The expected utility for the two players at target i can be computed as follows:

$$U_i^d(x_i) = x_i R_i^d + (1 - x_i) P_i^d \quad (1)$$

$$U_i^a(x_i) = x_i P_i^a + (1 - x_i) R_i^a \quad (2)$$

Adversarial behavioral models. Recent research in SSG has focused on modeling the attacker's bounded rationality and computing the defender's optimal strategy, assuming the attacker's response follows the given behavioral model. One leading class of behavioral models is Quantal Response (QR) (McFadden 1972; McKelvey and Palfrey 1995). In modeling the attacker's decision making, we consider a general discrete choice model of QR to capture behaviors of the attacker (Train 2009) in which the probability that the attacker chooses target i, $q_i(\mathbf{x})$, is predicted as follows:

$$q_i(x_i) = \frac{F_i(x_i)}{\sum_j F_j(x_j)} \tag{3}$$

where $F_i(x_i) : [0,1] \to \mathbb{R}^*$ is a positive and monotonically decreasing function of x_i at target i.

Uncertainty in Behavioral Modeling

Due to the behavioral uncertainty, we assume that the value of $F_i(x_i)$ in (3) is not perfectly known given x_i. Instead, $F_i(x_i)$ has known lower and upper bounds, $L_i(x_i) \leq F_i(x_i) \leq U_i(x_i)$ where $L_i(x_i), U_i(x_i) : [0,1] \to \mathbb{R}^*$ are positive functions of the defender coverage at target i. Denote by $\mathbf{I}(x_i) = [L_i(x_i), U_i(x_i)]$ the uncertainty interval, the interval size indicates the uncertainty level when modeling, which could be specified based on the available data for learning. We aim at computing the defender's optimal strategy by maximizing her utility under the worst case resulting from the behavioral uncertainty. The corresponding robust optimization problem is represented as follows:

$$\max_{\mathbf{x}\in\mathbf{X}} \min_{F_i(x_i)\in\mathbf{I}(x_i),\forall i} \sum\nolimits_i q_i(x_i) U_i^d(x_i) \tag{4}$$

Overall, the problem (4) is non-convex which is difficult to solve. We present our novel algorithm which efficiently solves the maximin problem (4) with a bound guarantee on its approximate solution. In short, there are three key ideas. First, we convert (4) into a single maximization problem via a non-linear conversion for fractional terms and dualty of the inner minimization in (4). Given a defender strategy $\mathbf{x}$, the objective of (4) remains a non-linear fractional function of $F_i(x_i)$, thus making the inner minimization problem in (4) non-linear and fractional. We introduce the following new variables: $y_i = q_i(x_i) = \frac{F_i(x_i)}{\sum_j F_j(x_j)}$ which is the attacking probability at target i and $z = \frac{1}{\sum_j F_j(x_j)}$ which is the normalization term in the attacking probabilities. By replacing $F_i(x_i)$ with the new variables and denote by $\mathbf{y} = \{y_i\}$, we can represent the inner minimization of (4) as the following linear minimization problem of the new variables $\mathbf{y}$ and z:

$$\min\nolimits_{\mathbf{y},z} \sum\nolimits_i y_i U_i^d(x_i) \tag{5}$$

$$\text{s.t.} \sum\nolimits_i y_i = 1 \tag{6}$$

$$L_i(x_i) z \leq y_i \leq U_i(x_i) z, \forall i. \tag{7}$$

where constraint (6) ensures the condition on the attacking probability distribution that $\sum_i q_i(x_i) = 1$ holds. In addition, constraint (7) is equivalent to the condition on the lower and upper bound of $F_i(x_i)$ that $F_i(x_i) \in [L_i(x_i), U_i(x_i)]$. As (5 – 7) is a linear minimization problem of $\mathbf{y}$ and z, its optimal solution is equivalent to the optimal solution of its duality which is a linear maximization problem. Therefore, we can merge this dual with the outer maximization of (4) to obtain a single maximization problem.

Given the resulting new maximization problem, we then apply a binary search to remove fractional terms. We explore extreme points of the feasible solution region and use piece-wise linear approximation to convert the resulting feasibility problem at each binary step into a MILP. Our new algorithm provides an $O\left(\epsilon + \frac{1}{K}\right)$-optimal solution of the maximin problem (4).

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

References

Aghassi, M., and Bertsimas, D. 2006. Robust game theory. *Mathematical Programming* 107(1-2):231–273.

Basilico, N.; Gatti, N.; and Amigoni, F. 2009. Leader-follower strategies for robotic patrolling in environments with arbitrary topologies. In *AAMAS*, 57–64.

Brown, M.; Haskell, W. B.; and Tambe, M. 2014. Addressing scalability and robustness in security games with multiple boundedly rational adversaries. In *GameSec*.

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

Korzhyk, D.; Conitzer, V.; and Parr, R. 2010. Complexity of computing optimal stackelberg strategies in security resource allocation games. In *AAAI*.

Lemieux, A. M. 2014. *Situational Prevention of Poaching*. Routledge.

Letchford, J., and Vorobeychik, Y. 2011. Computing randomized security strategies in networked domains. In *Applied Adversarial Reasoning and Risk Modeling*.

McFadden, D. 1972. Conditional logit analysis of qualitative choice behavior. Technical report.

McKelvey, R., and Palfrey, T. 1995. Quantal response equilibria for normal form games. *Games and economic behavior* 10(1):6–38.

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

Shieh, E.; An, B.; Yang, R.; Tambe, M.; Baldwin, C.; DiRenzo, J.; Maule, B.; and Meyer, G. 2012. Protect: A deployed game theoretic system to protect the ports of the united states. In *AAMAS*.

Tambe, M. 2011. *Security and Game Theory: Algorithms, Deployed Systems, Lessons Learned*. Cambridge University Press.

Train, K. E. 2009. *Discrete choice methods with simulation*. Cambridge university press.

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

Yang, R.; Ford, B.; Tambe, M.; and Lemieux, A. 2014. Adaptive resource allocation for wildlife protection against illegal poachers. In *AAMAS*.