Addressing Uncertainty in Stackelberg Games for Security: Models and Algorithms (Extended Abstract)

Zhengyu Yin
University of Southern California, Los Angeles, CA 90089
zhengyuy@usc.edu

ABSTRACT

There has been significant recent research interest in utilizing leader-follower Stackelberg game in security applications. Indeed, Stackelberg games are seen at many deployed applications [6]: ARMOR at Los Angeles International Airport, IRIS for Federal Air Marshals Service, GUARDS for the Transportation Security Administration, and TRUSTS for the Los Angeles Metro Rail System [3] (under evaluation). The foundational assumption for using Stackelberg games is that security forces (leaders), acting first, commit to a randomized strategy; while their adversaries (followers) choose their best response after surveillance of this randomized strategy.

Due to the adversarial environment and the nature of law enforcement activities, many types of uncertainty, such as execution, observation, and preference uncertainty, must be taken into account in game-theoretic modeling for practical security applications. To that end, focusing on *security games* I explicitly model the aforementioned uncertainty and present theoretical analysis and novel algorithms for computing robust solutions. Furthermore, as the cornerstone in providing real world evaluations of my robust solution techniques, I propose TRUSTS, a compact game-theoretic formulation, for fare evasion deterrence in the Los Angeles Metro Rail system. In my future research, I will extend TRUSTS to address real world uncertainty and evaluate the solutions within the LA Metro system.

Categories and Subject Descriptors

I.2.11 [**Artificial Intelligence**]: Distributed Artificial Intelligence

General Terms

Algorithms, Optimization

Keywords

Game Theory, Bayesian Stackelberg Games, Security

1. INTRODUCTION

In many of the security domains, e.g., airports, ports, transit systems, and other infrastructure, the security forces deploy security resources to provide protection and law enforcement against potential adversaries. With limited resources available, it is often impossible to cover all locations at all times. The use of game-theoretic concepts has allowed security forces to exert maximum leverage with limited resources. Indeed, game-theoretic solutions have been seen in many deployed applications: ARMOR at Los Angeles International Airport (LAX), IRIS for Federal Air Marshals Service, and GUARDS for the Transportation Security Administration [6], and TRUSTS for the Los Angeles Metro Rail System [3] (under evaluation). At the backbone of these applications is the leader-follower Stackelberg game model, where the leader (i.e., the security forces) acts first by committing to a (mixed) patrolling strategy, and the follower (i.e., the adversary) best-responds after observing the leader's choice perfectly.

Due to the adversarial environment and the nature of law enforcement activities, many types of uncertainty must be taken into account in game-theoretic modeling. There are currently three major challenges in handling uncertainty in Stackelberg games. First, there can be multiple follower types with entirely different preferences, often captured by the Bayesian extension of Stackelberg games. For example, the LA Metro system have tens of thousands of potential fare evaders daily, each may have a distinct intended trip and risk profile. Second, the follower may have imperfect observation and the leader may have imperfect execution. For example, an attacker at LAX may mistake a passing car as a security patrol; the police patrolling within the LA Metro system may fall behind a patrol schedule due to fare citations and criminal arrests. Third, there can be varying amounts of information available about the uncertainty (e.g., known vs. unknown distribution), depending on which different uncertainty models need to be considered.

My thesis addresses these three major challenges in handling uncertainty. Unlike previous work on Bayesian Stackelberg games [1, 4, 2] which assume perfect execution and observation, I explicitly model execution and observation uncertainty, presenting theoretical analysis and novel algorithms for computing robust solutions. In particular, I address three uncertainty models corresponding to different assumptions of the information available about the uncertainty. Furthermore, as the cornerstone of my future research on applying my robust solution techniques, I propose TRUSTS, a compact game-theoretic formulation for handling both a huge number of follower types and massive temporal and spatial constraints. In my future research, I will extend TRUSTS to address real world uncertainty and evaluate the solutions within the LA Metro system.

2. CONTRIBUTIONS

My thesis provides the following key contributions.

Addressing mixed uncertainty with known distribution: In the first contribution of my thesis, I consider scenarios in which the distribution of the uncertainty is known, providing the first unified approach to handling both continuous uncertainty (such as execution and observation error) and discrete uncertainty (such as fol-

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.

lower types) [9]. I show that via *sample average approximation*, continuous uncertainty such as payoff uncertainty, execution and observation error can be handled by solving a finite Bayesian Stackelberg game. I present a new algorithm for Bayesian Stackelberg games, called HUNTER, to scale up the number of types. HUNTER combines the following key ideas: i) efficient pruning via a best-first search; ii) a novel linear program for computing tight upper bounds for this search; iii) using Bender's decomposition for solving the upper bound linear program efficiently; iv) efficient inheritance of Bender's cuts from parent to child; v) an efficient heuristic branching rule. My experimental results suggest that HUNTER provides orders of magnitude speedups over the best existing methods [1, 4, 2]. Furthermore, under continuously distributed uncertainty, the HUNTER-based approach outperforms other existing robust solution methods.

Addressing bounded, distribution-free uncertainty: Often the precise distribution of continuous uncertainty such as execution and observation error is difficult to estimate. As my second contribution, I model distribution-free execution and observation uncertainty as a bounded hyper-rectangular, providing a robust optimization formulation, called RECON [7], to find risk-averse strategies for the leader. RECON assumes that nature chooses the noise realization to maximally reduce the leader's utility, and maximizes against this worst case. I provide a mixed-integer linear program (MILP) for solving the RECON formulation, and two novel heuristics that speed up the RECON MILP by orders of magnitude. My experimental result demonstrates the superiority of RECON under uncertainty where previous solutions [4, 5] perform poorly.

Addressing follower's observability uncertainty: In my previous uncertainty models, the follower is assumed to observe the leader's strategy (possibly imperfectly); yet, in many situations, the followers may act without any observation, essentially converting the game into a simultaneous-move game model. As the third contribution, my thesis addresses how a leader should compute her strategy given this fundamental uncertainty about the type of game faced (simultaneous-move vs. Stackelberg) [8]. Focusing on *security games*, I provide the following four key results. First, my thesis shows that the Nash equilibria in security games are interchangeable, thus alleviating the equilibrium selection problem. Second, resolving the leader's dilemma, it shows that under a natural restriction, any Stackelberg strategy is also a Nash equilibrium strategy; and furthermore, the solution is unique in a class of real world security games of which ARMOR is a key exemplar. Third, if a follower can attack multiple targets, many of these properties no longer hold. Fourth, my experimental results emphasize positive properties of games that do not fit the restrictions. This theoretical result has major implications for the real-world applications.

Compact formulation for fare evasion deterrence: In the fourth contribution of my thesis, I present a compact game-theoretic formulation, called TRUSTS, for scheduling randomized patrols for fare inspection in transit systems [3]. TRUSTS serves as a cornerstone of future research on applying my previous robust solution methods, facilitating future real world evaluations. TRUSTS models the problem as a zero-sum game, maximizing total revenue (total ticket sales plus penalties) against risk-neutral, perfectly rational followers. This problem differs from previously studied security games in that the leader strategies must satisfy massive temporal and spatial constraints; moreover, unlike in these counterterrorism-motivated applications, a large fraction of the ridership might realistically consider fare evasion, and so the number of followers is potentially huge. A third novelty in this work is deliberate simplification of leader strategies to reduce the likelihood of execution error due to human mistakes. I present an efficient algorithm for computing such patrol strategies and present experimental results using real world ridership data provided by the Los Angeles Sheriff's department (LASD). The LASD has begun trials of TRUSTS by deploying patrols according to my schedules and measuring the revenue recovered.

3. FUTURE WORK

Hybrid model of uncertainty: My contributions so far consider either completely distribution-free uncertainty (e.g., [7]) or uncertainty with perfectly known distribution (e.g., [9]). In real-world applications, it is often the case that we can estimate the uncertainty distribution accurately for some components but not for the others. For example, the leader can estimate the distribution of her execution error using the statistics collected from the past while she may have very limited idea about the follower's observation noise. Thus, in the future, I will investigate a hybrid model of continuous uncertainty that may consist of both components.

Addressing additional uncertainty in TRUSTS: I plan to extend TRUSTS to handle real world uncertainty such as preference, execution, and observation uncertainty considered previously. The extension of modeling followers' preference uncertainty (e.g., risk profiles) already voids the zero-sum assumption of the current TRUSTS formulation which holds only if the primary objective of the leader is to maximize total revenue and the followers are risk-neutral and perfectly rational. This general-sum extension of TRUSTS is a Bayesian Stackelberg game of potentially tens of thousands of types for a realistically sized problem, presenting significant computational challenge. In my future work, I plan to study possible decompositions of the resulting problem by exploiting the domain structure, and develop specialized algorithms based on HUNTER.

4. REFERENCES

[1] V. Conitzer and T. Sandholm. Computing the optimal strategy to commit to. In *EC*, 2006.

[2] M. Jain, C. Kiekintveld, and M. Tambe. Quality-bounded solutions for finite bayesian stackelberg games: Scaling up. In *AAMAS*, 2011.

[3] A. X. Jiang, Z. Yin, M. P. Johnson, C. Kiekintveld, K. Leyton-Brown, T. Sandholm, and M. Tambe. Towards optimal patrol strategies for fare inspection in transit systems. In *AAAI Spring Symposium on Game Theory for Security, Sustainability and Health*, 2012.

[4] P. Paruchuri, J. P. Pearce, J. Marecki, M. Tambe, F. Ordóñez, and S. Kraus. Playing games with security: An efficient exact algorithm for Bayesian Stackelberg games. In *AAMAS*, 2008.

[5] J. Pita, M. Jain, F. Ordóñez, M. Tambe, S. Kraus, and R. Magori-cohen. Effective solutions for real-world Stackelberg games: When agents must deal with human uncertainties. In *AAMAS*, 2009.

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

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

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

[9] Z. Yin and M. Tambe. A unified method for handling discrete and continuous uncertainty in bayesian stackelberg games. In *International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2012.