Computational Game Theory for Security: Progress and Challenges[*]

Milind Tambe, Albert Xin Jiang
Computer Science Department
University of Southern California
Los Angeles, CA 90089
{tambe, jiangx}@usc.edu

Bo An
School of Computer Engineering
Nanyang Technological University
Singapore, 639798
boan@ntu.edu.sg

Manish Jain
Department of Computer Science
Virginia Tech
Blacksburg, VA 24061
jmanish@cs.vt.edu

Abstract

The goal of this paper is to (re)introduce a real-world challenge problem for researchers in multiagent systems and beyond, where our collective efforts may have a significant impact on activities in the real-world. The challenge is in applying game theory for security: our goal is to not only introduce the research challenges for algorithmic and behavioral game theory in service of this problem, but also to provide initial exemplars of successes of deployed systems, and to challenges introduced by these deployments of computational game theory in the field. We also wish to provide an overview of key open research challenges and pointers to getting started in this research.

Introduction

Security is a critical concern around the world that arises in protecting our ports, airports, transportation or other critical national infrastructure, curtailing the illegal flow of drugs, weapons and money, suppressing urban crime, as well in protecting wildlife, fish and forests from poachers and smugglers; and it arises in problems ranging from physical to cyberphysical systems. In all of these problems, we have limited security resources which prevent full security coverage at all times; instead, limited security resources must be deployed intelligently taking into account differences in priorities of targets requiring security coverage, the responses of the adversaries to the security posture and potential uncertainty over the types, capabilities, knowledge and priorities of adversaries faced.

Game theory is well-suited to adversarial reasoning for security resource allocation and scheduling problems. Casting the problem as a Bayesian Stackelberg game, we have developed new algorithms for efficiently solving such games to provide randomized patrolling or inspection strategies. These algorithms have led to some initial successes in this challenge problem arena, leading to advances over previous approaches in security scheduling and allocation, e.g., by addressing key weaknesses of predictability of human schedulers. These algorithms are now deployed in multiple applications: ARMOR has been deployed at the Los Angeles International Airport (LAX) since 2007 to randomizes checkpoints on the roadways entering the airport and canine patrol routes within the airport terminals (Pita et al. 2008); IRIS, is a game-theoretic scheduler for randomized deployment of the US Federal Air Marshal Service (FAMS) requiring significant scale-up in underlying algorithms has been in use since 2009 (Tsai et al. 2009); and PROTECT, which requires further scale up is deployed for generating randomized patrol schedules for the US Coast Guard in Boston, New York, Los Angeles and other ports around the US (An et al. 2011b; Shieh et al. 2012; An et al. 2013b). Furthermore, TRUSTS is being evaluated for deterring fare evasion, suppressing urban crime and counter-terrorism within the Los Angeles Metro System (Yin et al. 2012; Jiang et al. 2013a), and GUARDS was earlier tested by the US Transportation Security Administration (TSA) for security inside the airport (Pita et al. 2011). These initial successes point the way to major future applications in a wide range of security arenas; with major research challenges in scaling up our game-theoretic algorithms, to addressing human adversaries' bounded rationality and uncertainties in action execution and observation, as well as in preference elicitation and multiagent learning.

This paper will provide pointers to our algorithms, key research challenges and information on getting started in this research. While initial research has made a start, a lot remains to be done; yet these are large-scale interdisciplinary research challenges that call upon researchers to work with researchers in other disciplines, be "on the ground" with domain experts, and examine real-world constraints and challenges that cannot be abstracted away. What this research is leading to is the very new area of computational game theory in the field. Together as an international community of multiagent researchers, we can accomplish more!

Deployed and Emerging Security Applications

The last several years have witnessed the successful application of multi-agent systems in allocating limited resources to protect critical infrastructures (Basilico, Gatti, and Amigoni 2009; Korzhyk, Conitzer, and Parr 2010; Jain et al. 2010b; Pita et al. 2011; An et al. 2011b; Tambe and An 2012; An et al. 2013b; Yin et al. 2012; Jiang et al. 2013a;

[*] This paper is an updated version of the paper appeared at the AAAI 2012 Spring Symposium on Game Theory for Security, Sustainability and Health (Tambe and An 2012).

An et al. 2013a; 2011a). The framework of Stackelberg games is well suited to formulate the strategic interaction in security domains in which there are usually two players: the security force (defender) commits to a security policy first and the attacker (e.g., terrorist, poacher and smuggler) conducts surveillance to learn the policy and then takes his best attacking action.[1] Stackelberg games have been widely used for modeling/reasoning complex security problems and a variety of algorithms have been proposed to efficiently compute the equilibrium strategy, i.e., defender's best way of utilizing her limited security resources (there is actually a special class of Stackelberg games that often gets used in these security domains, and this class is referred to as security games). In the rest of this section, we describe the application of the Stackelberg game framework in multiple significant security domains. We start with our first and by now the smallest-scale application; we discuss it as it is quite instructive to understand the overall problem.

ARMOR for Los Angeles International Airport

Los Angeles International Airport (LAX) is the largest destination airport in the United States and serves 60-70 million passengers per year. The LAX police use diverse measures to protect the airport, which include vehicular checkpoints and police units patrolling with canines. The eight different terminals at LAX have very different characteristics, like physical size, passenger loads, foot traffic or international versus domestic flights. Furthermore, the numbers of available vehicle checkpoints and canine units are limited by resource constraints. Thus it is challenging to optimally allocate these resources to improve their effectiveness while avoiding patterns in the scheduled deployments.

The ARMOR system (Assistant for Randomized Monitoring over Routes) focuses on two of the security measures at LAX (checkpoints and canine patrols) and optimizes security resource allocation using Bayesian Stackelberg games. Take the vehicle checkpoints model as an example. Assume that there are n roads, the police's strategy is placing $m < n$ checkpoints on these roads where m is the maximum number of checkpoints. The adversary may potentially choose to attack through one of these roads. ARMOR models different types of attackers with different payoff functions, representing different capabilities and preferences for the attacker. ARMOR uses DOBSS (Decomposed Optimal Bayesian Stackelberg Solver) to compute the defender's optimal strategy (Paruchuri et al. 2008). ARMOR has been successfully deployed since August 2007 at LAX to randomize checkpoints on the roadways entering the airport and canine patrol routes within the airport terminals (Pita et al. 2008).

IRIS for US Federal Air Marshals Service

The US Federal Air Marshals Service (FAMS) allocates air marshals to flights originating in and departing from the United States to dissuade potential aggressors and prevent an attack should one occur. Flights are of different importance based on a variety of factors such as the numbers of passengers, the population of source/destination, international flights from different countries, and special events that can change the risks for particular flights at certain times. Security resource allocation in this domain is significantly more challenging than for ARMOR: a limited number of FAMS need to be scheduled to cover thousands of commercial flights each day. Furthermore, these FAMS must be scheduled on tours of flights that obey various constraints (e.g., the time required to board, fly, and disembark). Therefore, we face significant computational challenge while generating the optimal scheduling policy that meets these scheduling constraints.

(a) PROTECT is being used in Boston (b) Extending PROTECT to NY

Figure 1: USCG boats patrolling the ports of Boston and NY

Figure 2: Protecting ferries with patrol boats

Against this background, the IRIS system (Intelligent Randomization In Scheduling) has been developed and has been deployed by FAMS since October 2009 to randomize schedules of air marshals on international flights. In IRIS, the targets are the set of n flights and the attacker could potentially choose to attack one of these flights. The FAMS can assign $m < n$ air marshals that may be assigned to protect these flights. Since the number of possible schedules exponentially increases with the number of flights and resources, DOBSS is no longer applicable to the FAMS domain. Instead, IRIS uses the much faster ASPEN algorithm (Jain et al. 2010a) to generate the schedule for thousands of commercial flights per day. IRIS also uses an attribute-based preference elicitation system to determine reward values for the Stackelberg game model.

PROTECT for US Coast Guard

The US Coast Guard's (USCG) mission includes maritime security of the US coasts, ports, and inland waterways; a security domain that faces increased risks due to threats such as terrorism and drug trafficking. Given a particular port and

[1] Or the attacker may be sufficiently deterred and dissuaded from attacking the protected target.

the variety of critical infrastructure that an adversary may attack within the port, USCG conducts patrols to protect this infrastructure; however, while the adversary has the opportunity to observe patrol patterns, limited security resources imply that USCG patrols cannot be at every location 24/7. To assist the USCG in allocating its patrolling resources, the PROTECT (Port Resilience Operational / Tactical Enforcement to Combat Terrorism) model is being designed to enhance maritime security and has been in use at the port of Boston since April 2011 (Figure 1). Similar to previous applications ARMOR and IRIS, PROTECT uses an attacker-defender Stackelberg game framework, with USCG as the defender against terrorist adversaries that conduct surveillance before potentially launching an attack.

PROTECT is currently deployed in the ports of Boston, New York, Los Angeles/Long Beach and several others (An et al. 2013b). Indeed the goal now is to deploy PROTECT at ports nationwide. Furthermore, beyond just port protection, PROTECT has been extended to protect ferry systems such as the Staten Island ferry in New York (Fang, Jiang, and Tambe 2013).

While PROTECT builds on previous work, it offers some key innovations. First, to improve PROTECT's efficiency, a compact representation of the defender's strategy space is used by exploiting equivalence and dominance. Second, the evaluation of PROTECT for the first time provides real-world data: (i) comparison of human-generated vs PROTECT security schedules, and (ii) results from an Adversarial Perspective Team's (human mock attackers) analysis.

GUARDS for US Transportation Security Agency

The US Transportation Security Administration (TSA) is tasked with protecting the nation's over 400 airports. To aid the TSA in scheduling resources to protect airports, a new application called GUARDS (Game-theoretic Unpredictable and Randomly Deployed Security) has been developed. While GUARDS also utilizes Stackelberg games as ARMOR and IRIS, GUARDS faces three key challenges (Pita et al. 2011): 1) reasoning about hundreds of heterogeneous security activities; 2) reasoning over diverse potential threats; and 3) developing a system designed for hundreds of end-users. To address those challenges, GUARDS created a new game-theoretic framework that allows for heterogeneous defender activities and compact modeling of a large number of threats and developed an efficient solution technique based on general-purpose Stackelberg game solvers. GUARDS was originally tested at an undisclosed airport and further results are awaited (Pita et al. 2011).

TRUSTS for Urban Security in Transit Systems

TRUSTS focuses on three major security challenges: deterring fare evasion, suppressing crime and counter-terrorism. Significant focus in TRUSTS has been on deterring fare evasion. Specifically, in some urban transit systems, including the Los Angeles Metro Rail system, passengers are legally required to purchase tickets before entering but are not physically forced to do so (Figure 4). Instead, patrol units move about through the transit system, inspecting tickets of passengers, who face fines for fare evasion. This setting yields the problem of computing optimal patrol strategies, to deter fare evasion and hence maximize revenue. The TRUSTS system (Tactical Randomization for Urban Security in Transit Systems) models the patrolling problem as a leader-follower Stackelberg game (Yin et al. 2012; Jiang et al. 2013a). Urban transit systems, however, present unique computational challenges since there are exponentially many possible patrol strategies, each subject to both the spatial and temporal constraints of travel within the transit network under consideration. To overcome this challenge, TRUSTS uses a compact representation which captures the spatial as well as temporal structure of the domain. The system has been evaluated using real-world ridership data from the Los Angeles Metro Rail system.

(a) Los Angeles Metro (b) Barrier-free entrance

Figure 3: TRUSTS for transit systems

One key finding from initial tests was that the schedules generated by officers were often interrupted. Interruptions occurred because in frequent interactions with the public, sometimes officers would get stopped by lost travelers, sometimes they would need to arrest someone. Such interruptions meant that the schedules now needed to be highly dynamic. To that end, a new generation of Stackelberg game based scheduling algorithms – using Markov Decision Problems – was designed. This led to schedules now being loaded onto smartphones and given to officers. The schedules are then automatically updated on the smartphone if interruptions occur (Luber et al. 2013).

Applications Focusing on Deterring Environmental and Urban Crime

A number of newer applications are focused on suppressing crime: both environmental crime and urban crime. One of those is protecting forests (Johnson et al. 2012), where we must protect a continuous forest area from extractors. Since the attacker's behavior (e.g., extracting important resources from the forest) could be effected by spatial considerations, it is critical for the defender to incorporate spatial considerations into her enforcement decisions (Albers 2010).

Another area of interesting is protecting endangered species. Endangered species poaching is reaching critical levels as the populations of these species plummet to unsustainable numbers. The global tiger population, for example, has dropped over 95% from the start of the 1900s and has resulted in three out of nine species extinctions. Depending on the area and animals poached, motivations for poaching range from profit to sustenance, with the former being more common when profitable species such as tigers, ele-

(a)

(b)

Figure 4: Computational game theory can play a role in the (a) global tiger initiative and (b) forest protection. Both figures are from http://commons.wikimedia.org.

phants, and rhinos are the targets. To counter poaching efforts and to rebuild the species' populations, countries have set up protected wildlife reserves and conservation agencies tasked with defending these large reserves. Because of the size of the reserves and the common lack of law enforcement resources, conservation agencies are at a significant disadvantage when it comes to deterring and capturing poachers. Agencies use patrolling as a primary method of securing the park. Due to their limited resources, however, patrol managers must carefully create patrols that account for many different variables (e.g., limited patrol units to send out, multiple locations that poachers can attack at varying distances to the outpost). Protection Assistant for Wildlife Security (PAWS) aims to assist conservation agencies in their critical role of patrol creation by predicting where poachers will attack and optimizing patrol routes to cover those areas.

Another emerging application domain is that of ensuring the sustainability of fish resources. Marine fisheries are acknowledged to be some of the most important food resources for countries around the world. As reported by World Wild Fund for Nature (WWF), cod are currently at risk from overfishing in the UK, Canada and most other Atlantic countries. Global cod catch has suffered a 70% drop over the last 30 years, and if this trend continues, the world's cod stocks will disappear in 15 years. Illegal, unreported, and unregulated (IUU) fishing is one of the major threats to the sustainability of ocean fish resources. As estimated by National Oceanic and Atmospheric Administration (NOAA), IUU fishing produces between 11 and 26 million tons of seafood annually, representing as much as 40 percent of the total catch in some fisheries. The driver behind IUU fishing is high economic profit and low chance of seizure. It is impossible to maintain a 24/7 presence to prevent IUU fishing everywhere due to the limited asset patrolling resources. Hence the allocation of the patrolling resources becomes a key challenge for security agencies like US Coast Guard.

Even with all of these applications, we have barely scratched the surface of possibilities in terms of potential applications for multiagent researchers for applying game theory for security.

Figure 5: US Coast Guard personnel on a mission to protect fisheries

Open Research Issues

While the deployed applications have advanced the state of the art, significant future research remains to be done. In the following, we highlight some key research challenges, including scalability, robustness, and human adversary modeling. The main point we want to make is that this research does not require access to classified information of any kind. Problems, solution approaches and datasets are well specified in the papers discussed below,

Scalability: The first research challenge is improving the scalability of our algorithms for solving Stackelberg (security) games. The strategy space of both the defender and the attacker in these games may exponentially increase with the number of security activities, attacks, and resources. As we scale up to larger domains, it is critical to develop newer algorithms that scale up significantly beyond the limits of the current state of the art of Bayesian Stackelberg solvers. Driven by the growing complexity of applications, a sequence of algorithms for solving security games have been developed including DOBSS (Paruchuri et al. 2008), ERASER (Jain et al. 2010b), ASPEN (Jain et al. 2010a). However, existing algorithms still cannot scale up to very large scale domains such as scheduling randomized checkpoints in cities. In such graph based security games, the strategy space of the defender grows exponentially with the number of available resources and the strategy space of the attacker grows exponentially with the size of the road network considered. The latest technique to schedule such checkpoints is based on a "double oracle approach" which does not require the enumeration of the entire strategy space for either of the players (Jain et al. 2011; Jain, Tambe, and Conitzer 2013).

Another approach for patrolling domains with spatiotemporal constraints is to compactly represent defender mixed strategies as fractional flows. This approach has recently been applied to efficiently compute fare-enforcement patrols in urban transit systems (Yin et al. 2012; Jiang et al. 2013b) and boat patrols for protecting ferries (Fang, Jiang, and Tambe 2013). An open problem is to find other types of security domains in which the strategy space can be compactly represented.

Robustness: The second challenge is improving solutions' robustness. Classical game theory solution concepts often make assumptions on the knowledge, rationality, and capability of players. Unfortunately, those assumptions could be wrong in real-world scenarios. Therefore, while computing the defender's optimal strategy, algorithms should take into account various uncertainties faced in the domain, including payoff noise (Kiekintveld, Marecki, and Tambe 2011), execution/observation error (Yin et al. 2011), uncertain capability (An et al. 2011c). For observation uncertainty, it is typically assumed that the attacker has perfect knowledge of the defender's randomized strategy or can learn the defender's strategy after conducting a fixed period of surveillance. In consideration of surveillance cost, these assumptions are clearly simplistic since attackers may act with partial knowledge of the defender's strategies and may dynamically decide whether to attack or conduct more surveillance. Security game models with limited observation (An et al. 2012; 2013a) have been proposed in which the attacker either makes limited number of observations or dynamically determines a place to stop surveillance. Since the belief state space exponentially increases with observation length, it is still computationally challenging to solve large games in consideration of limited observation.

Bounded Rationality: One required research direction with respect to robustness is addressing bounded rationality of human adversaries, which is a fundamental problem that can affect the performance of our game theoretic solutions. Recently, there has been some research on applying ideas (e.g., prospect theory (Kahneman and Tvesky 1979), and quantal response (McKelvey and Palfrey 1995)) from social science or behavioral game theory within security game algorithms (Yang et al. 2011; Pita et al. 2010; Nguyen et al. 2013; Yang et al. 2013). Previous work usually applies existing frameworks and sets the parameters of these frameworks by experimental tuning or learning. However, in real-world security domains, we may have very limited data, or may only have some limited information on the biases displayed by adversaries. Recently, monotonic maximin (Jiang et al. 2013a) was proposed as a robust solution concept to Stackelberg security games with boundedly rational adversaries. It tries to optimize defender utility against the worst-case monotonic adversary behavior, where monotonicity is the property that actions with higher expected utility is played with higher probability. An open research challenge is to combine such robust-optimization approaches with available behavior data. Furthermore, since real-world human adversaries are sometimes distributed coalitions of socially, culturally and cognitively-biased agents, acting behind a veil of uncertainty, we may need significant interdisciplinary research to build in social, cultural and coalitional biases into our adversary models.

In addition to the above research challenges, there are other on-going challenges such as preference elicitation for acquiring necessary domain knowledge in order to build game models and evaluation of the game theoretic applications (Taylor et al. 2010).

Resources for Starting This Research

Security is recognized as a world-wide challenge and game theory is an increasingly important paradigm for reasoning about complex security resource allocation. While the deployed game theoretic applications have provided a promising start, very significant amount of research remains to be done. These are large-scale interdisciplinary research challenges that call upon multiagent researchers to work with researchers in other disciplines, be "on the ground" with domain experts, and examine real-world constraints and challenges that cannot be abstracted away.

There are a number of resources (mostly online) for starting this research. The research papers related to game theory for security have been extensively published at AAMAS conference [2] and the reader can also find some papers from AAAI [3] and IJCAI [4] conferences. Additional resources:

- Key papers describing important algorithms and the deployed systems can also be found from a recently published book –*Security and Game Theory: Algorithms, Deployed Systems, Lessons Learned* (Tambe 2011).
- The details of those deployed systems can also be found at http://teamcore.usc.edu/projects/security/.
- From http://teamcore.usc.edu/projects/security/, the reader can also find a tutorial at UAI'2011 – Game Theory for Security: Lessons learned from deployed applications.

While we have focused on research conducted by our Teamcore group, there are a few other research groups that have started addressing challenges in security games (Basilico, Gatti, and Amigoni 2009; Dickerson et al. 2010; Korzhyk, Conitzer, and Parr 2011b; 2011a; Letchford and Vorobeychik 2013; 2012; Letchford and Conitzer 2013).

References

Albers, H. 2010. Spatial modeling of extraction and enforcement in developing country protected areas. *Resource and Energy Economics* 32(2):165–179.

An, B.; Jain, M.; Tambe, M.; and Kiekintveld, C. 2011a. Mixed-initiative optimization in security games: A preliminary report. In *Proc. of the AAAI Spring Symposium on Help Me Help You: Bridging the Gaps in Human-Agent Collaboration*, 8–11.

An, B.; Pita, J.; Shieh, E.; Tambe, M.; Kiekintveld, C.; and Marecki, J. 2011b. GUARDS and PROTECT: Next generation applications of security games. *SIGECOM* 10:31–34.

An, B.; Tambe, M.; Ordonez, F.; Shieh, E.; and Kiekintveld, C. 2011c. Refinement of strong stackelberg equilibria in security games. In *Proc. of the 25th Conference on Artificial Intelligence*, 587–593.

An, B.; Kempe, D.; Kiekintveld, C.; Shieh, E.; Singh, S.; Tambe, M.; and Vorobeychik, Y. 2012. Security games with limited surveillance. In *Proceedings of the 26th Conference on Artificial Intelligence (AAAI)*, 1241–1248.

An, B.; Brown, M.; Vorobeychik, Y.; and Tambe, M. 2013a. Security games with surveillance cost and optimal timing of attack execution. In *Proceedings of The 12th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 223–230.

An, B.; Ordóñez, F.; Tambe, M.; Shieh, E.; Yang, R.; Baldwin, C.; DiRenzo, J.; Moretti, K.; Maule, B.; and Meyer, G. 2013b. A deployed quantal response-based patrol planning system for the U.S. Coast Guard. *Interfaces* 43(5):400–420.

Basilico, N.; Gatti, N.; and Amigoni, F. 2009. Leader-follower strategies for robotic patrolling in environments with arbitrary topologies. In *Proc. of The 8th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 500–503.

Dickerson, J. P.; Simari, G. I.; Subrahmanian, V. S.; and Kraus, S. 2010. A graph-theoretic approach to protect static and moving

[2] www.aamas-conference.org.

[3] www.aaai.org/.

[4] ijcai.org/.

targets from adversaries. In *Proc. of The 9th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 299–306.

Fang, F.; Jiang, A.; and Tambe, M. 2013. Optimal patrol strategy for protecting moving targets with multiple mobile resources. In *AAMAS*.

Jain, M.; Kardes, E.; Kiekintveld, C.; Ordonez, F.; and Tambe, M. 2010a. Security games with arbitrary schedules: A branch and price approach. In *Proc. of The 24th AAAI Conference on Artificial Intelligence*, 792–797.

Jain, M.; Tsai, J.; Pita, J.; Kiekintveld, C.; Rathi, S.; Tambe, M.; and Ordonez, F. 2010b. Software assistants for randomized patrol planning for the lax airport police and the federal air marshal service. *Interfaces* 40:267–290.

Jain, M.; Korzhyk, D.; Vanek, O.; Pechoucek, M.; Conitzer, V.; and Tambe, M. 2011. A double oracle algorithm for zero-sum security games on graphs. In *Proc. of The 10th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*.

Jain, M.; Tambe, M.; and Conitzer, V. 2013. Security scheduling for real-world networks. In *AAMAS*.

Jiang, A. X.; Nguyen, T. H.; Tambe, M.; and Procaccia, A. D. 2013a. Monotonic maximin: A robust stackelberg solution against boundedly rational followers. In *Conference on Decision and Game Theory for Security (GameSec)*.

Jiang, A.; Yin, Z.; Kraus, S.; Zhang, C.; and Tambe, M. 2013b. Game-theoretic randomization for security patrolling with dynamic execution uncertainty. In *AAMAS*.

Johnson, M.; Fang, F.; Yang, R.; Tambe, M.; and Albers, H. 2012. Patrolling to maximize pristine forest area. In *Proc. of the AAAI Spring Symposium on Game Theory for Security, Sustainability and Health*.

Kahneman, D., and Tvesky, A. 1979. Prospect theory: An analysis of decision under risk. *Econometrica* 47(2):263–291.

Kiekintveld, C.; Marecki, J.; and Tambe, M. 2011. Approximation methods for infinite bayesian stackelberg games: modeling distributional uncertainty. In *Proc. of The 10th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*.

Korzhyk, D.; Conitzer, V.; and Parr, R. 2010. Complexity of computing optimal stackelberg strategies in security resource allocation games. In *Proc. of The 24th AAAI Conference on Artificial Intelligence*, 805–810.

Korzhyk, D.; Conitzer, V.; and Parr, R. 2011a. Security games with multiple attacker resources. In *Proc. of The International Joint Conference on Artificial Intelligence (IJCAI)*.

Korzhyk, D.; Conitzer, V.; and Parr, R. 2011b. Solving stackelberg games with uncertain observability. In *Proc. of The 10th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*.

Letchford, J., and Conitzer, V. 2013. Solving security games on graphs via marginal probabilities. In *Proceedings of the 27th AAAI Conference on Artificial Intelligence (AAAI)*.

Letchford, J., and Vorobeychik, Y. 2012. Computing optimal security strategies for interdependent assets. In *Proceedings of the conference on Uncertainty in Artificial Intelligence(UAI)*.

Letchford, J., and Vorobeychik, Y. 2013. Optimal interdiction of attack plans. In *International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*.

Luber, S.; Yin, Z.; Fave, F. D.; Jiang, A. X.; Tambe, M.; and Sullivan, J. P. 2013. Game-theoretic patrol strategies for transit systems: the trusts system and its mobile app (demonstration). In *International Conference on Autonomous Agents and Multiagent Systems (AAMAS)[Demonstrations Track]*.

McKelvey, R. D., and Palfrey, T. R. 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 *Conference on Artificial Intelligence (AAAI)*.

Paruchuri, P.; Pearce, J. P.; Marecki, J.; Tambe, M.; Ordonez, F.; and Kraus, S. 2008. Playing games with security: An efficient exact algorithm for Bayesian Stackelberg games. In *Proc. of The 7th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 895–902.

Pita, J.; Jain, M.; Western, C.; Portway, C.; Tambe, M.; Ordonez, F.; Kraus, S.; and Parachuri, P. 2008. Deployed ARMOR protection: The application of a game-theoretic model for security at the Los Angeles International Airport. In *AAMAS*, 125–132.

Pita, J.; Jain, M.; Tambe, M.; Ordóñez, F.; and Kraus, S. 2010. Robust solutions to stackelberg games: Addressing bounded rationality and limited observations in human cognition. *Artificial Intelligence* 174(15):1142–1171.

Pita, J.; Tambe, M.; Kiekintveld, C.; Cullen, S.; and Steigerwald, E. 2011. GUARDS - game theoretic security allocation on a national scale. In *Proc. of The 10th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*.

Shieh, E.; An, B.; Yang, R.; Tambe, M.; Baldwin, C.; DiRenzo, J.; Maule, B.; and Meyer, G. 2012. PROTECT: An application of computational game theory for the security of the ports of the United States. In *Proceedings of the 26th AAAI Conference on Artificial Intelligence (AAAI)*, 2173–2179.

Tambe, M., and An, B. 2012. Game theory for security: A real-world challenge problem for multiagent systems and beyond. In *Proc. of The AAAI Spring Symposium on Game Theory for Security, Sustainability and Health*, 69–74.

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

Taylor, M. E.; Kiekintveld, C.; Western, C.; and Tambe, M. 2010. A framework for evaluating deployed security systems: Is there a chink in your armor? *Informatica* 34:129–139.

Tsai, J.; Rathi, S.; Kiekintveld, C.; Ordonez, F.; and Tambe, M. 2009. IRIS: a tool for strategic security allocation in transportation networks. In *Proc. of The 8th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 37–44.

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.; Jiang, A. X.; Tambe, M.; and Ordonez, F. 2013. Scaling-up security games with boundedly rational adversaries: A cutting-plane approach. In *International Joint Conference on Artificial Intelligence (IJCAI)*.

Yin, Z.; Jain, M.; Tambe, M.; and Ordonez, F. 2011. Risk-averse strategies for security games with execution and observational uncertainty. In *Proc. of The 25th AAAI Conference on Artificial Intelligence (AAAI)*, 758–763.

Yin, Z.; Jiang, A.; Johnson, M.; Tambe, M.; Kiekintveld, C.; Leyton-Brown, K.; Sandholm, T.; and Sullivan, J. 2012. Trusts: Scheduling randomized patrols for fare inspection in transit systems. In *Conference on Innovative Applications of Artificial Intelligence (IAAI)*.