

Green Security Games: Apply Game Theory to Addressing Green Security Challenges

FEI FANG

University of Southern California

and

THANH H. NGUYEN

University of Southern California

In the past decade, game-theoretic applications have been successfully deployed in the real world to address security resource allocation challenges. Inspired by the success, researchers have begun focusing on applying game theory to green security domains such as protection of forests, fish, and wildlife, forming a stream of research on Green Security Games (GSGs). We provide an overview of recent advances in GSGs and list the challenges that remained open for future study.

Categories and Subject Descriptors: I.2.11 [**Distributed Artificial Intelligence**]: Multiagent systems—

General Terms: Security, Human factors, Algorithms

Additional Key Words and Phrases: Security, Game Theory, Conservation, Computational Sustainability

1. INTRODUCTION

Green security domains, including protecting the forest from illegal logging, protecting endangered wildlife from poaching, and protecting fish stocks from illegal fishing, are critical domains for environmental sustainability. For example, poaching has led to the population drop of tigers and other key species, presenting a significant threat to the whole ecosystem. Illegal logging is a major problem for many developing countries, with severe economic and environmental impacts. In these domains, law enforcement agencies often suffer from a lack of resources to combat illegal activities, and strategic allocation or scheduling of resources is in great need.

In the past decade, Stackelberg Security Games (SSGs) has been studied extensively for optimizing security resource allocation. In an SSG, the defender commits to a randomized strategy to protect a set of targets, and the attacker then responds by attacking a target with the highest expected utility. Previous work on SSGs focuses on designing models to represent complex real-world problems and developing efficient algorithms to compute the optimal strategy for the defender. Research on SSGs has led to many successfully deployed applications in infrastructure security domains, including ones for protecting airports, ports and flights [Tambe 2011; Pita et al. 2008; An et al. 2011; Fang et al. 2013].

Building on previous work on SSGs, researchers have started focusing on applying game theory to green security domains to optimize the allocation of limited

Authors' addresses: feifang@usc.edu, thanhhng@usc.edu

resources, leading to a stream of research on Green Security Games (GSGs). Similar to SSGs, there are two types of players in GSGs, the defender (e.g., the law enforcement agency) and the attacker (e.g., the poacher). Green security domains have several key differences when compared to infrastructure security domains. First of all, frequent and repeated attacks are involved. For example, poachers place snares frequently to poach wildlife. Therefore, if the problem is modeled as a game, it is no longer a one-shot game. In addition, the frequent attacks would bring in more attack data that can be exploited by the defender. The second main difference is in attackers' decision making. The attacks take place frequently, and it is impossible for the attacker to conduct long-term surveillance before each of the attacks. Also, due to frequent attacks and the relatively low cost of failure, the attacker will take less effort in planning the attacks and may be boundedly rational in their decision making. Third, green security domains often involve a large area in need of protection, and the spatial change (e.g., elevation change) within the area cannot be neglected. Fourth, the law enforcement agencies may need to form a team of different types of patrolling resources, ranging from local volunteers to police officers to NGO personnel.

These differences lead to a number of research challenges, including how to model the repeated attacks, how to represent the behavior model of attackers in GSGs, how to learn from data to improve the defender's performance, how to handle uncertainties in the game model, how to handle spatial constraints, and how to simultaneously find an optimal team of resources and the optimal allocation of the resources. Here we provide a brief overview of recent research advances to address these challenges in GSGs and outline several key challenges that remain open, pointing out possible directions for future research.

2. RESEARCH ADVANCES IN GSGS

2.1 Repeated Attacks and Deceptive Planning

[Yang et al. 2014] presents an initial effort to apply game theory to green security domains. To feature the repeated attacks in wildlife protection, it models the problem as a repeated Stackelberg game, i.e., it considers a multi-stage game where in each stage, attackers respond to a mixed strategy chosen by the defender. [Haskell et al. 2014] adopts this model for the problem of protecting fisheries from illegal fishing. [Fang et al. 2015] provides a formal definition of GSG, which generalizes the Stackelberg assumption. It captures the fact that the attackers may not be able to conduct long-term surveillance before each of their attacks, and there may be a delay in their understanding of the defender's strategy. This model enables deceptive planning, as the defender can exploit the delay and carefully plan for a sequence of strategies to be used to achieve a higher overall utility.

2.2 Behavior Model for Human Attackers

In GSGs, the attackers are often not perfectly rational utility maximizers. They are often boundedly rational due to limited time for planning the attacks, and it is necessary to incorporate behavior models to the game.

In previous work, several behavior models are proposed to describe the attacker's bounded rationality for SSGs [Pita et al. 2012; Yang et al. 2012; Nguyen et al. 2013]

and the SUQR (subjective utility quantal response) model has been demonstrated to have the best performance in human subject experiments so far [Nguyen et al. 2013]. The SUQR model proposes that the attacker evaluates key features of each target (such as defender coverage, adversary reward, and adversary penalty), and chooses more promising targets with higher probability. The parameters in SUQR model indicate the weights that the attacker may give to different features.

Motivated by the initial success of the SUQR in GSGs, significant efforts have been made to further improve the prediction accuracy, addressing key complexity when real-world green security domains are considered. To capture the complexity introduced by repeated interaction in GSGs, [Kar et al. 2015] proposes a new behavior model named SHARP that models the adversary’s adaptiveness by taking into account the success or failure of the adversary’s past actions. Also, SHARP reasons about the similarity between exposed and unexposed areas of the attack surface. Another complexity in GSGs is that the defender is often faced with many attackers. The Bayesian SUQR model [Yang et al. 2014] is proposed to handle multiple attackers who may have different parameter vectors. It assumes the parameter vectors of the group of attackers follow a Gaussian distribution. The latest model for describing the attacker’s behavior is CAPTURE [Nguyen et al. 2016]. CAPTURE model is built based on two key separate components: (1) the behavior component incorporates the dependence of the attackers’ behavior on their activities in the past; and (2) the detection component takes into account the defender’s imperfect detection of attack signs. The CAPTURE model is shown to be superior to SUQR in predicting the poachers’ behavior in wildlife protection based on the largest poaching dataset collected by rangers in the Queen Elizabeth National Park in Uganda over 12 years.

2.3 Learn to Play GSGs

Given the data available in GSGs, it is important to understand how the defender can exploit the data to improve their strategy. One major topic is how to learn the parameters in the attacker’s behavior models. [Yang et al. 2014] proposes a learning framework that can learn from a combination of abundant anonymous data (e.g., anonymous snares placed on the ground) and a few identified data (e.g., snares that are linked to individual poachers). It learns parameters from each anonymous data point using Maximum Likelihood Estimation (MLE) and then learns a Gaussian distribution for a population of attackers with the help of identified data points. [Haskell et al. 2014] also uses MLE to estimate the parameters from data. [Fang et al. 2015] proposes a Bayesian update based approach to estimate the parameter distribution for a population of attackers. Researchers have also worked on analyzing the theoretical aspects of learning in GSGs. [Sinha et al. 2016] analyzes learning the response function of the adversary based on the PAC model, and [Haghtalab et al. 2016] proposes an approach that learns the parameters in the attacker’s behavioral model by observing how the attacker responds to only three defender strategies.

Instead of learning the parameters in the behavior models, [Qian et al. 2016] propose a different solution framework called RMAB to solve GSGs by learning directly from previous defender-attacker interactions while considering partial observability. In particular, the problem is represented using a restless multi-armed bandit (RMAB) model to handle the limited observation challenge, i.e., the de-

fender does not have observations for targets (arms) they do not patrol (activate). In this RMAB model, the impact of the defender's patrols on the attackers' activities is modeled as a Markov process. Since the defender can learn the effect of patrols on the attackers' behavior from the historical data, the defender can predict attack activities at locations where they do not patrol.

2.4 Handle Uncertainties in GSGs

Prior game-theoretic research on GSGs assumes that payoff values of both the defender and the attacker are precisely estimated. However, in some green security domains, information on key domain features (e.g., animal density in wildlife protection) that contribute to the payoffs is not perfectly known, leading to the uncertainty over the payoff values. To address this challenge of payoff uncertainty, the ARROW algorithm is proposed to compute an optimal patrolling strategy for the defender which is robust against payoff uncertainty [Nguyen et al. 2015]. Essentially, ARROW focuses on computing Minimax Regret (MMR) strategies which handle uncertainty in both players' payoffs, given the presence of an attacker behavioral model. MMR is a robust solution method for handling uncertainty, attempting to find the solution which minimizes the maximum regret (i.e., utility loss of the defender) over a prior uncertainty set.

2.5 Handle Spatial/Practical Constraints

To develop a GSG-based application that can be used in the real world, it is important to consider the spatial constraints and other practical constraints. [Fang et al. 2016] handles the practical constraints in bringing GSG to the real world to combat poaching, and introduced an application PAWS (Protection Assistant for Wildlife Security) which has been deployed in Southeast Asia for wildlife conservation. PAWS provides the defender a set of suggested patrol routes that starts and ends from the base camp location within distance limit, and provide probabilities that each of the routes should be taken. To get a practical solution that can be used by patrollers for foot patrols, PAWS considers detailed elevation information and design patrol routes on a virtual street map that consists of terrain features which patrollers can easily follow. PAWS incorporates a modified version of ARROW and the cutting plane framework to handle the payoff uncertainty and scheduling constraints with boundedly rational attackers.

2.6 Team Formation

While it is important to investigate how to improve the tactics of conducting patrols in GSGs, there is another layer of optimization that could be involved, which is the team formation problem. The law enforcement agencies may need to team up different types of groups, from national police to local volunteers, each differing in their interdiction effectiveness, and with varying costs of deployment. [McCarthy et al. 2016] introduces a game model with this additional complexity in the domain of preventing illegal logging, and presents an algorithm FORTIFY that can simultaneously find an optimal team of resources and the optimal allocation of the resources.

3. FUTURE DIRECTIONS FOR GSGS

While efforts have been made to address the challenges in GSGs, there are a few challenges that are open for further investigation. The first challenge is scaling up the algorithms to handle large problems and provide a fine-grained solution. In previous work, the area in need of protection is often discretized into a grid, and each grid cell is treated as a target. Due to the complexity brought by the bounded rationality of attackers and uncertainty, the proposed algorithm cannot scale up to a large number of targets and therefore often fail to provide a fine-grained solution, leading to an important direction for future work. This challenge could potentially be addressed by exploiting the game structure, using abstraction and contraction, and applying hierarchical discretization.

Another direction of future work could be dealing with dynamic defender-attacker interactions in the presence of data. In domains such as wildlife protection, the players can partially observe the other players' actions and the observation data collected from informants and surveillance (e.g., camera traps) may lead to a change in their behavior in the future. It is important to understand how these factors affect the players' behavior and improve the way of collecting data for the defender (e.g., how to allocate the camera traps optimally to maximize the foot patrol efficiency).

PAWS is the first of a new wave of applications based on GSG research. With further research advances in GSGs, we expect more applications deployed in green security domains in the future.

REFERENCES

- AN, B., PITA, J., SHIEH, E., TAMBE, M., KIEKINTVELD, C., AND MARECKI, J. 2011. GUARDS and PROTECT: Next Generation Applications of Security Games. *ACM SIGecom Exchanges* 10, 1, 31–34.
- FANG, F., JIANG, A. X., AND TAMBE, M. 2013. Optimal Patrol Strategy for Protecting Moving Targets with Multiple Mobile Resources. In *AAMAS*.
- FANG, F., NGUYEN, T. H., PICKLES, R., LAM, W. Y., CLEMENTS, G. R., AN, B., SINGH, A., TAMBE, M., AND LEMIEUX, A. 2016. Deploying PAWS: Field Optimization of the Protection Assistant for Wildlife Security. In *Proceedings of the Twenty-Eighth Innovative Applications of Artificial Intelligence Conference (IAAI 2016)*.
- FANG, F., STONE, P., AND TAMBE, M. 2015. When Security Games Go Green: Designing Defender Strategies to Prevent Poaching and Illegal Fishing. In *International Joint Conference on Artificial Intelligence (IJCAI)*.
- HAGHTALAB, N., FANG, F., NGUYEN, T. H., SINHA, A., PROCACCIA, A. D., AND TAMBE, M. 2016. Three Strategies to Success : Learning Adversary Models in Security Games. In *IJCAI'16*.
- HASKELL, W. B., KAR, D., FANG, F., TAMBE, M., CHEUNG, S., AND DENICOLA, L. E. 2014. Robust protection of fisheries with COMPASS. In *IAAI*.
- KAR, D., FANG, F., FAVE, F. D., SINTOV, N., AND TAMBE, M. 2015. "A Game of Thrones": When Human Behavior Models Compete in Repeated Stackelberg Security Games. In *International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2015)*.
- MC CARTHY, S., TAMBE, M., KIEKINTVELD, C., GORE, M. L., AND KILLION, A. 2016. Preventing Illegal Logging: Simultaneous Optimization of Resource Teams and Tactics for Security. In *AAAI Computational Sustainability Track*.
- NGUYEN, T. H., FAVE, F. M. D., KAR, D., LAKSHMINARAYANAN, A. S., YADAV, A., TAMBE, M., AGMON, N., PLUMPTRE, A. J., DRICIRU, M., WANYAMA, F., AND RWETSIBA, A. 2015. Making the most of Our Regrets: Regret-based Solutions to Handle Payoff Uncertainty and Elicitation in Green Security Games. In *Conference on Decision and Game Theory for Security*.

- NGUYEN, T. H., SINHA, A., GHOLAMI, S., PLUMPTRE, A., JOPPA, L., TAMBE, M., DRICIRU, M., WANYAMA, F., RWETSIBA, A., CRITCHLOW, R., AND BEALE, C. 2016. CAPTURE: A new predictive anti-poaching tool for wildlife protection. In *15th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*.
- 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)*.
- PITA, J., JAIN, M., WESTERN, C., PORTWAY, C., TAMBE, M., ORDONEZ, F., KRAUS, S., AND PARUCHURI, P. 2008. Deployed ARMOR protection: The application of a game theoretic model for security at the Los Angeles International Airport. In *AAMAS*.
- PITA, J., JOHN, R., MAHESWARAN, R., TAMBE, M., AND KRAUS, S. 2012. A Robust Approach to Addressing Human Adversaries in Security Games. In *In ECAI*.
- QIAN, Y., ZHANG, C., KRISHNAMACHARI, B., AND TAMBE, M. 2016. Restless Poachers : Handling Exploration-Exploitation Tradeoffs in Security Domains. In *AAMAS'16*.
- SINHA, A., KAR, D., AND TAMBE, M. 2016. Learning Adversary Behavior in Security Games: A PAC Model Perspective. In *AAMAS'16*.
- TAMBE, M. 2011. *Security and Game Theory: Algorithms, Deployed Systems, Lessons Learned*. Cambridge University Press.
- YANG, R., FORD, B., TAMBE, M., AND LEMIEUX, A. 2014. Adaptive Resource Allocation for Wildlife Protection against Illegal Poachers. In *AAMAS*.
- YANG, R., ORDONEZ, F., AND TAMBE, M. 2012. Computing Optimal Strategy against Quantal Response in Security Games. In *AAMAS'12*.