Introduction to Green Security Games (Extended Abstract)

Fei Fang, Thanh Nguyen, Benjamin Ford, Nicole Sintov, Milind Tambe

University of Southern California, Los Angeles, CA, 90089

{feifang,thanhhng,benjamif,sintov,tambe}@usc.edu

1   Introduction

Conservation agencies around the world are tasked with protecting endangered wildlife from poaching. Despite their substantial efforts, however, species are continuing to be poached to critical status and, in some cases, extinction. In South Africa, rhino poaching has seen a recent escalation in frequency; while only 122 rhinos were poached in 2009, a record 1215 rhinos were poached in 2014 (approximately 1 rhino every eight hours)[the Rhino International, 2015]. To combat poaching, conservation agencies send well-trained rangers to patrol designated protected areas. However, these agencies have limited resources and are unable to provide 100% coverage to the entire area at all times. Thus, it is important that agencies make the most efficient use of their patrolling resources, and we introduce Green Security Games (GSGs) as a tool to aid agencies in designing effective patrols.

First introduced by [Von Stengel and Zamir, 2004] as a Leadership Game, Stackelberg Games have been applied in a variety of Security Game research (i.e., Stackelberg Security Games, or SSGs). In particular, the focus on randomization in Stackelberg Games lends itself to solving real-world security problems where defenders have limited resources, such as randomly allocating Federal Air Marshals to international flights [Tsai *et al.*, 2009]. However, the SSG model focuses on generating an optimal defender strategy against a single defender-attacker interaction (e.g., a single terrorist attack). For domains where attacks occur frequently, such as in wildlife conservation, another type of Security Game is needed that effectively models the repeated interactions between the defender and the attacker.

While still following the Leader-Follower paradigm of SSGs, GSGs have been developed as a way of applying Game Theory to assist wildlife conservation efforts, whether its to prevent illegal fishing [Haskell *et al.*, 2014], illegal logging [Johnson *et al.*, 2012], or wildlife poaching [Yang *et al.*, 2014]. GSGs are similar to SSGs except that, in GSGs, the game takes place over N rounds. In SSGs, once the attacker makes a decision, the game is over, but in GSGs, the attacker (e.g., the poacher) and defender have multiple rounds in which they can adapt to each other's choices in previous rounds. This multi-round feature of GSGs introduces some key research challenges that are being studied: (1) how can we incorporate the attacker's previous choices into our model of their behavior, in order to improve the defender's strategy, [Yang *et al.*, 2014; Kar *et al.*, 2015] and (2) how do we choose a strategy such that the long-term payoff (i.e., cumulative expected utility) is maximized [Fang *et al.*, 2015]? In addition to exploring these open research questions, we also discuss field tests of the Protection Assistant for Wildlife Security (PAWS) software in Uganda and Malaysia.

2   Green Security Games

In contrast to a Stackelberg Security Game (SSG), the interaction between the defender and attacker in a Green Security Game (GSG) takes places over many rounds. We describe the cyclical interaction as follows: (1) existing attack data is used to learn a behavior model for the attacker. (2) Based on this learned behavior model, the defender then calculates the best response. The best response takes the form of a mixed strategy - a probability distribution over feasible patrol routes. (3) The defender executes the mixed strategy; in the context of wildlife conservation, the execution phase can correspond to park rangers executing a patrol. (4) The attacker(s) choose a target(s) to attack, thus generating attack data; this phase can correspond to poachers placing wire snares in the park to poach wildlife. For the next round, this process begins again, except that the attack dataset now includes the attack data from the previous round.

Previous work in SSGs has studied the sub-optimal choices of humans in Security Games (referred to as bounded rationality) [Yang *et al.*, 2012; Nguyen *et al.*, 2013]. Using attack data from (single round) SSG human subject experiments, [Nguyen *et al.*, 2013] modeled humans' preferences for attacking targets via the Subjective Utility Quantal Response (SUQR) model. In reality, obtaining attack data for terrorist attacks is more difficult because attacks are relatively rare (in comparison to crime), and thus, it is more difficult to learn precise behavior models on real-world data. In the wildlife conservation domain, however, attacks happen much more frequently, and thus, this relative abundance of attack data enables more precise behavior models to be learned on real-world data. [Yang *et al.*, 2014] and [Kar *et al.*, 2015] both incorporate attack data from multiple rounds to improve the learned behavior models. While [Yang *et al.*, 2014] attempts to learn a distribution of heterogeneous attacker preferences in every round, [Kar *et al.*, 2015] demonstrates that it is better to use an adaptive utility model where the adversary's success or failures in each round will heavily influence his future

(a) PAWS patrol in QENP

(b) QENP ranger on patrol
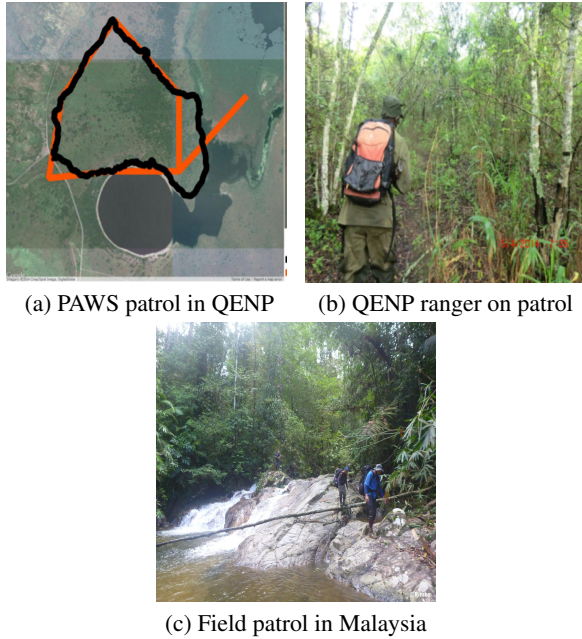

(c) Field patrol in Malaysia

Figure 1: Field trials for PAWS

round decisions.

As mentioned previously, the attackers in GSG domains (e.g., wildlife conservation) carry out frequent attacks, and, generally, attackers do not conduct extensive surveillance before performing an attack. As a result, the attackers' understanding of the defender strategy may not be up-to-date; the defender can exploit this delay through careful planning of mixed strategy changes throughout the rounds in the GSG. [Fang *et al.*, 2015] investigate how to design a sequence of defender strategies to maximize the long-term payoff for a multi-round GSG.

3 Field Trials

Originally developed in [Yang *et al.*, 2014], the Protection Assistant for Wildlife Security (PAWS) has successfully undergone field trials in Uganda's Queen Elizabeth National Park (QENP) and a protected area in Malaysia. For each field trial, the test area was converted into a grid, where each cell corresponded to a single target. Each target's attacker reward corresponded to the animal density in that cell. The defender's mixed strategy, in this context, corresponded to a set of patrol routes, where each route traveled through several targets in the grid. For the actual field test, these patrols were converted into a set of GPS waypoints. The park rangers then used a hand-held GPS to follow one of these waypoint sets as overall guidance for their patrol. An example patrol for QENP is shown in Figure 1a, where the orange line refers to the patrol generated by PAWS, and the black line is the patrol actually executed. Figure 1b shows a ranger during the execution of this patrol. Figure 1c shows a team of patrollers walking along a river during a field trial in Malaysia. Based off of feedback from these initial field trials, the development and

refinement of PAWS continues. Key new features of PAWS include taking into account various terrain information in designing patrol routes and dealing with payoff uncertainty in the game. When learning the parameters in behavioral models from previous data, terrain-related features will also be taken into consideration. The improved PAWS is currently planned for additional field trials in the Fall of 2015.

4 Acknowledgments

References

[Fang *et al.*, 2015] Fei Fang, Peter Stone, and Milind Tambe. When security games go green: Designing defender strategies to prevent poaching and illegal fishing. In *IJCAI*, 2015.

[Haskell *et al.*, 2014] William Haskell, Debarun Kar, Fei Fang, Milind Tambe, Sam Cheung, and Lt Elizabeth Denicola. Robust protection of fisheries with compass. In *IAAI*, 2014.

[Johnson *et al.*, 2012] Matthew P. Johnson, Fei Fang, , and Milind Tambe. Patrol strategies to maximize pristine forest area. In *AAAI*, 2012.

[Kar *et al.*, 2015] Debarun Kar, Fei Fang, Francesco Delle Fave, Nicole Sintov, and Milind Tambe. "a game of thrones": When human behavior models compete in repeated stackelberg security games. In *AAMAS*, 2015.

[Nguyen *et al.*, 2013] Thanh Hong Nguyen, Rong Yang, Amos Azaria, Sarit Kraus, and Milind Tambe. Analyzing the effectiveness of adversary modeling in security games. *In AAAI*, 2013.

[the Rhino International, 2015] Save the Rhino International. Poaching: The statistics, 2015. Accessed: 2015-07-05.

[Tsai *et al.*, 2009] Jason Tsai, Shyamsunder Rathi, Christopher Kiekintveld, Fernando Ordonez, and Milind Tambe. Iris - a tool for strategic security allocation in transportation networks. In *AAMAS*, 2009.

[Von Stengel and Zamir, 2004] Bernhard Von Stengel and Shmuel Zamir. Leadership with commitment to mixed strategies. 2004.

[Yang *et al.*, 2012] Rong Yang, Fernando Ordonez, and Milind Tambe. Computing optimal strategy against quantal response in security games. In *AAMAS*, 2012.

[Yang *et al.*, 2014] Rong Yang, Benjamin Ford, Milind Tambe, and Andrew Lemieux. Adaptive resource allocation for wildlife protection against illegal poachers. In *AAMAS*, pages 453–460, 2014.