Broken Signals in Security Games: Coordinating Patrollers and Sensors in the Real World[*]

Extended Abstract

Elizabeth Bondi
University of Southern California
Los Angeles, CA
bondi@usc.edu

Hoon Oh
Carnegie Mellon University
Pittsburgh, PA
hooh@andrew.cmu.edu

Haifeng Xu
Harvard University
Cambridge, MA
hx4ad@virginia.edu

Fei Fang
Carnegie Mellon University
Pittsburgh, PA
feifang@cmu.edu

Bistra Dilkina
University of Southern California
Los Angeles, CA
dilkina@usc.edu

Milind Tambe
University of Southern California
Los Angeles, CA
tambe@usc.edu

ABSTRACT

Mobile sensors, e.g., unmanned aerial vehicles (UAVs), are becoming increasingly important in security domains and can be used for tasks such as searching for poachers in conservation areas. Such mobile sensors augment human patrollers by assisting in surveillance and in signaling potentially deceptive information to adversaries, and their coordinated deployment could be modeled via the well-known security games framework. Unfortunately, real-world uncertainty in the sensor's detection of adversaries and adversaries' observation of the sensor's signals present major challenges in the sensors' use. This leads to significant detriments in security performance. We first discuss the current shortcomings in more detail, and then propose a novel game model that incorporates uncertainty with sensors. The defender strategy in this game model will consist of three interdependent stages: an allocation stage, a signaling stage, and a reaction stage.

KEYWORDS

security games; computational sustainability; uncertainty; sensors

ACM Reference Format:
Elizabeth Bondi, Hoon Oh, Haifeng Xu, Fei Fang, Bistra Dilkina, and Milind Tambe. 2019. Broken Signals in Security Games: Coordinating Patrollers and Sensors in the Real World. In *Proc. of the 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2019), Montreal, Canada, May 13–17, 2019,* IFAAMAS, 3 pages.

1 INTRODUCTION

In many real-world situations, there are not enough security resources, such as human patrollers, to protect all possible targets from attackers and prevent illegal activities. Security games have been used to model and solve strategic security resource allocation in these situations in the past decade for problems such as protecting airports, traffic enforcement, protecting elections, and protecting borders [3, 7, 9].

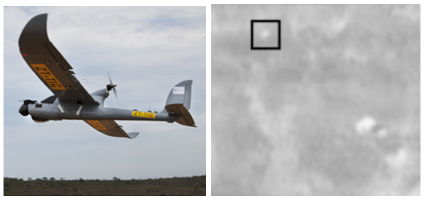

Figure 1: UAV monitoring area by capturing thermal infrared images of people (in black box).

Concurrently, mobile sensors such as unmanned aerial vehicles (UAVs or drones) have been introduced for security purposes with an increasing importance in domains such as traffic enforcement [8] and wildlife poaching prevention [5]. They are used to detect the actions of the adversaries, to assist human patrollers, and to signal potentially deceptive information to adversaries. The security game framework has been augmented and applied to the coordinated deployment of human patrollers and mobile sensors as well as strategic signaling [11].

Unfortunately, real-world circumstances inevitably involve uncertainty in both the sensors' detection of adversaries and adversaries' imperfect observation of sensors' signals, leading to challenges in successfully using sensors in security domains. As a key application involving both types of uncertainties, consider that to combat poaching, UAVs equipped with thermal infrared (heat-detecting) cameras are used to locate poachers at night when poaching typically occurs [1] and sometimes send warning signals to poachers through onboard lights for deterrence. Fig. 1 shows an example of a UAV used for this task, and a corresponding thermal infrared image. To assist the UAV crew in detecting poachers and animals in these videos automatically, a decision aid called SPOT was developed [2]. Although useful, detectors such as SPOT, [6], and [10] suffer from imperfect detection. Additionally, the presence of the UAV or signals may not be observed by the poacher, for example due to occlusions by trees.

Ignoring such uncertainties would result in significant detriments in security performance. Consider a sensor with a high false negative rate as an example. In this case, it could be beneficial for the human patroller to go and check a nearby location even if the

[*] Equal contribution by the first two authors.

Proc. of the 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2019), N. Agmon, M. E. Taylor, E. Elkind, M. Veloso (eds.), May 13–17, 2019, Montreal, Canada.

sensor in the location does not detect any adversary in order to confirm that there is no adversary there, rather than fully trusting the sensor. Fully trusting the sensors' capability of detecting adversaries leads to a wrong belief of the location of the adversary, and the efficiency of patrol can be even worse than not having any sensors. Similarly, when the uncertainty in the adversary's observation of the signal is high, the attacker may not be deterred even if the sensor sends out the signal indicating the presence of a patroller nearby. Not considering this uncertainty will lead to an overly optimistic estimation of the probability that the adversary will give up the attack. We aim to address this limitation and provide an efficient patrol plan that works in an environment with uncertainty.

2 MODEL

We consider a security game played between a defender and an attacker, who seeks to attack one target. The defender possesses k human patrollers and l sensors, and aims to protect N targets. Let $[N] = \{1, 2, ..., N\}$ denote the set of all targets. Let $U^{d/a}_{+/-}(i)$ be the defender/attacker (d/a) utility when the defender successfully protects/fails to protect ($+/-$) the attacked target i. By convention, we assume $U^d_+(i) \geq 0 > U^d_-(i)$ and $U^a_+(i) \leq 0 < U^a_-(i)$ for any $i \in [N]$. The underlying geographic structure of targets is captured by an undirected graph $G = (V, E)$. Mobile sensors cannot interdict an attack, though they can notify nearby patrollers to respond. If a target i is attacked, then we assume that a patroller at any neighboring target of i can move to i and successfully interdict the attack. Mobile sensors will send one of two signals – the quiet and warning signals to the attacker. The warning signal (lights on aboard the UAV) is used to warn the attacker off. We would like a model in which the adversary would stop the attack and run away upon seeing a warning signal.

2.1 Types of Uncertainty

Uncertainty is a crucial factor in automated applications of mobile sensors, yet has not been considered in previous work [11]. We consider two prominent types of uncertainties, motivated directly by the application of conservation drones. The first is *detection uncertainty*, i.e., the sensor could fail to detect a real attacker (false negative), or it could incorrectly classify something as an attacker (false positive) due to the inaccuracy of image recognition techniques [2, 6, 10]. We only consider false negatives in this work because the patrollers often have access to sensor videos, and the problem of false positives can be partly resolved by having a human in the loop. In contrast, verifying false negatives is harder due to various potential reasons, e.g., the poacher is simply hard to see as in Fig. 1. Also, deep learning algorithms can be tuned so there are more false negatives or positives depending on need.

The second type of uncertainty we consider is *observational uncertainty*, i.e., the attacker's imperfect detection of the sensors and the signals. When the attacker chooses one target to attack, he observes one of four possible signaling states at the target: (1) a patroller; (2) nothing; (3) a quiet signal (e.g., UAV only with no lights); (4) a warning signal. The existence of observational uncertainty means the true signaling state of the target may differ from the attacker's observation. Therefore, the attacker could observe nothing even when there is a warning signal.

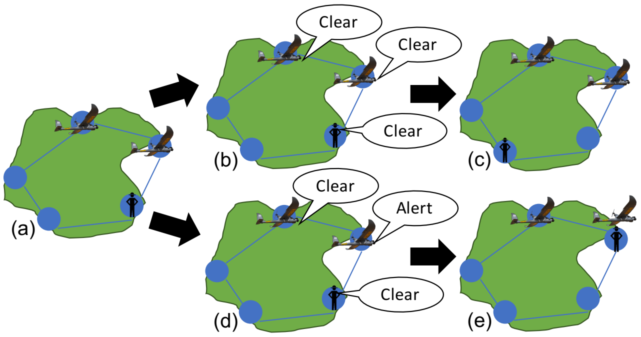

Figure 2: The reaction stage. (a) shows an initial allocation. In (b), no attacker is detected, so the patroller moves to the matched target (c). In (d), an attacker is detected, so the patroller responds (e).

2.2 Multistage Game Model

To facilitate incorporating uncertainty, we start with a novel three-stage game model: (1) *allocation stage* where (i) the defender places security resources (defender allocation stage), and (ii) the attacker chooses a target to attack based on the defender mixed strategy (attacker allocation stage); (2) *signaling stage* where the mobile sensors send signals based on detection (defender signaling stage); (3) *reaction stage* where (i) defender reacts to the sensor detection and relocates a human patroller to a nearby target (defender reaction stage), and (ii) the attacker chooses to deploy the attack or run away after the observation (attacker reaction stage). In practice, the defender signaling and reaction stages can happen simultaneously. In stage (3), the human defender moves from the original assigned location to a new location. If the attacker is detected by a sensor, nearby patroller(s) react by moving to the attacker's location to interdict. Unlike [11], if no sensors or patrollers detect the attacker, the defender still reacts by moving to another target (e.g., Fig. 2).

2.3 Using the Game Model

Uncertainty affects many aspects of the game model, such as the utilities, attacker behavior, and signaling and reaction strategy. Regarding attacker behavior, for instance, an attacker may stop attacking even when he observes a quiet signal because he is afraid of failing to observe a real warning signal. These different behaviors are not expected in the model thus far, so will need to be incorporated. In addition, we will need to modify the reaction stage to allow the patrollers to relocate to locations where sensors are placed rather than just empty locations. This provides an opportunity to check targets covered by drones but with false negative detection. In previous work [4, 11], if the sensor does not detect an attacker, the patroller does not do anything. However, due to the detection uncertainty, the defender has the incentive to *relocate* the human patrollers to nearby locations even if no detection is made.

3 ACKNOWLEDGEMENTS

This work was supported by Microsoft AI for Earth, NSF grants CCF-1522054 and IIS-1850477, and MURI W911NF-17-1-0370.

REFERENCES

[1] Air Shepherd. 2018. Air Shepherd: The Lindbergh Foundation. http://airshepherd.org. (2018). Accessed: 2018-08-01.

[2] Elizabeth Bondi, Fei Fang, Mark Hamilton, Debarun Kar, Donnabell Dmello, Jongmoo Choi, Robert Hannaford, Arvind Iyer, Lucas Joppa, Milind Tambe, and Ram Nevatia. 2018. SPOT Poachers in Action: Augmenting Conservation Drones with Automatic Detection in Near Real Time. In *IAAI*.

[3] Victor Bucarey, Carlos Casorrán, Óscar Figueroa, Karla Rosas, Hugo Navarrete, and Fernando Ordóñez. 2017. Building Real Stackelberg Security Games for Border Patrols. In *GameSec*.

[4] Xiaobo Ma, Yihui He, Xiapu Luo, Jianfeng Li, Mengchen Zhao, Bo An, and Xiaohong Guan. 2018. Camera Placement Based on Vehicle Traffic for Better City Security Surveillance. *IEEE Intelligent Systems* 33, 4 (2018), 49–61.

[5] Margarita Mulero-Pázmány, Roel Stolper, LD Van Essen, Juan J Negro, and Tyrell Sassen. 2014. Remotely piloted aircraft systems as a rhinoceros anti-poaching tool in Africa. *PloS one* 9, 1 (2014), e83873.

[6] Miguel A Olivares-Mendez, Changhong Fu, Philippe Ludivig, Tegawendé F Bissyandé, Somasundar Kannan, Maciej Zurad, Arun Annaiyan, Holger Voos, and Pascual Campoy. 2015. Towards an autonomous vision-based unmanned aerial system against wildlife poachers. *Sensors* (2015).

[7] Ariel Rosenfeld and Sarit Kraus. 2017. When security games hit traffic: Optimal traffic enforcement under one sided uncertainty. In *Proceedings of the 26th International Conference on Artificial Intelligence, IJCAI*.

[8] Ariel Rosenfeld, Oleg Maksimov, and Sarit Kraus. 2018. Optimal Cruiser-Drone Traffic Enforcement Under Energy Limitation.. In *IJCAI*.

[9] Milind Tambe. 2011. *Security and game theory: algorithms, deployed systems, lessons learned.* Cambridge University Press.

[10] Jan C van Gemert, Camiel R Verschoor, Pascal Mettes, Kitso Epema, Lian Pin Koh, Serge Wich, et al. 2014. Nature Conservation Drones for Automatic Localization and Counting of Animals. In *ECCV Workshops*.

[11] Haifeng Xu, Kai Wang, Phebe Vayanos, and Milind Tambe. 2018. Strategic coordination of human patrollers and mobile sensors with signaling for security games. In *AAAI*.