

To Signal or Not To Signal: Exploiting Uncertain Real-Time Information in Signaling Games for Security and Sustainability

Elizabeth Bondi,¹ Hoon Oh,² Haifeng Xu,³ Fei Fang,² Bistra Dilkina,⁴ Milind Tambe¹

¹Center for Research on Computation and Society, Harvard University, ebondi@g.harvard.edu, tambe@seas.harvard.edu

²Carnegie Mellon University, hooh@andrew.cmu.edu, feifang@cmu.edu

³University of Virginia, hx4ad@virginia.edu

⁴University of Southern California, dilkina@usc.edu

Abstract

Motivated by real-world deployment of drones for conservation, this paper advances the state-of-the-art in security games with signaling. The well-known defender-attacker security games framework can help in planning for such strategic deployments of sensors and human patrollers, and warning signals to ward off adversaries. However, we show that defenders can suffer significant losses when ignoring real-world uncertainties despite carefully planned security game strategies with signaling. In fact, defenders may perform worse than forgoing drones completely in this case. We address this shortcoming by proposing a novel game model that integrates signaling and sensor uncertainty; perhaps surprisingly, we show that defenders can still perform well via a signaling strategy that exploits uncertain real-time information. For example, even in the presence of uncertainty, the defender still has an informational advantage in knowing that she has or has not actually detected the attacker; and she can design a signaling scheme to “mislead” the attacker who is uncertain as to whether he has been detected. We provide theoretical results, a novel algorithm, scale-up techniques, and experimental results from simulation based on our ongoing deployment of a conservation drone system in South Africa.

1 Introduction

Conservation drones are currently deployed in South Africa to prevent wildlife poaching in national parks (Fig. 1). The drones, equipped with thermal infrared cameras, fly throughout the park at night when poaching typically occurs. Should anything suspicious be observed in the videos, nearby park rangers can prevent poaching, and a warning signal (e.g., drone lights) can be deployed for deterrence (Air Shepherd 2019). This requires a great deal of planning and coordination, as well as constant video monitoring. Rather than constant monitoring, we have recently worked with Air Shepherd to deploy an automatic detection system to locate humans and animals in these videos. Although an automatic detection system is helpful, its detections are uncertain. Potential false negative detections, in which the system fails to detect actual poachers, may lead to missed opportunities to




Figure 1: A drone and drone team member who are currently searching for poachers in a South African park at night.

deter or prevent poaching. This work is motivated by this ongoing, real-world deployment of drones for conservation.

Security challenges similar to those in conservation must be addressed around the world, from protecting large public gatherings such as marathons (Yin, An, and Jain 2014) to protecting cities. Security game models have been shown to be effective in many of these real-world domains (Tambe 2011; Bucarey et al. 2017). Recently, these models have begun to take into account real-time information, for example by using information from footprints when tracking poachers, or images from sensors (Wang et al. 2019; Basilio, De Nittis, and Gatti 2015). In particular, signaling based on real-time information, e.g., signaling to indicate the presence of law enforcement (Xu et al. 2018), has been introduced and established as a fundamental area of work.

Despite the rising interest in real-time information and signaling, unfortunately, security games literature has failed to consider uncertainty in sensing real-time information and signaling, hindering real-world applicability of the game models. Previously, only some types of uncertainty have been considered, such as uncertainty in the attacker’s observation of the defender’s strategy, attacker’s payoff values, or attacker’s rationality (Yin et al. 2011; Nguyen et al. 2014; Yang et al. 2011). However, there are fundamentally new insights when handling uncertainties w.r.t. real-time sensing and signaling, which we discuss at the end of this section.

We therefore focus on uncertainty in security games, in which real-time information comes from sensors that alert the defender when an attacker is detected and can also send warning signals to the attacker to deter the attack in real time. We consider both uncertainty in the sensor’s detection of adversaries (henceforth detection uncertainty) and uncer-

tainty in the adversaries’ observation of the sensor’s signals (henceforth observational uncertainty), and show that ignoring uncertainty hurts the defender’s expected utility. In our motivating domain of wildlife conservation with drones, automatic detection algorithms may make incorrect detections because humans in thermal infrared frames look similar to other objects (e.g., Fig. 1) and may even be occluded by other objects from the aerial perspective. The drone is also used to emit light to deter poachers, but such signals could sometimes be difficult for poachers to see in the wild, e.g., when trees block the sight.

We make contributions in (i) modeling, (ii) theoretical analysis, (iii) algorithmic design, and (iv) empirical evaluation. (i) We are the first to model uncertainty in sensing and signaling settings for security games. We introduce a novel reaction stage to the game model and construct a new signaling scheme, allowing the defender to mitigate the impact of uncertainty. In fact, this signaling scheme *exploits uncertain real-time information and the defender’s informational advantage*. For example, both the defender and attacker may know that there is detection uncertainty; however, the defender has an informational advantage in knowing that she has or has not actually detected the attacker, which she can exploit via a signaling scheme to “mislead” the attacker who is uncertain as to whether he has been detected. (ii) We provide several theoretical results on the impact of uncertainties, e.g., the loss due to ignoring observational uncertainty can be arbitrarily large, illustrating the need to handle uncertainty. (iii) To compute the defender’s optimal strategy given uncertainty, we develop a novel algorithm, GUARDSS, that not only uses six states to represent the type of protection a target has in a defender’s pure strategy but also uses a new matching technique in a branch-and-bound framework. (iv) We conduct extensive experiments on simulation based on our real-world deployment of a conservation drone system.

2 Related Work

Among the rich literature of Stackelberg security games (SSGs) (Tambe 2011; Bucarey et al. 2017), SSGs with real-time information have been studied recently. Some recent work in deception for cybersecurity, such as (Cooney et al. 2019; Thakoor et al. 2019), considers strategic signaling with boundedly rational attackers and attackers with different objectives and abilities, but no sensing is required to identify attackers; rather, the systems may interact with both normal and adversarial users. Some other work relies on human patrollers for real-time information (Zhang et al. 2019; Wang et al. 2019), and others rely on sensors that can notify the patroller when an opponent is detected (de Cote et al. 2013; Basilico, De Nittis, and Gatti 2015; De Nittis and Gatti 2018). Sensor placement (He et al. 2017) and drone patrolling (Rosenfeld, Maksimov, and Kraus 2018) have also been studied. Spatial and detection uncertainties in alarms are examined in (Basilico, De Nittis, and Gatti 2016; Basilico, De Nittis, and Gatti 2017). In all of these works, the sensors are only used to collect information, and do not actively and possibly deceptively disseminate information to the attacker. One work that does consider mobile sensors with detection and signaling capability is (Xu et al.

2018). However, it does not consider uncertainty in detection, which limits its capability in real-world settings. We add a new reaction stage and signaling strategy without detection, and compactly encode the different states that the defender resources can have at a target. Our model is therefore strictly more general than that in (Xu et al. 2018).

Our work is also related to multistage game models, e.g., defender-attacker-defender sequential games (DAD) (Brown et al. 2006; Alderson et al. 2011). In DAD, the defender and attacker take turns to commit to strategies, while in our game, the defender commits to a strategy of all stages at once. Extensive-form games (EFGs) also naturally model sequential games (Kroer et al. 2017; Brown and Sandholm 2017; Moravčík et al. 2017), and algorithms exist to efficiently solve the Stackelberg equilibrium in general two-player EFGs (Černý, Božanský, and Kiekintveld 2018; Cermak et al. 2016). However, GUARDSS is more scalable than the general EFG approach in this case (see Appendix¹).

3 Model

We consider a security game played between a defender and an attacker who seeks to attack one target. The defender has k human patrollers and l sensors to be allocated to targets in set $[N] = \{1, 2, \dots, N\}$. The sensor is the same as a drone in our motivation domain, and the attacker is the same as a poacher. 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 $G = (V, E)$ (e.g., Fig. 4). A patroller can move to any neighboring target and successfully interdict an attack at the target at no cost.

Sensors cannot interdict an attack, but they can notify nearby patrollers to respond and signal to deter the attacker. If the attacker is deterred by a signal (e.g., runs away), both players get utility 0. In practice, often one signal (σ_1 , e.g., illuminating the lights on the drone) is a warning that a patroller is nearby, while another signal (σ_0 , e.g., turning no lights on) indicates no patroller is nearby, although these may be used deceptively. Theoretically, (Kamenica and Gentzkow 2011) also showed two signals suffice (without uncertainty). We thus use two signals: σ_1 is a *strong signal* and σ_0 is a *weak signal*. When the attacker chooses one target to attack, he encounters one of four *signaling states*, based on the target either having a patroller, nothing, or a drone. The attacker may encounter: (1) a patroller and immediately get caught (state p); (2) nothing (state n); (3) a drone with signal σ_0 (state σ_0); (4) a drone with signal σ_1 (state σ_1). The attacker is caught immediately at state p, so there is no signal. Therefore, we omit p and let $\Omega = \{n, \sigma_0, \sigma_1\}$ be the set of signaling states.

3.1 Modeling Uncertainty

In this paper, we focus on two prominent uncertainties motivated directly by the use of conservation drones. The

¹bit.ly/aaai2020.signaluncertainty

first is the *detection uncertainty*, when there is a limitation in the sensor’s capability, e.g., a detection could be incorrect due to the inaccuracy of image detection techniques in the conservation domain (Bondi et al. 2020; 2018; Olivares-Mendez et al. 2015). We consider only false negative detection in this paper because patrollers often have access to sensor videos, so the problem of false positives can be partly resolved with a human in the loop. In contrast, verifying false negatives is harder, e.g., the attacker is easy to miss in the frame (Fig. 1) or is occluded. We therefore denote the false negative rate as γ for any sensor².

The second type of uncertainty we consider is the *observational uncertainty*, where the true signaling state of the target may differ from the attacker’s observation (e.g., a poacher may not be able to detect the drone’s signal). We use $\hat{\omega}$ to denote the attacker’s observed signaling state, and use ω to denote the true signaling state based on the defender signaling scheme. We introduce uncertainty matrix Π to capture observational uncertainty. The uncertainty matrix Π will contain the conditional probability $\Pr[\hat{\omega}|\omega]$ for all $\hat{\omega}, \omega \in \Omega$ to describe how likely the attacker will observe a signaling state $\hat{\omega}$ given the true signaling state is ω .

$$\Pi = \begin{bmatrix} \Pr[\hat{\omega} = n|n] & \Pr[\hat{\omega} = n|\sigma_0] & \Pr[\hat{\omega} = n|\sigma_1] \\ \Pr[\hat{\omega} = \sigma_0|n] & \Pr[\hat{\omega} = \sigma_0|\sigma_0] & \Pr[\hat{\omega} = \sigma_0|\sigma_1] \\ \Pr[\hat{\omega} = \sigma_1|n] & \Pr[\hat{\omega} = \sigma_1|\sigma_0] & \Pr[\hat{\omega} = \sigma_1|\sigma_1] \end{bmatrix}$$

Considering an arbitrary uncertainty matrix may unnecessarily complicate the problem, since some uncertainties never happen. We thus focus on a restricted class of uncertainty matrices that are natural in our domain.³ In our uncertainty model, we assume that a weak signal will never be observed as strong; moreover, n (the signaling state without any resource) will never be observed as strong or weak. As a result, the uncertainty matrix Π can be reduced to the following form, parameterized by κ, λ, μ , where $\kappa = \Pr[\hat{\omega} = n|\sigma_0]$, $\lambda = \Pr[\hat{\omega} = n|\sigma_1]$, $\mu = \Pr[\hat{\omega} = \sigma_0|\sigma_1]$:

$$\Pi_{\kappa\lambda\mu} = \begin{bmatrix} 1 & \kappa & \lambda \\ 0 & 1 - \kappa & \mu \\ 0 & 0 & 1 - \lambda - \mu \end{bmatrix}$$

As a result of this uncertainty, the attacker may not behave as expected. For example, if he knows that he has difficulty seeing the strong signal, he may decide to attack only when there is no drone, whereas typically we would expect him to attack on a weak signal. Therefore, let $\eta \in \{0, 1\}^3$ be the vector that depicts attacker behavior for each observation $\{n, \sigma_0, \sigma_1\} \in \Omega$, where 1 represents attacking, and 0 represents running away. So, $\eta = \mathbf{1}$ means an attacker will attack no matter what signaling state is observed, and $\eta = \mathbf{0}$ means an attacker will never attack.

3.2 Reaction Stage

Uncertainty motivates us to add an explicit reaction stage during which the defender can respond *or* re-allocate patrollers to check on extremely uncertain sensors or previously unprotected targets, for example. The timing of the

²False negative rate: $P(\text{no detection} | \text{poacher is present})$.

³Most results can be extended to general uncertainty matrices.

game is summarized in Fig. 2. In words, (i) the defender commits to a mixed strategy and then executes a pure strategy allocation; (ii) the attacker chooses a target to attack; (iii) the sensors detect the attacker with detection uncertainty; (iv) the sensors signal based on the signaling scheme; (v) *the defender re-allocates patrollers based on sensor detections and matching*; (vi) the attacker observes the signal with observational uncertainty; (vii) the attacker chooses to either continue the attack or run away. In (v), if a sensor detects the attacker, then nearby patroller(s) (if any) always go to that target, and the game ends; *or if no sensors or patrollers detect the attacker, the patroller moves to another target to check for the attacker*. The attacker reaction occurs after the defender reaction because the attacker reaction does not affect the defender reaction in the current model. In other words, there is no cost in reallocating the defender even if the attacker runs away, so the defender should begin moving right away.

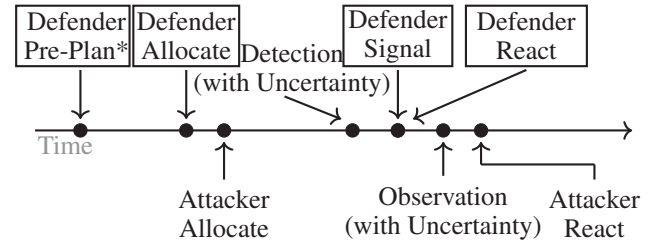


Figure 2: Game timing. Top and bottom are defender and attacker actions, respectively. *Defender fixes strategy offline.

3.3 Defender and Attacker Strategies

Defender Strategy: The strategy space consists of randomized resource allocation and re-allocation, and signaling. A deterministic resource allocation and re-allocation strategy (henceforth, a *defender pure strategy*) consists of allocating the patrollers to k targets, the sensors to l targets, and the neighboring target to which each patroller moves if no attackers are observed. Re-allocation can be equivalently thought of as matching each patroller’s original target to a neighboring target. A patroller goes to the matched target only if the attacker is not observed, and may respond to any nearby sensor detection, regardless of matching.

As a result of this rich structure, a pure strategy in the model needs to represent not only if the target is assigned a patroller (p), nothing (n), or a sensor (s), but also the allocation in neighboring targets. We compactly encode this pure strategy via 6 possible *allocation states* for each target. Let $\Theta = \{p, n+, n-, \bar{s}, s+, s-\}$ denote the set of all possible allocation states of an individual target. The target is assigned a patroller (p), nothing (n), or a sensor (s). If there is no patroller near a sensor (\bar{s}), then no one can respond to the sensor’s detection. If there is a nearby patroller, the target is either matched ($n+, s+$) or not matched ($n-, s-$). Therefore, each target is in one of the allocation states in Table 1. For example, $n+$ is the state of a target which was not allocated a patroller or sensor, but in the reaction stage has a

patroller from a neighboring target (“patroller matched”).

	Covered By:	Near Patroller?	Patroller Matched?	Protected Overall?
p	Patroller	N/A	N/A	Yes
n+	Nothing	Yes	Yes	Yes
n-	Nothing	N/A	No	No
\bar{s}	Sensor	No	N/A	No
s-	Sensor	Yes	No	Yes*
s+	Sensor	Yes	Yes	Yes

Table 1: Allocation State, *protected if sensor detects

Given Θ , a defender pure strategy can be compactly represented with an allocation state vector $\mathbf{e} \in \Theta^N$, in which $e_i \in \Theta$ denote the allocation state of a target $i \in [N]$. Let $\mathcal{E} \subseteq \Theta^N$ be the set of feasible allocation state vectors that corresponds to defender pure strategies. Note that not all vectors in Θ^N correspond to a feasible defender strategy due to the limited number of patrollers and sensors. A *defender mixed strategy* is thus a distribution over \mathcal{E} and can be described by $\{q_e\}_{e \in \mathcal{E}}$ where q_e is the probability of playing pure strategy $e \in \mathcal{E}$. Similarly, a defender mixed strategy can also be compactly represented by a marginal probability vector x , where x_i^θ represents the marginal probability that target i is in the allocation state $\theta \in \Theta$. This is similar to the coverage vector used in basic SSGs with schedules (Jain et al. 2010). We introduce the constraints that x needs to satisfy to be a valid mixed strategy in Section 5.

The defender also deploys a signaling process w.r.t. each target i . The defender’s signaling strategy can be specified by probabilities ψ_i^{s-} , ψ_i^{s+} , and $\psi_i^{\bar{s}}$. ψ_i^{s-} is the joint probability of allocation state s- and sending signal σ_0 together conditioned on the sensor detecting an attacker, i.e., $\Pr[s- \wedge \sigma_0 | \text{detected}]$. To be a valid signaling strategy, $\psi_i^{s-} \in [0, x_i^{s-}]$. Note that $x_i^{s-} - \psi_i^{s-}$ will be the joint probability of realized state s- and sending signal σ_1 , together conditioned on detection. The conditional probability of sending σ_0 given the target is in state s- and it is detected is ψ_i^{s-}/x_i^{s-} . We use the joint probability instead of the conditional probability as it results in linear terms for the optimal defender strategy. *Because of detection uncertainty, we add the option to signal without detecting the attacker.* Let $\varphi_i^\theta \in [0, x_i^\theta]$ be the joint probability of allocation state θ and sending signal σ_0 conditioned on the sensor not detecting an attacker, for all $\theta \in \{\bar{s}, s-, s+\}$. We use χ to denote the allocation, reaction, and signaling scheme, or *defender’s deployment strategy*: $\chi = (x, \psi, \varphi)$.

Attacker Strategy: Recall the attacker has the allocation and reaction stages. In the allocation stage, the attacker chooses a target to attack based on the defender deployment strategy χ . He will be caught if the target is at state p. When the attacker is not caught, he may observe any of the signaling states $\hat{\omega} \in \Omega$. Based on his observation, the attacker then has a choice in the reaction stage to run away or continue the attack. The attacker knows the defender mixed strategy χ when choosing a target to attack, and he can observe the realization of the target (with uncertainty) when choosing to attack or run away. Since this is a Stackelberg game and

the defender commits to allocation and signaling schemes, it suffices to consider only the attacker’s pure responses.

4 Why Do We Need to Handle Uncertainty

In this section, we prove several theoretical properties regarding how uncertainties affect the defender’s optimal strategy and utility. All formal proofs are deferred to the Appendix. Let $\chi^*(\gamma, \Pi)$ be the optimal allocation under detection uncertainty of γ and observational uncertainty Π . Let $\text{DefEU}(\chi, \gamma, \Pi)$ be the defender expected utility when the actual uncertainties are γ, Π and the defender’s deployment is χ . Let $\Pi_0 = \mathbf{I}$ denote no observational uncertainty. We assume in Propositions 1 and 2 and Theorem 1 that $\Pi = \Pi_0$ and analyze detection uncertainty, so omit for conciseness. We first show the loss due to ignoring detection uncertainty.

Proposition 1. *Let $\chi_0^* = \chi^*(0)$ be the defender optimal deployment when no uncertainties exist. There exist instances where $\text{DefEU}(\chi_0^*, \gamma) < \text{DefEU}(\chi^*(\gamma), \gamma)$ for some γ .*

In fact, $\text{DefEU}(\chi^*(\gamma), \gamma) - \text{DefEU}(\chi_0^*, \gamma) \geq \gamma \cdot \max_{i \in [N]} |U_-^d(i)|$

for some instance. If we ignore γ , we do not signal when we do not detect an attacker. Furthermore, the defender would never match a patroller to a target with a sensor (s+) in χ_0^* . Thus, if we ignore uncertainty, there can be a steep penalty; in contrast, with the optimal strategy considering uncertainty, if the false negative rate is high, we may match a patroller to a target to confirm the presence of an attacker. Given the attacker’s knowledge of the defender mixed strategy, the attacker is therefore more likely to run away.

Our next result (Theorem 1) shows that the defender expected utility is non-increasing as detection uncertainty γ increases. As a byproduct of the proof for Theorem 1, we also show that the optimal solution may change as detection uncertainty changes. This illustrates the necessity of an algorithm for dealing with detection uncertainties.

Theorem 1. *$\text{DefEU}(\chi^*(\gamma), \gamma) \geq \text{DefEU}(\chi^*(\gamma'), \gamma')$ for any $\gamma' > \gamma$ in any problem instance.*

Proposition 2. *$\chi^*(\gamma)$ differs from $\chi^*(\gamma')$ for any $\gamma' > \gamma$ when x_t^{s-} is nonzero for $\chi^*(\gamma')$, where target t is the attacker best responding target in $\chi^*(\gamma')$.*

The intuition underlying the proof of Theorem 1 is that if we have a drone with a low false negative rate, then we can simulate a drone with a high false negative rate by ignoring some of its detections. The optimal solution for drones with a low false negative rate cannot be worse than that for drones with a high false negative rate.

We now show several results for observational uncertainty. First, we show that the loss due to observational uncertainty can be arbitrarily large.

Proposition 3. *There exists Π such that the loss due to ignoring observational uncertainty is arbitrarily large. In other words, $\text{DefEU}(\chi^*(\gamma_0, \Pi), \gamma_0, \Pi) - \text{DefEU}(\chi^*(\gamma_0, \Pi_0), \gamma_0, \Pi) > M, \forall M > 0$.*

The original signaling strategy tries to ensure the attacker only attacks when he observes the weak signal, σ_0 , or nothing, n. However, with observational uncertainty, this may not be true because the true signal may be σ_1 , but the attacker

may have observed it mistakenly as σ_0 . Therefore, we need to enforce different attacker behaviors in order to obtain a better solution quality.

Now, we examine the attacker's behavior given a fixed deployment χ as observational uncertainty changes. Let (t, η) represent an attacker strategy of attacking target t and behaving according to η . Theorems 2 and 3 show that if we do not consider observational uncertainty, then the attacker behavior is more likely to converge to always attacking ($\eta = 1$) as observational uncertainty increases, where higher observational uncertainty means the attacker cannot distinguish between signaling states. Theorems 2 and 3 show that a deployment χ that does not consider observational uncertainty is more likely to result in this worst-case behavior of $\eta = 1$.

Theorem 2. *For any fixed deployment χ , if the attacker's best response is $(t, \mathbf{0})$ or $(t, \mathbf{1})$ at the Stackelberg equilibrium with Π_0 , then it stays as an equilibrium for any Π' .*

Note that $\eta = \mathbf{0}$ and $\eta = \mathbf{1}$ result in an action that is independent of the attacker's observation. Thus, no matter what the attacker observes, the attacker can obtain the same utility with $\eta = \mathbf{0}$ or $\eta = \mathbf{1}$. It's only left to show that the attacker cannot get strictly better utility in Π' with a different attacker behavior. Intuitively, Π_0 implies a perfect observation, thus the attacker cannot get better utility than the perfect observation. So, if $(t, \mathbf{1})$ or $(t, \mathbf{0})$ is a Stackelberg equilibrium, the defender can safely deploy the same strategy for any uncertainty matrix Π' , without any loss in her expected utility.

Even if $(t, \mathbf{0})$ or $(t, \mathbf{1})$ is not a best response with Π_0 , $(t, \mathbf{1})$ may still be a best response at high levels of uncertainty. First, we say a target t is a *weak-signal-attack target* if $\text{AttEU}(\sigma_0) \geq 0$ at t . Note that if $\text{AttEU}(\sigma_0) \geq 0$, then the attacker will either always attack at $\hat{\omega} = \sigma_0$, or is indifferent between attacking and running away. We say χ is a *weak-signal-attack deployment* if all targets are weak-signal-attack targets.

Theorem 3. *If $(t, \mathbf{1})$ is a best response for $\Pi_{\kappa\lambda\mu}$ and χ is a weak-signal-attack deployment, then $(t, \mathbf{1})$ is a best response for $\Pi_{\kappa'\lambda'\mu'}$ and χ for all $\kappa' \geq \kappa$, $\lambda' \geq \lambda$, $\mu' \geq \mu$.*

In our model of observational uncertainty, more uncertainty means that the attacker sees a weak signal more often. Further, the attacker always attacks when he observes a weak signal. Thus, if the attacker is always attacking with less uncertainty, he will only attack more often with more uncertainty. However, in order to obtain predictable attacker behavior, we need to show that a weak-signal-attack deployment always exists as an optimal solution. In other words, Theorem 3 holds if there is weak-signal-attack deployment, so we now have to show that such a deployment exists.

Proposition 4. *There always exists an optimal solution that is a weak-signal-attack deployment with Π_0 .*

The intuition behind the proof is that we can always decrease the probability of a weak signal such that we either do not send a weak signal, or the attacker attacks when he observes a weak signal. This holds optimally because when observational uncertainty is Π_0 , signals are interchangeable. To summarize, if the attacker behavior is $\mathbf{0}$ or $\mathbf{1}$, then the attacker behavior is independent of observational uncertainty. We may see this behavior emerge as uncertainty increases.

5 How to Handle Uncertainty

We provide a solution approach based on the well-known multiple LPs approach from (Conitzer and Sandholm 2006). In particular, for each target $t \in [N]$, we compute the optimal defender strategy given that the attacker's best response is t . Then, the optimal defender strategy is the mixed strategy that leads to the maximum defender expected utility among all $t \in [N]$. The problem is NP-hard without uncertainty (Xu et al. 2018), thus our ultimate goal is to develop an efficient algorithm to solve the problem. For expository purposes, we first focus on presenting the LP for detection uncertainty.

5.1 Detection Uncertainty

Using notation from Section 3.3, we first formulate each player's utility function by breaking it into three parts according to signaling states: 1) no sensor is allocated (states $n(+/-)$ and p , which we denote by $-s$); 2) sensor is allocated and sends σ_0 ; and 3) sensor is allocated and sends σ_1 .

1. $U_{-s}^{d/a}(i) = x_i^p \cdot U_+^{d/a}(i) + x_i^{n+} \cdot U_+^{d/a}(i) + x_i^{n-} \cdot U_-^{d/a}(i)$ is the expected defender/attacker utility of target i being attacked over states when i has no sensor ($p, n+, n-$).
2. $U_{\sigma_0}^{d/a}(i) = (1 - \gamma) \cdot [\psi_i^{s+} \cdot U_+^{d/a}(i) + \psi_i^{s-} \cdot U_+^{d/a}(i) + \psi_i^{\bar{s}} \cdot U_-^{d/a}(i)] + \gamma \cdot [\varphi_i^{s+} \cdot U_+^{d/a}(i) + \varphi_i^{s-} \cdot U_-^{d/a}(i) + \varphi_i^{\bar{s}} \cdot U_-^{d/a}(i)]$ is the defender/attacker expected utility when the attacker attacks target i and the defender signals σ_0 .
3. $U_{\sigma_1}^{d/a}(i) = (1 - \gamma) \cdot [(x_i^{s+} - \psi_i^{s+}) \cdot U_+^{d/a}(i) + (x_i^{s-} - \psi_i^{s-}) \cdot U_+^{d/a}(i) + (x_i^{\bar{s}} - \psi_i^{\bar{s}}) \cdot U_-^{d/a}(i)] + \gamma \cdot [(x_i^{s+} - \varphi_i^{s+}) \cdot U_+^{d/a}(i) + (x_i^{s-} - \varphi_i^{s-}) \cdot U_-^{d/a}(i) + (x_i^{\bar{s}} - \varphi_i^{\bar{s}}) \cdot U_-^{d/a}(i)]$

In words, 2) and 3) are the sum of expected utility on a detection and the sum of expected utility on no detection. In 3), in the no detection case, the defender exploits information asymmetry in signaling σ_1 . In particular, the defender knows that there is no detection, but in sending σ_1 to indicate a detection, relies on the uncertainty the attacker faces in determining if there was a detection. We are now ready to describe an (exponentially-large) linear program (LP) formulation for computing the optimal defender strategy assuming best attacker response t (not (t, η) since only detection uncertainty):

$$\max_{x, q, \psi, \varphi} U_{-s}^d(t) + U_{\sigma_0}^d(t) \quad (1)$$

$$\text{s.t.} \quad \sum_{e \in \mathcal{E}: e_i = \theta} q_e = x_i^\theta \quad \forall \theta \in \Theta, \forall i \in [N] \quad (2)$$

$$\sum_{e \in \mathcal{E}} q_e = 1 \quad (3)$$

$$q_e \geq 0 \quad \forall e \in \mathcal{E} \quad (4)$$

$$U_{\sigma_0}^a(i) \geq 0 \quad \forall i \neq t \quad (5)$$

$$U_{\sigma_1}^a(i) \leq 0 \quad \forall i \neq t \quad (6)$$

$$U_{-s}^a(t) + U_{\sigma_0}^a(t) \geq U_{-s}^a(i) + U_{\sigma_0}^a(i) \quad \forall i \neq t \quad (7)$$

$$0 \leq \psi_i^\theta \leq x_i^\theta \quad \forall \theta \in \{\bar{s}, s-, s+\}, \forall i \in [N] \quad (8)$$

$$0 \leq \varphi_i^\theta \leq x_i^\theta \quad \forall \theta \in \{\bar{s}, s-, s+\}, \forall i \in [N] \quad (9)$$

The objective function (1) maximizes defender expected utility. Since the attacker is running away when he observes

$\sigma_1, U_{\sigma_1}^d = 0$. Constraints (2)-(4) enforce that the randomized resource allocation is feasible (\mathcal{E} has exponential number of elements); (5)-(6) guarantee that σ_1, σ_0 result in the attacker best responses of running away and attacking⁴; (7) ensures the attacker expected utility at target t is bigger than at any other target i , thus t is attacker's best response; (8)-(9) ensure a feasible signaling scheme.

5.2 Acceleration via Branch and Price

We now describe the branch-and-price solution framework, which can be used for both uncertainty scenarios. There are two main challenges in efficiently solving the LP (1)-(9). First, the total number of possible q_e is $O(6^N)$. Second, we will need to solve N LPs (for each $t \in [N]$). Solving many of these large LPs is a significant barrier for scaling up. We therefore introduce Games with Uncertainty And Response to Detection with Signaling Solver (GUARDSS), which employs the branch-and-price framework. This framework is well-known for solving large-scale optimization programs, but the main challenges of applying this framework are to (1) design the efficient subroutine called the *slave problem* for solving each LP, and to (2) carefully design an upper bound for pruning LPs.

First, for one LP w.r.t. a specific t , to address the issue of the exponential size of set \mathcal{E} , we adopt the column generation technique. At a high level, we start by solving the LP for a small subset $\mathcal{E}' \subset \mathcal{E}$, and then search for a pure strategy $e \in \mathcal{E} \setminus \mathcal{E}'$ such that adding e to \mathcal{E}' improves the optimal objective value strictly. This procedure continues until convergence, i.e., no objective value improvement. The key component in this technique is an algorithm to search for the new pure strategy, which is a specially-crafted problem derived from LP duality and referred to as the *slave problem*. **Slave Problem:** Given different weights $\alpha_i^\theta \in \mathbb{R}$ for $\theta \in \Theta$, for each target i , solve the *weight maximization problem*:

$$\max_{e \in \mathcal{E}} \sum_{\theta \in \Theta} \sum_{i: e_i = \theta} \alpha_i^\theta \quad (10)$$

Note that $\{\alpha_i^\theta\}_{\theta \in \Theta}$ are the optimal dual variables for the previous LP constraint (2). We want to solve this without enumerating all of the elements in \mathcal{E} . Despite the added complexity compared to classic SSGs, in this section, we compactly represent this slave problem as a mixed integer linear program (MILP). To formulate the MILP, we introduce six binary vectors $\mathbf{v}^p, \mathbf{v}^{n+}, \mathbf{v}^{n-}, \mathbf{v}^s, \mathbf{v}^{s-}, \mathbf{v}^{s+} \in \{0, 1\}^N$ to encode for each target whether it is in each allocation state. For example, target i is at allocation state \bar{s} if and only if $v_i^{\bar{s}} = 1$. The main challenge then is to properly set up linear (in)equalities of these vectors to precisely capture their constraints and relations. The capacity for each resource type results in two constraints (number of patrollers and sensors):

$$\sum_{i \in [N]} v_i^p \leq k \quad (11)$$

$$\sum_{i \in [N]} (v_i^s + v_i^{s-} + v_i^{s+}) \leq l \quad (12)$$

Moreover, each target must be at one of these states:

$$v_i^p + v_i^{n-} + v_i^{n+} + v_i^s + v_i^{s-} + v_i^{s+} = 1 \quad \forall i \in [N] \quad (13)$$

⁴Although we minimize this behavior, we still model it.

Due to the reaction stage, we have to add constraints to specify (a) which targets have a patroller at a neighboring target; (b) which patroller goes to which nearby target if both sensors and patrollers do not detect the attacker. For (a), the non-zero entries of $A \cdot \mathbf{v}^p$ specify the targets with a patroller nearby, where A is the adjacency matrix of the underlying graph. Since three vectors encode the states requiring a nearby patroller, we have this constraint:

$$A \cdot \mathbf{v}^p \geq \mathbf{v}^{n+} + \mathbf{v}^{s-} + \mathbf{v}^{s+} \quad (14)$$

We ensure that a vertex with a patroller nearby cannot be \mathbf{v}^s :

$$A \cdot \mathbf{v}^p \leq \mathbf{v}^p + \mathbf{v}^{n+} + \mathbf{v}^{n-} + \mathbf{v}^{s-} + \mathbf{v}^{s+} \quad (15)$$

Constraint (b) means that patrollers must be “re-matched” to new vertices in the reaction stage. Specifically, targets in states p, n+, s+ must form a matching. To enforce this constraint, let G' be the directed version of G , i.e. for all $(i, j) \in E$ we have $(i, j), (j, i) \in E'$. We further introduce edge variables $y_{(i, j)} \in \{0, 1\}$ indicating whether the directed edge (i, j) is in the matching or not. The matching constraint can be expressed by the following linear constraints:

$$\sum_{(i, j) \in E': j \in [N]} y_{(i, j)} = v_i^p \quad \forall i \in [N] \quad (16)$$

$$v_j^{n+} + v_j^{s+} \geq y_{(i, j)} \quad \forall (i, j) \in E' \quad (17)$$

The resulting MILP for the slave problem is as follows.

$$\max_{\mathbf{v}, y} \sum_{\theta} \sum_i v_i^\theta \alpha_i^\theta \quad (18)$$

$$\text{s.t.} \quad (11) - (17) \quad (18)$$

$$\mathbf{v}^\theta \in \{0, 1\}^N \quad \forall \theta \in \Theta \quad (19)$$

$$y_{(i, j)} \in \{0, 1\} \quad \forall (i, j) \in E' \quad (20)$$

Second, to avoid solving LPs for all different targets $t \in [N]$, we use the branch and bound technique which finds an upper bound for each LP for pruning. The natural approach for finding an upper bound is to solve a relaxed LP corresponding to the original LP — in our case, essentially relax the original LP into its marginal space. As the set \mathcal{E} is exponentially large, we relax variables and constraints corresponding to \mathcal{E} in our LP. Concretely, we relax (2) - (4) into a polynomial number of variables and constraints. These variables and constraints are (18) - (20) with \mathbf{v}^θ replaced by \mathbf{x}^θ . We first use the relaxed LP to efficiently compute an upper bound for each LP. After solving each relaxed LP exactly, we solve original LPs chosen according to some heuristic order (typically the descending order of the relaxed optimal objective) using the column generation techniques, and we can safely prune out those LPs whose optimal relaxed value is less than the current largest achievable objective value. This process continues until no LP is left to solve, in which case the current largest objective value is optimal.

5.3 Detection and Observational Uncertainty

Finally, we briefly discuss the case with both uncertainties, as it can be solved in a similar way. Constraints (2)-(4) and (8)-(9) are the same. However, the remaining constraints must now account for attacker behavior, η . For example, the

utility functions $U_{\sigma_0}^{d/a}$ and $U_{\sigma_1}^{d/a}$ must change to incorporate attacker behaviors, and the objective function becomes that in (21) since the attacker may not run away when he observes σ_1 in the presence of observational uncertainty. Also, we add a constraint to ensure the attacker utilities are aligned with the attacker behavior $\eta \in \{0, 1\}^3$. These are primarily notational changes. We therefore provide the full LP for this case in the Appendix.

$$\max_{x, q, \psi, \varphi} U_{\sigma_0}^d(t) + U_{\sigma_1}^d(t) + U_{\sigma_0}^a(t) \quad (21)$$

6 Experiments

We generate random Watts-Strogatz graphs, which have small-world properties to describe more complex environments, such as roads connecting far-away nodes. For all tests, we average over 20 random graphs and include p -values. Utilities are randomly generated with a maximum absolute value of 1090 and based on the idea that the losses from undetected attacks are higher than the utility of catching adversaries (similar to (Xu et al. 2015)). This is realistic to the situation of preventing poaching, as animals are worth more for ecotourism than for sale on the black market as discussed in the Appendix. Additionally, we see that if we test on a set of utilities that is slightly different from the original input, the defender’s utility does not vary greatly. Fig. 3a-3b show timing tests run on a cluster with Intel(R) Xeon(R) CPU E5-2683 v4 @ 2.1 GHz with at most 16 GB RAM. We set the number of patrollers to be $k = \sqrt{N/2}$ and the number of drones to be $l = 2N/3 - k$. As shown, the full LP scales up to graphs of $N = 14$ only and exceeds the cutoff time limit of 3600s for all $N = 16$ graphs. Branch and price scales up to $N = 80$ and runs out of time for larger games, and a warm-up enhancement that greedily select an initial set of \mathcal{E} further improves scalability and solves 13/40 graphs within cutoff time at $N = 90$ and $N = 100$. This is sufficient for middle-scale real-world problems, with further scalability being an interesting direction for future work. The heuristics provide the same solution as the full LP in most of the instances tested.

Next, we show the loss due to ignoring uncertainty empirically. In Figs. 3c-3d we compare $\text{DefEU}(\chi^*(\gamma, \Pi), \gamma, \Pi)$ computed by GUARDSS and $\text{DefEU}(\chi^*(0, \Pi_0), \gamma, \Pi)$, the defender expected utility when ignoring uncertainty for graphs with $N = 10$, $k = 1$, $l = 3$. We consider only one type of uncertainty at a time (e.g., $\gamma = 0$ when varying observational uncertainty). For detection uncertainty, GUARDSS’s defender expected utility only decreases by 12%, whereas ignoring uncertainty decreases by 210% when γ varies from 0 to 0.9 ($p \leq 1.421e-03$ for $\gamma \geq 0.2$ in Fig. 3c)⁵. Some initial analysis shows that it is robust in most of the cases when we slightly under- or overestimate γ (e.g., the differences in defender expected utility are typically within 5-6% when the estimate of gamma is off by 0.1 or 0.2), but further investigation on dealing with such uncertainty over uncertainty would be an interesting direction for future work. For observational uncertainty, GUARDSS’s defender

expected utility only decreases by 1%, whereas ignoring uncertainty decreases by 18% as the observational uncertainty, parameterized by κ ($\lambda = \frac{\kappa}{2}$, and $\mu = \frac{\kappa}{2}$) varies from 0 to 0.9 ($p \leq 0.058$ for $\kappa \geq 0.4$ in Fig. 3d).

We also observe that when ignoring detection uncertainty, the attacker’s best response is typically a target with a sensor, which implies that the attacker is taking advantage of the defender’s ignorance of uncertainty. In fact, there is a statistically significant ($p = 1.52e-08$) difference in the mean probability of a sensor at the attacker’s best response when ignoring uncertainty (0.68) versus GUARDSS (0.19).

How does the defender avoid these challenges and achieve such a small performance drop with GUARDSS when facing uncertainty? Statistics of the resulting defender strategy as well as Fig. 3e indicate that the defender *exploits the uncertain real-time information and the information asymmetry*, including (a) frequently but not always sending patrollers to check important targets when there is no detection; (b) sending strong signals more frequently than the probability that the patroller will visit the target (either due to response to detection or planned reallocation in the case of no detection), leveraging the informational advantage in which the attacker does not know whether he is detected or whether a patroller is matched; (c) using different signaling schemes with and without detection, leveraging the information advantage that the attacker does not know whether he is detected. In the GUARDSS strategies in Figs. 3c-3d, the mean probability of the attacker’s best response target being at state $s-$ (with sensor but without a matched patroller) is 0.04, versus 0.43 when ignoring uncertainty ($p = 2.70e-09$), indicating point (a). If we call the strong signal sent when there is no detection a *fake signal*, Fig. 3e shows that the probability of the strong signal an attacker observes is a fake signal is non-zero and increases in a non-linear fashion, indicating points (b) and (c). Also, note that the strong signal is used with nonzero probability on average on targets with a nonzero probability of having a drone present.

Despite considering uncertainty, sensors may be less valuable at a high level of uncertainty. In Fig. 3f, the defender expected utility is influenced by the number of drones and uncertainty in size $N = 15$ graphs. In Fig. 3g, drones are better than an extra patroller at $\gamma = 0.3$ ($p \leq 6.661e-02$), but at $\gamma = 0.8$, patrollers are better than drones ($p \leq 1.727e-07$).

7 Conservation Drones

We have deployed a drone in South Africa, equipped with a thermal camera and detection system (Air Shepherd 2019). A photo of the drone team in South Africa currently is included in Fig. 1 (center). To ease the challenges faced by these operators in coordination of drones with imperfect sensors and patrollers, we apply GUARDSS and show that it provides positive results in simulation to support future potential deployment. To facilitate the most realistic simulation possible, we utilize example poaching hotspots in a real park. We cannot provide the exact coordinates in order to protect wildlife, but we selected points based on geospatial features, and selected utilities to reflect the fact that the reward and penalty of the attackers are impacted by animal presence, price, and distance to several park features used

⁵% change once normalized by largest defender/attacker utility.

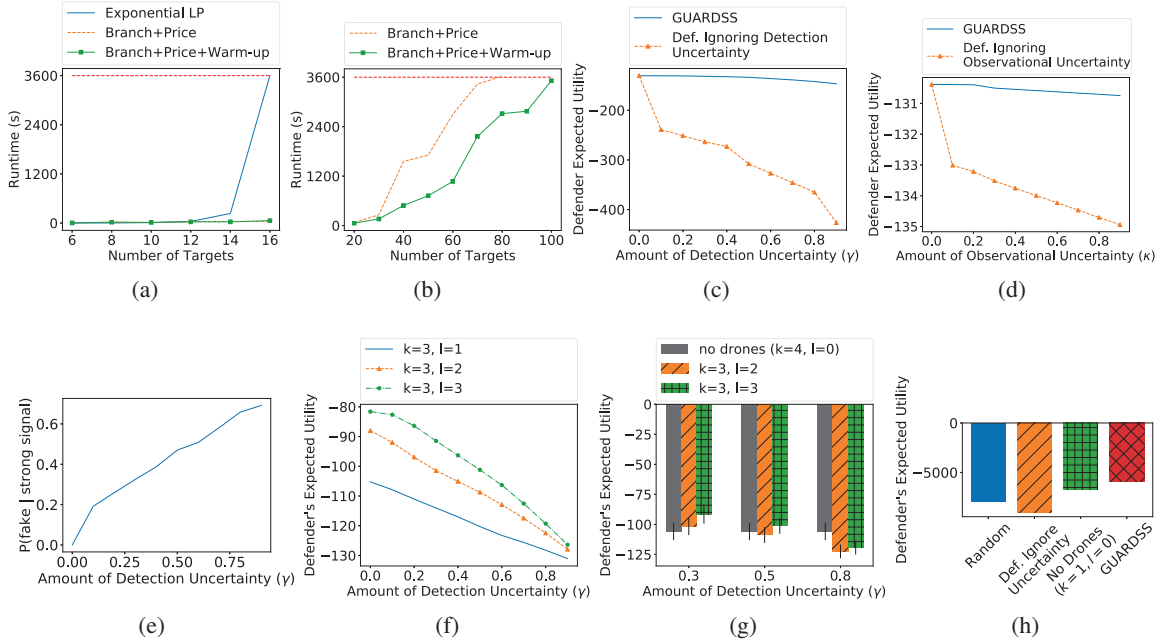


Figure 3: Experimental results. Figs. 3a-3b compare multiple LPs approach (Exponential LP) with GUARDSS branch-and-price and heuristic method. Figs. 3c-3d show defender expected utility when amount of detection uncertainty γ and observational uncertainty vary. Defender expected utility decreases much more when uncertainties are ignored. Fig. 3e shows the informational advantage of the defender as uncertainty increases. Figs. 3f-3g show that in the presence of a high false negative rate, extra patrollers may be more useful than drones. Fig. 3h contains the results from the case study, where GUARDSS performs best.

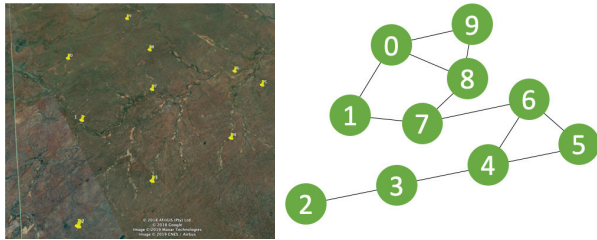


Figure 4: A park in Google Maps with potential poaching hotspots and the resulting graph (edges for $< 5\text{ km}</math>).$

in (Gholami et al. 2018). The targets are shown in Fig. 4 (left). Any targets within 5 km are connected via edges in the graph, as park rangers could cover 5km for response. The resulting graph is shown in Fig. 4 (right). The utilities are included in the Appendix along with further details. For the simulation, we use 3 drones and 1 patroller. In the “no drones” scenario only, there are 0 drones and 1 patroller. We use $\gamma = 0.3$ for detection uncertainty and no observational uncertainty (see the Appendix for results with other γ). These details are directly input to GUARDSS, and then a mixed strategy is determined to cover the park. Fig. 3h shows the defender expected utility in this park using GUARDSS with and without uncertainty, and several baselines. A negative defender expected utility indicates that animals were lost, so a higher positive number is

ideal. Therefore, we perform better with GUARDSS than using a random allocation, ignoring uncertainty, or forgoing drones. In fact, *ignoring uncertainty is worse than forgoing drones completely*. For varying γ (see Appendix), the gap between ignoring detection uncertainty and GUARDSS increases as γ increases, and the gap between the no drones case and GUARDSS decreases as γ increases, showing a similar trend to Fig. 3g. However, in all cases, the results emphasize the importance of correctly optimizing to get value from drones even with uncertainty.

8 Conclusion

The loss due to ignoring uncertainty can be high such that sensors are no longer useful. Nevertheless, by carefully accounting for uncertainty, uncertain information can still be exploited via a novel reaction stage and signaling even upon no detection. In this case, despite being aware of uncertainty, the attacker does not know whether he was detected, nor whether a patroller will respond in the reaction stage. Our results illustrate that the defender can exploit this informational advantage even with uncertain information. Thriving under this uncertainty makes real-world deployment of GUARDSS promising, as shown through simulation.

9 Acknowledgements

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

We would also like to thank Air Shepherd for their valuable insights and collaboration.

References

- Air Shepherd. 2019. <http://airshepherd.org>.
- Alderson, D. L.; Brown, G. G.; Carlyle, W. M.; and Wood, R. K. 2011. Solving defender-attacker-defender models for infrastructure defense. Technical report, Naval Postgraduate School.
- Basilico, N.; De Nittis, G.; and Gatti, N. 2015. A security game model for environment protection in the presence of an alarm system. In *GameSec*.
- Basilico, N.; De Nittis, G.; and Gatti, N. 2016. A security game combining patrolling and alarm-triggered responses under spatial and detection uncertainties. In *AAAI*.
- Basilico, N.; De Nittis, G.; and Gatti, N. 2017. Adversarial patrolling with spatially uncertain alarm signals. *Artificial Intelligence*.
- Bondi, E.; Fang, F.; Hamilton, M.; Kar, D.; Dmello, D.; Choi, J.; Hannaford, R.; Iyer, A.; Joppa, L.; Tambe, M.; and Nevatia, R. 2018. Spot poachers in action: Augmenting conservation drones with automatic detection in near real time. In *IAAI*.
- Bondi, E.; Jain, R.; Aggrawal, P.; Anand, S.; Hannaford, R.; Kapoor, A.; Piavis, J.; Shah, S.; Joppa, L.; Dilkina, B.; and Tambe, M. 2020. BirdSai: A dataset for detection and tracking in aerial thermal infrared videos. In *WACV*.
- Brown, N., and Sandholm, T. 2017. Superhuman AI for heads-up no-limit poker: Libratus beats top professionals. *Science*.
- Brown, G.; Carlyle, M.; Salmerón, J.; and Wood, K. 2006. Defending critical infrastructure. *Interfaces*.
- Bucarey, V.; Casorrán, C.; Figueroa, Ó.; Rosas, K.; Navarrete, H.; and Ordóñez, F. 2017. Building real stackelberg security games for border patrols. In *GameSec*.
- Cermak, J.; Bosansky, B.; Durkota, K.; Lisy, V.; and Kiekintveld, C. 2016. Using correlated strategies for computing stackelberg equilibria in extensive-form games. In *AAAI*.
- Černý, J.; Bojanský, B.; and Kiekintveld, C. 2018. Incremental Strategy Generation for Stackelberg Equilibria in Extensive-Form Games. In *EC*.
- Conitzer, V., and Sandholm, T. 2006. Computing the Optimal Strategy to Commit to. In *EC*.
- Cooney, S.; Wang, K.; Bondi, E.; Nguyen, T.; Vayanos, P.; Winetrobe, H.; Cranford, E. A.; Gonzalez, C.; Lebiere, C.; and Tambe, M. 2019. Learning to signal in the goldilocks zone: Improving adversary compliance in security games. In *ECML PKDD (Research Track)*.
- de Cote, E.; Stranders, R.; Basilico, N.; Gatti, N.; and Jennings, N. 2013. Introducing alarms in adversarial patrolling games. In *AAMAS*.
- De Nittis, G., and Gatti, N. 2018. Facing Multiple Attacks in Adversarial Patrolling Games with Alarmed Targets. *arXiv preprint arXiv:1806.07111*.
- Gholami, S.; Mc Carthy, S.; Dilkina, B.; Plumtre, A.; Tambe, M.; Driciru, M.; Wanyama, F.; Rwetsiba, A.; Nsubaga, M.; Mabonga, J.; et al. 2018. Adversary models account for imperfect crime data: Forecasting and planning against real-world poachers. In *AAMAS*.
- He, Y.; Ma, X.; Luo, X.; Li, J.; Zhao, M.; An, B.; and Guan, X. 2017. Vehicle Traffic Driven Camera Placement for Better Metropolis Security Surveillance. *arXiv preprint arXiv:1705.08508*.
- Jain, M.; Kardes, E.; Kiekintveld, C.; Ordóñez, F.; and Tambe, M. 2010. Security games with arbitrary schedules: A branch and price approach. In *AAAI*.
- Kamenica, E., and Gentzkow, M. 2011. Bayesian persuasion. *American Economic Review*.
- Kroer, C.; Waugh, K.; Kilinc-Karzan, F.; and Sandholm, T. 2017. Theoretical and practical advances on smoothing for extensive-form games. *EC*.
- Moravčík, M.; Schmid, M.; Burch, N.; Lisy, V.; Morrill, D.; Bard, N.; Davis, T.; Waugh, K.; Johanson, M.; and Bowling, M. 2017. Deepstack: Expert-level artificial intelligence in heads-up no-limit poker. *Science*.
- Nguyen, T. H.; Yadav, A.; An, B.; Tambe, M.; and Boutilier, C. 2014. Regret-Based Optimization and Preference Elicitation for Stackelberg Security Games with Uncertainty. In *AAAI*.
- Olivares-Mendez, M. A.; Fu, C.; Ludivig, P.; Bissyandé, T. F.; Kannan, S.; Zurad, M.; Annaiyan, A.; Voos, H.; and Campoy, P. 2015. Towards an autonomous vision-based unmanned aerial system against wildlife poachers. *Sensors*.
- Rosenfeld, A.; Maksimov, O.; and Kraus, S. 2018. Optimal cruiser-drone traffic enforcement under energy limitation. In *IJCAI*.
- Tambe, M. 2011. *Security and game theory: algorithms, deployed systems, lessons learned*. Cambridge University Press.
- Thakoor, O.; Tambe, M.; Vayanos, P.; Xu, H.; Kiekintveld, C.; and Fang, F. 2019. Cyber camouflage games for strategic deception. In *GameSec*, 525–541. Springer.
- Wang, Y.; Shi, Z. R.; Yu, L.; Wu, Y.; Singh, R.; Joppa, L.; and Fang, F. 2019. Deep reinforcement learning for green security games with real-time information. In *AAAI*.
- Xu, H.; Rabinovich, Z.; Dughmi, S.; and Tambe, M. 2015. Exploring Information Asymmetry in Two-Stage Security Games. In *AAAI*, 1057–1063.
- Xu, H.; Wang, K.; Vayanos, P.; and Tambe, M. 2018. Strategic coordination of human patrollers and mobile sensors with signaling for security games. In *AAAI*.
- Yang, R.; Kiekintveld, C.; Ordóñez, F.; Tambe, M.; and John, R. 2011. Improving resource allocation strategy against human adversaries in security games. In *IJCAI*.
- Yin, Y.; An, B.; and Jain, M. 2014. Game-theoretic resource allocation for protecting large public events. In *AAAI*.
- Yin, Z.; Jain, M.; Tambe, M.; and Ordóñez, F. 2011. Risk-averse strategies for security games with execution and observational uncertainty. In *AAAI*.
- Zhang, Y.; Guo, Q.; An, B.; Tran-Thanh, L.; and Jennings, N. R. 2019. Optimal interdiction of urban criminals with the aid of real-time information. In *AAAI*.