

Visionary Security: Using Uncertain Real-Time Information in Signaling Games

Elizabeth Bondi
 Harvard University
 ebondi@g.harvard.edu

1 Introduction

In important domains from natural resource conservation to public safety, real-time information is becoming increasingly important. Strategic deployment of security cameras and mobile sensors such as drones can provide real-time updates on illegal activities. To help plan for such strategic deployments of sensors and human patrollers, as well as warning signals to ward off adversaries, the defender-attacker security games framework can be used. [Zhang *et al.*, 2019] has shown that real-time data (e.g., human view from a helicopter) may be used in conjunction with security game models to interdict criminals. Other recent work relies on real-time information from sensors that can notify the patroller when an opponent is detected [Basilico *et al.*, 2017; Xu *et al.*, 2018]. Despite considering real-time information in all cases, these works do not consider the combined situation of uncertainty in real-time information in addition to strategically signaling to adversaries. In this thesis, we will not only address this gap, but also improve the overall security result by considering security game models and computer vision algorithms together.

A major aspect of this work is in applying it to real-world challenges, such as conservation. Although it applies to many environmental challenges, such as protecting forests and avoiding illegal mining, we will focus particularly on reducing poaching of endangered wildlife as an example. To reduce poaching, human patrollers typically search for snares and poaching activity as they patrol, as well as intervene if poaching activity is found. Drones are useful patrolling aids due to their ability to cover additional ground, but they must interpret their environments, notify nearby human patrollers for intervention, and send potentially deceptive signals to the adversary to deter poaching. Rather than treating these as separate tasks, models must coordinate to handle challenges found in real-world conservation scenarios (Fig. 1). We will determine the success of this work both in simulated experiments and through work with conservation agencies such as Air Shepherd to implement the system in the real world.

2 Proposed Research Plan

Our goal is to consider the “end-to-end” process of image recognition of adversary behavior to strategically counteracting such behavior via game theoretic reasoning. We will (i)

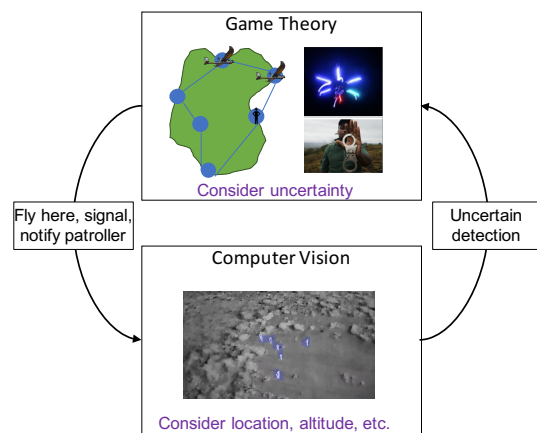


Figure 1: Considering uncertainty during high-level planning (game theory) may improve performance; considering the high-level plan may help improve detection (vision).

interpret imagery automatically through the use of computer vision, and (ii) develop strategies that counteract adversarial behavior via game theoretic reasoning. This will be done holistically to ensure real-world problems (e.g., uncertainty) are addressed by the system.

2.1 Detection in Imagery

One major aspect of this thesis is interpreting the images captured. We have already completed a great deal of this work, particularly in detection in thermal infrared videos, which can be used for nighttime surveillance. Automatic detection in thermal infrared videos captured aboard UAVs is extremely difficult since (i) the varying altitude of the UAV can sometimes lead to extremely small humans and animals, (ii) the motion of the UAV makes stabilization, and consequently human and animal motion detection, difficult, and (iii) the thermal infrared sensor itself leads to lower resolution, single-band images, much different from typical RGB images. Because thermal infrared imagery is different from the photos used to train algorithms like Faster RCNN [Ren *et al.*, 2015], labeled thermal infrared imagery is required to use these models for our detection. As a result, we developed VIOLA [Bondi *et al.*, 2017], an application that assists in labeling objects of interest, such as wildlife and poachers, in thermal




Figure 2: Example of providing a real-time security recommendation based on an image detection.

infrared imagery. After labeling 70 videos of varying altitude and resolution over the course of 6 months, we produced 39,380 labeled frames and approximately 180,000 individual poacher and animal labels on those frames. With this dataset, we developed SPOT [Bondi *et al.*, 2018b], the first (to our knowledge) aerial thermal detector for wildlife and poachers.

We evaluated SPOT based on both historical videos captured at different altitudes and a test run by Air Shepherd in the field (Fig. 2). We perform better than Air Shepherd’s current application, EyeSpy, in both precision and recall for large-sized poachers and animals, and in the field test video. For example, for the historical video containing large poachers, SPOT achieves 0.3977 precision and 0.0188 recall, whereas EyeSpy achieves 0.0052 precision and 0.0159 recall. Additionally, SPOT significantly reduces the burden on human operators since it is fully automated, whereas EyeSpy requires expert tuning of 14 parameters. We also use simulated data to augment the current dataset using AirSim-W [Bondi *et al.*, 2018a], a simulator for UAVs, by converting the simulated RGB images into thermal infrared images with simple physical models. With simulated data, SPOT achieves 0.7799 precision and 0.0374 recall on the same historical video.

2.2 Incorporating Real-Time Data into Game Theoretic Framework

In [Bondi *et al.*, 2019], our goal is to incorporate real-time data, such as our drone video detections, into a game theoretic framework. In the conservation application and other security applications, real-time data may be beneficial to help defenders find the attackers, and it may allow the defenders to cover more territory by utilizing deceptive signaling when anything is observed [Xu *et al.*, 2018]. In other words, our detections guide whether to signal. For example, if a poacher is detected far from a park ranger, we may have to strategically decide whether to signal depending on the park ranger location and the detection location, bearing in mind that the detections are subject to detection uncertainty. A demonstration of this concept is shown in Fig. 2, where a security recommendation (e.g., signal and go to this location) is provided after SPOT makes a detection.

2.3 Future Work

We would like to develop a new detection model to improve detection of small objects and take motion into account, potentially through the use of a tracking algorithm. We also plan to consider combining the two models directly to improve detections for decision-making purposes (e.g., [Wilder *et al.*, 2019]), and conduct field tests of the system.

Acknowledgements

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

References

- [Basilico *et al.*, 2017] Nicola Basilico, Giuseppe De Nittis, and Nicola Gatti. Adversarial patrolling with spatially uncertain alarm signals. *Artificial Intelligence*, 2017.
- [Bondi *et al.*, 2017] Elizabeth Bondi, Fei Fang, Mark Hamilton, Debarun Kar, Donnell Dmello, Jongmoo Choi, Robert Hannaford, Arvind Iyer, Lucas Joppa, Milind Tambe, and Ram Nevatia. Viola: Video labeling application for security domains. In *GameSec*, 2017.
- [Bondi *et al.*, 2018a] Elizabeth Bondi, Debadepta Dey, Ashish Kapoor, Jim Piavis, Shital Shah, Fei Fang, Bistra Dilikina, Robert Hannaford, Arvind Iyer, Lucas Joppa, and Milind Tambe. Airsim-w: A simulation environment for wildlife conservation with uavs. In *ACM COMPASS*, 2018.
- [Bondi *et al.*, 2018b] Elizabeth Bondi, Fei Fang, Mark Hamilton, Debarun Kar, Donnell Dmello, Jongmoo Choi, Robert Hannaford, Arvind Iyer, Lucas Joppa, Milind Tambe, and Ram Nevatia. Spot poachers in action: Augmenting conservation drones with automatic detection in near real time. In *IAAI*, 2018.
- [Bondi *et al.*, 2019] Elizabeth Bondi, Hoon Oh, Haifeng Xu, Fei Fang, Bistra Dilikina, and Milind Tambe. Broken signals in security games: Coordinating mobile patrollers and sensors in the real world. *AAMAS Extended Abstract*, 2019.
- [Ren *et al.*, 2015] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *NIPS*, 2015.
- [Wilder *et al.*, 2019] Bryan Wilder, Bistra Dilikina, and Milind Tambe. Melding the data-decisions pipeline: Decision-focused learning for combinatorial optimization. *AAAI*, 2019.
- [Xu *et al.*, 2018] Haifeng Xu, Kai Wang, Phebe Vayanos, and Milind Tambe. Strategic coordination of human patrollers and mobile sensors with signaling for security games. *AAAI*, 2018.
- [Zhang *et al.*, 2019] Youzhi Zhang, Qingyu Guo, Bo An, Long Tran-Thanh, and Nicholas R Jennings. Optimal interdiction of urban criminals with the aid of real-time information. In *AAAI*, 2019.