Vision for Decisions: Utilizing Uncertain Real-Time Information and Signaling for Conservation

Doctoral Consortium

Elizabeth Bondi
Harvard University
Cambridge, MA
ebondi@g.harvard.edu

ABSTRACT

Recent advances in fields such as computer vision and natural language processing have created new opportunities for developing agents that can automatically interpret their environment. Concurrently, advances in artificial intelligence have made the coordination of many such agents possible. However, there is little work considering both the low-level reasoning that allows agents to interpret their environment, such as deep learning techniques, and the high-level reasoning that coordinates such agents. By considering both together, we can better handle real-world scenarios. We will describe a real-world deployment of conservation drones to illustrate this point.

KEYWORDS

security games; computational sustainability; uncertainty; sensors; unmanned aerial vehicles; conservation

ACM Reference Format:
Elizabeth Bondi. 2020. Vision for Decisions: Utilizing Uncertain Real-Time Information and Signaling for Conservation. In *Proc. of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2020), Auckland, New Zealand, May 9–13, 2020,* IFAAMAS, 3 pages.

1 INTRODUCTION

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

In important domains like conservation and public safety, the real-time information provided by these sensors (e.g., drones) is becoming increasingly important. To help plan for 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. [9] has shown that real-time data 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 [1, 8]. However, these works do not consider the combined situation of uncertainty in real-time information in addition to strategically signaling to adversaries, which may lead to huge losses to defenders (e.g., park rangers). In this thesis, we will address this gap by (i) interpreting imagery automatically through the use of computer vision, and (ii) improve strategic, real-world decision-making given these interpretations, particularly by developing strategies that counteract adversarial behavior via game theoretic reasoning.

Figure 1: A conservation drone currently used for real-world deployment.

2 INTERPRETING SENSOR DATA

Our first goal is to interpret the sensors' data as accurately and robustly as possible. Despite recent advances in computer vision allowing agents to interpret their environment, automatic detection in thermal infrared videos captured aboard UAVs is still difficult since (i) the varying altitude of the drone can sometimes lead to extremely small humans and animals, (ii) the motion of the drone makes 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 [7], labeled thermal infrared imagery is required to use these models for our detection. As a result, we developed VIOLA [3], an application that assists in labeling objects of interest, such as wildlife and poachers, in thermal infrared imagery. After labeling 48 videos of varying altitude and resolution,

---

Proc. of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2020), B. An, N. Yorke-Smith, A. El Fallah Seghrouchni, G. Sukthankar (eds.), May 9–13, 2020, Auckland, New Zealand.

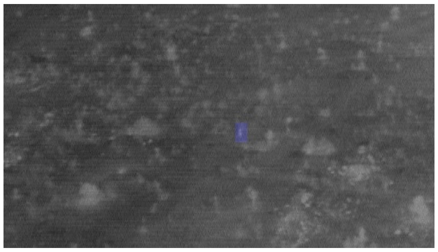

Figure 2: Example of human walking amongst trees, which is typically harder due to background and size.

We recommend that you turn on your lights and respond to this location.

Figure 3: Example of providing a real-time security recommendation based on an image detection.

we produced approximately 62,000 labeled frames. We then developed SPOT [4], the first (to our knowledge) aerial thermal detector for wildlife and poachers. *SPOT has recently been deployed to help ease the burden on park rangers in a real-world park*, whose name is hidden to protect wildlife. A conservation drone used during deployment is shown in Fig. 1.

We continued to improve this work by augmenting our data with synthetic data [2]. By using these synthetic data with labeled real data and domain adaptation [6], we achieve 0.459 mAP, whereas without these synthetic data, we achieve 0.438 mAP. An example of primarily successful detections is shown in Fig. 3. However, there are also more difficult images, as shown in Fig. 2, which are easy to miss, as other objects also look like.

3 USING UNCERTAIN SENSOR INFORMATION STRATEGICALLY IN A MULTIAGENT SYSTEM

Our goal, however, is to build a practical system that can overcome limitations in our interpretations of sensor imagery. With conservation drones, real-time data may help park rangers find poachers. As we also mentioned, the drones can be used to emit warning signals, such as turning lights on aboard the drone. This may allow park rangers to protect more animals by utilizing deceptive signaling when a poacher is observed. For example, if a poacher is detected far from a park ranger, we may want to strategically signal that the park ranger will respond, even though they may not. However, given real-world uncertainty, we may also want to strategically signal just in case there was a false negative detection, and furthermore, we may want to have a human to confirm there was no false negative detection. The detection result therefore influences the signaling and response. A demonstration of a signaling recommendation after a true positive detection by SPOT is shown in Fig. 3.

We are the first to model this uncertainty in sensing and signaling settings for security games [5]. We show that ignoring uncertainty leads to large losses for the defender (park ranger). We therefore (i) introduce a novel reaction stage to the game model, in which a park ranger could visit another target if nothing is observed by the drones or park rangers at first, and (ii) construct a new signaling scheme which includes signaling when there is no detection. Both of these allow the defender (park ranger) to mitigate the impact of uncertainty by (i) utilizing the chance to double check targets with drones just in case there are many false negatives expected, and (ii) signaling to make sure false negatives are not always accidentally skipped. In fact, this signaling scheme *exploits uncertain real-time information and the defender's informational advantage.* For example, both the defender (park ranger) and attacker (poacher) 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. Through this multiagent system, we are mitigating the effects of uncertainty in our interpretations of the imagery from the sensor, and in so doing, we are creating a better system.

4 FUTURE WORK

The main focus of upcoming work is on (i) improving robustness and performance in image detection, (ii) incorporating real-time data into game theoretic frameworks via a combined reinforcement learning and optimization framework, and (iii) working with real-world experts to analyze and improve the deployment of such systems for a wide range of conservation domains, from conducting animal surveys with conservation drones to reducing the negative effects of invasive species. In all cases, these will help further the goal of creating complete, practical multiagent systems with realistic sensors that can be used to alleviate challenges in conservation.

5 ACKNOWLEDGEMENTS

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

REFERENCES

[1] Nicola Basilico, Giuseppe De Nittis, and Nicola Gatti. 2017. Adversarial patrolling with spatially uncertain alarm signals. *Artificial Intelligence* (2017).
[2] Elizabeth Bondi, Debadeepta Dey, Ashish Kapoor, Jim Piavis, Shital Shah, Fei Fang, Bistra Dilkina, Robert Hannaford, Arvind Iyer, Lucas Joppa, and Milind Tambe. 2018. AirSim-W: A Simulation Environment for Wildlife Conservation with UAVs. In *ACM COMPASS*.
[3] Elizabeth Bondi, Fei Fang, Mark Hamilton, Debarun Kar, Donnabell Dmello, Jongmoo Choi, Robert Hannaford, Arvind Iyer, Lucas Joppa, Milind Tambe, and Ram Nevatia. 2017. VIOLA: Video Labeling Application for Security Domains. In *GameSec*.
[4] 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*.

[5] Elizabeth Bondi, Hoon Oh, Haifeng Xu, Fei Fang, Bistra Dilkina, and Milind Tambe. 2019. Broken Signals in Security Games: Coordinating Mobile Patrollers and Sensors in the Real World. AAMAS Extended Abstract.
[6] Yuhua Chen, Wen Li, Christos Sakaridis, Dengxin Dai, and Luc Van Gool. 2018. Domain Adaptive Faster R-CNN for Object Detection in the Wild. In *Computer Vision and Pattern Recognition (CVPR)*.
[7] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. 2015. Faster R-CNN: Towards real-time object detection with region proposal networks. In *NIPS*.
[8] Haifeng Xu, Kai Wang, Phebe Vayanos, and Milind Tambe. 2018. Strategic coordination of human patrollers and mobile sensors with signaling for security games. AAAI.
[9] Youzhi Zhang, Qingyu Guo, Bo An, Long Tran-Thanh, and Nicholas R Jennings. 2019. Optimal Interdiction of Urban Criminals with the Aid of Real-Time Information. In *AAAI*.