Using Game Theory in Real Time in the Real World: A Conservation Case Study

Demonstration

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
hxu@seas.harvard.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

In the real world, real-time data are now widely available, especially in security domains. Security cameras, aerial imagery, and even social media keep defenders informed when protecting important events, locations, and people. Further, advances in artificial intelligence have led to tools that can interpret these data automatically. Game theoretic models, for example, have shown great success in security. However, most of them ignore real-time information. In this paper, we demonstrate the potential to use real-time information from imagery to better inform our decisions in game theoretic models for security. As a concrete example, a conservation group called Air Shepherd uses conservation drones equipped with thermal infrared cameras to locate poachers at night and alert park rangers. They have also used lights aboard the drones, or signaled, to warn poachers of their presence, which often deters the poachers. We propose a system that (i) allocates drones and humans strategically throughout a protected area, (ii) detects poachers in the thermal infrared videos recorded by the conservation drones flying through the protected area in the predetermined location, and (iii) recommends moving to the location and/or signaling to the poacher that a patroller is nearby depending on real-time detections. **View the demonstration: http://bit.ly/aamas19-demo-bondi-et-al.**

KEYWORDS

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

ACM Reference Format:

Elizabeth Bondi, Hoon Oh, Haifeng Xu, Fei Fang, Bistra Dilkina, and Milind Tambe. 2019. Using Game Theory in Real Time in the Real World: A Conservation Case Study. 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 conservation, we can optimize limited conservation security resources to protect endangered wildlife and forests using the green security games (GSG) framework. Specifically, the GSG framework has shown success in the task of allocating patrollers in a protected area [5, 7]. The patrollers can log any signs of poaching they may encounter while patrolling, which may be useful in determining future patrols. However, by incorporating real-time data, we may be able to do more for current patrols. Consider using conservation drones during patrols, for example. They can make real-time detections of poachers and notify patrollers of the detection location, and they can send signals to notify poachers that a patroller is nearby. Although these signals may deter poachers, eventually the poachers will likely attack if warning signals are always sent without any response from human patrollers. To maintain a deterrence effect, it is necessary to signal truthfully at least some of the time. A game theoretic model can be used to determine when signaling should be done, as shown in [8]. It may even be used to plan for false negative (i.e., missed) image detections [4]. In fact, in this demonstration, we present a system to make recommendations based on real-time image detection with uncertainty in the domain of conservation.

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.

2 SPOT: USING REAL-TIME INFORMATION

Thermal infrared cameras are used aboard these conservation drones in order to detect poachers at night when poaching typically occurs. However, aerial thermal infrared imagery is quite different from the eye-level, visible spectrum photos used to train deep learning algorithms like Faster RCNN. Therefore, VIOLA [3] was used to label objects of interest, such as wildlife and poachers, in historical thermal infrared imagery from Air Shepherd. SPOT [2] was developed by training Faster RCNN on these data, and was the first (to our knowledge) aerial thermal detector for wildlife and poachers.

To evaluate SPOT's performance, precision and recall were measured for historical videos and a field test run by Air Shepherd in the field. SPOT outperformed Air Shepherd's previous application in both precision and recall for large-sized poachers and animals, and in the field test video. By adding simulated data generated using AirSim-W [1], a simulator for UAVs, SPOT achieves 0.7799 precision and 0.0374 recall on the large-size poacher historical video, as opposed to the previous algorithm which achieved only 0.0052 precision and 0.0159 recall. It is important to note that there is still ample detection uncertainty using SPOT with real and simulated data, which we address in our game theoretic model [4].

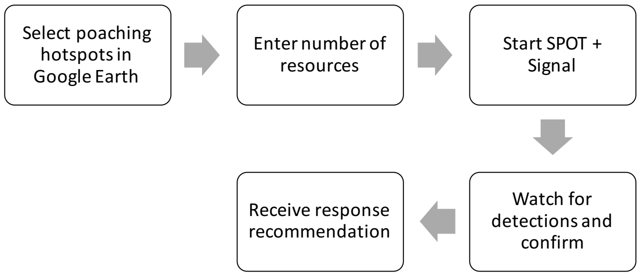

Figure 1: Steps in demonstration.

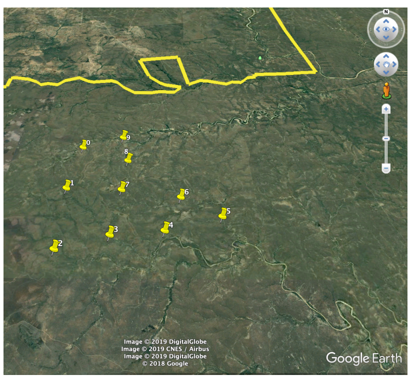

Figure 2: Google Earth snapshot showing potential poaching hotspots in a protected area in South Africa.

Figure 3: After detecting and confirming a poacher, provide recommendation, which in this case is sending a signal.

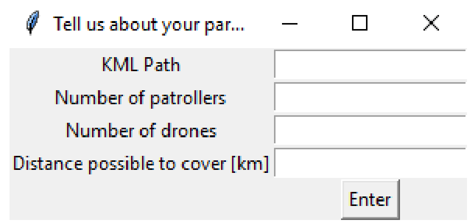

Figure 4: Read in the poaching hotspots and gather information about the protected area's security resources.

3 DEMONSTRATION

We demonstrate a software tool that could be deployed in the real world to detect poachers and respond to poaching incidents strategically, thereby easing the burden on Air Shepherd and park rangers. The demonstration will consist of the steps shown in Fig. 1 that integrate SPOT and game theoretic models to provide real-time detections in images and real-time recommendations to users.

First, the user can select poaching hotspots in the protected area. The poaching hotspots could be chosen based on a predictive model for poaching incidents, such as [6], or from domain expert knowledge. In our example, we are using potential hotspots in a protected area in South Africa. This is shown in Fig. 2. Users would be able to select their own hotspots.

Next, the user would enter details about the resources available to them, such as the number of human and drone patrollers and the approximate distance that could be traversed to respond to a poaching incident. This determines if a patroller could respond if a poacher was detected nearby, or if signaling was the only option, for example. These details are directly input to the game theoretic model, and then a randomized (mixed) strategy is determined to cover the park. At that point, a (pure) strategy is selected for the current night based on the mixed strategy. In the real world, the option to select from a previously computed mixed strategy would be provided, but for the purposes of the demonstration, the strategy will be recomputed and a pure strategy will be selected each time we use this software tool. We will then report this to the user.

Then, in the field, the drone(s) would start flying in the location to which it was assigned, and video would be transmitted to the base station in real time. In the demonstration, we have several historical videos associated with various potential hotspots in the park. We then start the SPOT detection system with the signaling and response recommendation enabled. The user is also prompted for other options of the SPOT system, such as whether the video assigns warm objects white or black in the grayscale images, whether there is a border or any text (e.g., altitude, time) on the video feed, whether sound should be made when there is a detection, etc.

SPOT then alerts the user when a poacher is detected. It first asks whether the detection is true (i.e., if it is really a poacher), and if so, a recommendation is made based on the pure strategy and the location of the drone. For example, if there is a human patroller near the drone, they should respond and a signal should likely also be sent notifying the poachers that someone is coming. If no human patroller is within the distance in which they could reasonably respond, it may be beneficial to send a deceptive signal to deter the poaching. An example of the detections and the recommendation are shown in Fig. 4. In assembling this system, we combine real-time information from thermal infrared cameras and SPOT with a game theoretic model to provide real-time security recommendations.

4 ACKNOWLEDGEMENTS

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

REFERENCES

[1] 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*.

[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] Elizabeth Bondi, Fei Fang, Debarun Kar, Venil Noronha, Donnabell Dmello, Milind Tambe, Arvind Iyer, and Robert Hannaford. 2017. VIOLA: Video Labeling Application for Security Domains. In *GameSec*.

[4] 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 *AAMAS Extended Abstract*.

[5] Fei Fang, Thanh Hong Nguyen, Rob Pickles, Wai Y Lam, Gopalasamy R Clements, Bo An, Amandeep Singh, Milind Tambe, and Andrew Lemieux. 2016. Deploying PAWS: Field Optimization of the Protection Assistant for Wildlife Security.. In *AAAI*. 3966–3973.

[6] Shahrzad Gholami, Benjamin Ford, Fei Fang, Andrew Plumptre, Milind Tambe, Margaret Driciru, Fred Wanyama, Aggrey Rwetsiba, Mustapha Nsubaga, and Joshua Mabonga. 2017. Taking it for a Test Drive: A Hybrid Spatio-temporal Model for Wildlife Poaching Prediction Evaluated through a Controlled Field Test. In *ECML PKDD*.

[7] Binru Wang, Yuan Zhang, and Sheng Zhong. 2017. On Repeated Stackelberg Security Game with the Cooperative Human Behavior Modelfor Wildlife Protection. In *AAMAS*.

[8] Haifeng Xu, Kai Wang, Phebe Vayanos, and Milind Tambe. 2018. Strategic coordination of human patrollers and mobile sensors with signaling for security games. In *Thirty-Second AAAI Conference on Artificial Intelligence*. AAAI.