Bridging the Gap Between High-Level Reasoning in Strategic Agent Coordination and Low-Level Agent Development

Doctoral Consortium

Elizabeth Bondi
University of Southern California
bondi@usc.edu

ABSTRACT

Recent advances in fields such as computer vision and natural language processing have paved the way for developing agents capable of automatically interpreting their surrounding 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, for example by planning at a high level with low-level uncertainty in mind, or even by improving low-level processing by using high-level reasoning to place the agent in the best scenario for success.

KEYWORDS

security games; computational sustainability; uncertainty; sensors

ACM Reference Format:
Elizabeth Bondi. 2019. Bridging the Gap Between High-Level Reasoning in Strategic Agent Coordination and Low-Level Agent Development. 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

Using artificial intelligence and deep learning algorithms in the real world has recently shown great promise and success (e.g., [1, 3, 7, 8, 14]). However, much of this work considers only high- *or* low-level reasoning. For example, [8, 14] use high-level reasoning, specifically security game models, to protect ports and national parks. These strategies largely ignore low-level processing of real-world data, whereas [3] processes data from unmanned aerial vehicles (UAVs or drones), but does not reason about strategic coordination of UAVs.

In this work, we specifically consider the detection of adversary behavior in imagery, such as UAV video, to strategically counteracting such behavior via game theoretic reasoning. Traditionally in artificial intelligence, adversary behavior modeling does not include computer vision processing, whereas computer vision research traditionally does not extend to strategic game theoretic reasoning. By considering both components together, we can (i) create agents that interpret videos automatically, and (ii) develop strategies that counteract adversarial behavior and consider real-world challenges.

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 can search for snares and poaching activity as they patrol, as well as intervene if poaching activity is found. UAVs are useful patrolling aids due to their ability to cover additional ground, but for easy use, we must give them the ability to interpret their environments, notify nearby human patrollers for intervention, and send potentially deceptive signals to the adversary to deter poaching. To accomplish these tasks strategically, we require both high- and low-level reasoning, for example in allocating UAVs and patrollers throughout a national park, and in detecting poachers in UAV videos. Rather than treating these as separate tasks, the different levels must coordinate to handle the challenges found in real-world conservation scenarios (Fig. 1). We will also work with conservation agencies such as Air Shepherd [12] to implement the system in the real world.

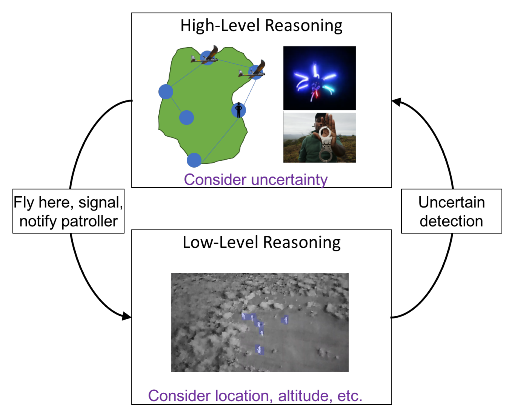

Figure 1: Conservation coordination: Considering uncertainty during high-level planning may improve overall performance; high-level planning may improve detection.

2 RELATED WORK

Combining high- and low-level reasoning is not a new idea, but has recently begun to receive more thought. In [9], multi-agent systems are composed of computer vision algorithm agents that each solve the respective task in some cases. There have also been promising results in end-to-end learning from high-dimensional sensory inputs to action outputs, with a famous example being [2]. In both cases, the context is slightly different than our proposal, which is to make better decisions based on imagery, and vice versa.

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.

3 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. To this end, we will (i) create agents to interpret imagery automatically, and (ii) use game theoretic reasoning to strategically counteract adversarial behavior. This will be done holistically to ensure real-world challenges (e.g., uncertainty) are addressed by the system.

3.1 Creating the Agent: Detection in Imagery

We first discuss the agent's low-level reasoning, specifically the computer vision task of interpreting captured images. 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 difficult since the UAV's altitude variation can lead to small humans and animals, UAV motion makes human and animal motion detection difficult, especially if requiring stabilization, and the thermal infrared sensor produces lower resolution, single-band images. Because thermal infrared imagery is different from the photos used to train algorithms like Faster RCNN [11], labeled thermal infrared imagery is required to use these models for our detection. As a result, we developed VIOLA [4], an application that assists in labeling objects of interest, such as wildlife and poachers, in thermal infrared imagery. After labeling 70 videos of varying altitude and resolution over the course of 6 months, we produced about 39,380 labeled frames and 180,000 individual poacher and animal labels on those frames. With this dataset, we developed SPOT [5], an aerial, thermal detector for wildlife and poachers. SPOT is based on Faster RCNN, which, after training, locates and classifies objects of interest as poachers or animals.

We evaluated SPOT with both historical videos and a test run by Air Shepherd in the field (Fig. 2) [5]. We compared SPOT to Air Shepherd's current application, EyeSpy, which requires expert tuning of 14 parameters at run time. In addition to reducing the burden on human operators by removing this tuning at run time, we also perform better in both precision and recall in most cases. 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. We also used simulated data to augment the current dataset using AirSim-W [3], a simulator for UAVs. SPOT with simulated data achieves 0.7799 precision and 0.0374 recall on the same historical video.

3.2 Towards Incorporating Real-Time Image Data into the Game Theoretic Framework

We now consider integrating the high- and low-level reasoning involved in game theoretic and image detection models. In terms of high-level reasoning, we need to consider where to deploy UAVs and human patrollers, as well as how to respond to the observation of an adversary. When a UAV observes an adversary, responses include sending a (deceptive) signal and sending a human patroller to interdict. In [6], we propose a game theoretic model for this scenario while also incorporating a major challenge in real-world deployment: uncertainty. As we already saw from the results in [3], detection uncertainty is present in SPOT. When SPOT is considered as part of this wider system, ignoring this uncertainty could unfortunately lead to significant losses to the defender. For example, false negative detections (i.e., when the object is not detected at all even though it is truly present) would mislead the defender and therefore prevent the defender from interdicting an attack. Therefore, we should keep low-level uncertainty in mind in the high-level game theoretic model. To do this, one improvement is to have human patrollers move to neighboring locations, even if they have UAVs already. This allows patrollers to double check for adversaries, especially if the UAVs have high false negative rates [6]. In Fig. 2, SPOT makes a possibly uncertain detection, and then the game theoretic model provides the response for that scenario.

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

3.3 Future Work

To improve the agent, a new detection model will be developed that works better for small objects, takes advantage of motion, and is tailored to thermal infrared imagery. We would also like to further consider human-agent teamwork. Although agent teams are discussed in [10, 13], human team members are not explicitly considered. To further facilitate teamwork between humans and agents, we will expand upon our current work by considering patroller and poacher routes and considering that there could be attacks at more than one time. Finally, we will also consider combining the two models directly using decision-focused learning [15].

4 CONCLUSION

We aim to bridge the gap between high- and low-level reasoning in the development of security agents and systems. One major application of this work is in conservation. By considering both computer vision techniques in agents, as well as game theoretic reasoning for coordinating agents, we can make the full system work better in real-world tasks.

5 ACKNOWLEDGEMENTS

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

REFERENCES

[1] Sara Beery, Grant Van Horn, and Pietro Perona. 2018. Recognition in terra incognita. In *Proceedings of the European Conference on Computer Vision (ECCV)*. 456–473.

[2] Mariusz Bojarski, Davide Del Testa, Daniel Dworakowski, Bernhard Firner, Beat Flepp, Prasoon Goyal, Lawrence D Jackel, Mathew Monfort, Urs Muller, Jiakai Zhang, et al. 2016. End to end learning for self-driving cars. *arXiv preprint arXiv:1604.07316* (2016).

[3] 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*.

[4] 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*.

[5] 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*.

[6] 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. Proceedings of the 19th International Conference on Autonomous Agents and Multiagent Systems. International Foundation for Autonomous Agents and Multiagent Systems, 2019 (Extended Abstract).

[7] Rob Critchlow, Andrew J Plumptre, Margaret Driciru, Aggrey Rwetsiba, Emma J Stokes, Charles Tumwesigye, Fred Wanyama, and CM Beale. 2015. Spatiotemporal trends of illegal activities from ranger-collected data in a Ugandan national park. *Conservation Biology* 29, 5 (2015), 1458–1470.

[8] 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.

[9] Alfonso González-Briones, Gabriel Villarrubia, Juan F De Paz, and Juan M Corchado. 2018. A multi-agent system for the classification of gender and age from images. *Computer Vision and Image Understanding* (2018).

[10] Jonathan Grizou, Samuel Barrett, Manuel Lopes, and Peter Stone. 2016. Collaboration in Ad Hoc Teamwork: Ambiguous Tasks, Roles, and Communication. In *AAMAS Adaptive Learning Agents (ALA) Workshop*.

[11] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. [n. d.]. Faster R-CNN: Towards real-time object detection with region proposal networks. In *NIPS*.

[12] Air Shepherd. 2018. Air Shepherd: The Lindbergh Foundation. http://airshepherd.org. (2018). Accessed: 2018-09-01.

[13] Milind Tambe. 1997. Towards flexible teamwork. *Journal of artificial intelligence research* 7 (1997), 83–124.

[14] Milind Tambe. 2011. *Security and Game Theory: Algorithms, Deployed Systems, Lessons Learned*. Cambridge University Press.

[15] Bryan Wilder, Bistra Dilkina, and Milind Tambe. 2019. Melding the Data-Decisions Pipeline: Decision-Focused Learning for Combinatorial Optimization. *AAAI* (2019).