

PAWS: Adaptive Game-theoretic Patrolling for Wildlife Protection

(Demonstration)

Benjamin Ford, Debarun Kar, Francesco M. Delle Fave, Rong Yang, Milind Tambe
University of Southern California, Los Angeles, CA, US
{benjamif,dkar,dellefav,yangrong,tambe}@usc.edu

Categories and Subject Descriptors

I.2.11 [Artificial Intelligence]: Distributed Artificial Intelligence - multiagent systems

1. INTRODUCTION

Endangered species around the world are in danger of extinction from poaching. From the start of the 20th century, the African rhino population has dropped over 98% [4] and the global tiger population has dropped over 95% [5], resulting in multiple species extinctions in both groups. Species extinctions have negative consequences on local ecosystems, economies, and communities. To protect these species, countries have set up conservation agencies and national parks, such as Uganda's Queen Elizabeth National Park (QENP). However, a common lack of funding for these agencies results in a lack of law enforcement resources to protect these large, rural areas. As an example of the scale of disparity, one wildlife crime study in 2007 reported an actual coverage density of one ranger per 167 square kilometers [2]. Because of the hazards involved (e.g., armed poachers, wild animals), rangers patrol in groups, further increasing the amount of area they are responsible for patrolling.

Security game research has typically been concerned with combating terrorism, and this field has indeed benefited from a range of successfully deployed applications [1, 6]. These applications have enabled security agencies to make more efficient use of their limited resources. In this previous research, adversary data has been absent during the development of these solutions, and thus, it has been difficult to make accurate adversary behavior models during algorithm development. In a domain such as wildlife crime, interactions with the adversary are frequent and repeated, thus enabling conservation agencies to collect data. This presence of data enables security game researchers to begin developing algorithms that incorporate this data into, potentially, more accurate behavior models and consequently better security solutions.

Developed in conjunction with staff at QENP, the Protection Assistant for Wildlife Security (PAWS) generates optimized defender strategies for use by park rangers [7]. Due to the repeated nature of wildlife crime, PAWS is able to lever-

Appears in: *Alessio Lomuscio, Paul Scerri, Ana Bazzan, and Michael Huhns (eds.), Proceedings of the 13th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2014), May 5-9, 2014, Paris, France.*
Copyright © 2014, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.




Figure 1: PAWS system data flow

age crime event data - a previously unrealized capability in security games research. Thus, PAWS implements a novel adaptive algorithm that processes crime event data, builds multiple human behavior models, and, based on those models, predicts where adversaries will attack next. These predictions are then used to generate a patrol strategy for the rangers (i.e., a set of patrol waypoints) that can be viewed on a GPS unit.

Against this background, the demonstration presented in this paper introduces two contributions. First, we present the PAWS system which incorporates the algorithm in [7] into a scheduling system and a GPS visualizer. Second, we present a software interface to run a number of human subject experiments (HSE) to evaluate and improve the efficacy of PAWS before its deployment in QENP. By conducting these HSEs, we can (i) test the PAWS algorithms with repeated interactions with humans, thus providing a more realistic testing environment than in its previous simulations; (ii) generate data that can be used to initialize PAWS's human behavior models for deployment, and (iii) compare the current PAWS algorithms' performance to alternatives and determine if additional improvements are needed prior to deployment. To provide proper context for the presentation, this paper also presents a brief overview of the PAWS system data flow and its adaptive algorithms.

The demonstration will engage audience members by having them participate in the HSEs and using the GPS unit to visualize a patrol schedule in QENP.

2. PAWS SYSTEM

The PAWS system is composed of two main components: a centralized scheduler and a GPS unit.

2.1 PAWS Patrol Generation

The key component of PAWS is the scheduler [7]. The scheduler formulates the wildlife crime problem as a Bayesian Stackelberg game with infinite types, where the leaders are the rangers and the followers are the poachers. The poachers' behavior is modeled using a Subjective Utility Quantal Response model [3]. This enables PAWS to account for hu-

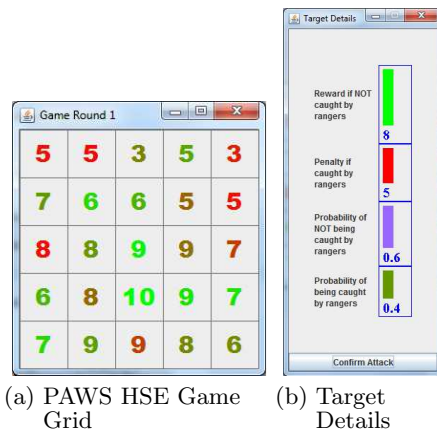


Figure 2: PAWS HSE Screenshots

mans’ bounded rationality and obtain more accurate predictions of where poachers will attack next.

PAWS’s adaptability stems from its PAWS-Learn algorithm (see [7] for details) which, as a first in Stackelberg Security games, enables PAWS to continuously update the SUQR behavior model parameters with incoming data. Once new data is received from ranger patrols, PAWS will re-evaluate the poachers’ behavior model and output a new mixed strategy for the defender from which ranger patrols can be sampled.

The scheduling component of PAWS runs on a server and uses patrol data (consisting of suspected attack events, their locations, and, if available, the responsible poacher) as input into the previously described algorithm. As described in the next section, this patrol data is created at the end of a patrol and is stored in a database on the server.

2.2 GPS Visualization

Once patrols are generated by the PAWS algorithms, they are downloaded onto GPS units as a set of waypoints. One example of the GPS unit is shown in Figure 3. Rangers in QENP are currently using these GPS units to conduct their patrols. As a part of their patrol duties, rangers will collect data on any suspicious activities or poachers encountered. Once the rangers return from their patrol, this data will be uploaded to PAWS and thus completing an iteration of the patrol generation, execution, and data acquisition cycle illustrated in Figure 1.

3. THE SIMULATION GAME

Demo: <https://www.youtube.com/watch?v=zpzVbGd1kf4>
 Before deploying PAWS, we will conduct a series of HSEs to test PAWS’s algorithms’ performance in both predicting human behavior and adapting to behavior patterns. During our demonstration, the audience members will be able to directly participate in these HSEs. First, they will be presented a high-level description of the wildlife crime problem, details on how their participation can help wildlife rangers more effectively combat poaching, and a set of instructions on how to play the HSE game. Second, they will be asked to participate in a game, whereby they need to choose a cell within a game grid (see Figure 2(a)). This grid represents a wildlife area to protect and each cell is a location that can, potentially, be attacked by one or more poachers.

In essence, each cell represents a target within a Stack-

elberg Security Game (see [6]). Hence, it is characterized by a reward value, denoted by the number, and a probability of coverage, denoted by the color of the number: red colors are a high probability of coverage, green colors are a low probability of coverage, and yellow colors are almost equally likely to be covered as not covered. If a player clicks on a cell, it will display the cell’s complete payoff information (i.e., if the player attacks, the reward they will receive if successful or penalty they will receive if they fail) and defender coverage probabilities. Hence, in this experiment, players have full knowledge of the defender’s mixed strategy (i.e., coverage probabilities for each target) as required by a Stackelberg game formulation. If the player decides to attack that target, they can click on a “Confirm Attack” button. After players make their attack choices, their choices will be logged, and they will be informed of their success or failure. Since we are comparing the performance of different algorithms, including PAWS (e.g., adaptive and robust algorithms), players will be playing a game in which coverage probabilities may vary greatly between two game types.

4. THE GPS VISUALIZER

Following participation in the game and concluding the demonstration, audience members will be able to interact with a model of GPS (Fig: 3) that is currently in use in QENP. Here, we will demonstrate how generated patrol strategies will be displayed on a GPS for rangers to use in the field.



Figure 3: QENP GPS Model

5. ACKNOWLEDGMENTS

This research was supported by MURI Grant W911NF-11-1-03. We thank IBM for the fellowship for Rong Yang.

6. REFERENCES

- [1] V. Conitzer. Computing game-theoretic solutions and applications to security. In *AAAI*, pages 2106–2112, 2012.
- [2] T. Holmern, M. J., and R. E. Local law enforcement and illegal bushmeat hunting outside the serengeti national park, tanzania. *Environmental Conservation*, 34:55–63, 2007.
- [3] T. H. Nguyen, R. Yang, A. Azaria, S. Kraus, and M. Tambe. Analyzing the effectiveness of adversary modeling in security games. In *AAAI*, 2013.
- [4] J. R. Platt. How the western black rhino went extinct. blog, November 2013.
- [5] G. T. I. Secretariat. Global tiger recovery program implementation plan: 2013-14. Report, The World Bank, Washington, D.C., 2013.
- [6] M. Tambe. *Security and Game Theory: Algorithms, Deployed Systems, Lessons Learned*. Cambridge University Press, New York, NY, 2011.
- [7] R. Yang, B. Ford, M. Tambe, and A. Lemieux. Adaptive resource allocation for wildlife protection against illegal poachers. In *AAMAS*, 2014.