

PAWS: Game Theory Based Protection Assistant for Wildlife Security

This chapter introduces Protection Assistant for Wildlife Security (PAWS) (Yang, Ford, Tambe, & Lemieux, 2014) as a joint effort done by computer scientists, conservation researchers, and conservation practitioners from two nongovernmental organizations—Panthera, and Rimba. PAWS is a game theory based application to assist conservation agency officials in planning wildlife ranger patrols to prevent wildlife crime. Reducing risks to people and wildlife from wildlife crime ideally includes the combined effort of practitioners worldwide and researchers in many different disciplines; PAWS demonstrates the positive impact that research in computational game theory, an important topic within the field of Artificial Intelligence (AI), can have in assisting wildlife conservation agencies to prevent poaching. In recent deployment efforts, patrol planners mentioned that the routes generated by PAWS came close to an actual planner's routes, a promising sign that PAWS can suggest feasible routes and help reduce the significant burden of patrol planning.

Poaching directly threatens some species' survival. For example, tigers, along with many other endangered species, are in danger of extinction because of poaching risks (Global Tiger Initiative Secretariat, 2013; National Wildlife Refuge Association, 2015; Montesh, 2013). The global population of tigers has dropped over 95% since the 1900s, resulting in 3 out of 9 species going extinct (Global Tiger Initiative Secretariat, 2013) in part to poaching. In 2015, South African rhino poaching reached a rate of approximately 1 death every 8 hours (Save the Rhino International, 2015). Species extinction can destroy ecosystems and weaken the communities and economies that depend on those ecosystems (Global Tiger Initiative Secretariat, 2013). In some cases, such as with the illegal rhino and tiger trades, poachers can be part of well-funded organized crime groups. Many other poachers, however, are not part of organized crime syndicates but still threaten species with over-hunting by snare poaching. Regardless of the scale, wildlife poaching poses negative risks to wildlife and people.

Worldwide, patrols are the most widespread method to combat and prevent wildlife poaching. However, many conservation agencies suffer from a lack of law enforcement resources to conduct these patrols over what are typically vast uninhabited protected areas. One wildlife crime study reported an actual coverage density of 1 ranger per 167 square kilometers (Holmern, Muya, & Roskaft, 2007) while current law enforcement density statistics for New York City show approximately 28 officers per square kilometer (e.g., 34,500 officers over a total land and sea area of 1,213 square kilometers) (City of New York, 2013). It would be impossible to adequately protect the entire protected areas at the same density as in urban areas, increasingly the importance and necessity of planning efficient patrols. However, current patrols in many protected or other conservation areas may not make the most efficient use of their limited patrolling resources.

PAWS generates a set of strategically randomized patrol routes, based on a game-theoretic analysis, such that a conservation agency can choose to execute any of these patrols with a given probability. This game-theoretic approach provides a valuable degree of unpredictability, and while this by itself would be useful for assisting agencies with the time-consuming process of patrol route planning, PAWS also formalizes and incorporates two important concepts in patrol planning:

1. Building models of human behavior to better predict where poachers will attack and how they will react to any executed patrol routes.
2. Incorporating domain features such as terrain information into the game-theoretic analysis so that patrollers can easily execute any of the patrols generated by PAWS.

PAWS was first tested at Uganda's Queen Elizabeth National Park (QENP) where poaching is believed to be the most harmful and frequent illegal activity in the park. Now PAWS is regularly deployed in a Southeast Asia protected area in support of tiger conservation. For the security of animals and patrollers, no latitude/longitude information is presented about this site.

Applying game theoretic analysis to poaching

Game theory is the study of strategic decision-making and, more specifically, the conflict and cooperation between intelligent decision-makers (Myerson, 1997). In combating poaching, there is a strategic interaction between the conservation agency (or wildlife ranger or other patrollers) and the poachers. In this game, there are two types of players with conflicting interests: the defender (i.e., conservation agency or patrollers) and the poachers. Each player wants to take action(s) so that the outcome is in their best interest. For example, poachers want to place snares and capture wildlife without having their snares confiscated by patrollers, and on the other hand, patrollers want to find snares before they capture any animals. To that end, each player needs to reason about his or her opponent's potential actions and play intelligently. For example, if the patroller always takes the same patrol route every day, the poachers will always be able to avoid the patroller and will always succeed in placing snares in the unpatrolled locations. Thus, it is in the defender's best interest to play unpredictably (i.e., patrol somewhat randomly). Instead of simply "rolling a die" and choosing which patrol route to take in a uniformly random way, a defender should choose patrol routes that visit more important locations (e.g., areas with higher animal densities of key species) more often. Which patrol routes to consider and how to randomly choose among these patrol routes is called the defender strategy or the patrol strategy. Ideally, for the defender, poachers would be deterred from locations with high animal density since they are often patrolled, and the poachers would be reluctant to place snares in areas with low animal density since the chance of successfully capturing an animal in those areas is low.

Even though this ideal case may not be achievable given the defender's limited patrolling resources, it does not mean a game theoretic approach cannot have a positive impact. Indeed, game-theory based decision support systems have been successfully deployed in the real-world to protect critical infrastructure such as airports (Pita et al., 2008), airline flights (Tsai, Rathi, Kiekintveld, Ordonez, Tambe, 2009), seaports (Shieh et al., 2012), and metro trains (Yin, Jiang, Johnson, Kiekintveld, & Leyton-Brown, 2012). In each of these cases, even though the defender also had a limited amount of patrolling resources, these works aided defenders in more efficiently and effectively allocating their resources (Tambe, 2011).

PAWS was inspired by these successes and was the first of a new wave of proposed applications in the subarea called Green Security Games (GSGs) (Fang, Stone, & Tambe, 2015; Kar, Fang, Delle Fave, Sintov, & Tambe, 2015), which focus on resource allocation and scheduling problem in domains such as protecting forest, wildlife, and fisheries. PAWS provides quantitative analysis of the game between the defenders and poachers and calculates the optimal patrol strategy.

Modeling human behavior to create optimal patrol strategies

In generating an optimal patrol strategy (i.e., one that offers the best chance of stopping poaching attacks), modeling poachers' behavior is important. Without an accurate idea of how the poacher is currently planning their attacks and how they will react to the defender's patrols, it is difficult to generate a patrol strategy that performs well in practice. Because there is a large number of poachers and poaching data potentially available, we can use this crime data to build a model of the poachers' behavior. In the case of poaching, crime data can be anonymous, or identified if it can be linked to confessed adversaries. Although the latter type of data provides rich information about individual adversaries, it is sparse and hard to collect. Indeed, it is very difficult to catch poachers on site since it requires the patrollers to be at the same location as poachers at the same time. The majority of collected data is evidence on crimes committed by anonymous adversaries (e.g., a snare found on a trail). Compared to identified data, anonymous data provides no information about the characteristics of the adversary that committed the crime and therefore cannot be used to build accurate behavioral models on the individual level. The open questions here are then, how do we utilize both types of data to build and learn a better model of the large population of criminals and moreover, how does the learned model help better predict future crime events and thus help law enforcement officials to improve their patrols?

Domain feature modeling

To bring PAWS from theory to real-world deployment, we capture and integrate important domain properties of wildlife crime into PAWS. By doing so, PAWS can generate optimal patrol strategies that accurately reflect many of the factors that influence the creation of patrols. First, we incorporate terrain information (e.g., elevation, ridgelines) and use a novel hierarchical modeling approach to build a virtual street map of the protected area. This virtual street map allows PAWS to scale-up to patrol large areas while simultaneously providing fine-grained guidance. Essentially, the street map connects the whole protected area through easy-to-follow route segments, where each segment could be a ridgeline, stream, river bank, or other geographic feature. The rationale for this approach is threefold:

1. Animals use these terrain features and are most easily trapped here.
2. Poachers use these features for trapping and moving about in general.
3. Patrollers can move along these features more efficiently than traversing across ridges.

In other words, animals, poachers, and patrollers all use these features while moving. In addition to incorporating terrain information, PAWS also accounts for incomplete information regarding animal density. Finally, PAWS accounts for real-world patrolling constraints such as time limits on patrolling and needing to start and end at a base camp. In the rest of the chapter we discuss the related work of PAWS and provide a detailed

description of the wildlife poaching problem domain. We will then present an overview of the PAWS system and explain how it works, in detail, with respect to game-theoretic analysis, human behavior modeling, and domain feature modeling.

The Genesis of PAWS from Synthesizing Conservation, Computer science, and Criminology

Poaching is increasingly studied by criminologists (Montesh, 2013; Pires, 2012) and geographic information systems (GIS) experts (Hamisi, 2008; Ouko, 2013) in addition to conservation practitioners (Wato, Wahungu, & Okello, 2006). A variety of methods are used to identify critical points in the poaching system, such as GIS analysis and interviews with apprehended poachers. In spite of all these efforts, returns on applied research investment can be low because of a lack of law enforcement resources (Hamisi, 2008; Pires, 2012).

Conducting patrols is an important way to combat poaching. In recent years, data collection and aggregation software such as MIST (Stokes, 2010) and SMART (Smart Collaboration, 2015) have enabled conservation managers to more effectively coordinate their protection efforts. These tools are developed to help conservation managers record data and analyze patrols retrospectively. However, these works do not create patrol routes or identify targets to protect; the creation of patrols is still done by an experienced patrol manager. It is well known that humans have an extremely difficult time generating feasible schedules that are also unpredictable (Wagenaar, 1972). In contrast, PAWS builds on concepts and models from game theory, in particular, security games and provides an automated approach that generates efficient and randomized patrol schedules.

Research on security games focuses on overcoming the security and conservation agencies' challenge of limited law enforcement resources. In optimizing security resource allocation, previous work on Stackelberg Security Games (SSGs) has led to many successfully deployed applications to improve the security of airports, ports and flights (Fang, Jiang, & Tambe, 2013; Pita et al., 2008; Tsai, Rathi, Kiekintveld, Ordonez, & Tambe, 2009). Based on the early work on SSGs, recent work has focused on GSGs (Fang, Stone, & Tambe, 2015; Kar et al., 2015) for domains such as protecting forest, wildlife, and fisheries resources (Haskell et al., 2014; Qian, Haskell, Jiang, & Tambe, 2014). Unlike a standard SSG where the attack is assumed to be cross sectional, interactions between the defender and their adversary in GSG are repeated. For example, in the security domain of wildlife protection poachers conduct illegal activities (e.g., place snares in a protected area as shown in Figure 10.1(c)) frequently and repeatedly). Although some GSG research provides conceptual advances in integrating learning and planning (Fang et al., 2015), PAWS is the first GSG application to combat wildlife crime. PAWS models the interaction between the patroller (i.e., defender) and the poacher (i.e., attacker) as a basic GSG. Every few months, collected poaching data is analyzed, behavioral models of the poacher are updated and reparameterized. Improved patrols are generated.

Describing the poaching domain to create patrols that prevent poaching

Wire snaring, shown in Figure 10.1(c), is one of the main techniques used by poachers. Poachers can set and leave snares unattended and when they think an animal has been captured they will come back and kill the animal. Snares can be designed to target many species. The main targets of snares in Uganda's QENP are hippo, kobs, and antelope. In Southeast Asia where PAWS is now regularly deployed the main targets are elephants, tigers, and Sambar deer (i.e., a key prey species of tigers).

Law enforcement officers, such as park rangers, or patrollers from NGOs and other government agencies are responsible for patrolling a given protected area. One of their primary goals is to reduce successful poaching activities as much as possible so as to protect wildlife in that geographic area. During a given patrol patrollers will typically search for signs of illegal activity inside the protected area, confiscate any poaching equipment found, and apprehend any persons inside the area illegally (e.g., poachers). During their patrol all their findings are recorded in a log book. In most cases, if patrollers find wire snares, they will not find the poacher that set them because the poachers usually leave after setting the snares. If patrollers do encounter and apprehend poachers, however, patrollers are sometimes able to generate confessions from the poachers as to where they set their snares. Thus, among all the records of wire snares found by patrollers, most of them are anonymous. Identified data points, when a poacher is captured and divulges where they placed snares, are inherently more useful as they can be used to obtain a complete behavioral model that can better predict where future poachers will place their traps. If recording and monitoring tools such as SMART and MIST are in use in a protected area, collected data are uploaded to a database after the patrollers return to the outpost. Patrollers usually patrol in teams, and there can be multiple teams of patrollers conducting patrols simultaneously, each team taking a different patrol route.



(a) QENP in Uganda (b) Outline of QENP (c) A caught poacher holding up a snare (photo by Andrew Lemieux).

Figure 10.1. Snare poaching in Queen Elizabeth National Park

Poachers can conduct surveillance on patrollers' activities and patrol patterns. Wildlife patrollers are well aware that some neighboring villagers will inform poachers of when patrollers begin their patrols and where they are patrolling (Moreto, 2013). For any number of reasons, such as changes that impact animal migration habits, patrollers may change their patrolling patterns. Poachers, in turn, continually conduct surveillance on the patrollers' changing patrol strategy and adapt their poaching strategies accordingly.

Within a given protected area, different locations may have different importance for animals and people. Locations of high animal density are more important and attractive to poachers. In the case of QENP, areas that contain fresh water (e.g., watering holes, lakes) are known to be high-risk areas for poaching (Montesh, 2013; Moreto, 2013; Wato et al., 2006). In areas with more hilly and complex terrains, such as tropical forests in Southeast Asia, areas along ridgelines or streams serve as natural conduits for wildlife. Despite the available information on animal ecology and ethology from natural science studies, there are still too many areas for patrollers to patrol, and it is a huge cognitive burden to simultaneously account for animal density factors, physical distance constraints and base camp locations while also ensuring that the new patrols are unpredictable. PAWS aims to aid patrol managers by generating an optimal strategy that accounts for all of these factors; PAWS will generate a strategy that enables patrollers to effectively cover these numerous areas with their limited resources.

The PAWS Model

PAWS is based on game theoretic analyses (Figure 10.2). Input data includes the following information: contour lines that describe elevation, terrain information such as lakes and drainage, base camp locations, previous patrol observations (e.g., animal signs, human activities), and previous patrol tracks. Based on the input data, we estimate current animal population distribution. It is necessary to estimate this distribution because individual sightings of animals are not likely to be spatially representative of the population. To estimate the distribution for tigers, the species of interest in recent deployments, we use Just Another Gibbs Sampler (JAGS) (Plummer, 2003) to produce a posterior predictive density raster derived from a spatially explicit capture-recapture analysis conducted in a Bayesian framework.

Because wildlife patrols and poaching attacks happen frequently, we model this recurrent behavior as a repeated game. PAWS first builds the game model, models the poachers' behavior using wildlife crime data and then calculates the optimal patrol strategy according to the game model, behavior model, and additional patrolling constraints (e.g., contour lines, terrain information). When patrollers execute the PAWS patrols over a period of time, they will collect more crime data and those data are fed into PAWS as inputs for the next set of generated PAWS patrols. In the following sections, we focus on how the game model is built, how the poachers' behavior is modeled, and how the optimal patrol strategy is generated.

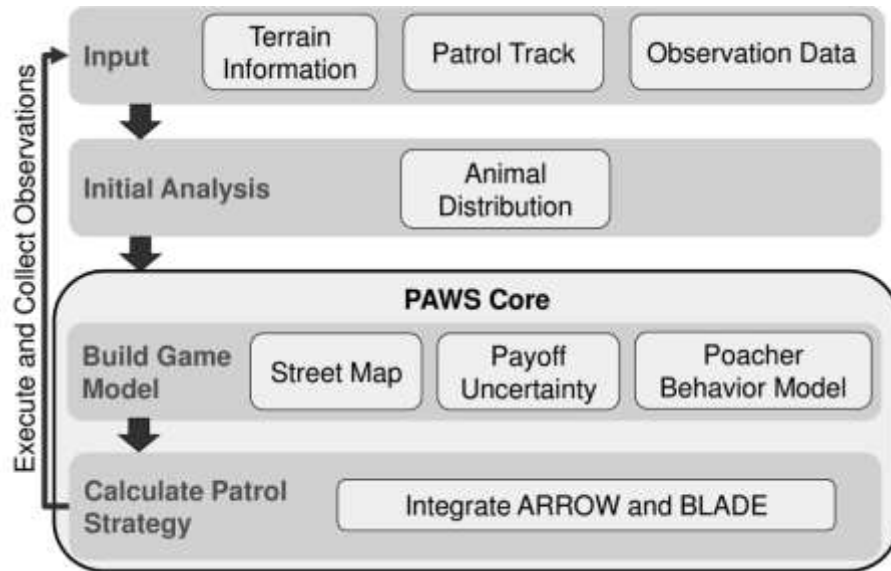


Figure 10.2. PAWS overview.

The basis of game-theoretic analysis in PAWS

The problem of combating poaching can be seen as a game with two types of players, the defender (i.e., the conservation agency's wildlife patrollers) and the attacker (i.e., poachers). The players have conflicting interests and there is a strategic interaction between the players. The defender must account for the attacker's actions and vice versa. To reduce successful poaching activities so as to protect wildlife in the area, the defender conducts randomized patrols against poachers while balancing the priorities of different locations with different animal densities. If patrols are not randomized and are deterministic instead (i.e., no randomization), poachers are able to exploit this predictability and thus circumvent any efforts by the defender. To decide which patrol routes need to be considered and how often each of them should be taken so as to provide maximum protection to wildlife, PAWS builds an evidence-based mathematical game model.

In PAWS, a protected area is discretized into a grid (e.g., 1km by 1km), where each grid cell is viewed as a target location (i.e., targets) for poachers. The defender tries to protect these T targets from poachers by optimally allocating a set of R patrolling resources (i.e., teams of patrollers, where each team can take a different patrol route). Note that R is typically much less than T ; there are typically many more targets to protect than there are resources available to the defender. By executing a patrol route, patrollers can protect the targets along that route. The assignment of resources to patrol routes (i.e., sending out patrollers to patrol) is called the defender's pure strategy. Since our goal is to randomize patrol routes, the defender can choose to execute any of the patrol routes with a certain probability; this probability distribution over pure strategies (i.e., patrol routes) is called the defender's mixed strategy. We compactly represent the defender's mixed strategy as a coverage vector $c = \langle c_i \rangle$ where c_i is the coverage probability for target i (i.e., the probability that target i is protected) (Korzhyk, Conitzer, & Parr, 2010). After committing to a mixed strategy the defender will then randomly select a pure strategy to execute (i.e., send patrollers out on

patrol). The adversary observes the defender's mixed strategy through surveillance and then chooses a target to attack (i.e., places snares at a target). We provide an illustrative example of the game model in Figure 10.3.

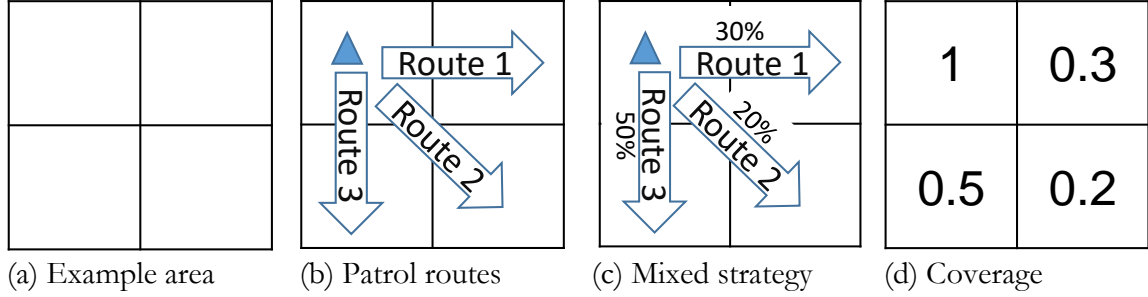


Figure 10.3. Illustrative example for PAWS game model. Figure 10.3(a) shows an example area that is discretized into four cells, i.e., four targets. Figure 10.3(b) shows three possible patrol routes for the example area, each of which starts from the upper left cell (base camp location) and protects two targets. Figure 10.3(c) shows a mixed strategy for the example area when there is only one team of patrollers. The compact representation of the mixed strategy in Figure 10.3(c) is shown in Figure 10.3(d).

Each target is associated with payoff values. Higher animal density implies higher payoffs, which indicate separate reward and penalty values for the defender and the poachers. If a poacher places snares in target i , and i is protected by the patroller (i.e., the executed patrol goes through location i), the defender gets a reward $U_{r,i}^d$, and the poacher receives a penalty $U_{p,i}^a$. Conversely, if target i is not protected, the defender gets a penalty $U_{r,i}^d$ and the poacher receives a reward $U_{p,i}^a$.

Given a defender mixed strategy (compactly represented by $c = \langle c_i \rangle$), the probability that target i is protected by patrollers is c_i , and the probability that target i is not protected is $1 - c_i$. Thus, the poacher's expected payoff (also called expected utility) when he chooses target i is

$$U_i^a = c_i U_{p,i}^a + (1 - c_i) U_{r,i}^a$$

On the other hand, the defender's expected utility with respect to target i is

$$U_i^d = c_i U_{r,i}^d + (1 - c_i) U_{p,i}^d$$

Each player in the game seeks to maximize their expected utility; defenders want to stop attacks as much as possible, and attackers want to successfully attack as much as possible. In PAWS, the game is zero-sum, $U_{r,i}^d = -U_{p,i}^a$, $U_{p,i}^d = -U_{r,i}^a$. In other words, whatever the poacher gains by successfully poaching is also how much the defender loses. Since placing snares in a location with high animal density has a higher chance to lead to a successful capture of an animal (and is thus more rewarding to the poacher), the poacher's reward $U_{r,i}^a$ is determined by animal density. When there are multiple poachers, the defender's expected utility is the overall expected utility against all the poachers (see Table 1 for notions).

Table 1: Notations used in this chapter

Notation	Meaning
T	Number of targets
R	Number of defender resouces (teams of patrollers)
$U_{r,i}^d$	Reward for defender if target i is selected by poacher and is protected
$U_{p,i}^a$	Penalty for poacher if target i is selected by poacher and is protected
$U_{p,i}^d$	Penalty for defender if target i is selected by poacher and is not protected
$U_{r,i}^a$	Reward for poacher if target i is selected by poacher and is not protected
U_i^a	Poacher's expected utility of selecting target i if strategy c is played by defender
U_i^d	Defender's expected utility of playing strategy c if target i is selected by poacher
c_i	Coverage probability on target i . Probability that target i is protected by patroller
ω	Parameter of the SUQR model. $\omega = (\omega_1, \omega_2, \omega_3)$ where $\omega_1, \omega_2, \omega_3$ are the coefficients for $c_i, U_{r,i}^a$ and $U_{p,i}^a$ respectively

Modeling human behavior for PAWS

In SSGs, the adversary's behavior model represents how he will choose a target to attack in response to the defender's mixed strategy. Past work has often assumed that the adversary will act with perfect rationality and will always choose the target with the highest expected utility for him (Pita et al., 2008). However, it is well known that humans do not act in such a way (Camerer, 2003), and bounded rationality models seek to model how humans act (i.e., not always choosing the target with highest expected utility). PAWS is the first deployed application that uses the Subjective Utility Quantal Response model (SUQR), which is a bounded rationality model. SUQR models the adversary's probabilistic response (i.e., with some probability, attack a target) to the defender's strategy and was shown to perform the best in human subject experiments when compared with other bounded rationality models (Nguyen, Yang, Azaria, Kraus, & Tambe, 2013). SUQR suggest adversaries evaluate targets based on a linear combination of multiple observable features: the probability that a target will be protected and thus the probability that the adversary will be caught, the reward for a successful attack at a target for the adversary, and the penalty for being captured at a target for the adversary. The value of $\omega_1 c_i + \omega_2 U_{r,i}^a + \omega_3 U_{p,i}^a$ is denoted as the subjective utility

where $\omega_1, \omega_2, \omega_3$ are the parameters indicating the importance of the aforementioned features. The intuition of the SUQR model is that the adversary will, with higher probability, attack targets with higher subjective utility.

Incorporating learning into the behavioral model

Next, we discuss our method for incorporating past crime data to model the poachers' behavior. Each poacher in the population is modeled with the SUQR model, and we use the crime data to learn the values of each poacher's SUQR model's parameters $\omega = (\omega_1, \omega_2, \omega_3)$. Because we are modeling each poacher's preferences individually and we need to create our strategies with a population of poachers in mind, we compute a distribution of the poacher population's SUQR model parameters. We assume the distribution is normally distributed. To learn this distribution, we must incorporate collected crime data. As discussed previously, we have two types of data that are collected by patrollers: identified data and anonymous data. When a poacher is captured, the patrollers question the poacher and get him to confess to any previous attacks that he committed in the protected area. As such, an identified data point can consist of attacks from multiple rounds and can be used to learn an accurate (SUQR) model of the poacher's behavior. Unfortunately, this type of data is sparse, and we thus cannot rely on using only identified data to construct the entire distribution. On the other hand, there are abundant anonymous data. One potential approach is to assume that a new poacher committed this single attack. Because we only have a single attack data point on this "new poacher," it is impossible to learn that poacher's model in the same way we do with identified data. While anonymous data provides a noisy estimation of an individual poacher's behavioral model, it gives a sufficiently accurate measurement of the crime distribution of the poacher population due to its abundance.

Because each type of data has its strengths or weaknesses, we present the PAWS-Learn algorithm that combines both types of data. By doing so, we can obtain a more accurate distribution than if we just used one type or the other. More detailed explanations of the PAWS-Learn algorithm can be found in Yang et al. (2014).

PAWS-Learn

The PAWS-Learn algorithm combines both identified and anonymous data together to obtain a more accurate distribution of the poachers' behavior than if just one type were used. First, a distribution, $f(\omega)'$, is computed using the identified data, and this computation is done via Maximum Likelihood Estimation (MLE). To summarize this process in the form of a question: given the identified data, what distribution $f(\omega)'$ is most likely (as defined by a likelihood function)? Next, we compute the proportion of anonymous crimes committed at each target. Finally, we compute a new distribution, $f(\omega)$, such that the mean squared error (MSE) is minimized between the identified data distribution ($f(\omega)'$) and the proportion of anonymous crimes. This new distribution, $f(\omega)$, is then used by PAWS to compute a new defender strategy for the next round. Combining the two types of data (i.e., identified, anonymous) in a simulation helps the model learn the distribution of poacher behavior faster

than using one or the other type of data (Figure 10.4). As can be seen, by combining both types of data with the PAWS-Learn algorithm, we can learn a more accurate model of poacher behavior and thus generate strategies according to that model.

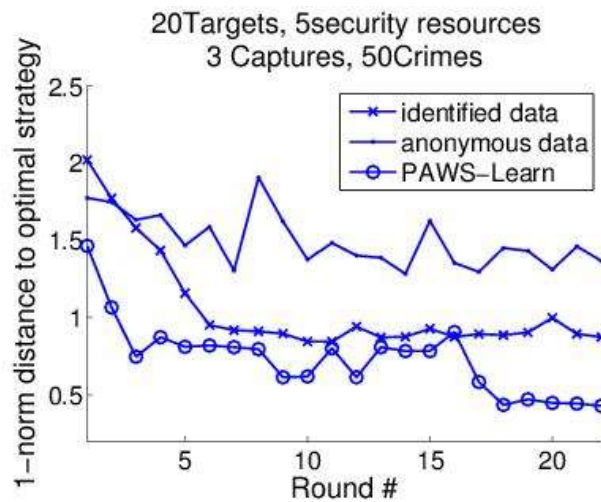


Figure 10.4. Strategy convergence simulation. The X-Axis shows the number of rounds that have elapsed (i.e., the number of interactions between the defender and attacker, and also the number of times a new strategy has been generated). The Y-Axis compares the generated defender mixed strategy in that round to an “optimal” strategy which is generated from the true distribution of poacher behavior (which, outside of a simulation, isn't known); lower values on the axis indicate that the generated strategy is closer to the optimal strategy, indicating that the learned distribution is of higher quality. The first two lines in the figure correspond to learning a distribution with only one type of data available (only identified or anonymous data, respectively) whereas the last line corresponds to the PAWS-Learn algorithm (both data types are combined).

Domain feature modeling

PAWS aims to generate optimal patrol strategies that accurately reflect the many factors that influence the creation of patrols. To that end, providing the game-theoretic analysis and modeling the poacher behavior is not adequate. A key missing piece is to capture and integrate important domain features of wildlife crime into PAWS. In this section, we discuss how domain features are incorporated, which leads to the regular deployment of PAWS in the field.

Terrain information. The first important domain feature is the terrain information. The critical importance of topographic information was ignored at the beginning when PAWS was first proposed (Yang et al., 2014) and was identified by patrollers during first tests of PAWS in the tropical forests in Southeast Asia. Topography can affect patrollers’ speed in key ways. For example, lakes are inaccessible for foot patrols. Not considering such information may lead to the failure of completing the patrol route. Also, changes in elevation require extra patrol effort and extreme changes may stop the patrollers from following a route. In addition, in areas with complex terrain, it is necessary to focus on terrain features such as ridgelines and streams when planning routes for three reasons: (a) they are important conduits for certain mammal species such as tigers; (b) hence, poachers use these features

for trapping and moving about in general; and (c) patrollers find it easier to move around here than on slopes. To incorporate terrain information we use a hierarchical modeling approach to build a virtual “street map” of a protected area. This virtual street map allows PAWS to scale-up to patrol large areas while providing fine-grained guidance. We first discretize a protected area into 1km by 1km grid cells and treat every grid cell as a target. We further discretize the grid cells into 50m by 50m raster pieces and describe the topographic information such as elevation in 50m scale. The virtual street map is built in the terms of raster pieces while aided by the grid cells in this abstraction as described below. With this hierarchical modeling, the model keeps a small number of targets and reduces the number of patrol routes while allowing for details at the 50m scale. The street map is a graph consisting of nodes and edges, where the set of nodes is a small subset of the raster pieces and edges are sequences of raster pieces linking the nodes. We denote nodes as Key Access Points (KAPs) and edges as route segments. The street map not only helps scalability but also allows us to focus patrolling on preferred terrain features such as ridgelines. The street map is built in three steps

1. Determine the accessibility type for each raster piece.
2. Define KAPs.
3. Find route segments to link the KAPs.

Patrolling constraint. In practice, total patrolling time is limited and the patrollers can only move to nearby areas. Also, it is common that there is a base location where patrol routes start from and end with. Such patrolling constraints should be taken into account when designing the patrol routes. Given the street map, these constraints can be easily incorporated to the model by restricting defender’s pure strategy to be a patrol route on the street map, starting from the base camp, walking along route segments and ending with base camp, with its total distance satisfying the patrol distance limit.

Uncertainty in animal distribution. PAWS models a zero-sum game and the reward for the attacker, and the penalty for the defender, is ultimately dictated by animal distribution. However, key domain features such as animal density that contribute to the payoffs are difficult to estimate precisely, leading to uncertainties in the payoff in the game model. Not considering such uncertainty may lead to high degradation in patrol quality. PAWS considers such uncertainty and calculates an optimal patrol strategy that is robust against the uncertainty. In PAWS, the algorithm for calculating the optimal strategy is based on algorithms from the rich security game literature. Specifically, we integrate an algorithm ARROW (Nguyen et al., 2015) with another algorithm BLADE (Yang, Jiang, Tambe, Ordonez, 2013) to fit the need of the problem: (a) we must generate patrol routes over the street map over the entire conservation area region while; (b) simultaneously addressing payoff uncertainty; and (c) bounded rationality of the adversary. To delineate the algorithm in detail is beyond the scope of this chapter; interested readers may explore the aforementioned references.

Discussion

Antipoaching patrols, while essential for combating wildlife poaching, are difficult to create and introduce a large cognitive burden on patrol managers due to the large number of

factors involved. In addition, it is challenging to create patrols that are randomized, but without doing so, patrols run the risk of becoming predictable and will thus be less effective at deterring and stopping poachers. Resolving these problems requires the efforts of people in many different disciplines; this chapter describes the innovative SSG based application PAWS, the result of joint efforts among computer scientists, conservationists, and criminologists.

PAWS seeks to aid patrol managers by addressing the challenges of creating effective and efficient wildlife patrols; PAWS patrols are intelligently randomized and incorporate many important factors such as terrain information, poacher behavior, and time constraints. Although PAWS cannot completely replace skilled patrol managers' experience and intuition, PAWS provide sufficient aid to conservation agencies in making intelligent decisions. That being said, PAWS represents a novel approach to generating patrols that highlights the importance of: (a) modeling human behavior, so as to better predict where poachers will attack and how they will react to patrols; (b) remaining unpredictable to poachers; (c) modeling the strategic interaction between rangers and poachers.

PAWS poses multiple implications for conservation practice. First, PAWS brings in game-theoretic perspective in designing and analyzing defender's patrol strategy. This proactive information about our adversaries is beneficial because they react to different policies designed to reduce risks to biodiversity and livelihoods and change their behavior accordingly. The overall approach may be applied to diverse conservation policies given the common strategic interaction between conservation policy and noncompliant individuals or groups..

The model emphasizes the importance of human behavior in understanding the effects of conservation policies designed to reduce risk. Human beings are not perfectly rational in most cases and their behavior can affect the effectiveness of conservation policy. PAWS adopts the SUQR model for modeling poachers' bounded rationality and this model may have more general applications in conservation problems. In addition, the method for learning the parameters in SUQR also has broader applications where data is available.

PAWS is the first application of GSG-based research for protecting wildlife and conserving global biodiversity. PAWS has been tested on grassy plains in QENP in Uganda and went through a significant evolution to be deployed in a Southeast Asia protected area with complex terrains; PAWS has proven capable at functioning in areas with diverse terrain, ecosystems, species and human cultures. Based on this success, PAWS shows promise for future deployments to different sites through collaborations with different organizations. It should be noted that future deployments of PAWS will require initial investment by each site and collaborating organization. For past deployments, this investment took the form of aggregating past patrol data, past poaching data, and terrain information for the purpose of training the PAWS models.

PAWS shows the potential that algorithms and techniques in computer science can have in combatting poaching. Indeed, PAWS focuses on one facet of this large, complex problem, namely how to make the most efficient use of foot patrols. AI-based solutions hold promise for addressing additional facets of this issue, such as investigating and interfering with the trafficking of illegal wildlife products.

References

- Camerer, C. (2003). *Behavioral game theory: Experiments in strategic interaction*. Princeton University Press.
- City of New York (2013). NYPD website- frequently asked questions. *NYPD*. Retrieved from http://www.nyc.gov/html/nypd/html/faq/faq_police.shtml
- Fang, F., Jiang, A. X., & Tambe, M. (2013). Optimal patrol strategy for protecting moving targets with multiple mobile resources. In *Proceedings of the 2013 international conference on Autonomous agents and multi-agent systems*. International Foundation for Autonomous Agents and Multiagent Systems.
- Fang, F., Stone, P., & Tambe, M. (2015). When security games go green: Designing defender strategies to prevent poaching and illegal fishing. In *International Joint Conference on Artificial Intelligence (IJCAI)*.
- Hamisi, M. (2008). *Identification and mapping risk areas for zebra poaching: A case of Tarangire National Park, Tanzania* (Doctoral dissertation, Thesis, ITC).
- Haskell, W. B., Kar, D., Fang, F., Tambe, M., Cheung, S., & Denicola, E. (2014). Robust Protection of Fisheries with COMPASS. In *Proceedings of the Twenty-Sixth Annual Conference on Innovative Applications of Artificial Intelligence* (pp. 2978-2983).
- Holmern, T., Muya, J., & Røskaft, E. (2007). Local law enforcement and illegal bushmeat hunting outside the Serengeti National Park, Tanzania. *Environmental Conservation*, 34, 55-63.
- Kar, D., Fang, F., Delle Fave, F., Sintov, N., & Tambe, M. (2015). A game of thrones: when human behavior models compete in repeated Stackelberg security games. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems*. International Foundation for Autonomous Agents and Multiagent Systems.
- Korzhyk, D., Conitzer, V., & Parr, R. (2010). Complexity of Computing Optimal Stackelberg Strategies in Security Resource Allocation Games. In *Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence* (pp. 805-810).
- Montesh, M. (2013). Rhino poaching: A new form of organized crime. Technical report, College of Law Research and Innovation Committee of the University of South Africa, vol, 27, 1-23.
- Moreto, W. (2013). *To conserve and protect: Examining law enforcement ranger culture and operations in Queen Elizabeth National Park, Uganda*. Doctoral dissertation, Rutgers University-Graduate School-Newark.
- Myerson, R. B. (2013). *Game theory*. Cambridge, MA: Harvard University Press.
- National Wildlife Refuge Association. (2015). *Global anti-poaching act seeks to restrict illicit wildlife trade*. Retrieved from <http://refugeassociation.org/2015/07/global-anti-poaching-act-seeks-to-restrict-illicit-wildlife-trade/>
- Nguyen, T. H., Delle Fave, F. M., Kar, D., Lakshminarayanan, A. S., Yadav, A., Tambe, M., ... & Rwetsiba, A. (2015). Making the most of our regrets: Regret-based solutions to handle payoff uncertainty and elicitation in green security games. In Baras, J. S., Katz, J., & Altman, E. (Eds.). In *Decision and Game Theory for Security* (pp. 170-191). Springer International Publishing.

- Nguyen, T. H., Yang, R., Azaria, A., Kraus, S., & Tambe, M. (2013). Analyzing the Effectiveness of Adversary Modeling in Security Games. In *Proceedings of the Twenty-Seventh AAAI Conference on Artificial Intelligence* (pp. 718-724).
- Ouko, E. (2013). *Where, when, and why are there elephant poaching hotspots in Kenya*. Doctoral dissertation. Netherlands: University of Twente.
- Pires, S. F. (2012). *The illegal parrot trade in the neo-tropics: The relationship between poaching and illicit pet markets*. Doctoral dissertation. Newark, NJ: Rutgers University.
- Pita, J., Jain, M., Marecki, J., Ordóñez, F., Portway, C., Tambe, M., ... & Kraus, S. (2008). Deployed ARMOR protection: the application of a game theoretic model for security at the Los Angeles International Airport. In *Proceedings of the 7th international joint conference on Autonomous agents and multiagent systems: industrial track* (pp. 125-132). International Foundation for Autonomous Agents and Multiagent Systems.
- Plummer, M. (2003). JAGS: A program for analysis of Bayesian graphical models using Gibbs sampling. In *Proceedings of the 3rd international workshop on distributed statistical computing*. Wien, Austria: Technische Universität Wien.
- Qian, Y., Haskell, W. B., Jiang, A. X., & Tambe, M. (2014). Online planning for optimal protector strategies in resource conservation games. In *Proceedings of the 2014 international conference on Autonomous agents and multi-agent systems* (pp. 733-740). International Foundation for Autonomous Agents and Multiagent Systems.
- Save the Rhino International. (2015). *Poaching statistics*. Retrieved from https://www.savetherhino.org/rhino_info/poaching_statistics
- Secretariat, G.T.I. (2013). *Global tiger recovery program implementation plan: 2013-14*. Report, The World Bank: Washington, D.C. Retrieved from <http://globaltigerinitiative.org/publication/global-tiger-recovery-program-implementation-plan-2013-14/>
- Shieh, E., An, B., Yang, R., Tambe, M., Baldwin, C., DiRenzo, J., ... & Meyer, G. (2012). Protect: A deployed game theoretic system to protect the ports of the united states. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems - Volume 1* (pp. 13-20). International Foundation for Autonomous Agents and Multiagent Systems.
- Smart Collaboration (2015). Retrieved from <http://www.smartconservationsoftware.org/>
- Stokes, E. J. (2010). Improving effectiveness of protection efforts in tiger source sites: developing a framework for law enforcement monitoring using MIST. *Integrative Zoology*, 5, 363-377.
- Tambe, M. (2011). *Security and game theory: Algorithms, deployed systems, lessons learned*. Cambridge, United Kingdom: Cambridge University Press.
- Tsai, J., Rathi, S., Kiekintveld, C., Ordóñez, F., Tambe, M. (2009). IRIS - a tool for strategic security allocation in transportation networks. In *Proceedings of 8th International Conference on Autonomous Agents and Multiagent Systems - Industry Track* (pp. 1327-1334). International Foundation for Autonomous Agents and Multiagent Systems.
- Wagenaar, W. A. (1972). Generation of random sequences by human subjects: A critical survey of literature. *Psychological Bulletin*, 77, 65.
- Wato, Y. A., Wahungu, G. M., & Okello, M. M. (2006). Correlates of wildlife snaring patterns in Tsavo West National Park, Kenya. *Biological Conservation*, 132, 500-509.
- Yang, R., Ford, B., Tambe, M., & Lemieux, A. (2014). Adaptive resource allocation for wildlife protection against illegal poachers. In *Proceedings of the 2014 international*

- conference on Autonomous agents and multi-agent systems* (pp.453-460). International Foundation for Autonomous Agents and Multiagent Systems.
- Yang, R., Jiang, A. X., Tambe, M., & Ordonez, F. (2013). Scaling-up security games with boundedly rational adversaries: A cutting-plane approach. In *Proceedings of the Twenty-Third International Joint Conference on Artificial Intelligence* (pp. 404-410).
- Yin, Z., Jiang, A. X., Tambe, M., Kiekintveld, C., Leyton-Brown, K., Sandholm, T., & Sullivan, J. P. (2012). TRUSTS: Scheduling randomized patrols for fare inspection in transit systems using game theory. *AI Magazine*, 33, 59.