

PAWS — A Deployed Game-Theoretic Application to Combat Poaching

Fei Fang, Thanh H. Nguyen, Robert Pickles, Wai Y. Lam, Gopalasamy R. Clements, Bo An, Amandeep Singh, Brian C. Schwedock, Milind Tambe, Andrew Lemieux

■ Poaching is considered a major driver for the population drop of key species such as tigers, elephants, and rhinos, which can be detrimental to whole ecosystems. While conducting foot patrols is the most commonly used approach in many countries to prevent poaching, such patrols often do not make the best use of the limited patrolling resources.

This article presents PAWS, a game-theoretic application deployed in Southeast Asia for optimizing foot patrols to combat poaching. In this article, we report on the significant evolution of PAWS from a proposed decision aid introduced in 2014 to a regularly deployed application. We outline key technical advances that lead to PAWS's regular deployment: (1) incorporating complex topographic features, for example, ridgelines, in generating patrol routes; (2) handling uncertainties in species distribution (game-theoretic payoffs); (3) ensuring scalability for patrolling large-scale conservation areas with fine-grained guidance; and (4) handling complex patrol scheduling constraints.

Poaching is a serious threat to wildlife conservation and can lead to the extinction of species and destruction of ecosystems. For example, poaching is considered a major driver (Chapron et al. 2008) of why tigers are now found in less than 7 percent of their historical range (Sanderson et al. 2006), with three out of nine tiger subspecies already extinct (IUCN 2015). As a result, efforts have been made by law enforcement agencies in many countries to protect endangered animals from poaching. The most direct and commonly used approach is conducting foot patrols. However, given their limited human resources and the vast area in need of protection, improving the efficiency of patrols remains a major challenge.

Game theory has become a well-established paradigm for addressing complex resource allocation and patrolling problems in security and sustainability domains. Models and algorithms have been proposed and studied extensively in the past decade, forming the general area of security games (Tambe 2011). Furthermore, several security-game-based decision support systems have previously been successfully deployed in protecting critical infrastructure such as airports, ports, and metro trains (Pita et al. 2008; Shieh et al. 2012; Yin et al. 2012). Inspired by the success of these deployments, researchers have begun applying game theory to generating effective patrol strategies in green security domains such as protecting wildlife (Yang et al. 2014; Fang, Stone, and Tambe 2015), preventing overfishing (Haskell et al. 2014, Qian et al. 2014), and illegal logging (Johnson, Fang, and Tambe 2012).

Among these prior works, a novel emerging application called PAWS (protection assistant for wildlife security) (Yang et al. 2014) was introduced as a game-theoretic decision aid to optimize the use of human patrol resources to combat poaching. PAWS was the first of a new wave of proposed applications in the subarea now called *green security games* (Fang, Stone, and Tambe 2015; Kar et al. 2015). Specifically, PAWS solves a repeated Stackelberg security game, where the patrollers (defenders) conduct randomized patrols against poachers (attackers) while balancing the priorities of different locations with different animal densities. Despite its promise, the initial PAWS effort did not test the concept in the field.

This article reports on PAWS's significant evolution over the last two years from a proposed decision aid to a regularly deployed application. We report on the innovations made in PAWS and lessons learned from the first tests in Uganda in spring 2014, through its continued evolution since then, to current deployments in Southeast Asia and plans for future worldwide deployment. In this process, we have worked closely with two nongovernment organizations (Panthera and Rimba) and incorporated extensive feedback from professional patrolling teams. Indeed, the first tests revealed key shortcomings in PAWS's initial algorithms and assumptions (we will henceforth refer to the initial version of PAWS as PAWS-Initial, and to the version after our enhancement as PAWS). First, a major limitation was that PAWS-Initial ignored topographic information. Second, PAWS-Initial assumed animal density and relevant problem features at different locations to be known, ignoring the uncertainty. Third, PAWS-Initial could not scale to provide detailed patrol routes in large conservation areas. Finally, PAWS-Initial failed to consider patrol-scheduling constraints.

In this article, we outline novel research advances that remedy the aforementioned limitations, making it possible to deploy PAWS on a regular basis. First, we incorporate elevation information and land fea-

tures and use a novel hierarchical modeling approach to building a virtual street map of the conservation area. This virtual street map helps scale-up while providing fine-grained guidance and is an innovation that would be useful in many other domains requiring patrolling of large areas. Essentially, the street map connects the whole conservation area through easy-to-follow route segments such as ridgelines, streams, and river banks. The rationale for this comes from the fact that animals, poachers, and patrollers all use these features while moving. To address the second and third limitations, we build on the street map concept with a novel algorithm that uniquely synthesizes two threads of prior work in the security games literature; specifically, the new PAWS algorithm handles payoff uncertainty using the concept of minimax regret (Nguyen et al. 2015), while simultaneously ensuring scalability — using our street maps — through the cutting plane framework (Yang et al. 2013). Finally, we incorporate into PAWS the ability to address constraints such as patrol time limits and starting and ending at the base camp. In the final part of the article, we provide detailed information about the regular deployment of PAWS.

Background and Related Work

Criminologists have worked on the problem of combating poaching, from policy design to illegal trade prevention (Lemieux 2014). Geographic information systems (GIS) experts (Hamisi 2008) and wildlife management staff (Wato, Wahungu, and Okello 2006) have carefully studied the identification of poaching hotspots. In recent years, software tools such as SMART² and MIST (Stokes 2010) have been developed to help conservation managers record data and analyze patrols retrospectively. We work on a complementary problem of optimizing the patrol planning of limited security staff in conservation areas.

In optimizing security resource allocation, previous work on Stackelberg security games (SSGs) has led to many successfully deployed applications for the security of airports, ports, and flights (Pita et al. 2008; Fang, Jiang, and Tambe 2013). Based on the early work on SSGs, recent work has focused on green security games (Kar et al. 2015), providing conceptual advances in integrating learning and planning (Fang, Stone, and Tambe 2015) and the first application to wildlife security, PAWS-Initial. PAWS-Initial (Yang et al. 2014) models the interaction between the patroller (defender) and the poacher (attacker) who places snares in the conservation area (see figure 1) as a basic green security game, that is, a repeated SSG, where every few months, poaching data is analyzed, and a new SSG is set up enabling improved patrolling strategies. The deployed version of PAWS adopts this framework.

We provide a brief review of SSGs, using PAWS as a

key example. In SSGs, the defender protects T targets from an adversary by optimally allocating a set of R resources ($R < T$) (Pita et al. 2008). In PAWS, the defender discretizes the conservation area into a grid, where each grid cell is viewed as a target for poachers, to be protected by a set of patrollers. The defender's pure strategy is an assignment of the resources to targets. The defender can choose a mixed strategy, which is a probability distribution over pure strategies. The defender strategy can be compactly represented as a coverage vector $c = \langle c_i \rangle$ where c_i is the coverage probability, that is, the probability that a defender resource is assigned to be at target i (Korzhyk, Conitzer, and Parr 2010). The adversary observes the defender's mixed strategy through surveillance and then attacks a target. An attack could refer to the poacher, a snare, or some other aspect facilitating poaching (for example, poaching camp). Each target is associated with payoff values that indicate the reward and penalty for the players. If the adversary attacks target i , and i is protected by the defender, the defender gets reward $U_{r,i}^d$ and the adversary receives penalty $U_{p,i}^a$. Conversely, if not protected, the defender gets penalty $U_{p,i}^d$ and the adversary receives reward $U_{r,i}^a$. $U_{r,i}^a$ is usually decided by animal density — higher animal density implies higher payoffs. Given a defender strategy c and the penalty and reward values, we can calculate the players' expected utilities U_i^a and U_i^d when target i is attacked accordingly.

In SSGs, the adversary's behavior model decides his response to the defender's mixed strategy. Past work has often assumed that the adversary is perfectly rational, choosing a single target with the highest expected utility (Pita et al. 2008). PAWS is the first deployed application that relaxes this assumption in favor of a bounded rationality model called SUQR, which models the adversary's stochastic response to defender's strategy (Nguyen et al. 2013). SUQR was shown to perform the best in human subject experiments when compared with other models. SUQR predicts the adversary's probability of attacking i based on a linear combination of three key features at the targets, including the coverage probability c_i , the attacker's reward $U_{r,i}^a$ and penalty $U_{p,i}^a$. A set of parameters (w_1, w_2, w_3) are used for combining the features. They indicate the importance of the features and can be learned from data.

First Tests and Feedback

We first tested PAWS-Initial (Yang et al. 2014) at Uganda's Queen Elizabeth National Park (QENP) for three days. Subsequently, with the collaboration of Panthera and Rimba, we started working in forests in Malaysia in September 2014¹. These protected forests are home to endangered animals such as the Malayan tiger and Asian elephant but are threatened by poachers. One key difference of this site compared to



Figure 1. Snares Found by Patrollers.

QENP is that there are large changes in elevation, and the terrain is much more complex. The first four-day patrol in Malaysia was conducted in November 2014. For this test, we set the value of $U_{r,i}^a$ (input for PAWS-Initial) by aggregating observation data recorded during April 2014–September 2014. We first set the importance value of each cell as a weighted sum of the observed counts of different types of animals and different human activities. We then dilute the importance value of available cells to nearby cells by applying a 5 by 5 Gaussian kernel for a convolution operation so as to estimate the value for cells with no data recorded.

These initial tests revealed four areas of shortcomings, which restricted PAWS-Initial from being used regularly and widely. The first limitation, which was surprising given that it has received no attention in previous work on security games, is the critical importance of topographic information that was ignored in PAWS-Initial. Topography can affect patrollers' speed in key ways. For example, lakes are inaccessible for foot patrols. Not considering such information may lead to failure to complete the patrol route. Figure 2 shows one patrol route during the test in Uganda. The suggested route (orange straight line) goes across the water body (lower right part of figure), and hence the patrollers decided to walk along the water body (black line). Also, changes in elevation require extra patrol effort, and extreme changes may stop the patrollers from following a route. For example, in figure 3a, PAWS-Initial planned a route on a 1 kilometer by 1 kilometer grid (straight lines), and suggested that the patrollers walk to the north area (row 1, column 3) from the south side (row 2, column 3). However, such movement was extremely difficult because of the changes in elevation. So patrollers decided to head toward the




Figure 2. One Patrol Route During the Test in Uganda.

northwest area as the elevation change is more gentle. In addition, it is necessary to focus on terrain features such as ridgelines and streams (figure 3b) when planning routes for three reasons:

First, they are important conduits for certain mammal species such as tigers; hence, second, poachers use these features for trapping and moving about in general; and third, patrollers find it easier to move around here than on slopes. Figure 4a shows a prominent ridgeline.

The second limitation is that PAWS-Initial assumes the payoff values of the targets — for example, $U_{r,i}^a$ — are known and fixed. In the domain of wildlife protection, there can be uncertainties due to animal movement and seasonal changes. Thus, considering payoff uncertainty is necessary for optimizing patrol strategy.

The third limitation is that PAWS-Initial cannot scale to provide detailed patrol routes in large conservation areas, which is necessary for successful deployment. Detailed routes require fine-grained discretization, which leads to an exponential number of routes in total.

The fourth limitation is that PAWS-Initial considers covering individual grid cells, but not feasible routes. In practice, the total patrolling time is limited, and the patrollers can move to nearby areas. A patrol strategy for implementation should be in the form of a distribution over feasible patrol routes satisfying these constraints. Without taking these scheduling (routing) constraints into account, the optimal coverage probabilities calculated by PAWS-Initial may not be implementable. Figure 4b shows an example area that is discretized into four cells and the base camp is located in the upper left cell. There

are three available patrol routes, each protecting two targets. The coverage probabilities shown in figure 4c cannot be achieved by a randomization over the three routes because the coverage of the upper left cell (Target 1) should be no less than the overall coverage of the remaining three cells since all routes start from the base camp.

PAWS Overview

Figure 5 provides an overview of the deployed version of PAWS. PAWS first takes the input data and estimates the animal distribution and human activity distribution. Based on this information, an SSG-based game model is built, and the patrol strategy is calculated. In wildlife protection, there is repeated interaction between patrollers and poachers. When patrollers execute the patrol strategy generated by PAWS over a period (for example, three months), more information is collected and can become part of the input in the next round.

PAWS provides significant innovations in addressing the aforementioned limitations of PAWS-Initial. In building the game model, PAWS uses a novel hierarchical modeling approach to building a virtual street map, while incorporating detailed topographic information. PAWS models the poachers bounded rationality as described by the SUQR model and considers uncertainty in payoff values. In calculating the patrol strategy, PAWS uses the ARROW (Nguyen et al. 2015) algorithm to deal with payoff uncertainty and adopts cutting plane approach and column generation to address the scalability issue introduced by scheduling constraints.

Input and Initial Analysis

The input information includes contour lines that describe the elevation, terrain information such as lakes and drainage, base camp locations, previous observations (animals and human activities), as well as previous patrol tracks. However, the point detections of the presence of animal and human activity are not likely to be spatially representative. As such, it is necessary to predict the animal and human activity distribution over the entire study area. To this end, we used (1) JAGS (Plummer 2003) to produce a posterior predictive density raster for tigers (as a target species) derived from a spatially explicit capture-recapture analysis conducted in a Bayesian framework; and (2) MaxEnt (Phillips, Anderson, and Schapire 2006) to create a raster of predicted human activity distribution based on meaningful geographical covariates (for example, distance to water, slope, elevation) in a maximum entropy modeling framework.

Build Game Model

Based on the input information and the estimated distribution, we build a game model abstracting the

strategic interaction between the patroller and the poacher as an SSG. Building a game model involves defender action modeling, adversary action modeling, and payoff modeling. We will discuss all three parts but emphasize defender action modeling since this is one of the major challenges to bring PAWS to a regularly deployed application. Given the topographic information, modeling defender actions in PAWS is far more complex than any other previous security game domain.

Defender Action Modeling

Based on the feedback from the first tests, we aim to provide detailed guidance to the patrollers. If we use a fine-grained grid and treat every fine-grained grid cell as a target, computing the optimal patrolling strategy is exceptionally computationally challenging due to the large number of targets and the exponential number of patrol routes. Therefore, a key novelty of PAWS is to provide a hierarchical modeling solution, the first such model in security game research. This hierarchical modeling approach allows us to attain a good compromise between scaling up and providing detailed guidance. This approach would be applicable in many other domains for large open area patrolling where security games are applicable, not only other green security games applications, but others including patrolling of large warehouse areas or large open campuses by robots or UAVs.

More specifically, we leverage insights from hierarchical abstraction for heuristic search such as path planning (Botea, Müller, and Schaeffer 2004) and apply two levels of discretization to the conservation area. We first discretize the conservation area into 1 kilometer by 1 kilometer grid cells and treat every grid cell as a target. We further discretize the grid cells into 50 meters by 50 meters raster pieces and describe the topographic information such as elevation in 50-meter scale. The defender actions are patrol routes defined over a virtual “street map” — which is built in 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 50-meter 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, and (3) find route segments to link the KAPs.

In the first step, we check the accessibility type of every raster piece. For example, raster pieces in a lake are inaccessible, whereas raster pieces on ridgelines or

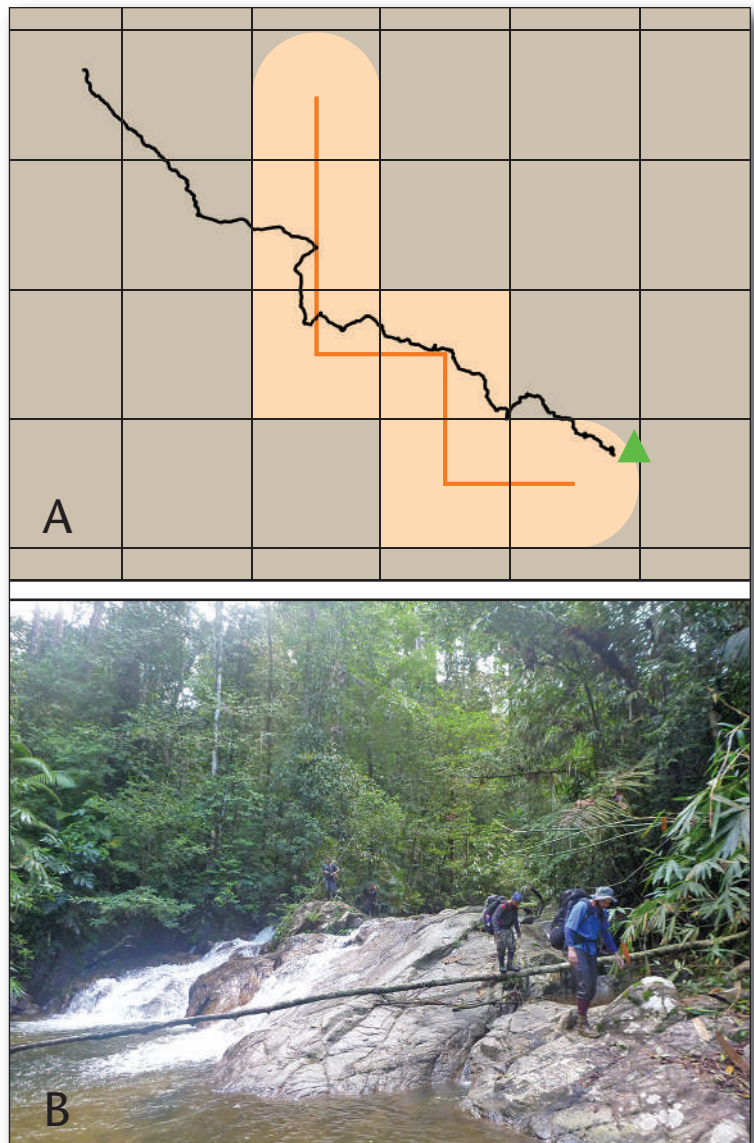


Figure 3. First Four-Day Patrol in Malaysia.

Figure 3a shows one suggested route (orange straight lines) and the actual patrol track (black line). Figure 3b shows the patrollers walking along the stream during the patrol.

previous patrol tracks are easily accessible. Ridgelines and valley lines are inferred from the contour lines using existing approaches in hydrology (Tarboton, Bras, and Rodriguez-Iturbe 2007).

The second step is to define a set of KAPs, through which patrols will be routed. We want to build the street map in such a way that each grid cell can be reached. So we first choose raster pieces that can serve as entries and exits for the grid cells as KAPs, that is, the ones that are on the boundary of grid cells and are easily accessible according to the accessibility type calculated in the first step. When there are

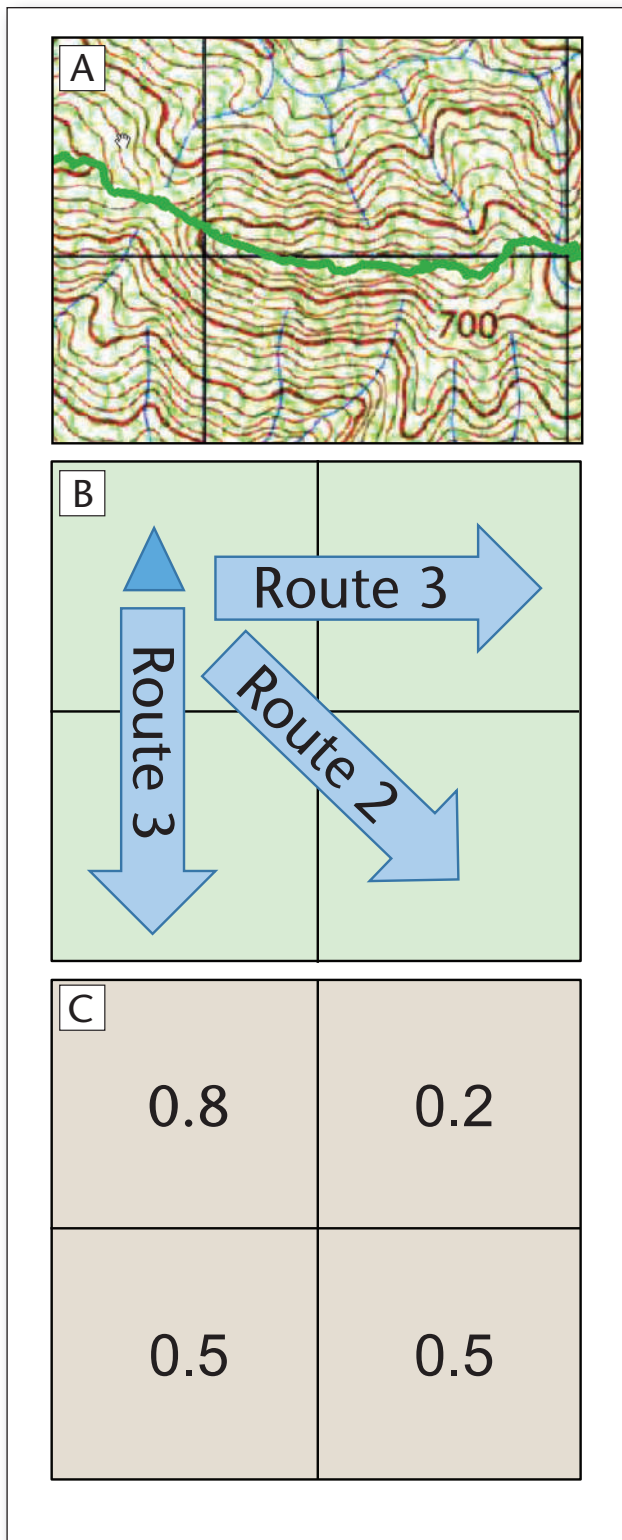


Figure 4. Illustrative Examples.
a. Ridgeline. b. Feasible routes. c. Coverage

multiple adjacent raster pieces that are all easily accessible, we add the midpoint as a KAP. If there are no easily accessible raster pieces on one side of the boundary, we choose the raster piece with the lowest slope as a KAP. In addition, we consider existing base camps as KAPs as they are key points in planning the patroller's route. We choose additional KAPs to ensure KAPs on the boundary of adjacent cells are paired. Figure 6 shows identified KAPs and easily accessible pieces (black and gray raster pieces respectively).

The last step is to find route segments to connect the KAPs. Instead of inefficiently finding route segments to connect each pair of KAPs on the map globally, we find route segments locally for each pair of KAPs within the same grid cell, which is sufficient to connect all the KAPs. When finding the route segment, we design a distance measure that estimates the actual patrol effort and also gives high priority to the preferred terrain features. The effort needed for three-dimensional movement can be interpreted as the equivalent distance on flat terrain. For example, for gentle slopes, equivalent "flat-terrain" distance is obtained by adding eight kilometers for every one kilometer of elevation ascent according to Naismith's rule (Thompson 2011). In PAWS, we apply Naismith's rule with Langmuir corrections (Langmuir 1995) for gentle slopes ($< 20^\circ$) and apply Tobler's hiking speed function (Tobler 1993) for steep slopes ($\geq 20^\circ$). Very steep slopes ($> 30^\circ$) are not allowed. We penalize not walking on preferred terrain features by adding extra distance. Given the distance measure, the route segment is defined as the shortest distance path linking two KAPs within the grid cell.

The defender's pure strategy is defined as 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 (all measured as the distance on flat terrain). The patroller confiscates the snares along the route and thus protects the grid cells. More specifically, if the patroller walks along a route segment that covers a sufficiently large portion (for example, 50 percent of animal distribution) of a grid cell, the cell is considered to be protected. The defender's goal is to find an optimal mixed patrol strategy — a probability distribution over patrol routes.

Poacher Action Modeling and Payoff Modeling

The poacher's actions are defined over the grid cells to aid scalability. In this game, we assume the poacher can observe the defender's mixed strategy and then chooses one target (a grid cell) and places snares in this target. Following earlier work, the poacher in this game is assumed to be boundedly rational, and his actions can be described by the SUQR model.

Each target is associated with payoff values indicating the reward and penalty for the patrollers and the poachers. As mentioned earlier, PAWS models a zero-sum game and the reward for the attacker (and

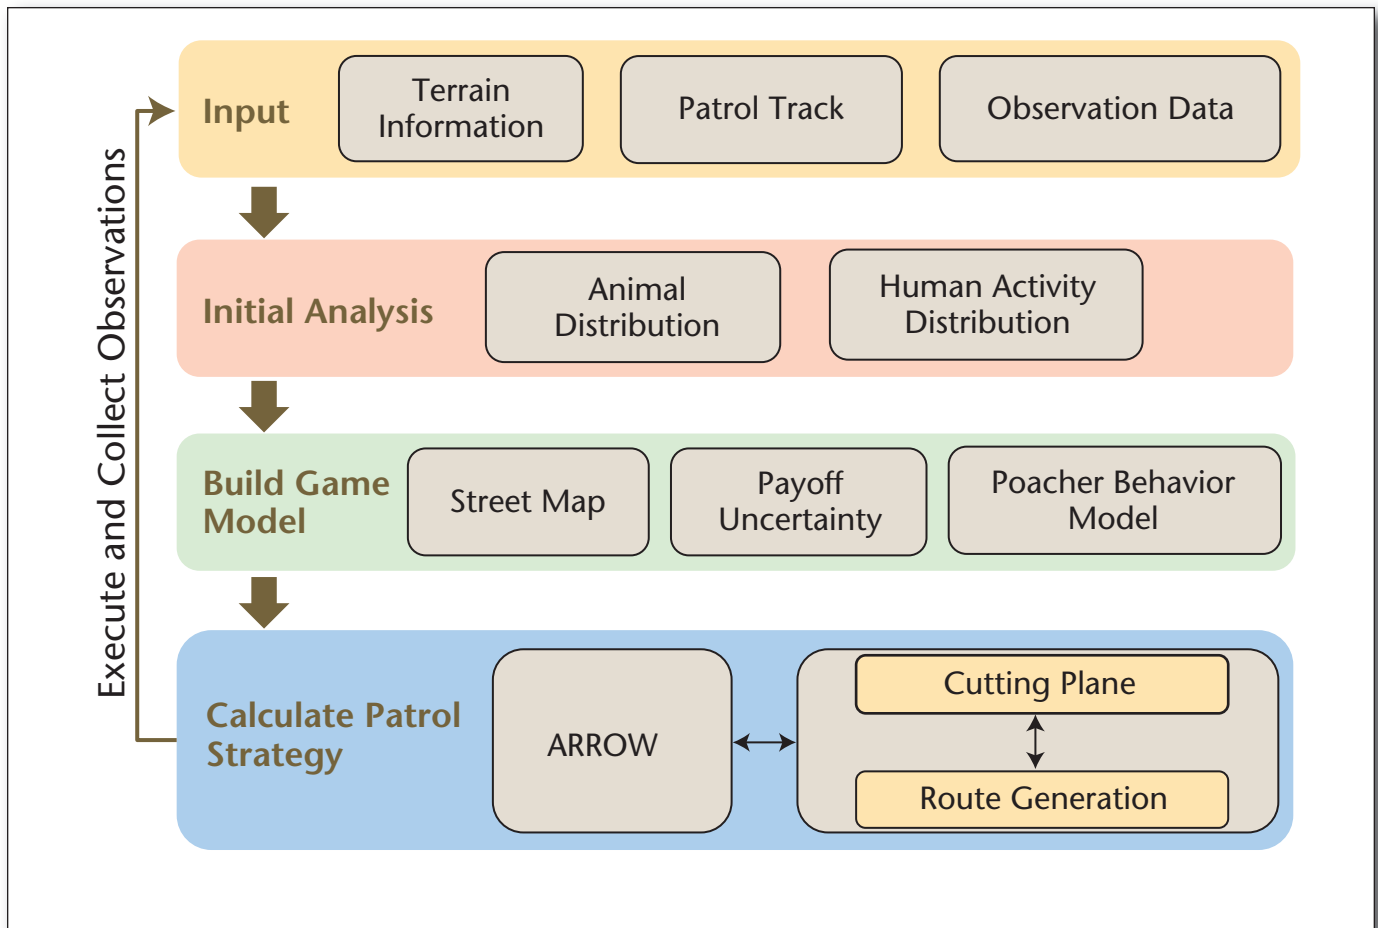


Figure 5. PAWS Overview.

the penalty for the defender) is decided by the animal distribution. However, in this game model, we need to handle uncertainty in the players' payoff values since key domain features, such as animal density, that contribute to the payoffs are difficult to precisely estimate. In addition, seasonal or dynamic animal migration may lead to payoffs to become uncertain in the next season. We use intervals to represent payoff uncertainty in PAWS; the payoffs are known to lie within a certain interval whereas the exact values are unknown. Interval uncertainty is, in fact, a well-known concept to capture uncertainty in security games (Nguyen et al. 2014, 2015). We determine the size of the payoff intervals at each grid cell based on patrollers' patrol efforts at that cell. If the patrollers patrol a cell more frequently, there is less uncertainty in the players' payoffs at that target and thus a smaller size of the payoff intervals.

Calculate Patrol Strategy

We build on algorithms from the rich security game literature to optimize the defender strategy. However, we find that no existing algorithm directly fits our

needs as we need an algorithm that can scale up to the size of the domain of interest, where (1) we must generate patrol routes over the street map over the entire conservation area region, while (2) simultaneously addressing payoff uncertainty and (3) bounded rationality of the adversary. While the ARROW (Nguyen et al. 2015) algorithm allows us to address (2) and (3) together, it cannot handle scale-up over the street map. Indeed, while the (virtual) street map is of tremendous value in scaling up as discussed earlier, scaling up given all possible routes (approximately equal to 1012 routes) on the street map is still a massive research challenge. We, therefore, integrate ARROW with another algorithm BLADE (Yang et al. 2013) for addressing the scalability issue, resulting in a novel algorithm that can handle all the three aforementioned challenges. The new algorithm is outlined in figure 7. In the following, we explain how ARROW and BLADE are adapted and integrated.

ARROW attempts to compute a strategy that is robust to payoff uncertainty given that poachers' responses follow SUQR. The concept of minimizing maximum regret is a well-known concept in AI for




Figure 6. KAPs (Black) for 2 by 2 Grid Cells.

decision making under uncertainty (Wang and Bouilrier 2003). ARROW uses the solution concept of behavioral minimax regret to provide the strategy that minimizes regret or utility loss for the patrollers in the presence of payoff uncertainty and bounded rational attackers. In small-scale domains, ARROW could be provided all the routes (the defender's pure strategies), on the basis of which it would calculate the PAWS solution — a distribution over the routes. Unfortunately, in large-scale domains like ours, enumerating all the routes is infeasible. We must, therefore, turn to an approach of incremental solution generation, which is where it interfaces with the BLADE framework.

More specifically, for scalability reasons, ARROW first generates the robust strategy for the patrollers in the form of coverage probabilities over the grid cells without consideration of any routes. Similar to BLADE, a separation oracle is then called to check if the coverage vector is implementable. If it is implementable, the oracle returns a probability distribution over patrol routes that implements the coverage vector, which is the desired PAWS solution. If it is not implementable (see figure 4c for an example of a coverage vector that is not implementable) the oracle returns a constraint (cutting plane) that informs ARROW why it is not. For the example in figure 4b and 4c, if ARROW generates a vector as shown in figure 4c, the constraint returned could be

$$c_1 \geq \sum_{i=2}^4 c_i$$

since all implementable coverage vectors should satisfy this constraint. This constraint helps ARROW refine its solution. The process repeats until the coverage vector generated by ARROW is implementable.

As described in BLADE (Yang et al. 2013), to avoid enumerating all the feasible routes to check whether the coverage vector is implementable, the separation oracle iteratively generates routes until it has just enough routes (usually after a small number of iterations) to match the coverage vector probabilities or get the constraint (cutting plane). At each iteration of route generation (shown in the bottommost box in figure 7), the new route is optimized to cover targets of high value. However, we cannot directly use any existing algorithm to find the optimal route at each iteration due to the presence of our street map. But we note similarities to the well-studied orienteering problem (Vansteenwegen, Souffriau, and Oudheusden 2011) and exploit the insight of the S-algorithm for orienteering (Tsiligiridis 1984).

In particular, in this bottommost box in figure 7, to ensure each route returned is of high quality, we run a local search over a large number of routes and return the one with the highest total value. In every iteration, we start from the base KAP and choose which KAP to visit next through a weighted random selection. The next KAP to be visited can be any KAP on the map, and we assume the patroller will take the shortest path from the current KAP to the next KAP. The weight of each candidate KAP is proportional to the ratio of the additional target value that can be accrued and distance from the current KAP. We set the lower bound of the weight to be a small positive value to make sure every feasible route can be chosen with positive probability. The process continues until the patroller has to go back to the base to meet the patrol distance limit constraint. Given a large number of such routes, our algorithm returns a route close to the optimal solution.

Integrating all these algorithms, PAWS calculates the patrol strategy consisting of a set of patrol routes and the corresponding probability for taking them.

Addressing Additional Practical Challenges

We have introduced the technical innovations that lead to PAWS's deployment. In addition to these innovations, we have addressed a number of practical constraints to make the strategy suggested by PAWS easier to follow by human patrollers. In this section, we summarize these challenges and our solutions to them.

First, mountaintops should be considered as key points in the patrol route. In PAWS, we require that the patrollers always move between KAPs, which are located at the boundary of the grid cells or the base

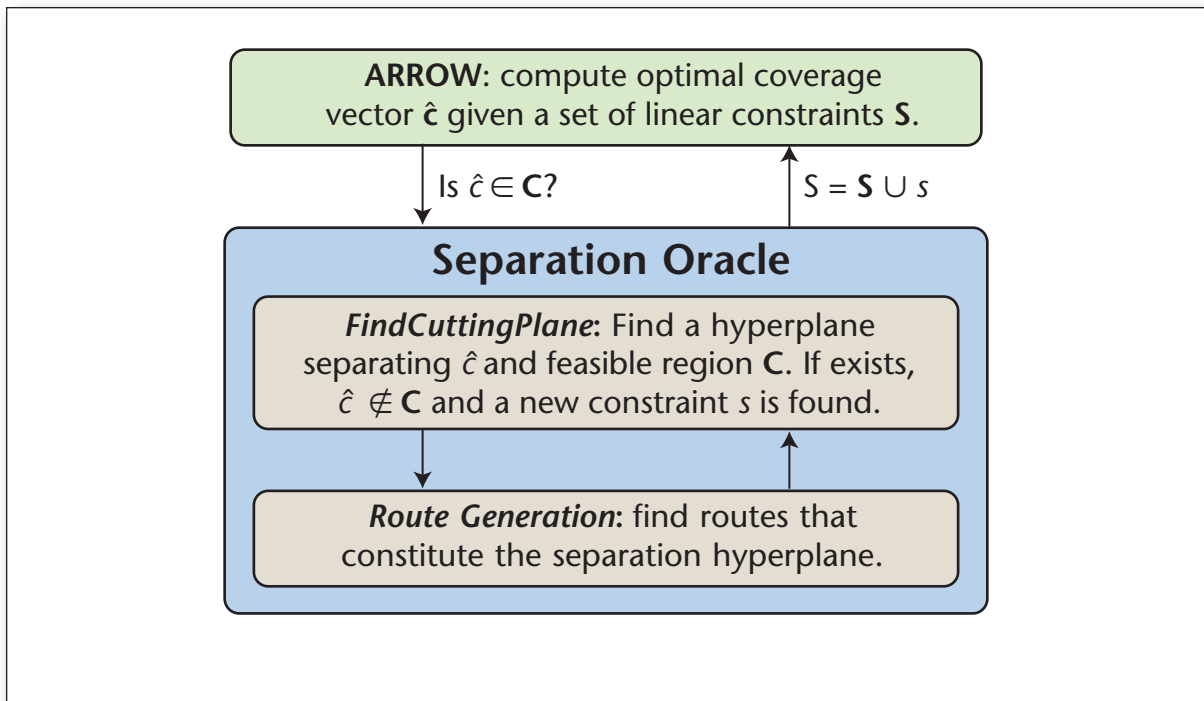


Figure 7. New Integrated Algorithm.

camps. Therefore, in some suggested patrol routes, the patroller is asked to go to a mountaintop and go downhill for a short distance and then backtrack. However, this kind of short downhill followed by returning uphill will annoy the patrollers, and they will naturally ignore the downhill part. We address this problem by considering mountaintops as KAPs when building the street map. With these additional KAPs, patrollers are not forced to take the short downhill unless necessary (that is, when the short downhill covers areas with high animal density).

Second, there is a limit on working time in addition to the limit on walking distance. It takes less time for the patroller to go to an area and then backtrack than to take a loop even if the walking distance is the same. The reason is the patrollers need to spend the time to record what they observe, including animal signs and human activity signs. If the patrollers walk along the same ridgeline twice in a day, they only need to record the signs once. Therefore, in designing the patrol routes, we should consider the total working time in addition to the total distance. This is implemented in PAWS by adding additional constraints in route generation.

Third, not all terrain features should be treated in the same way. In building the street map, we give preference to terrain features like ridgelines by designing a distance measure that penalizes not walking along the terrain feature. However, how much priority should be given to the terrain features depends on the cost of the alternative routes, or how

much easier it is compared to taking other routes. On the one hand, in a very hilly region of the area where there are large elevation changes, the patrollers would highly prefer the terrain features as it is much easier to walk along them than taking an alternative route. On the other hand, if the elevation change in the region is small, the effort of taking a ridgeline for unit distance is comparable to that of taking an alternative route. To differentiate these different cases, we use secondary derivatives to check how important the ridgeline is. Instead of penalizing not walking along the terrain features, we can use a discount factor for taking the preferred route, and assign a higher discount factor for terrain features with a higher (regarding absolute value) secondary derivative.

Finally, additional factors such as slope should be considered when evaluating the walking effort. In the distance measure introduced in the previous section, elevation change and terrain features have been considered, but there are other factors that contribute to the walking effort. For example, walking along the contour line will lead to zero elevation change along the way, but the effort needed highly depends on the slope. Walking along the hillside of a steep slope takes much more effort than walking on the flat terrain. Therefore, in the distance measure, we penalize walking along the hillside and assign a higher penalty factor for a higher slope.

Trading Off Exploration and Exploitation

Building on the PAWS framework, we provide a vari-

Average # of Reachable Raster Pieces	9066.67
Average # of Reachable Grid Cells (Targets)	22.67
Average # of Reachable KAPs	194.33

Table 1. Problem Scale for PAWS Patrols.

Average Trip Length	4.67 Days
Average Number of Patrollers	5
Average Patrol Time Per Day	4.48 hours
Average Patrol Distance Per Day	9.29 km

Table 2. Basic Information of PAWS Patrols.

ation of PAWS (denoted as PAWS-EvE) that offers the option of assigning a probability range for selecting an explorative route. Explorative routes are those that cover a significant portion of previously unpatrolled land, while exploitative routes are those that cover a significant portion of land previously patrolled.

The major advantage of selecting exploitative routes is that patrollers are familiar with those patrol routes. Having experience with the routes, patrollers also require less effort on following the routes as they would with unfamiliar territory, enabling them to better cover the area they patrol. Further, the explorative routes may require more effort on checking the map, finding water refill points, and figuring out the best way around unexpected obstacles. Some of these tasks require experienced patrollers and additional equipment, which are not available for every patrol. However, if patrollers were to only take exploitative routes, then poachers would easily be able to observe such a strategy and then focus on targeting other areas. With the objective of PAWS to minimize poaching activity, it is necessary to also take explorative routes. Therefore, offering the option of setting a probability range for selecting an explorative route would be helpful for practical use. The range is set before generating the patrols based on these practical concerns.

To implement the functionality of PAWS-EvE, we modify the previously mentioned separation oracle by introducing two sets of routes, the explorative set and the exploitative set. Two new user-defined parameters θ and δ are introduced and we add two new constraints to make sure the probability of assigning a patrol route from the explorative set lies within a δ margin of θ . Also, in route generation, we

check if adding a route from the explorative set would improve the solution, and then we also check if adding a route from the exploitative set would improve the solution.

Deployment and Evaluation

PAWS patrols are now regularly deployed at a conservation area in Malaysia. This section provides details about the deployment and both subjective and objective evaluations of PAWS patrols.

PAWS patrol aims to conduct daily patrols from base camps. Before the patrol starts, PAWS generates the patrol strategy starting from the base camp selected by the patrol team leader. The patrol distance limit considered by PAWS is 10 kilometer per day (equivalent flat terrain). As shown in table 1, this leads to about 9000 raster pieces to be considered. Thus, it is impossible to consider each raster piece as a separate target or consider all possible routes over the raster pieces. With the two-level discretization and the street map, the problem scale is reduced, with $8.57(=194.33/22.67)$ KAPs and 80 route segments in each grid cell on average, making the problem manageable. The strategy generated by PAWS is a set of suggested routes associated with probabilities and the average number of suggested routes associated with probability > 0.001 is 12.

Each PAWS patrol lasts for 4–5 days and is executed by a team of 3–7 patrollers. The patrol planner will make plans based on the strategy generated by PAWS. After reaching the base camp, patrollers execute daily patrols, guided by PAWS's patrol routes. Table 2 provides a summary of basic statistics about the patrols. During the patrol, the patrollers are equipped with a printed map, a hand-held GPS, and data-recording booklet. They detect animal and human activity signs and record them with detailed comments and photos. After the patrol, the data manager will put all the information into a database, including patrol tracks recorded by the hand-held GPS, and the observations recorded in the logbook.

Figure 8 shows various types of signs found during the patrols. Table 3 summarizes all the observations. These observations show that there is a serious ongoing threat from the poachers. Column 2 shows results for all PAWS patrols. Column 3 shows results for explorative PAWS patrols, the (partial) patrol routes, which go across areas where the patrollers have never been before. To better understand the numbers, we show in column 4 the statistics about early-stage non-PAWS patrols in this conservation area, which were deployed for a tiger survey. Although it is not a fair comparison as the objectives of the non-PAWS patrols and PAWS patrols are different, comparing columns 2 and 3 with column 4 indicates that PAWS patrols are effective in finding human activity signs and animal signs. Finding the human activity signs is important to identify hotspots of poaching activity,



Figure 8. Various Signs Recorded During PAWS Patrols.

Patrol Type	All PAWS Patrol	Explorative PAWS Patrol	Previous Patrol for Tiger Survey
Total Distance (kilometers)	130.11	20.1	624.75
Average Number of Human Activity Signs per kilometer	0.86	1.09	0.57
Average Number of Animal Signs per kilometer	0.41	0.44	0.18

Table 3. Summary of Observations.

and patrollers' presence will deter the poachers. Animals signs are not a direct evaluation of PAWS patrols, but they indicate that PAWS patrols prioritize areas with higher animal density. Finding these signs is aligned with the goal of PAWS — combat poaching to save animals — and thus is a proof for the effec-

tiveness of PAWS. Comparing column 3 with column 2, we find the average number of observations made along the explorative routes is comparable to and even higher than that of all PAWS patrol routes. The observations on explorative routes are important as they lead to a better understanding of the unex-




Figure 9. One Daily PAWS Patrol Route in August 2015.

plored area. These results show that PAWS can guide the patrollers toward hotspots of poaching activity and provide valuable suggestions to the patrol planners.

Along the way of PAWS deployment, we have received feedback from patrol planners and patrollers. The patrol planners mentioned that the top routes in PAWS solution (routes with the highest probability) come close to an actual planner's routes, which shows PAWS can suggest feasible routes and potentially reduce the burden of the planning effort. As we deploy PAWS in the future at other sites, the cumulative human planners' effort saved by using PAWS will be a considerable amount. In addition, patrollers commented that PAWS was able to guide them toward poaching hotspots. The fact that they found multiple human signs along the explorative PAWS patrol routes makes them believe that PAWS is good at finding good ridgelines that are taken by animals and humans. Patrollers and patrol planners also agree that PAWS generates detailed suggested routes, which can guide the actual patrol. Patrollers commented that the suggested routes were mostly along the ridgeline, which is easier to follow, compared with the routes from the first trial by PAWS-Initial. Figure 9 shows one suggested route (orange line) and the actual patrol track (black line) during PAWS patrol in August 2015 (shown on a 1 kilometer by 1 kilometer grid). Due to the precision of the contour lines we get, we provide a 50-meter buffer zone (light orange polygon) around the suggested route (orange/light-gray line). The patrollers started from the base camp (the green or shaded triangle) and headed to the southeast. The patrollers mostly followed PAWS's suggested route, indicating that the route generated by PAWS is easy to follow (contrast

with PAWS-Initial as shown in figure 3a). Finally, the power of randomization in the PAWS solution can be expected in the long term.

Lessons Learned

During the development and deployment process, we faced several challenges, and here we outline some lessons learned.

First, firsthand immersion in the security environment of concern is critical to understanding the context and accelerating the development process. The authors (from USC and NTU) intentionally went for patrols in the forest with the local patrolling team to familiarize themselves with the area. The firsthand experience confirmed the importance of ridgelines, as several human and animal signs were found along the way, and also confirmed that extreme changes in elevation require a considerable extra effort of the patrollers. This gave us the insight for building the street map.

Second, visualizing the solution is important for communication and technology adaptation. When we communicate with domain experts and human planners, we need to effectively convey the game-theoretic strategy generated by PAWS, which is a probability distribution over routes. We first visualize the routes with probability > 0.01 using ArcGIS so that they can be shown on the topographic map and the animal distribution map. Then for each route, we provide detailed information that can assist the human planners' decision making. We not only provide basic statistics such as probability to be taken and total distance, but also estimate the difficulty level for patrol, predict the probability of finding animals and human signs, and provide an elevation chart that shows how the elevation changes along the route. Such information can help planners' understanding of the strategy, and also help the planner assign patrol routes to the appropriate team of patrollers, as some patrollers may be good at long-distance walking with flat terrain while others would prefer short-distance hiking with high elevation change.

Third, minimizing the need for extra equipment/effort would further ease PAWS future deployment, that is, patrollers would prefer having a single hand-held device for collecting patrol data and displaying suggested patrol routes. If PAWS routes could be embedded in the software that is already in use for collecting data in many conservation areas, for example, SMART, it would reduce the effort required of planners. This is one direction for future development.

Summary

PAWS is a first deployed green security game application to optimize human patrol resources to combat

poaching. We provided key research advances to enable this deployment; this has provided a practical benefit to patrol planners and patrollers. The deployment of PAWS patrols will continue at the site in Malaysia. Panthera has seen the utility of PAWS, and we are taking steps to expand PAWS to its other sites. This future expansion and maintenance of PAWS will be taken over by Armorway,³ a security games company (starting in spring 2016); Armorway has significant experience in supporting security-games-based software deployments.

Acknowledgement

This research was supported by MURI Grant W911NF-11-1-0332. We also thank our partners in the field who made these tests possible and subcontract from Cornell University for NSF Grant CCF-1522054.

Notes

1. For the security of animals and patrollers, no latitude/longitude information is presented in this article.
2. The Spatial Monitoring and Reporting Tool (SMART), www.smartconservationsoftware.org.
3. The company has now changed its name. See avataai.com.

References

- Botea, A.; Müller, M.; and Schaeffer, J. 2004. Near Optimal Hierarchical Path-Finding. *Journal of Game Development* 1(1): 7–28.
- Chapron, G.; Miquelle, D. G.; Lambert, A.; Goodrich, J. M.; Legendre, S.; and Clobert, J. 2008. The Impact on Tigers of Poaching Versus Prey Depletion. *Journal of Applied Ecology* 45(6): 1667–1674.
- Fang, F.; Jiang, A. X.; and 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*, 957–964. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems.
- Fang, F.; Stone, P.; and Tambe, M. 2015. When Security Games Go Green: Designing Defender Strategies to Prevent Poaching and Illegal Fishing. In *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence (IJCAI)*. Palo Alto, CA: AAAI Press.
- Hamisi, M. 2008. Identification and Mapping Risk Areas for Zebra Poaching: A Case of Tarangire National Park, Tanzania. Ph.D. Dissertation, Thesis, ITC Faculty of Geo-Information Science and Earth Observation of the University of Twente, Enschede, The Netherlands.
- Haskell, W.; Kar, D.; Fang, F.; Tambe, M.; Cheung, S.; and Denicola, E. 2014. Robust Protection of Fisheries with Compass. In *Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence and the Twenty-Sixth Innovative Applications of Artificial Intelligence Conference*, 2978–2983. Palo Alto, CA: AAAI Press.
- IUCN. 2015. IUCN Red List of Threatened Species. Version 2015.2. Gland, Switzerland: International Union for Conservation of Nature. (www.iucnredlist.org)
- Johnson, M. P.; Fang, F.; and Tambe, M. 2012. Patrol Strategies to Maximize Pristine Forest Area. In *Proceedings of the Twenty-Sixth AAAI Conference on Artificial Intelligence (AAAI)*, 295–301. Palo Alto, CA: AAAI Press.
- Kar, D.; Fang, F.; Fave, F. D.; Sintov, N.; and Tambe, M. 2015. A Game of Thrones: When Human Behavior Models Compete in Repeated Stackelberg Security Games. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems*. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems.
- Korzhyk, D.; Conitzer, V.; and Parr, R. 2010. Complexity of Computing Optimal Stackelberg Strategies in Security Resource Allocation Games. In *Proceedings of the 24th National Conference on Artificial Intelligence*, 805–810. Palo Alto, CA: AAAI Press.
- Langmuir, E. 1995. *Mountaineering and Leadership: A Handbook for Mountaineers and Hillwalking Leaders in the British Isles*. Manchester, UK: Mountain Leader Training Board.
- Lemieux, A. M., ed. 2014. *Situational Prevention of Poaching*. Crime Science Series. New York: Taylor & Francis Group / Routledge.
- Nguyen, T. H.; Fave, F. M. D.; Kar, D.; Lakshminarayanan, A. S.; Yadav, A.; Tambe, M.; Agmon, N.; Plumtore, A. J.; Driciru, M.; Wanyama, E.; and Rwetsiba, A. 2015. Making the Most of Our Regrets: Regret-Based Solutions to Handle Pay-off Uncertainty and Elicitation in Green Security Games. In *Decision and Game Theory for Security: 6th International Conference*. Berlin: Springer.
- Nguyen, T. H.; Yadav, A.; An, B.; Tambe, M.; and Boutilier, C. 2014. Regret-Based Optimization and Preference Elicitation for Stackelberg Security Games with Uncertainty. In *Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence*. Palo Alto, CA: AAAI Press.
- Nguyen, T. H.; Yang, R.; Azaria, A.; Kraus, S.; and Tambe, M. 2013. Analyzing the Effectiveness of Adversary Modeling in Security Games. In *Proceedings of the Twenty-Seventh AAAI Conference on Artificial Intelligence*. Palo Alto, CA: AAAI Press.
- Phillips, S. J.; Anderson, R. P.; and Schapire, R. E. 2006. Maximum Entropy Modeling of Species Geographic Distributions. *Ecological Modelling* 190(3–4): 231–259.
- Pita, J.; Jain, M.; Marecki, J.; Ordóñez, F.; Portway, C.; Tambe, M.; Western, C.; Paruchuri, P.; and 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 Conference on Autonomous Agents and Multiagent Systems: Industrial Track, AAMAS 2008*, 125–132. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems.
- Plummer, M. 2003. JAGS: A Program for Analysis of Bayesian Graphical Models Using Gibbs Sampling. Paper presented at the 3rd International Workshop on Distributed Statistical Computing, March 20–22, Vienna, Austria.
- Qian, Y.; Haskell, W. B.; Jiang, A. X.; and Tambe, M. 2014. Online Planning for Optimal Protector Strategies in Resource Conservation Games. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems (AAMAS'14)*. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems.
- Sanderson, E.; Forrest, J.; Loucks, C.; Ginsberg, J.; Dinersteiner, E.; Seidensticker, J.; Leimgruber, P.; Songer, M.; Heydlauff, A.; O'Brien, T.; Bryja, G.; Klenzendorf, S.; and Wikramanayake, E. 2006. Setting Priorities for the Conservation and Recovery of Wild Tigers: 2005–2015. In *Tigers of the*

World: *The Science, Politics, and Conservation of Panthera tigris*, chapter 9, 143–161, ed. R. Tilson and P. J. Nyhus. Amsterdam, The Netherlands: Elsevier.

Shieh, E.; An, B.; Yang, R.; Tambe, M.; Baldwin, C.; DiRenzo, J.; Maule, B.; and 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, AAMAS'12, 13–20. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems.

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(4): 363–377.

Tambe, M. 2011. *Security and Game Theory: Algorithms, Deployed Systems, Lessons Learned*. Cambridge, UK: Cambridge University Press.

Tarboton, D. G.; Bras, R. L.; and Rodriguez-Iturbe, I. 2007. On the Extraction of Channel Networks from Digital Elevation Data. *Hydrologic Processes* 5(1): 81–100.

Thompson, S. 2011. *Unjustifiable Risk?: The Story of British Climbing*. Cumbria, UK: Cicerone Press.

Tobler, W. 1993. Three Presentations on Geographical Analysis and Modeling. Nonisotropic Geographic Modeling: Speculations on the Geometry of Geography, and Global Spatial Analysis. Technical Report 93-1, National center for Geographic Information and Analysis. Santa Barbara, CA: University of California, Santa Barbara.

Tsiligiridis, T. 1984. Heuristic Methods Applied to Orienteering. *The Journal of the Operational Research Society* 35(9): 797–809.

Vansteenwegen, P.; Souffriau, W.; and Oudheusden, D. V. 2011. The Orienteering Problem: A Survey. *European Journal of Operational Research* 209(1): 1–10.

Wang, T., and Boutilier, C. 2003. Incremental Utility Elicitation with the Minimax Regret Decision Criterion. In *Proceedings of the Eighteenth International Joint Conference on Artificial Intelligence*, 309–318. San Francisco: Morgan Kaufmann Publishers.

Wato, Y. A.; Wahungu, G. M.; and Okello, M. M. 2006. Correlates of Wildlife Snaring Patterns in Tsavo West National Park, Kenya. *Biological Conservation* 132(4): 500–509.

Yang, R.; Ford, B.; Tambe, M.; and Lemieux, A. 2014. Adaptive Resource Allocation for Wildlife Protection Against Illegal Poachers. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems.

Yang, R.; Jiang, A. X.; Tambe, M.; and Ordéñez, F. 2013. Scaling-Up Security Games with Boundedly Rational Adversaries: A Cutting-Plane Approach. In *Proceedings of the 23rd International Joint Conference on Artificial Intelligence*. Palo Alto, CA: AAAI Press.

Yin, Z.; Jiang, A. X.; Johnson, M. P.; Kiekintveld, C.; Leyton-Brown, K.; Sandholm, T.; Tambe, M.; and Sullivan, J. P. 2012. TRUSTS: Scheduling Randomized Patrols for Fare Inspection in Transit Systems. In *Proceedings of the Twenty-Sixth AAAI Conference on Artificial Intelligence and the Twenty-Fourth Conference on Innovative Applications of Artificial Intelligence*, 2348–2355. Palo Alto, CA: AAAI Press.

Fei Fang is a postdoctoral fellow at the Center for Research

on Computation and Society (CRCS) at Harvard University and an adjunct assistant professor at Carnegie Mellon University. Her research lies in the field of artificial intelligence and multiagent systems, focusing on computational game theory with applications to security and sustainability domains.

Thanh H. Nguyen is a postdoctoral researcher at the University of Michigan. Her research is in the field of artificial intelligence, multiagent systems, and machine learning, focusing on game theory with connections to optimization, behavioral modeling, and operations research.

Robert Pickles is a monitoring specialist and the head of monitoring in the Tiger Program in Panthera. He is also a postdoctoral research fellow at the University of Trent. He gained his Ph.D. from the University of Kent and the Zoological Society of London in 2010.

Wai Y. Lam is the technical advisor to Rimba's Project Harimau Selamanya, as well as a Tiger Program scholar with Panthera and a master of science student at Universiti Malaysia Terengganu. Her work involves research on spatiotemporal patterns of people and wildlife in protected areas for antipoaching efforts, camera trap studies, and development of conservation technology in database management.

Gopalasamy R. Clements is an associate professor at the Kenyir Research Institute in Universiti Malaysia Terengganu. He is also the cofounder of Rimba (rimbaresearch.org), a nonprofit research group conducting research on threatened species and ecosystems in Malaysia. His project, Harimau Selamanya, aims to improve protection of three large carnivores in the state of Terengganu, Peninsular Malaysia.

Bo An is an assistant professor at the School of Computer Science and Engineering of the Nanyang Technological University. His research interests include artificial intelligence, multiagent systems, computational game theory, and optimization.

Amandeep Singh is a doctoral candidate at the OID Department at Wharton School of Business. His research interests lie in the areas of digital advertising, crowdfunding, and econometric methods.

Brian C. Schwedock is an undergraduate student at the University of Southern California studying computer engineering and computer science. His research with Teamcore USC includes statistical analysis and security game models.

Milind Tambe is founding codirector of CAIS, the USC Center for AI for Society, and Helen N. and Emmett H. Jones Professor in Engineering at the University of Southern California (USC). He is a fellow of AAAI and ACM, as well as recipient of the ACM/SIGART Autonomous Agents Research Award.

Andrew Lemieux is a researcher at the Netherlands Institute for the Study of Crime and Law Enforcement. His main areas of interest are the spatial and temporal distribution of crime, the use of technology to improve law enforcement operations, and the utility of wildlife crime analysis for data-driven antipoaching operations in Africa and Asia.