POLICE PRACTICE AND RESEARCH, 2018
VOL. 19, NO. 1, 61–82
https://doi.org/10.1080/15614263.2017.1295243

Routledge
Taylor & Francis Group

Check for updates

Optimal patrol routes: interdicting and pursuing rhino poachers

Timothy C. Haas[a] and Sam M. Ferreira[b]

[a]Lubar School of Business, University of Wisconsin-Milwaukee, Milwaukee, WI, USA; [b]Scientific Services, SANParks, Skukuza, South Africa

ABSTRACT
Interdiction patrolling (also known as hot spots policing) is a proactive police operation that depends on good planning. And the pursuit of fleeing suspects is a challenging operation that all police forces engage in on a regular basis. We give software tools that improve the effectiveness of both. We demonstrate the effectiveness of our tools by applying them to the case of rhino poaching in wildlife reserves. An interdiction patrol pre-emptively seeks out poaching parties within a reserve. Upon picking up the trail of a poaching party, a pursuit team needs a pursuit strategy that optimizes their chances of an interception. Our interdiction patrol route tool identifies optimal interdiction patrol routes to walk. This tool is based on a Stackelberg game and represents different types of poachers and their utilities. These utilities are updated (learned) each time new information on poaching party routes is added to the database. Our second tool provides a real-time sequence of team member moves that maximizes the team's chances of apprehending a pursued party. We show that our tools perform well when applied to simulated and real data.

ARTICLE HISTORY
Received 26 February 2016
Accepted 8 February 2017

KEYWORDS
Interdiction patrols; pursuit strategies; anti-poaching; security games

Introduction

We describe two data-based, analytical software tools that can help police forces plan more effective interdiction patrols and conduct more successful pursuits of fleeing suspects. A well-planned interdiction patrol route can place officers at the scene of an attempted crime. Doing so increases the chance of the crime being deterred or at the very least, the perpetrators being apprehended. An interdiction patrol is also referred to as a 'combat patrol,' 'night patrol,' 'security patrol,' or 'reconnaissance patrol' (Frenzen, 2010; Flynn, 2014). Our interdiction patrol route tool can make police efforts to intercept shipments of smuggled drugs or illegal immigrants more successful. This tool can also be used to break up burglary or auto theft rings by increasing the chances that a police patrol will make contact with the would-be thieves at the location of their next heist. In this sense, our tool can be viewed as an analytical form of 'hot spots policing' (National Institute of Justice, 2016). In general, interdiction patrols provide denial-of-entry and property protection (Law Library, 2016). Further examples include border patrolling, airport security, and building security.

Our pursuit strategy tool can be used to track-down escaped fugitives or suspects fleeing a crime scene. Responding to such a *fleeing felon* call means that the responding officers need to conduct a pursuit in real-time. Doing so is challenging in two ways: predicting where the suspect is headed, and

CONTACT Timothy C. Haas haas@uwm.edu

coordinating pursuit members so as to maximize their chances of apprehending the pursued. Such pursuits, also called *tactical tracking* are increasingly frequent and unfortunately are often focused on pursuing a fugitive who is armed, particularly dangerous, and versed in traveling through woodlands to avoid detection. Such pursuits often begin in urban settings but quickly enter woodlands as the fugitive attempts to shake his pursuers (Patten, 2015). Several recent pursuits of fugitives suspected of shooting police officers have highlighted the need for improved training and techniques for this particular type of police operation. A growing number of tracking schools have opened in the United States to provide such training to a wide spectrum of police forces including the FBI, ATF, USFS, NPS, BLM, several state patrols, and the BLM (Patten, 2015).

In general, running interdiction patrols and pursuing suspects are two activities that are of increasing importance to any police force in any country (Adrian, 2000; Schroeder, 2016). And the use of analytical techniques to conduct such patrols and pursuits is seen by many police chiefs to be increasingly useful to them (Police Executive Research Forum, 2014). Our two tools therefore, can significantly improve the overall effectiveness of every country's police operations.

We show the effectiveness of our tools by applying them to the policing of wildlife reserves. Such reserves are increasingly being relied on to protect biodiversity. And biodiversity is considered to be a world-wide, nonrenewable resource (United States Geological Survey [USGS], 1997). Wildlife trafficking is a major threat to it (Santos, Satchabut, & Vigo Trauco, 2011; Xie, 2015). Several species face increased extinction risks associated with the illegal trade of products derived from them (Coghlan, 2015; Rosen & Smith, 2010). Much of the challenge associates with wildlife included in the trade goods of transnational organized crime syndicates (WWF/Dalberg, 2012). Authorities, however, only have local jurisdiction to enforce laws (Haas & Ferreira, 2015).

Such illegal trade in wildlife often begins with poaching within a protected area and falls to the protected area's security staff to stem it. Hence, wildlife managers increasingly need to be knowledgeable of, and involved in the policing of wildlife trafficking crime (Dudley, Stolton, & Elliott, 2013). The tools provided in this article can make poaching interdiction patrols and poaching party pursuits more efficient by maximizing their effectiveness under law enforcement budget constraints. Running efficient anti-poaching operations becomes more important as the ratio of protected area size to its law enforcement budget increases.

Anti-poaching security aims to enforce anti-poaching laws that protect animals situated in public wildlife reserves. Anti-poaching security for private reserves aims to protect the owner's legal right to the protection of their private property – in this case, their animals. As such, anti-poaching security can be viewed as protection of public or private property from theft or destruction. Law enforcement in national parks is a growing problem and has been receiving increasing attention and resources (Frank, 2016). Such law enforcement is charged with not only public safety, but protection of the park's resources including its animals, forests, and bodies of water.

But the anti-poaching policing of a large outdoor area that has no clear border, let alone a border barrier, presents a number of challenges not usually encountered in other policing settings such as providing security for a building or a well-fenced airport. These challenges include typically low numbers of police (hereafter *rangers*) per square kilometer, modest enforcement budgets, the need to provide areal rather than strictly boundary protection, highly stealthy attackers, lack of intrusion alarm systems, lack of roads, presence of dangerous wildlife, and little control over the location of the mobile assets (animals) that are to be protected.

To address these challenges, we describe two tools for helping a protected area's rangers intercept and pursue poaching parties. Our first tool delivers optimal interdiction patrol routes while our second tool aids rangers engaged in an active pursuit of a recently-detected poaching party. Of course these tools may be used opportunistically: using the interdiction patrol route tool, rangers begin a patrol, encounter fresh spoor from a poaching party, and decide to abandon the interdiction patrol to give chase using the pursuit strategy tool.

Interdiction patrol route tool

Call a square region that is to be protected against poaching parties through interdiction patrols, the *patrol region*. Our first tool produces an interdiction patrol route. This route is computed based on past behavior of poaching parties, current knowledge of animal locations, and a collection of candidate patrol routes. This route is not computed as a response to a just-received report of an incursion. Rather, this route is used pre-emptively and pro-actively by the anti-poaching commander.

Our approach directly updates the poachers' utilities based on the recent record of those poacher interdiction patrol routes that did and did not stop a poaching event from happening. The tool based on this approach first computes a set of potential patrol routes and a probability associated with each one. The algorithm then randomly selects one of these patrol routes for execution using the associated probabilities. We call such a set of patrol routes and their associated probabilities an *interdiction patrol strategy*.

Pursuit strategy tool

Our second tool generates a sequence of recommended locations for a team of rangers to walk as they pursue a poaching party. Taken together, these recommended moves form a pursuit strategy that maximizes the chance of an interception. As opposed to the interdiction patrol route tool, above, this tool is used in a re-active mode: the tool begins to produce a pursuit strategy only when triggered by ranger reports of either recent poaching party spoor or gunshots from a poaching party's hunting rifle.

Note that pursuit of a poaching party is more complex than the classic pursuit-evasion game (see Raboin, Kuter, & Nau, 2012) as the poaching party, before the shooting of an animal, is not simply trying to avoid capture, but rather has the additional goal of trying to poach an animal. It is also important to note that any computer support for a pursuit needs to be provided to the pursuing rangers in real-time.

This article is structured as follows. First, we briefly review previous work. Then, to set context, we describe the situation surrounding the poaching of rhinos in Kruger National Park (KNP), South Africa. Next, we describe our interdiction patrol tool and give a simulated example of its use followed by an example of it being applied to an unnamed private rhino reserve. Then, we describe our pursuit strategy tool and give examples of its use that includes a continuation of the private reserve example. Following that, we discuss the place our tools occupy in the larger context of human-wildlife conflict and finally, we conclude.

Previous work

Several methods have been developed to meet the challenge of poacher interdiction. Park, Serra, Snitch, and Subrahmanian (2016) provide a data-based model that supports finding optimal ranger patrol routes and drone flight plans that minimize the number of rhinos predicted to be vulnerable to poaching (see also Park, Serra, & Subrahmanian, 2015). These authors represent poacher behavior with a regression model of the probability of poachers attacking a cell within a gridded map of the protected area that employs the predictor variables of *distance to nearest building from the current cell, distance to nearest vegetation from the current cell, distance to a freeway from the current cell, number of rhinos in the current cell, elevation of the current cell, elevation range in the current cell,* and *number of cells reachable by a poacher from the current cell*. A subset of these variables is used to compute an estimate of the probability of a rhino visiting a cell in the grid using several data analytic algorithms, namely Support Vector Machines, K-Nearest Neighbors, and Adaboost (see Park et al., 2016). These two models are then fitted to data on poaching events and rhino locations taken from a protected area in South Africa. Finally, ranger patrol routes and drone flight plans are found that maximize the probability of thwarting poacher attacks on rhinos.

The optimization procedure of Park et al. (2016) models anti-poaching units as agents who seek to maximize their profit function. Such anti-poaching units are composed of a team of rangers who may or may not have the ability to fly a drone over a chosen flight plan. Poachers are not modeled as utility-maximizing agents nor are they modeled as having any knowledge of previous ranger-drone unit patrol routes. Patrol routes are deterministically found as solutions to the optimization problem. As the authors state, 'Poachers are not considered agents because we assume they behave in accordance with the behavioral model we have learned about them' (Park et al., 2016).

Yang, Ford, Tambe, and Lemieux (2014) use a Bayesian Stackelberg game to model poacher and anti-poaching unit interactions. These authors use a multivariate normal tri-variate vector of stochastic parameters for the weights in the poachers' per-cell utility function. This function consists of the probability that the cell will be defended, the utility to the poachers of attacking the cell when it is defended, and the utility to the poachers of attacking the cell when it is undefended. The mean vector and covariance matrix of this multivariate normal distribution are updated (learned) each time a new set of arrested poachers is interviewed, and when a new set of carcasses and/or snares has been discovered. An infinite number of poacher types are indexed by the values of the tri-variate parameter vector making this game a Bayesian one. Interdiction patrols are found by sampling from the mixed strategies that are the solutions to this game. Yang et al. (2014) model wildlife locations statically through a single observation of their spatial density. This density function in-part determines a poacher's utility for attacking a particular cell.

Protecting African rhinos from poaching

African rhino species epitomize the challenges of human-wildlife conflict. The number of rhinos killed increased exponentially from 2008 (Knight, 2013) although this increase had slowed by the end of 2015. South Africa is home to 90% of white and 36% black rhinos in the world with 8365–9337 white and 313–453 black rhinos living in KNP. Of the remaining white rhinos, 189 live in other national parks, 2742–3743 in provincial reserves and 5693 are owned by private individuals (AfRSG, 2016 Unpublished data[1]). Of other black rhinos comprising three subspecies, 274 reside in other national parks, 828–880 in provincial reserves, and 419 on private or communal properties (AfRSG, 2016 Unpublished data[2]).

Protectors of such biodiversity are charged with implementing rhino security tactics. During 2014, KNP, the epicenter of rhino poaching, had a protection force that was comprised of 400 field rangers, 22 section rangers, 4 regional rangers, 15 special rangers, 12 investigators, 2 pilots, 140 personnel in security services and 1 commanding officer. Section rangers are in charge of sections within KNP and direct field ranger patrols that systematically cover each section. Field rangers search for any sign of poacher activity and include observation posts as well as information through various types of detection technology and local informants to direct patrols and effort. SANParks security personnel respond to incursions and incidences of poaching using helicopters and special rangers supported by tracking dogs (SANParks, Unpublished data[3]).

Poachers enter KNP through various means including being dropped off after entering through tourist gates and coming into the park across boundaries and fences. Before 2014, most incursions were from Mozambique, but by the end of 2014, 65% of incidences were from the west into KNP. At that time, a poaching party was typically comprised of three people each with different responsibilities – a shooter, a water carrier, and a horn and axe carrier. Poachers who hunt at night and during full-moon phases associate with higher incidences of poaching. Incidences tend to cluster. Typically poachers, after a successful hunt, seek to find the shortest and quickest route out of the park, usually on foot. During 2014, KNP experienced an estimated 4413 person-incursions with some individuals entering multiple times. These incursions were of such duration that at any time, there were usually about 12 to 15 poacher parties in the Park. Rangers experienced 137 contacts with poachers and saw poachers 118 times. A total of 765 poachers were thus exposed to rangers during 2014 of which 202 were arrested – a 4.6% chance of being arrested when illegally entering the Park (SANParks, Unpublished data[4]).

The size of Kruger and the magnitude of the poacher onslaught provide key challenges in protecting rhinos. Protectors could thus benefit from pro-active approaches. When authorities face enforcement requirements over large areas or with limited data on criminal activities, unpredictability of law enforcement responses may be an effective way to leverage limited resources. For instance, authorities at the Los Angeles International Airport use Bayesian Stackelberg games (Pita, Jain, Tambe, Ordóñez, & Kraus, 2010) to provide protection (Trejo, Clempner, & Poznyak, 2015) while the United States Coast Guard uses Bayesian Stackelberg games to protect fishery stocks in the Gulf of Mexico from poaching by Mexican fishing boats (Brown, Haskell, & Tambe, 2014).

Effective interdiction and pursuit strategies are important enablers that could provide rhinos a chance in the face of an increasing onslaught by poachers (Ferreira et al., 2015). Although authorities acknowledge several causes of poaching storms (Conrad, 2012), interim measures that enable protectors to respond pro-actively carry high priority because crime disruption, creating opportunities of more equitable benefit sharing with stakeholders, and long-term sustainability interventions have long time-lags and carry some amount of uncertainty as to their effectiveness in curbing rhino poaching. Poacher interdiction patrols and poacher pursuit strategies are part of the suite of activities that allows authorities to effectively protect rhinos and thereby buy time to implement policies that target the ultimate cause of poaching.

Interdiction patrol route tool

There are n_y types of poachers. A patrol route is defined to be a particular path walked by a team of rangers over a contiguous period of time. An attacker route is a particular path walked by a poaching party over a contiguous period of time.

A Stackelberg game

A Stackelberg game (Pita et al., 2010) is played by two players: a *leader* and a *follower*. This game is sequential and consists of only one period: the collection of events wherein each player makes exactly one move. What the player does on his/her move is called a *strategy*. In a Stackelberg game, the leader moves first. The follower observes this strategy and then moves by implementing his/her strategy. The leader knows that the follower will observe what strategy the leader implements and takes this into consideration when selecting his/her strategy. When the game is over, each player receives a payoff or *utility* that depends on the particular set of strategies implemented by the players during the game's period.

In a Stackelberg game formulation of an anti-poaching patrol route planning problem, the protection force plays the role of the leader or *defender* while poaching parties play follower or *attacker* roles with routes corresponding to strategies. Here, we consider only finite games: the only possible defender routes are $\sigma_i,\ i = 1, \ldots, n_d$, and the only possible attacker routes are $\sigma_j,\ i = 1, \ldots, n_a$ where n_d and n_a are positive, finite integers.

A *mixed strategy* consists of two parts: a probability distribution across the leader's possible strategies; and a random draw of one of these strategies according to this probability distribution. If the leader follows a mixed strategy, the follower can only observe the probability distribution – not the particular strategy that the leader ultimately draws from this distribution. A Stackelberg game is said to be *Bayesian* if there is more than one follower type of which the defender has knowledge of only up to a probability distribution over these types. For a finite number of follower types, the probability distribution is denoted p_i, $i = 1, \ldots, n_y$.

Say that each player selects a strategy that they will implement when their turn to make a move comes around. Call this particular set of strategies assigned to players, a *player-strategy combination*. A solution to any game is a player-strategy combination for which no player would prefer to change his/her chosen strategy because doing so would reduce his/her payoff. This player-strategy combination identifies the *Nash equilibrium* point in the space of all possible player-strategy combinations. One way to solve a Stackelberg game is to find a mixed strategy probability distribution for the leader that maximizes his/her expected utility. This is the standard approach to solving Stackelberg games applied to security domains (Pita et al., 2010).

Parameterization of ranger and poacher utilities

Let s_a, and s_d be the sensor radius of the poaching party, and ranger team, respectively. Let D_{ij} be the smallest distance between a line segment in patrol route σ_i and the closest line segment in attacker route σ_j. Let $L_i = 1$ if evidence of live rhinos is observed closer than s_d units to at least one line segment in patrol route σ_i or is closer than s_a units to at least one line segment in attacker route σ_i. Such evidence includes visual contact or spoor. Let $K_i = 1$ if at least one rhino carcass that is no more than one week old is closer than s_d to at least one line segment in patrol route σ_i.

Let the poaching party's utility of patrol route σ_i and attacker route σ_j be linear in the above functions as follows. The attacker's utility function is a cost function to the defender, $C_{lij} = \alpha_{l1} L_j + \alpha_{l2} D_{ij}$ where α_{l1} and α_{l2}, $l = 1, \ldots, n_y$ are parameters to be estimated ('C' for *costs*). Similarly, the defender's utility function is $R_{ij} = \beta_1 K_i + \beta_2 L_i + \beta_3 D_{ij}$ where the values of β_1, β_2, and β_3 are specified by the anti-poaching commander ('R' for *return*).

Algorithm

Say that an interdiction patrol route is needed for the next patrol period. Denote this time period with t_f. The algorithm consists of three steps: (1) updating the poacher utilities matrix; (2) finding the optimal mixed strategy probability distribution; and (3) sampling once from this distribution to select a patrol route to walk.

Finding the optimal interdiction route

Find the optimal mixed patrolling strategy for time period t_f by solving a Mixed Integer Quadratic Program (MIQP) (Pita et al., 2010, Equations (1–7)) using an updated poacher utilities matrix, $C^{(u)}$ (see below). At time t, these equations are:

$$\max_{\mathbf{x}_t, \mathbf{q}_{tl}, a_{tl}} \sum_{l=1}^{n_y} \sum_{i=1}^{n_d} \sum_{j=1}^{n_a} p_l R_{lij} x_{ti} q_{tlj} \tag{1}$$

s.t.

$$\sum_i x_{ti} = 1, \tag{2}$$

$$\sum_j q_{tlj} = 1 \quad \forall l, \tag{3}$$

$$0 \leq \left(a_{tl} - \sum_i C_{lij}^{(u)} x_{ti} \right) \leq (1 - q_{tlj}) M \quad \forall l \text{ and } j, \tag{4}$$

$$0 \leq x_{ti} \leq 1 \quad \forall i, \tag{5}$$

and

$$q_{tlj} \in \{0, 1\} \quad \forall l \text{ and } j. \tag{6}$$

A characteristic of this Mixed Integer Nonlinear Program (MINLP) is that regions in $\mathbf{x}_t$-space are indexed by $\mathbf{q}_{tl}$ vectors. In other words, a new $\mathbf{q}_{tl}$ vector is only optimal for the attacker if $\mathbf{x}_t$ is already a member of the associated $\mathbf{x}_t$-space region. Hence, such an $\mathbf{x}_t$ vector needs to be found before the objective function can be evaluated at all – otherwise, the constraint defined by (4) will be violated.

The objective function is discontinuous at the boundaries of these regions. Any search algorithm that uses gradient information will have a difficult time moving across such boundaries.

Here, we solve the MINLP given in (1)–(6) by maximizing the value of the Stackelberg game objective function (1) over the variables in $\mathbf{x}_t$, alone as the $\mathbf{q}_{tl}$ and $\mathbf{a}_t$ vectors are ultimately functions of the $\mathbf{x}_t$ vector.

Desirable solutions are those for which the mixed strategy has high variance, i.e., has high entropy (the value of $-\sum_{i=1}^{n_d} p_i \ln p_i$ is large and close to its maximum of $\ln n_d$). Note that at the solution, a prediction is provided of which route the attacker will follow.

Updating the poacher utilities matrix

A poaching event is identified by the presence of a carcass. Forensic analysis provides an estimate of when the rhino was shot. This estimated time is used for the time of the poaching event.

An interdiction patrol route that would have stopped a past poaching event is found by selecting a route that is at the same time as the poaching event's time, and contains a route segment that is very close to the associated carcass. Such a route is called a *thwart route* – so named because if rangers had walked such a route, they would have thwarted that poaching event.

For each of the most recent m patrol times, find updated poacher utilities, $C^{(u)}$ such that the thwart route, $\sigma^{(k_t)}$ has high probability within the mixed strategy found by solving the Stackelberg game at time t. The intent of this approach is to adjust the parameters that determine C so that the values $x_{k_t}^{(t)}$, $t = 1, \ldots, m$ are as large as possible given the constraints. In other words, use of this updated poachers' utility matrix would have resulted in more successful interdiction patrols.

The variables that are actually searched over in order to find the maximized probabilities are α_{l1} and α_{l2}, $l = 1, \ldots, n_y$. This problem may be cast as a Nonlinear Program (NLP) as follows.

$$\max_{\alpha_{l1}, \alpha_{l2}, l=1,\ldots,n_y} \quad \sum_{t=1}^{m} x_{k_t}^{(t)} \tag{7}$$

where $x_{k_t}^{(t)}$ is the thwart route's probability within the mixed strategy found by solving the Stackelberg game at time t.

Computing a solution to the NLP defined by (7) involves solving all m Stackelberg games whenever the objective function is evaluated at a new point in the solution space.

Multiple criteria optimization

This NLP actually expresses a multi-objective function in that the end-goal is to simultaneously maximize $x_{t_{k_t}}^{(t)}$ for $t = 1, \ldots, m$. A Pareto optimal solution may be found however, by re-expressing the multiple objectives as a sum as is done in (7). Such optimality is justified as follows.

A generalization of the classic constrained optimization problem is having n objective functions instead of just one. Let $\mathbf{x}$ be the vector of independent variables, $f_i(\mathbf{x})$, $i = 1, \ldots, n$ be the objective functions to be maximized, and $\mathbf{f}(\mathbf{x})$ the vector of these functions. A desirable characteristic of the solution to a multiple criteria optimization problem is that it be efficient or equivalently, be Pareto optimal. The solution $\mathbf{x}_s$ is efficient if there does not exist another $\mathbf{x}$ such that $f_i(\mathbf{x}) \geq f_i(\mathbf{x}_s)\ \forall i$ and $f_i(\mathbf{x}) > f_i(\mathbf{x}_s)$ for at least one i (Mockus, 1989, p. 110). This condition is met if $W^T\mathbf{f}(\mathbf{x}_s)$ is the maximum of $W^T\mathbf{f}(\mathbf{x})$ where $W = (w_, \ldots, w_n)^T$, $w_i > 0$, and $\sum w_i = 1$ (Marler & Arora, 2010; Steuer, 1986, p. 167). Here, we consider all patrol times to be equally important and hence implicitly set $w_i = 1/m$, $i = 1, \ldots, m$.

Generating sets of patrol routes

Candidate patrol routes are needed as input so that the algorithm may compute a probability distribution over them. We describe three ways to specify/generate such routes: ranger-specified, rhino-shadowing, and area-covering. We allow for the possibility of ranger-specified routes because the anti-poaching commander may have identified several interdiction patrol routes due to past experience, known idiosyncrasies of the terrain, or constraints not represented in the MINLP.

We describe the remaining two route-generating algorithms in the following Sections.

Rhino-shadowing patrol routes

A simple approach to generating patrol routes is to have rangers walk routes that keep them close to live rhinos. One algorithm for randomly generating such rhino-shadowing routes is as follows.

(1) Randomly select five live rhino locations within the patrol region.
(2) Let the first waypoint be the live rhino location that is the furthest south and the furthest west.
(3) Select the next waypoint to be the live rhino location that is closest to the first waypoint.
(4) Repeat Step 3 three times.

Execute this algorithm repeatedly until five different patrol routes are generated.

Area-covering patrol routes

Define an area-covering route to be a route that, when walked, results in the sensor circle of the ranger who is walking sweeping out most of the area of the patrol region. We consider the six area-covering routes of Figure 1. The list of waypoints for a route in Figure 1 consists of the route's sequential list of turning points.

As described above, the defender's utility is computed for each of these six routes based on how close the route takes rangers past live rhinos, rhino carcasses, and how close the route is to previously recorded attacker routes in the patrol region. Because these are area-covering routes, these utilities will be very similar to each other – leading to a nearly uniform discrete probability distribution across these routes. Again, this distribution forms the defender's mixed strategy. This nearly-uniform distribution is desirable because over repeated plays of the Stackelberg security game, poachers will see a high-entropy frequency distribution of routes chosen by the anti-poaching commander for his/her ranger team to walk. Such a high-entropy distribution makes it very difficult for the poachers to accurately predict what route the rangers will take next and hence makes it difficult for the poachers to confidently select a ranger-avoiding attacker route.

Poaching party (attacker) routes

Routes are copied-in from a database of previous attacker routes that have been documented by rangers following poaching party spoor and recording such spoor locations along the route with their hand-held GPS devices, e.g., Garmin Radios Integrated with Navigation for the Outdoors (RINO) 650t handheld GPS radios (Hass, 2015).

Experience suggests that poachers often prefer to return repeatedly to the same location until it is either is drained of targets or they are driven out by intensive patrolling of the area (Wasser, 2015). To allow the tool to account for this observed behavior, those utility parameters associated with attacker routes that resulted in a poaching event are initialized to high values.

Example with simulated routes

A simple example of computing an optimal interdiction patrol route is displayed in Figure 2. Consider a patrol region that contains a portion of an unnamed private reserve in South Africa. The enclosed rectangle has a diagonal length of about 20 km.

This example has two observed attacker routes (At1 and At2) and four potential interdiction patrol routes (IP1–IP4). There are three live rhinos, and two carcasses. These carcasses indicate poaching events: the one on the left occurred at time 2015.4, and the one on the right at time 2015.49. The interdiction patrol routes all begin and end at the same location. There are two types of attackers: inexperienced and experienced.

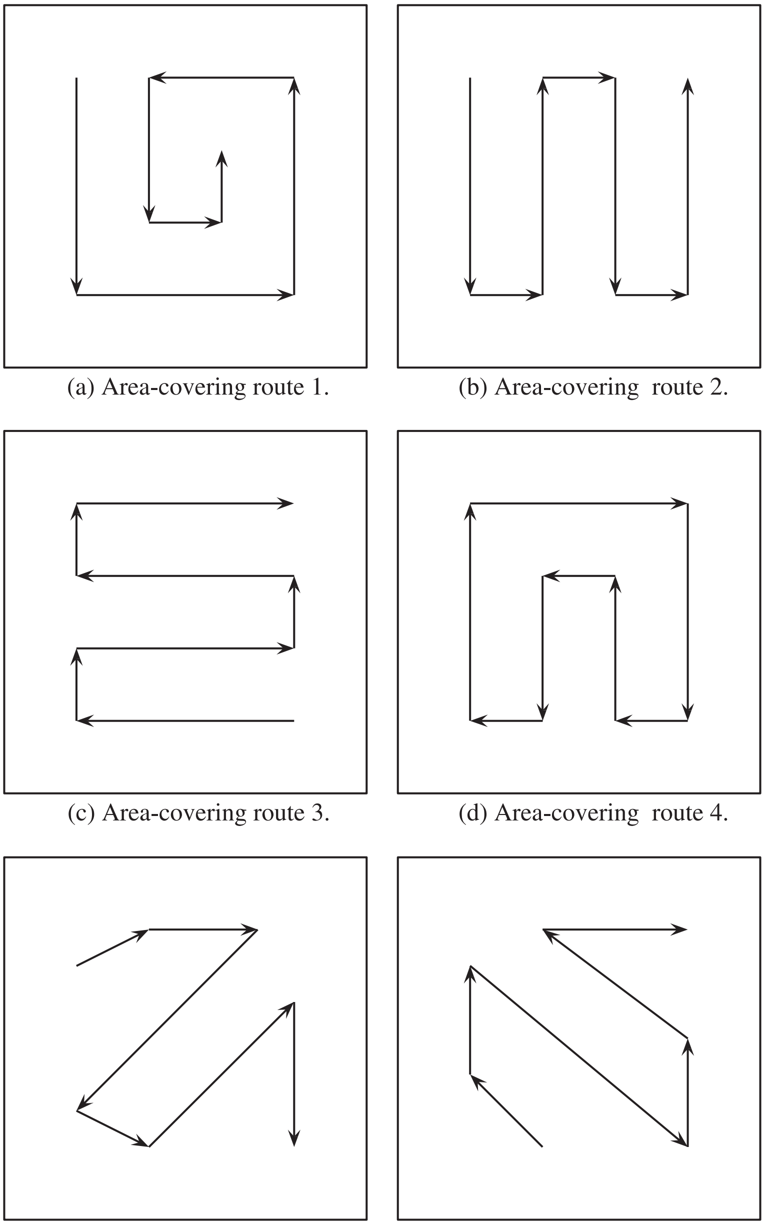

(a) Area-covering route 1.
(b) Area-covering route 2.
(c) Area-covering route 3.
(d) Area-covering route 4.
(e) Area-covering route 5.
(f) Area-covering route 6.

Figure 1. Patrol region and the six different area-covering patrol routes.

Interdiction patrol routes IP1 and IP2 had been considered in the past, but are not currently being considered. Rather, only routes IP3 and IP4 are under current consideration.

The optimal route based on expert judgement alone

The original poachers' utility matrix is based on the expert judgement of anti-poaching staff who are of the opinion that poachers place higher priority on avoiding rangers than on walking a route that brings them close to many rhinos. Using this utility matrix, interdiction patrol route IP4 has the highest probability of being selected (specifically, 0.174, 0.174, 0.199, and 0.450 for interdiction patrol routes 1 through 4, respectively). This is because it is believed that the poachers will choose attacker route At1 so as to avoid rangers.

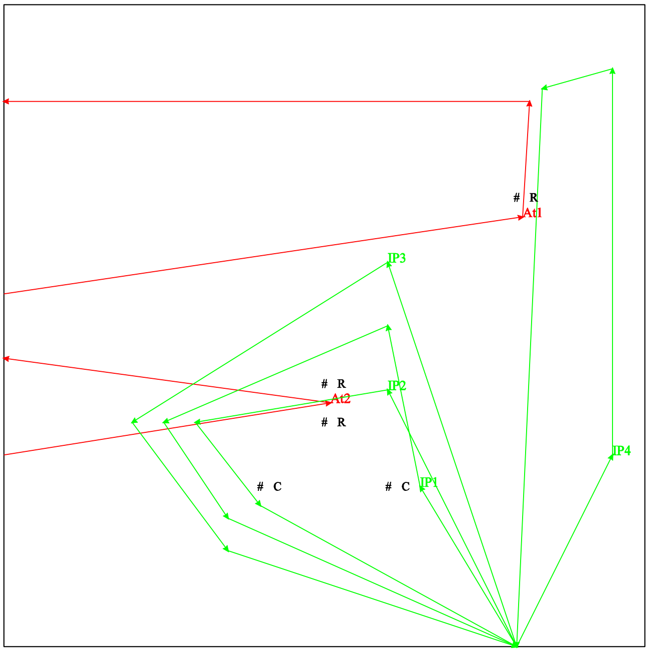

Figure 2. Optimal interdiction patrol route example.
Notes: #C symbols refer to previous carcasses, and #R symbols refer to expected live rhinos.

The optimal route based on past behavior of poachers

Say that the anti-poaching commander suspects the poachers' utility matrix is out of date and decides to update it using the recent past record of poaching events and all interdiction patrol routes that had been previously considered. Here, these consist of the two carcasses and interdiction patrol routes IP1 and IP2.

Running the update algorithm described, above yields the interdiction patrol route IP2 as a thwart route at time 1 and IP1 as a thwart route at time 2. The algorithm adjusts the relative weights poachers place on being caught vs. walking by many rhinos until the total chance the defender will select a thwart route is maximized. Doing so results in interdiction patrol route IP2 having the highest probability of being selected (0.300, 0.324, 0.275, and 0.100 for interdiction patrol routes 1 through 4, respectively) because now, based on the two poaching events, it is believed that the poachers care little about crossing paths with rangers and are most concerned with selecting an attacker route that will bring them close to many rhinos.

This example shows that using a poachers' utility matrix updated by recent poacher behavior can produce a different interdiction patrol route mixed strategy than the strategy arrived at using the original poachers' utility matrix. The reason is that although $C^{(u)}$ is only an estimate of the poachers' true utility matrix, it at least captures information on attacker route preferences gleaned from recent observations of what routes these attackers chose.

Hardware/software for implementing the optimal interdiction patrol route tool

One way to implement this tool is as follows.

(1) Run Garmin BaseStation and Garmin BaseCamp software (Garmin, 2015) on a computer housed within the anti-poaching operations center. Also run the above JAVA program, **id** on this computer.

(2) BaseStation software enables RINO 650t handheld radios to transmit and receive text and GPS eXchange Format (GPX) files between radios and with the operations center. GPX files follow a text-based, open format (topografix, 2015).

(3) The operations center computer also has access to an updated list of previous attacker routes, σ_j, $j = 1, \ldots, n_a$; a list of rhino carcasses (K_i) with associated poaching event times; and an updated list of live rhino locations (L_i). Such data is ideally updated every three or four days.

(4) When the anti-poaching commander decides to order an interdiction patrol, the operations center computer first solves for the optimal route and outputs a GPX file of its list of waypoints. Then, this file is loaded into BaseCamp and from there into each patrol team member's RINO 650t handheld radio.

(5) The anti-poaching commander follows the progress of the interdiction patrol through text messages sent by patrol team members from their Rhino 650t radios back to the BaseCamp system.

A real-world example

Figure 3 contains a plot of the exterior boundary of a real but unnamed reserve (hereafter, the *reserve*). This plot also contains carcass locations and attacker routes that were digitized from the provided attacker track reports. An additional attacker route was created for each carcass. Each of these created routes has an entry point that is the closest boundary point to the carcass; a straight path to the carcass; and then a straight path back to the entry point.

A set of candidate patrol routes was created. Each of these patrol routes traces a square that is 1.5 km on a side for a total patrol length of six kilometers (see the green squares in Figure 3). Patrol routes are separated by 200 m as it was assumed that rangers are able to sense attackers up to 100 m away (their *sensor range*). These routes are designed so that a ranger team may begin to walk the route from their parked vehicle and return to that vehicle in less than 5 h.

Attacker reward function

The attacker's reward function consists of two components: a score based on the highest rhino density along an attacker route; and a score based on how close an attacker route's entry point is to the preferred entry point of this attacker type. Five locations evenly-spaced around the reserve's boundary were specified as these preferred entry points. Because the reserve is randomly attacked from entry points roughly aligned with centers of population density of reserve-bordering towns, it is reasonable to identify attacker type by the location they would likely use to enter the reserve. These preferred entry points could be moved around to improve the effectiveness of chosen patrol routes. Because this grid of potential patrol routes is new, no component representing the preference of avoiding known patrol routes is included in the attacker's reward function.

As time-stamped data on live rhino locations was not available, carcass locations were spatially dithered and then used as rhino (target) locations for time values previous to their respective shot-time.

To keep computation manageable, in lieu of all attacker routes, only 20 randomly chosen attacker routes were employed when computing game solutions.

Ranger reward function

The rangers' reward function also consists of two components: a score based on how many times a patrol route crosses attacker routes; and a score based on the largest rhino density along a patrol route.

Patrol route selection procedure

It is assumed that during each month from 2013 through 2015, the reserve was attacked from each of these five attacker types. To counter these attacks, five patrol routes were selected from the set of candidate patrol routes as follows. First, a set of five patrol route probability distributions was computed

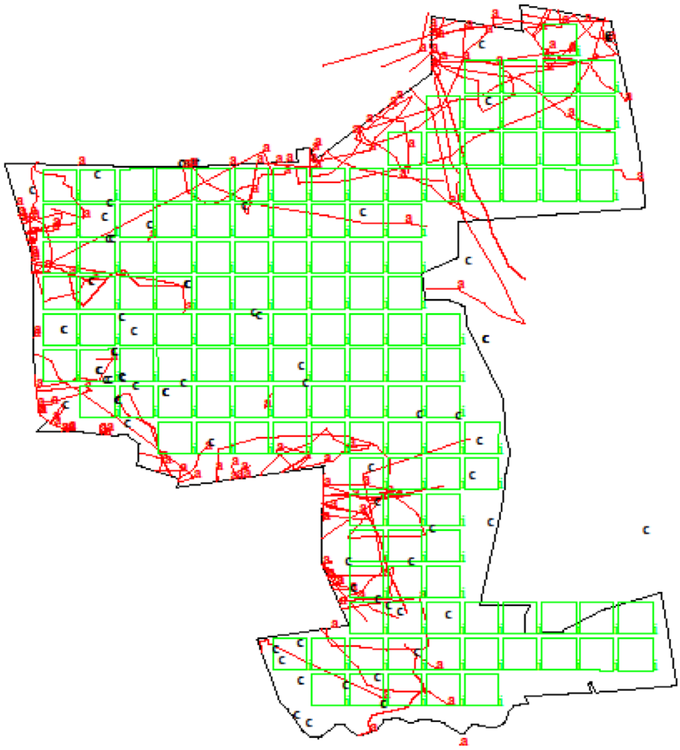

Figure 3. Reserve boundary.
Notes: The symbol, 'C' indicates a carcass location, the symbol 'a' indicates an attacker route, and the symbol 'i' indicates a candidate interdiction patrol route (located in the lower-right corner of each patrol route's square).

during the first month by computing solutions to five associated Stackelberg games. Then, one patrol route was drawn from each of these distributions to form the set of five patrol routes to walk for that month. In total then, five different patrol routes were walked during each of 36 months. And, there were 71 rhinos shot during these 36 months.

Because each patrol route distribution was the solution to a Stackelberg game, each of these patrol routes maximized the chance that the patrolling rangers would interdict an attacker that was getting ready to shoot a rhino – called here an *active attacker*.

Finally, the closest distance between any of these patrol routes and the location of a rhino kill during that month was found. The measure of effectiveness of a patrol route planning tool is defined here to be the median of these distances across all months.

Results

Use of this Stackelberg game approach resulted in a median distance to an active attacker of 533 m. In other words, if each of these patrol routes had been walked during their associated month, half the time rangers would have been closer than 533 m to an active attacker.

This game-theoretic interdiction patrol route selection strategy is compared to the strategy of always walking patrol routes that are close to the most recent kill. This is a standard strategy in anti-poaching work and is manifest in several forms most notably sending out a patrol in the direction of a gunshot. This strategy was implemented here by walking once a month a single patrol route that is closest to the most recent carcass. Note that the strategy of walking randomly selected patrol routes is not a realistic comparison strategy because it is not likely that a real-world anti-poaching commander would approve of such a strategy.

Figure 4. Carcass locations connected in the sequence of their shot times.

The median distance to an active attacker is 2788 m using this strategy of always walking a patrol route that is closest to the most recent carcass. This is about five times worse than that of the game-theoretic strategy.

Based on these results, an interdiction patrol route planning tool that assumes poachers have knowledge of live rhino locations produces effective interdiction patrol routes. Therefore, having at least predictions of live rhino locations at each patrol time appears to be an important component of such a tool.

A component of the attacker's reward function was evaluated that gave high value to previous kills being close to a contemplated route. This component represents a strategy by which attackers intentionally target rhinos that are close to their most recent kill. This component did not perform well and was not pursued. An explanation for why this component was not helpful can be had by examining Figure 4 which contains a plot of carcass locations connected by lines in the sequence of their shot-time. If attackers were re-visiting a location over and over again until it was depleted of rhinos, there would be mostly short lines connecting carcasses together since the location of a poached rhino would usually be close to the location of the next poached rhino. This is not the case for this reserve. Either intentionally or because many different poaching parties are engaging in random, uncoordinated attacks, the sequence of rhino-kill locations through time appears to be nearly random although there is some tendency to shoot rhinos who are closer to the border than those located deeper in the interior – possibly because doing so allows an attacker to spend less time in the reserve.

Needed improvements

The present version of the algorithm does not account for moon phase as it affects the poaching party's ability to hunt rhinos.

Pursuit strategy tool

Let the group of rangers sent to pursue a poaching party be called the *pursuit team*. Let the member of this team in charge of the team's actions be called the *pursuit team's commander*.

Applying pursuit-evasion games to rhino poaching

Sensor limitations

Raboin et al. (2012) give an algorithm for computing in real-time a sequence of moves by n pursuers that reduces their uncertainty about an evader's location over a finite tracking period. This algorithm allows the pursuers to have sensors with only limited range, obstacles that block their line of sight, and the possibility of extended periods of no communication between themselves. The authors assume sensor ranges are in the form of circles of radii r_{sens} centered on each pursuer.

In KNP and surrounding private nature reserves however, the range of visual or other sensing methods by a pursuit team member at a particular location is a function of time-of-day, moon size, and the moon's position. And, poachers often prefer to walk on rocks to reduce their spoor signature. Further, because spoor is rarely discovered soon enough to ensure any chance of audio-visual contact, almost all pursuits depend on the tracking skill of team members assisted by wet- or dry-skin dogs, depending on whether the pursuit team commander believes poachers are walking in streams in order to reduce their spoor signature or not. In other words, sensor range is a function of the terrain around a pursuit team member and what type of sensor system that individual is using, e.g., dog type, night vision system, or man-launchable drone. These effects on a team member's sensor range could be represented by replacing the pursuer-centered sensor range disc of Raboin et al. (2012) with a polygon enclosing the team member that is a function of time-of-day; moon position and size; and local terrain features such as rocky sub-regions, gullies, or stream reaches.

Poachers' utility function

Raboin et al. (2012) assume that the evader has only one goal: avoid capture. Because of this assumption, these authors do not explicitly model the evader's utility function, but rather simply assume the evader always selects a strategy that minimizes the pursuers' utility. This is equivalent to playing zero-sum game against an evader who always chooses the best-response to the pursuers' strategy.

But before they shoot a rhino, poachers have two goals: poach rhinos, and avoid capture. From their apparent willingness to risk death in a shoot-out with a pursuit team, we assume here that poachers see poaching a rhino as twice as important as avoiding capture. We model this utility function explicitly as follows. First, we assume the poaching party always knows the locations of all n pursuit team members that are after them – and the locations of all rhinos that are within the patrol region. The poaching party's goals are met by minimizing the distance between them and the nearest target (a rhino) while also maximizing the distance between themselves and the nearest pursuer.

Encirclement of the poaching party

Once a poaching party's next location is predicted, a pursuit strategy may be found. Here, we assume that the pursuit team wishes to surround the poaching party so as to cut off avenues of escape. To this end, while the poaching party is not within spoor and/or audio-visual range, team members select trajectories that will place them all equidistant from the poaching party's predicted location and for which the smallest separation distance between team members is d_s meters. Once spoor and/or audio-visual contact has been made, team members select trajectories such that they continue to be equidistant from the poaching party, but with separation distance between team members reduced by 50% of its current value. Then, for each move for which team members stay within spoor and/or audio-visual contact of the poaching party, the minimum separation distance between team members is again reduced by 50%. Doing so draws this circle tighter around the poaching party. If spoor and/or audio-visual contact is lost, the minimum distance between team members is reset back to d_s.

Algorithm

Let t_s be the planning horizon time step. Say that poachers move with velocity v_p and pursuit team members with velocity v_r ('p' for *poacher* and 'r' for *ranger*). Then, the distances that can be covered by poachers, and team members are $d_p = t_s \times v_p$, and $d_r = t_s \times v_r$, respectively. Let o_p and t_0 be the team commander's best guess of the poaching party's most recent location and when they were there, respectively. These values are based on the pursuit team's analysis of spoor and/or audio-visual contact with the poaching party. Let o_{ri}, $i = 1, \ldots, n$ be the current locations of the pursuit team's members, and l_{hi}, $i = 1, \ldots, m$ the locations of the m live rhinos in the patrol region. The following three steps are repeated until the poaching party is apprehended.

Step 1:

Define the poaching party's utility of a move to location l_p from location o_p when $|l_p - o_p| < d_p$ to be

$$u(l_p) = \frac{1}{3}\left\{ \min_{i \in \{1,\ldots,n\}} |l_p - o_{ri}| \right\} - \frac{2}{3}\left\{ \min_{i \in \{1,\ldots,m\}} |l_p - l_{hi}| \right\}. \tag{8}$$

Compute a prediction of the poaching party's next location, l_p^* by solving

$$l_p^* = \operatorname{argmax}_{l_p} \{u(l_p)\}. \tag{9}$$

If pursuit team members discover a recently poached rhino, they report this event with a short text message to their pursuit team commander via their RINO 650t radios and the poaching party's utility function is modified by dropping the last term in (8) for all subsequent poacher location prediction computations.

Step 2:

Based on the value of l_p^*, new, optimal locations for team members, l_{ri}, $i = 1, \ldots, n$ are computed by minimizing

$$f_{\text{team}}(l_{ri}, \ldots, l_{rn}) = 0.5\left\{ \sum_{i=1}^{n} |l_{ri} - l_p^*|/n \right\} + 0.3\left\{ \max_{i \in \{1,\ldots,n\}} |l_{ri} - l_p^*| - \min_{i \in \{1,\ldots,n\}} |l_{ri} - l_p^*| \right\} + 0.2\left| \max_{i \neq j}\left\{ |l_{ri} - l_{rj}| \right\} - \frac{d_s}{2^d} \right| \tag{10}$$

under the constraint $|l_{ri} - o_{ri}| < d_r$, $i = 1, \ldots n$.

This objective function is the weighted sum of (a) the average distance between the poaching party and a team member; (b) the difference in separation distances between the poaching party and the furthest and closest team members to that party; and (c) the absolute difference between the maximum inter-team member separation distance and the desired maximum inter-team member separation distance. The first of these terms minimizes the distance between the poaching party and the team as a whole. The second term leads to each team member being about the same distance from the poaching party. The third term in conjunction with the second term, results in team members gradually encircling the poaching party as long as spoor and/or audio-visual contact is maintained (called here the team's *entrapment circle*) while at the same time, gradually shrinking the radius of this circle. In the third term, d is the number of most recent contiguous update steps for which spoor and/or audio-visual contact with the poaching party has been maintained. The value of d is reset to zero if spoor and/or audio-visual contact with the poaching party is lost.

Step 3:

Pursuit team members move to their new locations and then re-estimate o_p and t_0.

A review of the literature (see Klein and Suri (2012)) reveals a prominent conclusion that occurs in many theoretical and simulation results: use of only one pursuer leads to a high probability of no capture, whereas use of at least two pursuers greatly increases the likelihood of a capture. Our pursuit algorithm is no exception.

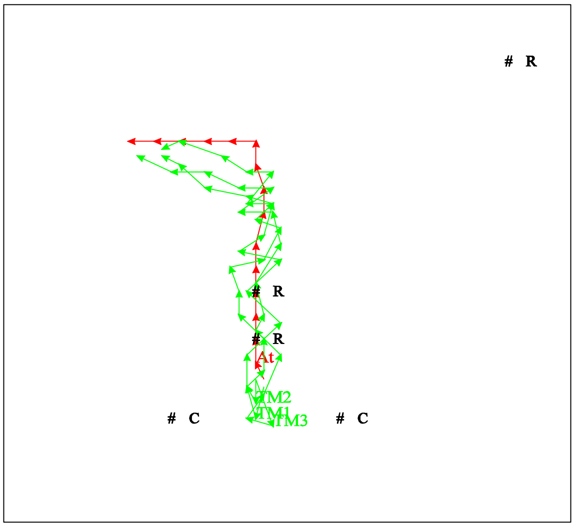

Figure 5. Optimal pursuit patrol strategy example.
Notes: The red-colored path is the poaching party's route. The green paths are the routes taken by the three member pursuit team.

Example

Consider a situation similar to the first interdiction patrol route example, above. Say that a poaching party has been detected at a location just south of several rhinos (Figure 5). The poaching party realizes they are being pursued and head north for several kilometers, poach a rhino on the way, and then abruptly turn west and make a run for the reserve's border. They hope that their abrupt turn will shake the pursuit team off their trail.

The optimal pursuit strategy tool is run on a two-hour time interval. The poaching party makes random changes to their course every time interval. The pursuit team is slightly faster than the poaching party and apprehends the poaching party after 16 location-updates (Figure 5). In this example, the pursuit team never loses spoor and/or audio-visual contact with the poaching party.

Real-world example continued

To assess the effectiveness of this pursuit strategy tool in the real world, it is applied to each observed attacker route contained in the Real-World Example section, above. These routes, indicated in red in Figure 3, were not the result of active pursuits but do represent actual attacker behavior under the threat of being pursued. If a carcass along an attacker route has a shot-date that coincides with the route's date, it is assumed that the rhino was shot at the point in the attackers' route's that is closest to that carcass. It is further assumed that a team of three rangers begins a pursuit at the first point of an attacker route. In other words, routine border patrolling is so intense that as soon as an entry occurs, a pursuit is initiated. Hence, this assumption is liberal as it gives the rangers the most opportunity to capture attackers.

As an aside, this latter assumption also highlights why interdiction patrolling is critical: for large reserves, it may be too expensive to have enough resources on the ground to assure a pursuit being initiated as soon as a poaching party enters the reserve.

Both attackers and rangers have a 1.0 km sensor range but apprehension is declared only if a team member is closer than half this distance to the pursued poaching party. Differing from the previous pursuit example, rangers are assumed to travel at the same speed as attackers: 1.0 km over an update interval. Rangers perform such updates at each point in the attackers' route. The data-set contains 69 attacker routes that consist of four or more points. These routes were acquired by rangers skilled in tracking. This tracking activity however, is typically performed the following day of a night raid, long after the poachers have either escaped or have been arrested. Locations along such an attacker route are irregularly observed, i.e., an attacker may have travelled much further than 1.0 km or much less than 1.0 km between two points in a route. Our example ignores this irregularity: at each update, we have rangers computing a predicted attacker location assuming that attackers have had time to move exactly 1.0 km. Doing so makes our results conservative.

Under these assumptions, our pursuit strategy tool allows rangers to apprehend the poaching party in 42 of these routes, a 61% success rate. Poachers manage to acquire rhino horn during 66 of these attacks. The median distance between the poaching party and the closest team member at the end of the pursuit (successful or not) is 469 m. This compares to the 533 m median distance to an active attacker achieved through interdiction patrolling that was computed in the first part of this real-world example.

This example suggests that (a) our pursuit strategy tool can be effective at apprehending attackers who follow routes that are similar to observed ones, and (b) large expenditures may be needed to create a security policy based on pursuit alone that is as effective as a policy based on interdiction patrolling.

Hardware/software to implement the optimal pursuit strategy tool

Because of the adaptive nature of this tool, its implementation requires a linked and real-time software system that ideally, is local to the pursuit team. The proposed system would operate as follows.

(1) A field ranger who has just picked up a poaching party's trail sends the spoor's GPS coordinates via his/her RINO 650t back to the operations center. Using the short text message feature of the RINO 650t, this ranger includes information on whether they need a wet-skin dog, a dry-skin dog, or no dog at all.

A pursuit team is formed and sent to the poaching party's reported location. This location is also downloaded into BaseCamp running on a mobile device carried by the team's commander. This particular device has sufficient capacity to run three software systems: BaseStation, BaseCamp, and **id**. One such device is the *Bobcat*, a ruggedized Microsoft tablet computer (about 1.0 kg, see Xplore Technologies (2014)).

(2) The pursuit team's commander transfers the poaching party's location from BaseStation to BaseCamp and from there into **id** where a set of new locations for team members is computed. These new locations are relayed via BaseCamp and BaseStation to pursuit team member RINO 650t radios.

(3) Pursuit team members move as close as they are able to these locations, communicate with each other, and send their new locations and any poaching party contact reports back to BaseStation running on the team commander's mobile device.

(4) A new set of pursuit team member locations is computed and sent back to pursuit team members.

(5) The previous two steps are repeated until the pursuit team closes on the poaching party and apprehends them.

Discussion

Police and law enforcement professionals generally seek to prevent crimes (Arrington, 2007), but face challenges when enforcement focuses on large areas typified by crimes such as wildlife trafficking. In this case, authorities require more officers, public participation and/or smarter law enforcement. Large protected areas like Kruger National Park have experienced escalating rhino poaching since 2008 (Knight, 2013). Authorities have responded with smart, technologically-enabled integrated approaches (Child, 2012; Knight, Balfour, & Emslie, 2013). Our interdiction and pursuit tools provide one such smart approach to aid authorities in policing large protected areas.

Using the most recent set of observations on poacher behavior, our interdiction patrol route tool updates a Bayesian Stackelberg game in order to identify optimal poacher interdiction patrol routes that acknowledge poacher knowledge of past patrol routes. When this model is an accurate representation of poacher decision-making, the interdiction patrol route that the tool selects is likely to result in an encounter with a poaching party. The tool does, however, assume that poachers know the location of their wildlife targets, e.g., rhinos.

Our interdiction patrol route tool is data-intensive in that it needs recent data on poacher and rhino locations. Close cooperation between a protected area's conservation staff and security staff is needed to support the maintenance of such an inter-departmental but secure database.

The algorithm that Yang et al. (2014) use for their mean vector and covariance matrix, updates the poachers' utility function without explicitly taking into account the impact on the leader's utility function that would-be successful interdiction patrol routes would have produced. In contrast, our algorithm for updating the poachers' utility function explicitly improves the chances that particular past interdiction patrol routes would have been successful at thwarting poaching events.

Our approach differs from that of Park et al. (2016) in that we model the interactions between poachers, rhinos, and anti-poaching units as a Bayesian Stackelberg game played between the poachers and the anti-poaching units. In this game, we model poachers as being utility maximizing agents who have knowledge of previous patrol routes walked by anti-poaching units. By employing a mixed strategy (random selection of a route from a set of candidate routes), our approach has anti-poaching units choosing patrol routes that, due to their randomness, are difficult to predict by the poachers. The use of a mixed strategy is seen by many as a key ingredient in modern policing and crime fighting (Chen, Cheng, & Wise, 2015).

Both Park et al. (2016) and Yang et al. (2014) represent the protected area as a grid of cells. We avoid this discretization and instead work directly with the continuously-valued GPS coordinate system. This has the advantage of allowing poacher and ranger routes to be directly and precisely represented.

The pursuit strategy tool, much like a missile's target-tracking system (e.g., Zarchan, 2013), depends on fast, local, and frequent updates of the poaching party's position and predicted heading based on a model of the poachers' utility maximization function. When this model is an accurate representation of the poachers' decision-making, the tool should provide a series of moves for the pursuit team members that result in the poaching party being apprehended. Use of GPS radios (Hass, 2015) allows our tool to be used in areas not served by a cellular phone network. This tool depends on the team commander being able to quickly calculate the team's next set of positions every few hours. Thus, reliable mobile computing capability is required.

The Garmin technology described above may be expanded to dog radio collars. Doing so would allow the locations of dogs to be maintained by either tool. These collars can transmit the event of barking which could indicate a dog has located a poaching party. Airplanes and/or helicopters could be incorporated into either tool. These airborne team members would have the ability to be placed almost anywhere for purposes of observation – as long as they are within radio contact of the team's commander. The database used by our tools could be expanded to hold information on particular poaching party members and middlemen who sponsor poaching raids. This latter group forms the lower echelons of the criminal syndicate that is running the poaching operation. Such data could be used in separate efforts to disrupt the international wildlife trade operation itself (e.g., Haas & Ferreira, 2015).

Several drivers including long traditions of trade, inelastic demand, high profit potential, inadequate law enforcement, human wildlife conflict disincentives and unclear property rights have created the current storm of wildlife poaching (Conrad, 2012). And in particular, these factors have made wildlife attractive to organized crime as part of their illegal inventory (Coghlan, 2015; Rosen & Smith, 2010).

Authorities in South Africa, the epicenter of rhino poaching, have embraced an integrated approach to curb the present rhino poaching onslaught. This integrated approach is built on the four pillars of *biological management interventions, long-term sustainability interventions, compulsory anti-poaching interventions,* and *game-changing interventions*. Biological management interventions, the backbone of rhino conservation success (e.g., Knight et al., 2013 and Knight, Emslie, Smart, & Balfour, 2015), seek to use strategic removals of rhinos from areas of high poaching risk and use these to create various kinds of rhino strongholds. Long-term sustainability interventions seek to develop legal trade systems that could meet demand for the use of rhino horn in Asian countries (e.g., Biggs, Courchamp, Martin, & Possingham, 2013; Child, 2012; Di Minin, et al. 2015). Compulsory interventions aim to protect rhinos by using zonal approaches with technology-led, intelligence-based, and rhino-guardian protection techniques (Haas & Ferreira, 2016). Complimenting these rhino security measures are game-changing interventions that seek to integrate disrupting organized crime (Haas & Ferreira, 2015) with more equitable sharing with local stakeholders of ecosystem service benefits (see Child (2012)).

Although our tools primarily support the pillar of compulsory anti-poaching interventions, they may also provide support to other pillars. For instance, once a model of rhino spatial behavior is added to our tools, they could be used by the conservation staff of rhino sanctuaries, strongholds or guardian zones to conduct surveillance of the rhinos themselves. In this case, our tools would contribute to the biological interventions pillar. By defining leader and follower strategies to be those safe houses that middlemen (those who sponsor poaching parties) choose to hide in within settlements that abut a protected area, our interdiction patrol route tool could be used by law enforcement to unpredictably target areas for police raids. In this case, our interdiction patrol route tool would be supporting the game-changing intervention pillar.

Conclusions

We have presented software systems that can identify more successful interdiction patrols, and manage more successful police pursuits. Both of our tools learn in real time how the preferences and resultant strategies of criminals change through their interactions with targets and police forces. Data-based learning of likely criminal behaviors is central to increasing interdiction patrol strategy effectiveness and to improving the success rate of police pursuits. These two operations determine to a large part, the overall effectiveness of police forces the world over. Our interdiction patrol tool can be used immediately to patrol urban environments since its use of a predefined route list automatically respects constraints on routes imposed by the urban environment, e.g., the street grid and/or the locations of subway stations.

We have applied these tools to the task of launching successful patrols for interdicting and/or pursuing wildlife poachers who invade protected areas. In doing so, we have highlighted the challenge of policing large scales – people or space. Our example application to rhino poaching epitomizes this challenge. Our results indicate that our tools could improve wildlife crime prevention and increase arrest rates for crimes that have already taken place. Advertising such an increased poacher arrest rate could serve as deterrent for would-be poachers thinking of entering a protected area.

Our interdiction patrol route tool uses game theory to introduce unpredictable anti-poaching tactics, a key element in successful policing of wildlife poaching crime. The application of these tools spans at least three of the four pillars of integrated responses that authorities are using to curb rhino poaching including compulsory anti-poaching, biological management, and game-changing interventions. Such tools and approaches could aid in addressing several other wildlife trafficking challenges.

Both of our tools could be improved with the inclusion of a model of the target wildlife's spatial behavior. For such a model specific to rhinos, see Rachlow, Kie, and Berger (1999). Poachers have a spatial behavior too in that they often return to the areas where they were successful before – a behavior shared by many other types of criminals (Johnson, 2014). Our interdiction patrol route tool exploits such predictable behavior through its learning mechanism.

The interdiction patrol route tool needs to incorporate the effects of time-of-day, moonphase, and terrain-type on the utility of a candidate route to the attackers (a cost to the defender). This tool depends on the acquisition of recent routes taken by attackers, and up-to-the-minute information on target locations. Such data may be difficult to acquire in real time.

The effect of terrain on the utility of a new location to a fleeing suspect needs to be added to the pursuit strategy tool's algorithm. This tool needs reliable and fast communication between pursuit team members and the ability to periodically detect the location of the pursued party via some form of remote sensing. As with the interdiction patrol route tool, real-time acquisition of such data may be difficult. The pursuit strategy tool's implementation with hand-held radios means that team members need to stay in radio contact with each other. Although RINO 650t radios have a range of up to 32 km, this restriction may be difficult to achieve in certain terrains.

All software and files necessary for running the examples given herein are available at Haas (2016).

Notes

1. Mike Knight, Chairman, African Rhino Specialist Group, mike.knight@sanparks.org.
2. See note 1 above.
3. SANParks, Maj-Gen Johan Jooste, johan.jooste@sanparks.org.
4. See note 3 above.

Disclosure statement

No potential conflict of interest was reported by the authors.

Funding

This work was supported by the World Wildlife Fund [ZA2321].

Notes on contributors

Timothy C. Haas earned a PhD in Statistics from Colorado State University and has been at the Lubar School of Business, University of Wisconsin-Milwaukee since 1990. Professor Haas has developed semi-parametric methods for prediction of nonstationary spatio-temporal processes, algorithms for the redesign of monitoring networks, Bayesian network models of aspen stand survival, forestry ranger decision-making, and integrated models of human-wildlife conflict. This work has been published in the Journal of the American Statistical Association, Forest Science, Atmospheric Environment, Environmetrics, AI Applications, Stochastic Environmental Research and Risk Assessment, and Security Informatics. In addition, Professor Haas has published two books with Wiley on ecosystem management.

Sam M. Ferreira received a PhD in Zoology with a focus on restoration and community ecology from the University of Pretoria. Sam has coordinated the Seal Research Program at Marion Island and the Richards Bay Dune Forest Restoration Research Program both under the auspices of the University of Pretoria. In addition, he has worked for the Department of Conservation in New Zealand as Conservancy Advisory Scientist on projects including marine reserve planning, dolphin and bird research, and alien mammal eradications from islands. Sam completed post-docs at the Department of Entomology and Zoology at the University of Pretoria on elephant temporal dynamics, and the second at the Department of Statistics of the University of Auckland on modeling large mammal dynamics. Sam is currently a SANParks Large Mammal Ecologist and is a South African Council for Natural Scientific Professions registered Professional Scientist in Ecological Science.

References

Adrian, R. (2000). Gaining ground through tactical tracking. *Police: The Law Enforcement Magazine*, August 1. Retrieved from http://www.policemag.com/channel/patrol/articles/2000/08/gaining-ground-through-tactical-tracking.aspx

Arrington, R. L. (2007). *Crime prevention: The law enforcement officer's practical guide*, p. 33. Sudbury, MA: Jones and Bartlett Publishers.

Biggs, D., Courchamp, F., Martin, R., & Possingham, H. P. (2013). Legal trade of Africa's rhino horns. *Science, 339*, 1038–1039.

Brown, M., Haskell, W. B., & Tambe, M. (2014). Addressing scalability and robustness in security games with multiple boundedly rational adversaries. In R. Poovendran & W. Saad (Eds.), *Decision and game theory for security*. Lecture notes in computer science, vol. 8840 (pp. 23–42). New York, NY: Springer-Verlag.

Chen, H., Cheng, T., Wise, S. (2015). *Designing daily patrol routes for policing based on ant colony algorithm*. ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences, Volume II-4/W2 (2015), pp. 103–109. International Workshop on Spatiotemporal Computing, 13–15 July 2015, Fairfax, VA, USA.

Child, B. (2012). The sustainable use approach could save South Africa's rhinos. *South African Journal of Science, 108*, 21–25.

Coghlan, A. (2015). UN puts wildlife crime on a par with drug and people trafficking. *New Scientist*, August 3. Retrieved from https://www.newscientist.com/article/dn27995-un-puts-wildlife-crime-on-a-par-with-drug-and-people-trafficking/

Conrad, K. (2012). Trade bans: A perfect storm for poaching? *Tropical Conservation Science, 5*, 245–254.

Di Minin, E., Laitila, J., Montesino-Pouzols, F., Leader-Williams, N., Slotow, R., Goodman, P. S., … Moilanen, A. (2015). Identification of policies for a sustainable legal trade in rhinoceros horn based on population projection and socioeconomic models. *Conservation Biology, 29*, 545–555.

Dudley, N., Stolton, S., & Elliott, W. (2013). Editorial: Wildlife crime poses unique challenges to protected areas. *PARKS, 19*, 7–12. Retrieved from https://cmsdata.iucn.org/downloads/parks_19_1_editorial.pdf

Ferreira, S. M., Greaver, C., Knight, G. A., Knight, M. H., Smit, I. P. J., & Pienaar, D. (2015). Disruption of rhino demography by poachers may lead to population declines in Kruger National Park, South Africa. *Pubic Library of Science ONE, 10*, e0127783.

Flynn, M. (2014). There and back again: On the diffusion of immigration detention. *Journal on Migration and Human Security, 2*, 165–197, page 166.

Frank, R. (2016). Major challenges face the National Park Service in its next century. *LegalPlanet*, BerkeleyLaw, UCLA Law, August 25. Retrieved from http://legal-planet.org/2016/08/25/major-challenges-face-the-national-park-service-in-its-next-century/

Frenzen, N. (2010). U.S. Migrant interdiction practices in international and territorial waters. In B. Ryan & V. Mitsilegas (Eds.), *Extraterritorial immigration patrol: Legal challenges*, page 390. Leiden: Nijhoff Publishers.

Garmin. (2015). What is the BaseStation feature and how can it be used? *Garmin, Inc*. Retrieved from https://support.garmin.com/support/searchSupport/case.faces?caseId={c05dac80-80fa-11e2-65d0-000000000000}

Haas, T. C. (2016). *Optimal interdiction patrol routes and pursuit strategies*. Retrieved from www4.uwm.edu/people/haas/interdict/

Hass, C. (2015). Garmin Rino 650t review: GPS navigator and radio in one. *IndefinitelyWild*. Retrieved from http://indefinitelywild.gizmodo.com/garmin-rino-650t-review-gps-navigator-and-radio-in-one-1739798572

Haas, T. C., & Ferreira, S. M. (2015). Federated databases and actionable intelligence: Using social network analysis to disrupt transnational wildlife trafficking criminal networks. *Security Informatics, 4*(2), 1–14. doi: 10.1186/s13388-015-0018-8

Haas, T. C., & Ferreira, S. M. (2016). Conservation risks: When will rhinos be extinct? *IEEE Transactions on Cybernetics, 46*, 1721–1734. Retrieved from http://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=7236914

Johnson, S. D. (2014). How do offenders choose where to offend? Perspectives from animal foraging. *Legal and Criminological Psychology, 19*, 193–210.

Klein, K., & Suri, S. (2012). Catch me if you can: Pursuit and capture in polygonal environments with obstacles. In *Proceedings of the 26th Conference on Artificial Intelligence* (pp. 2010–2016). Toronto.

Knight, M. H. (2013). African rhino specialist group report. *Pachyderm, 53*, 7–24.

Knight, M. H., Balfour, D., & Emslie, R. H. (2013). Biodiversity management plan for the black rhinoceros (*Diceros bicornis*) in South Africa 2011–2020. *Government Gazette (South Africa), 36096*, 5–76.

Knight, M. H., Emslie, R. H., Smart, R., & Balfour, D. (2015). *Biodiversity management plan for the white rhinoceros (Ceratotherium simum) in South Africa 2015–2020*. Pretoria: Department of Environmental Affairs.

Law Library. (2016). *Police & law enforcement: Types of police*. University of Iowa Law Library. Retrieved from http://libguides.law.uiowa.edu/c.php?g=103220&p=669182

Marler, R. T., & Arora, J. S. (2010). The weighted sum method for multi-objective optimization: New insights. *Structural and Multidisciplinary Optimization, 41*, 853–862.

Mockus, J. (1989). *Bayesian approach to global optimization*. Dordrecht: Kluwer Academic Publishers.

National Institute of Justice. (2016). *Practice profile: Hot spots policing*. United States Department of Justice. Retrieved from https://www.crimesolutions.gov/PracticeDetails.aspx?ID=8

Park, N., Serra, E., Snitch, T., & Subrahmanian, V. S. (2016). APE: A data-driven, behavioral model based anti-poaching engine. *IEEE Transactions on Computational Social Systems, 2*(2). Retrieved from http://ieeexplore.ieee.org/document/7407518/

Park, N., Serra, E., & Subrahmanian, V. S. (2015). Saving rhinos with predictive analytics. *IEEE Intelligent Systems, 30*, 86–88. doi: 10.1109/MIS.2015.62

Patten, P. (2015). Tactical tracking: Ancient strategies for modern manhunts. *S.W.A.T. Magazine*, August. Retrieved from www.woodlandoperations.com/SWAT_aug15_Tracking_wpermission.pdf

Pita, J., Jain, M., Tambe, M., Ordóñez, F., & Kraus, S. (2010). Robust solutions to Stackelberg games: Addressing bounded rationality and limited observations in human cognition. *Artificial Intelligence, 174*, 1142–1171.

Police Executive Research Forum. (2014). *Future trends in policing*. Washington, DC: Office of Community Oriented Policing Services, U.S. Department of Justice. Retrieved from http://www.policeforum.org/assets/docs/Free_Online_Documents/Leadership/future%20trends%20in%20policing%202014.pdf

Raboin, E., Kuter, U., & Nau, D. (2012). Generating strategies for multi-agent pursuit- evasion games in partially observable euclidean space. *AAMAS'12 Proceedings of the 11th International Conference on Autonomous Agents and Multi-age Systems* (Vol. 3, pp. 1201–1202). Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems.

Rachlow, J. L., Kie, J. G., & Berger, J. (1999). Territoriality and spatial patterns of white rhinoceros in Matobo National Park, Zimbabwe. *African Journal of Ecology, 37*, 295–304.

Rosen, G. E., & Smith, K. F. (2010). Summarizing the evidence on the illegal international trade in wildlife. *EcoHealth, 7*, 24–32.

Santos, A., Satchabut, T., & Vigo Trauco, G. (2011). Do wildlife trade bans enhance or undermine conservation efforts? *Applied Biodiversity Perspective Series, 1*(3), 1–15.

Schroeder, R. D. (2016). Holding the line in the 21st century, U.S. Customs and Border Protection. Retrieved from https://www.cbp.gov/sites/default/files/documents/Holding%20the%20Line_TRILOGY.pdf

Steuer, R. E. (1986). *Multiple criteria optimization*. New York, NY: Wiley.

Technologies, Xplore (2014). The Bobcat fully rugged Windows tablet PC. *Xplore Technologies*. Retrieved from http://www.xploretech.com/products/bobcat

topografix. (2015). *GPX 1.1 developer's manual*. Retrieved from http://www.topografix.com/gpx.asp

Trejo, K. K., Clempner, J. B., & Poznyak, A. S. (2015). A Stackelberg security game with random strategies based on the extraproximal theoretic approach. *Engineering Applications of Artificial Intelligence, 37*, 145–153.

United States Geological Survey. (1997). Recent highlights – Natural resources. *Factsheet* FS-010-97. Retrieved from www.usgs/gov/themes/FS-010-97/

Wasser, S. (2015). Where is all the ivory from? Using forensic science and elephant DNA to stop poachers. *Elsevier SciTech Connect*. Retrieved from http://scitechconnect.elsevier.com/ivory-dna-poachers/

WWF/Dalberg. (2012). *Fighting illicit wildlife trafficking: A consultation with governments*. Gland: WWF International.

Xie, K. (2015). Crime gone wild: The dangers of the international illegal wildlife trade. *Harvard International Review, 36*(4). Retrieved from http://hir.harvard.edu/crime-gone-wild-the-dangers-of-the-international-illegal-wildlife-trade/

Yang, R., Ford, B., Tambe, M., & Lemieux, A. (2014). Adaptive resource allocation for wildlife protection against illegal poachers. In A. Lomuscio, P. Scerri, A. Bassan, & M. Huhns (Eds.), *Proceedings of the 13th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2014)*, May 5–9, Paris.

Zarchan, P. (2013). *Tactical and strategic missile guidance* (6th ed.). Reston, VA: American Institute of Aeronautics and Astronautics, 1026 pages.