

Patrol Strategies to Maximize Pristine Forest Area

Matthew P. Johnson, Fei Fang, and Milind Tambe

Department of Computer Science
University of Southern California
Los Angeles, CA 90089, USA

Abstract

Illegal extraction of forest resources is fought, in many developing countries, by patrols that try to make this activity less profitable, using the threat of confiscation. With a limited budget, officials will try to distribute the patrols throughout the forest intelligently, in order to most effectively limit extraction. Prior work in forest economics has formalized this as a Stackelberg game, one very different in character from the discrete Stackelberg problem settings previously studied in the multiagent literature. Specifically, the leader wishes to minimize the distance by which a profit-maximizing extractor will trespass into the forest—or to maximize the radius of the remaining “pristine” forest area. The follower’s cost-benefit analysis of potential trespass distances is affected by the likelihood of being caught and suffering confiscation.

In this paper, we give a near-optimal patrol allocation algorithm and a $1/2$ -approximation algorithm, the latter of which is more efficient and yields simpler, more practical patrol allocations. Our simulations indicate that these algorithms substantially outperform existing heuristic allocations.

Introduction

Illegal extraction of fuelwood and other natural resources from forests is a problem confronted by officials in many developing countries, with only limited success (MacKinnon et al. 1986; Dixon and Sherman 1990; Clarke, Reed, and Shrestha 1993; Robinson 2008). To cite just two examples, Tanzania’s Kibaha Ruvu Forest Reserves are “under constant pressure from the illegal production of charcoal to supply markets in nearby Dar es Salaam,”¹ and illegal logging is reportedly “decimating” the rosewood of Cambodia’s Central Cardamom Protected Forest (see Fig. 1). In many cases, forest land covers a large area, which the local people may freely visit. Rather than protecting the forest by denying extractors entry to it, therefore, protective measures take the form of patrols throughout the forest, seeking to observe and hence deter illegal extraction activity (Lober 1992). With a limited budget, a patrol strategy will seek to distribute the patrols throughout the forest, in order to minimize the resulting amount of extraction that occurs or protect as much of the forest as possible.

Copyright © 2012, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

¹<http://www.tfcg.org/ruvu.html>



Figure 1: “A truck loaded with illegally cut rosewood passes through Russey Chrum Village...in the Central Cardamom Protected Forest.” Photo from (Boyle Dec 21 2011).

We follow (Albers 2010) in posing this problem as a Stackelberg game in which the policymaker or *leader* publicly chooses a (mixed) patrol strategy; in response, the extractor or *follower* then chooses whether or not to extract, or to what degree. The problem we study is of computing optimal leader strategies in such a game. The potential extraction-preventing benefits of patrols are twofold: extraction is prevented directly, when catching would-be extractors in the act, and also indirectly, through deterrence. As in other Stackelberg application settings, here the followers are presumed likely to learn the leader’s chosen strategy—the patrol personnel are often observed by the (many) villagers, who can communicate with one another over time. The leader wishes to arrange the potential troublemaker’s environment so as to render his choice of engaging in this behavior as expensive for him as possible.² More precisely, given the continuous nature of this setting, the leader seeks a patrol allocation that *minimizes* the distance by which the profit-maximizing follower will trespass into and extract from the forest.

Background. Economists have studied the relationship generally between enforcement policy for protecting natural resources and the resulting incentives for neighbors of the protected area (Milliman 1986; Robinson 2008; Sanchirico and Wilen 2001). Our point of departure in this paper is the forest protection model of (Albers 2010) (see also (Robin-

²By convention, we refer to leader as *she* and follower as *he*.

son, Albers, and Williams 2008; 2011)), in which a circular forest is surrounded by villages (hence potential extractors); the task is to distribute the probability density of patrols within the forest area; the objective is to minimize the distance by which rational extractors will trespass into the forest and hence (since nearby villagers will extract as a function of this distance (Hofer et al. 2000)) or maximize the resulting amount of *pristine* forestland. We assume extractors are rational, risk-neutral economic actors, who choose a trespass distance in order to maximize expected profits (Becker and Landes 1974). In the language of mathematical morphology (Soille 2004), the pristine forest area due to a given patrol strategy (or no patrols) will be an *erosion* $F \ominus B$ of the forest F by a shape B , where B is a circle whose radius equals the trespass distance. The *erosion* is the set of points reached by the center of B as it moves about inside of F , which in this case is the pristine area.

We strengthen this model in several ways, permitting spatial variation in patrol density, multiple patrol units, and convex polygon-shaped forests. As has been observed (Albers 2010), exogenous legal restrictions on patrol strategies, such as requiring *homogenous* patrols, can degrade protection performance (MacKinnon et al. 1986; Hall and Rodgers 1992). Unlike the existing work on this model, we bring to bear algorithmic analysis on the problem. Specifically, we show that while certain simple allocations can perform arbitrarily badly compared to the optimal, provably approximate or near-optimal allocations can be found efficiently.

The forest patrol problem we study here is an instance of the leader-follower Stackelberg game model, which has been the topic of much recent research, applied to a number of real-world security domains, including the Los Angeles International Airport (Paruchuri et al. 2008), the Federal Air Marshals Service (Tsai et al. 2009), and the Transportation Security Administration (Pita et al. 2011). The forest patrol setting differs from the settings of these previous works, most crucially in that it is essentially continuous rather than discrete, both spatially and in terms of player actions. In the existing problems there are a finite number of discrete locations to protect (e.g., modeled as nodes of a graph), whereas ideally the entire forest area would be protected from extraction. The spatial continuity of our problem setting permits a very different approach, in which we solve for optimal or approximate probability distributions over the region using efficient, combinatorial algorithms, without the use of general-purpose solvers. (Of course, the continuous space could be discretized by superimposing a grid on it, but such an approach would be inefficient due to the geometric density.) Once we have computed a distribution over patrol locations, selecting locations from the distribution is straightforward. As such, our primary focus is on choosing and computing this distribution on patrol density over the forest region.

Contributions. We give a full analysis of the problem of maximizing pristine forest radius. Our main contributions are efficient near-optimal and $1/2$ -approximation algorithms for this problem (both with additive error ϵ due to binary search), the latter of which has the advantage of both greater efficiency and more practical, easier to implement solutions.

Our results generalize a) from one to multiple patrol units, and b) from circular forests to convex polygon forests with symmetric patrols. Simulations indicate that our algorithms substantially outperform baseline strategies.

Problem Setting

In this section we present a version of the forest model of (Albers 2010) and formulate a corresponding optimization problem. Villagers are distributed about the forest perimeter (see Fig. 2), which is initially assumed to be a circular region, of radius 1.

Follower strategies. An extractor’s action is to choose some distance d to walk into the forest, on a ray from a point on the perimeter towards the forest center, before reversing and returning to his starting point, extracting on the return trip. We assume the extractors trespass *towards the forest center* because they will naturally wish to avoid one another as much as possible. In this case, the extractor only extracts—and hence is in danger of getting caught doing so—on the return trip, approaching the forest boundary.

Extractors gain a benefit if *never caught* and incur a cost, based on a convex increasing cost function $C(d)$ and a concave increasing benefit function $B(d)$, for trespassing distance d into the forest. If caught, the extractor’s benefit becomes 0 (the extracted resources are confiscated) but the cost does not change since the extractor has at this point already traveled distance d




Figure 2: The forest, with pristine area shaded.

into the forest and must exit the forest, despite having been caught (there is no positive punishment beyond the confiscation itself and being prevented from performing further extraction while exiting). Thus a given patrol strategy s will reduce the extractor’s *expected* benefit for an incursion of distance d from $B(d)$ to some value $B_s(d)$. Equivalently, once the patrol strategy s is fixed, the choice of d determines a cost $C(d)$ and an expected benefit $B_s(d)$.

The rational extractor will choose a distance trespass d that maximizes his expected profit $B_s(d) - C(d)$, which is the d for which the curves $b_s(d)$ and $c(d)$ intersect (see Fig. 3(a)) and his marginal profit $b_s(d) - c(d)$ equals zero. (Or the infimum of such d if the intersection extends over an interval, which can be justified by slight perturbations.) Beyond this point, the marginal cost of extraction outweighs the marginal benefit. We emphasize that the extractor’s strategy (the value d) is chosen offline, in advance, with the knowledge of the patrol allocation $\phi(\cdot)$. The extractor acquires no new information online that can affect his decision-making: patrols are invisible to the trespasser until they catch him, at which point he is forced to exit. For a sufficiently fast-growing cost function relative to the benefit function, there will be a *natural core* of pristine forest even with no patrols

at all (Albers 2010), which we assume.

Notation. $b(\cdot), c(\cdot), \phi(\cdot)$ are the (marginal) benefit, cost, and capture probability density functions, respectively, and are the derivatives of $B(\cdot), C(\cdot), \Phi(\cdot)$, the corresponding cumulative functions. $p(\cdot) = b(\cdot) - c(\cdot)$ and $P(\cdot) = B(\cdot) - C(\cdot)$ are the corresponding (net) profits. d_s for $s \in \{n, o, r\}$ is the *trespass distance* under no patrols, the optimal patrol, and the best ring patrol, respectively. $r_s = 1 - d_s$ is the radius of the pristine forest area under some patrol strategy s . Similarly for $b_s(\cdot), B_s(\cdot), p_s(\cdot), P_s(\cdot)$, which, under patrols, are benefits and profits in expectation. $\hat{r}_s = r_s - r_n = d_n - d_s$ is the *pristine radius increase* or *trespass distance decrease* under patrol strategy s .

Detection models. An extractor is detected if he comes within some distance $\Delta \ll 1$ of the patrol. Under our time model, the patrol units move much less quickly than the extractors, and so patrols can be modeled as stationary from the extractor’s point of view. Therefore, if e.g. $\phi(\cdot)$ is constant (for a single patrol unit) over the region R (of size $|R|$), then the probability of detection for an extraction path of length d is proportional to ϕd , specifically $\phi d 2\Delta / |R|$, where the total area within distance Δ of the length- d walk is approximated as $d \cdot 2\Delta$. That is, probabilities are added rather than “multiplied” due to stationarity. (We assume the patrol unit is *not* visible to the extractor.) The model described here also covers settings in which the amount spent at a location determines the sensing range Δ there. For notational convenience, we drop Δ and $|R|$ throughout the paper, assuming normalization as appropriate.

Leader strategies. The leader has a budget $E \in [0, 1]$ specifying a bound on the total detection probability mass, for the patrol’s presence at a given location, that can be distributed across the region. The task is to choose an allocation in order to minimize the rational extractor’s resulting trespass distance d_s . Due to forest symmetry and the fact that extractors’ decisions are uncoordinated, the problem is essentially one-dimensional. Specifically, the leader strategy specifies a patrol probability density $\phi(x)$ for each $x \in [0, 1]$, which density is reflected symmetrically about the circle of radius x .

Definition 1. Let $OPT(I)$ be the optimal solution value of a problem instance I , and let $ALG(I)$ be the solution value computed by a given algorithm. An algorithm for a maximization problem is a c -approximation (with $c < 1$) with additive error ϵ if, for every problem instance I , we have $ALG(I) \geq c \cdot OPT(I) - \epsilon$.

Minimizing d_s and maximizing r_s are equivalent in terms of optimal solutions; for approximations, we optimize for maximizing r_s . Note that a c -approximation for the pristine radius increase $\hat{r}_s = r_s - r_n$ implies a c -approximation for radius r_s .

Numerical precision issues. Our two algorithms involve binary search, numerical integration, and root-finding, all of which can introduce errors in subtle ways. Both algorithms’ binary search outer loops introduce an additive error ϵ to the resulting pristine radius. When $b(\cdot)$ and $c(\cdot)$ are not both

polynomials, Algorithm 1 performs a numerical integration whose error effectively reduces the algorithm’s budget by ϵ , with the effect that the ϵ error bound is in comparison to the optimal radius achievable with budget $E - \epsilon$. Finally, both algorithms’ binary search loops perform a root-finding step. To simplify the analysis and presentation, we assume that root-finding can be performed exactly in constant time.

Patrol Allocations

Let the *patrol zone* be the region of the forest assigned nonzero patrol density. We note three patrol allocation strategies that have been proposed in the past:

- **Homogeneous:** Patrol density distributed uniformly over the entire region.
- **Boundary:** Patrol density distributed uniformly over a ring (of some negligible width w) at the forest boundary.
- **Ring:** Patrol density distributed uniformly over a ring (of negligible width w) concentric with the forest.

Boundary patrols can be superior to homogenous patrols, since homogeneous patrols waste enforcement on the natural core (Albers 2010). It is interesting to note that this is not always so. Suppose w is very small and the homogenous-induced core radius is less than $1 - d$ for some trespass distance d satisfying $w < 1/2 < d \leq 1$. With homogenous patrols, we will have $\Phi(d) = E/\pi \cdot d$. With boundary patrols, however, this probability for any $d \geq w$ will be $\frac{E}{\pi - \pi(1-w)^2} \cdot w = E/\pi \cdot \frac{w}{1-(1-w)^2}$, which approaches $\frac{E}{2\pi}$ as $w \rightarrow 0$. In this case, homogeneous patrols will actually outperform boundary patrols. Intuitively, this is because patrols in the interior will “intersect” more boundary-to-center trespass rays than a patrol on the boundary will. Unfortunately, both boundary and homogeneous patrols can perform arbitrarily badly.

Proposition 1. *The approximation ratios of homogeneous and boundary patrols are both 0.*

Proof. To see this, consider the following example. Let the (non-normalized) forest radius R be extremely large, with $B(x) = x$ for $x \leq R$ (and 0 thereafter) and $C(x) = 0$ for all x . Then the rational extractor will trespass by distance R , and the natural core is empty.

Let the budget be fixed at $E = 1$, and first consider homogenous patrols. This implies a constant patrol density of $\frac{1}{\pi R^2}$. For each possible trespass distance $0 \leq x \leq R$, the cumulative expected profit is $(1 - \Phi(x))B(x) = (1 - \frac{x}{\pi R^2})x$, which is strictly increasing for all $0 \leq x \leq R$.

A width-1 boundary patrol (recall $R \gg 1$) with budget 1 will place a density of $\frac{1}{\pi(R^2 - (R-1)^2)} = \frac{1}{\pi(2R-1)}$ about the perimeter. This yields an expected cumulative profit of $(1 - \Phi(x))B(x) = (1 - \frac{\min\{x, 1\}}{\pi(2R-1)})x$, which again is strictly increasing $0 \leq x \leq R$.

Thus in both cases the rational extractor will continue to trespass all the way to the forest center, yielding a pristine area radius of 0, although a strictly positive pristine radius is achievable. Consider, e.g., allocating constant density over a radius-1 disk at the forest center, which would imply a




Figure 3: The shaded regions in the latter two subfigures correspond to the reduction in marginal benefits *within the patrol zone*. Note that the marginal benefits under patrols are also reduced *following* the patrol zone.

density of $1/\pi$ there. Then the expected cumulative profit for $R-1 \leq x \leq R$ will be $(1 - \Phi(x))B(x) = (1 - \frac{x-(R-1)}{\pi}) \cdot x$, which is strictly *decreasing* for such x . Thus the rational extractor will never enter this disk, and so the pristine area radius is 1. \square

Instead, our optimal patrol will be of the following sort:

- **Band:** The shape of the patrol zone is a band, i.e., the set difference of two circles, both concentric with the forest.

The cumulative profit for trespassing by distance x under some patrol strategy s is:

$$\begin{aligned} P_s(x) &= B_s(x) - C(x) \\ &= (1 - \Phi_s(x)) \cdot B(x) - C(x) \\ &= P(x) - \Phi_s(x)B(x) \end{aligned}$$

and so the marginal profit is:

$$\begin{aligned} p_s(x) &= dP_s(x)/dx \\ &= p(x) - \Phi_s(x)b(x) - \phi_s(x)B(x) \end{aligned} \quad (1)$$

Observe that $\phi(\cdot)$ influences the extractor's expected benefit function, and hence his cost-benefit analysis, in *two* ways. First, the probability of successfully traveling distance x (on the return trip) is reduced by the cumulative probability $\Phi(x)$ of capture up to that point, and so the marginal profit at point x is reduced from $p(x)$ by amount $\Phi(x)b(x)$. Second, being caught at point x with probability density $\phi(x)$ means losing the full benefit accrued so far, which further reduces the marginal profit at this point by $\phi(x)B(x)$.

Lemma 1. *Without loss of generality, we may assume:*

1. *The patrol zone of $\phi_o(\cdot)$ is (d_o, e_o) for some e_o .*
2. *$\phi_o(x)$ at each point $x \in (d_o, e_o)$ is the smallest possible value providing no positive profit at point x , i.e., that density yielding $b_o(x) = c(x)$.*
3. *$b_o(x) \leq c(x)$ and $\phi_o(x) = 0$ for $x \geq e_o$.*

Proof. (1) is provable by a modification to the following argument proving (2). Consider an allocation $\phi_o(\cdot)$ that successfully stops the extractor at d_o but which violates the stated property, at some particular level of discretization. That is, partition the interval (d_o, e_o) into n equal sized subintervals, numbered x_1, \dots, x_n . For this discretization, we write $B_o(x_i) = \sum_{j=1}^{i-1} b_o(j)$ and $\Phi_o(x_i) = \sum_{j=1}^{i-1} \phi_o(i)$

(omitting coefficients). Let x_i be the first such subinterval for which $p_o(x_i) < 0$, and let x_i^+ be shorthand for $x_i + 1$. In this case (see Eq. 1) we have $p(x_i) - \Phi_o(x_i)b(x_i) - \phi_o(x_i)B(x_i) < 0$. We correct this by subtracting a value δ from $\phi_o(x_i)$ to bring about equality, and adding δ to $\phi_o(x_i^+)$.

The marginal profit of step x_i is then 0 (by construction), and that of step x_i^+ is only lower than it was before, so there is no immediate payoff to walking from x_i to x_i^+ or $x_i + 2$. Clearly $\Phi_o(x_i + 2)$ is unchanged. Finally, we verify that the expected total profit of walking to position $x_i + 2$ is unchanged. This benefit is affected by the changes to both $\phi_o(x_i)$ and $\phi_o(x_i^+)$. First, $\delta B(x_i)$ is added to $b_o(x_i)$ by subtracting δ from $\phi_o(x_i)$; second, $b_o(x_i^+)$ becomes

$$\begin{aligned} &b(x_i^+)(1 - (\Phi_o(x_i^+) - \delta)) - (\phi_o(x_i^+) + \delta) \cdot B(x_i^+) \\ &= b(x_i^+)(1 - \Phi_o(x_i^+)) + b(x_i^+)\delta - \phi_o(x_i^+)B(x_i^+) - \delta B(x_i^+) \\ &= b(x_i^+)(1 - \Phi_o(x_i^+)) - \phi_o(x_i^+)B(x_i^+) + b(x_i^+)\delta - \delta B(x_i^+) \\ &= b_o(x_i^+) - \delta B(x_i) \end{aligned}$$

Thus, since these two changes cancel out and there was no incentive for walking from x_i *past* $x_i + 2$ prior to the modification, this remains true, and so the extractor will walk no farther than he did before the modification. We repeat this modification iteratively for all earliest adjacent violations (x_i, x_i^+) , and for discretization precisions n . Since outer rings of circular (or, more generally, convex) forests have greater circumference, each such operation of moving patrol density forward only lowers the patrol's total cost.

Finally, (3) follows from $\phi_o(\cdot)$ being a band that stops the extractor at position d_o . \square

Remark. We emphasize that the lemma implies that in an optimal allocation the patrols will, perhaps counterintuitively, occur within the pristine area of the forest, and so the rational trespasser will actually never encounter the patrols. We remark that under the resulting allocation, patrol density will decline monotonically with distance into the forest. Intuitively, the reason for this is that as distance into the forest grows, there is a smaller and smaller remaining marginal profit $p(x)$ that we need to compensate for with threat of confiscation, and yet the *magnitude* of the potential confiscation $B(x)$ grows only larger. We also remark that this point e_o will occur strictly before d_n , because for large enough $x \in (d_o, d_n)$, $p(x)$ will be small enough to ensure

Algorithm 1 Computing the optimal allocation(b, c, E, ϵ)

```
1:  $(d_1, d_2) \leftarrow (0, d_n)$ 
   binary search:
2: while  $d_1 < d_2 - \epsilon$  do
3:    $d \leftarrow (d_1 + d_2)/2$ 
4:    $\phi(x) \triangleq \frac{b(x) \cdot (C(x) + P(d)) - B(x)c(x)}{B(x)^2}$ 
5:    $e \leftarrow x$  s.t.  $d \leq x \leq d_n$  and  $\phi(x) = 0$ 
6:    $E(d) \leftarrow \int_d^e 2\pi(1-x)\phi(x)dx$  (within error  $+\epsilon$ )
7:    $\{d_2 \leftarrow d, \phi_2 \leftarrow \phi\}$  if  $E(d) \leq E$  else  $d_1 \leftarrow d$ 
8: end while
9: return  $(d_2, \phi_2)$ 
```

that $\Phi_o(e_o)b(x) < c(x)$, rendering additional patrol density within $[e_o, d_n]$ superfluous.

Theorem 1. *Algorithm 1 provides an allocation whose pristine radius (obtained using budget E) is within ϵ of the optimal pristine radius obtainable using budget $E - \epsilon$.*

Proof. We assume the properties stated by Lemma 1. Observe that for $x \leq d_o$, $b_o(x) = b(x)$; for $x \geq e_o$, $b_o(x)$ is determined only by $b(x)$ and the cumulative capture probability, i.e., $b_o(x) = (1 - \Phi_o(x)) \cdot b(x)$. e_o is the point at which $\phi_o(x) = 0$ and $(1 - \Phi_o(x)) \cdot b(x) - c(x) = 0$. Now we compute $\phi_o(\cdot)$. Setting Eq. 1 to 0 yields:

$$\phi_o(x) = \frac{p(x) - \Phi_o(x)b(x)}{B(x)} \quad (2)$$

The solution to this standard-form first-order differential equation (recall that $\Phi_o(x) = \int_{d_o}^x \phi_o(y)dy$, and note that $\Phi_o(\cdot)$ depends on the value d_o) is:

$$\Phi_o(x) = e^{-\int R(x)dx} \cdot \left(\int Q(x) \cdot e^{\int R(x)dx} dx + K \right)$$

where $R(x) = \frac{b(x)}{B(x)}$, $Q(x) = \frac{p(x)}{B(x)}$, and K is a constant. Since $\int R(x)dx = \int \frac{b(x)}{B(x)}dx = \ln B(x)$, we have $e^{\int R(x)dx} = B(x)$. Therefore,

$$\begin{aligned} \int Q(x) \cdot e^{\int R(x)dx} dx &= \int \frac{p(x)}{B(x)} \cdot B(x)dx \\ &= \int p(x)dx = P(x) \end{aligned}$$

and, based on initial condition $\Phi_o(d_o) = 0$,

$$K = - \int Q(x) \cdot e^{\int R(x)dx} dx \Big|_{d_o} = -P(d_o)$$

Since $\phi_o(x) = \Phi'_o(x)$, this yields:

$$\begin{aligned} \Phi_o(x) &= \frac{P(x) - P(d_o)}{B(x)} \\ \phi_o(x) &= \frac{b(x) \cdot (C(x) + P(d_o)) - B(x)c(x)}{B(x)^2} \end{aligned}$$

Then an allocation of $\phi_o(x)$ for $x \in (d_o, e_o)$ will provide the optimal trespass distance d_o for the budget $E(d_o) =$

Algorithm 2 Computing the best ring patrol(b, c, E, ϵ)

```
1:  $(d_1, d_2) \leftarrow (\epsilon/2, d_n)$ 
   binary search:
2: while  $d_1 < d_2 - \epsilon/4$  do
3:    $d \leftarrow (d_1 + d_2)/2$ ,  $a \leftarrow d - \epsilon/2$ 
4:    $\phi \leftarrow E/(\pi((1-d)\epsilon + \epsilon^2/4))$ ,  $\Phi \leftarrow \phi \cdot \epsilon/2$ 
5:    $f \leftarrow x$  s.t.  $(1 - \Phi)b(x) = c(x)$ 
6:    $\Delta P_1 \leftarrow P(d) - P(a) - \Phi B(d)$ 
7:    $\Delta P_2 \leftarrow P(f) - P(d) + \Phi \cdot (B(d) - B(f))$ 
8:    $\{d_2 \leftarrow d, \phi_2 \leftarrow \phi\}$  if  $\Delta P_1 + \Delta P_2 < 0$  else  $d_1 \leftarrow d$ 
9: end while
10: return  $(d_2, \phi_2)$ 
```

$\int_{d_o}^{e_o} 2\pi(1-x)\phi_o(x)dx$. We search for the smallest feasible d_o by binary search, with error ϵ . The integral in line 6 is computed within *positive* error ϵ (in some time $\tau(\epsilon)$ depending on $b(\cdot)$ and $c(\cdot)$), by adding $\epsilon/2$ to a $\pm\epsilon/2$ -error approximation, which effectively reduces budget E to $E - \epsilon$. \square

Remark. If $b(\cdot)$ and $c(\cdot)$ are polynomials, then the guarantee can be made relative to the same budget E , since in this case $\phi(x)$ is a rational function and so $E(d_o)$ is solvable analytically, by the method of partial fractions.

The varying-density allocation of Algorithm 1 may be difficult or impractical to implement; moreover, each iteration of the binary search loop requires an expensive iterative approximation parameterized by $1/\epsilon$, if the integral in line 6 is not solvable analytically. Now we present a more efficient algorithm that produces easier-to-implement allocations. Assuming the intersection of $b(\cdot)$ and $c(\cdot)$ can be found in constant time, Algorithm 2 runs in time $O(\log 1/\epsilon)$.

Definition 2. *Let $R_\epsilon(r)$ indicate the interval $(r, r + \epsilon)$. Let an ϵ -ring of radius r be a patrol of constant density over $R_\epsilon(r)$, rotated about the circle.*

Lemma 2. *The $\epsilon/2$ -ring given by Algorithm 2 increases the pristine radius by within $3/4\epsilon$ of the largest possible increase achievable with an $\epsilon/2$ -ring (and budget $E - \epsilon$).*

Proof. For a given possible value r (set $d = 1 - r$), an $\epsilon/2$ -ring of radius r and interval $(a, d) = (d - \epsilon/2, d)$ has $\phi = E/(\pi r \epsilon + \epsilon^2/4)$ for x within the ring and $\Phi = \phi \cdot \epsilon/2$ for x beyond the ring. Consider an extractor choosing between the stopping points of a and some point $f > d$ beyond the ring. In order to deter the extractor from continuing from a to any such f , it must be the case that the resulting change in expected profit is negative. That is, we must have $\Delta P_1 + \Delta P_2 < 0$, where these are the profit changes during the ring $P_r(x)|_a^d$ and after the ring $P_r(x)|_d^f$, respectively. Because $\Phi_r(a) = 0$ and $\Phi_r(d) = \Phi_r(f) = \Phi$, we have:

$$\begin{aligned} \Delta P_1 &= P(d) - P(a) - \Phi_r(d)B(d) + \Phi_r(a)B(a) \\ &= P(d) - P(a) - \Phi B(d) \\ \Delta P_2 &= P(f) - P(d) - \Phi_r(f)B(f) + \Phi_r(d)B(d) \\ &= P(f) - P(d) + \Phi \cdot (B(d) - B(f)) \end{aligned}$$

We do binary search for the smallest value d for which $\Delta P_1 + \Delta P_2 < 0$, to within error $\epsilon/4$. This means the chosen ring radius r is within $\epsilon/4$ of the largest $\epsilon/2$ -ring radius

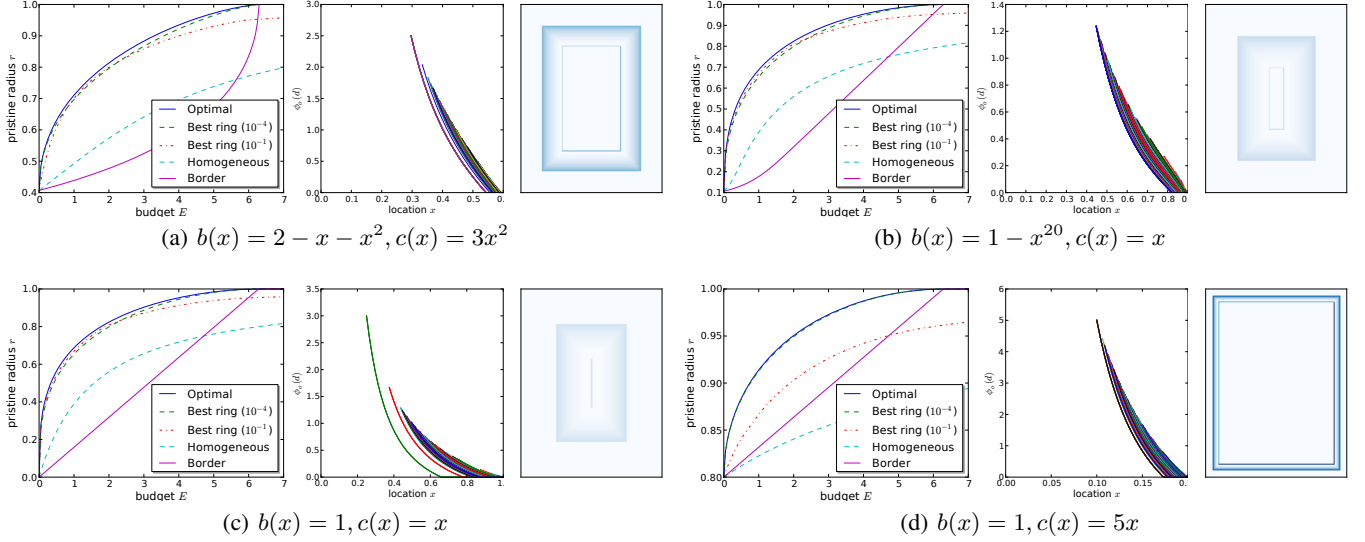


Figure 4: Patrol strategy effectiveness for sample $b(\cdot), c(\cdot)$ functions.

for which $d - \epsilon/2$ is a better stopping point than anywhere beyond the ring.

This does not preclude the possibility of there being a smaller-radius ϵ -ring that would induce a smaller trespass distance, since perhaps the trespasser could be induced to stop within the ring itself. But the possible improvement to obtained by decreasing the radius is of course bounded by the width of the ring— $\epsilon/2$. Thus the total error in trespass distance compared to that of the best possible ϵ -ring patrol is at most $3/4\epsilon$. \square

Lemma 3. *There exists an $\epsilon/2$ -ring that increases the pristine radius by at least $1/2\hat{r}_o - \epsilon/4$.*

Proof. As x grows from d_o to d_n , $b_o(x)$ falls monotonically while $c(x)$ grows, and $B_o(x)$ and $\Phi_o(x)$ both grow monotonically (because $b(x) \geq 0$ and $\phi_o(x) \geq 0$). Thus by Eq. 2, $\phi_o(x)$ falls monotonically over (d_o, d_n) .

Now consider the location $d_r = (d_o + d_n)/2 + \epsilon/4$ and the $\epsilon/2$ -ring of radius $r_r = 1 - d_r$, for which $\phi_r(\cdot) = \frac{1}{\epsilon/2} \int_{d_o}^{d_n} \phi_o(x) dx$ within $R_{\epsilon/2}(r_r)$ and 0 elsewhere. The true cost of patrol density $\phi(x)$ at location x rotated about the circle is $2\pi(1-x)\phi(x)$. Because $\phi_o(\cdot)$ is monotonic decreasing, we have $\int_{d_o}^{d_r} \phi_o(x) dx \geq \int_{d_r}^{d_n} \phi_o(x) dx$. Thus $\phi_r(\cdot)$ will be only cheaper than $\phi_o(\cdot)$, and so also in budget.

We claim that the rational extractor will be deterred from crossing the constructed $\epsilon/2$ -ring, which means its pristine radius increase is at least $(r_o - r_n)/2 - \epsilon/4$. Indeed, we know by definition that $P_o(d_o) > P_o(x)$ for all $x \in (d_o, d_n)$, and in particular for all $x \in [d_r, d_n]$ (recall $d_o < d_r < d_n$). With the constructed ring patrol and for such x , we have $P_r(x) \leq P_o(x)$, because $\Phi_r(x) \geq \Phi_o(x)$. Finally, we have $P_o(d_o) = P_r(d_o)$, since for small enough ϵ (specifically, $\epsilon/2 < \hat{r}_o/2$), the ring will not begin until after d_o . Combining inequalities, we obtain: $P_r(x) \leq P_o(x) < P_o(d_o) = P_r(d_o)$. Hence d_o

in particular is a better stopping point for the extractor than any point $x \in [d_r, d_n]$, and so the extractor's best stopping point will lie somewhere within $[d_r, r_r)$. \square

Combining the two preceding lemmas, we then immediately obtain that Algorithm 2 is a $1/2$ -approximation:

Theorem 2. *The allocation of Algorithm 2 increases the pristine radius by at least $(1/2\hat{r}_o - \epsilon/4) - 3/4\epsilon = 1/2\hat{r}_o - \epsilon$.*

We note that the approximation ratio is tight. To see this, problem instances can be constructed satisfying the following: $c(x) = 0$ and $b(x)$ is constant (and small) over the interval (d_o, d_n) (which meets an empty natural core, i.e. $d_n = 1$), and E is very small and hence (d_o, d_n) is very narrow. In this case, $\Phi_o(x)$ grows very slowly over the patrol region, and $\phi_o(x)$ declines very slowly over it. In the extreme case, the weight of $\phi_o(x)$'s probability mass to the right of d_r approaches the weight to the left.

Algorithmic extensions

Multiple patrol units. We can extend from one to multiple patrol units, weighted equally or unequally. Given k patrol units, each given budget $E_i \in [0, 1]$, with $E = \sum E_i$, we partition the forest into k sectors, each of angle $2\pi E_i/E$. We run one of our algorithms above, with budget E . (Observe that for both algorithms, having $E > 1$ is equivalent to having $E = 1$ on some smaller slice of the circle.) Then we choose positions for patrol unit i within sector i , i.e., proportional to $\phi(\cdot)$, which is reflected about the sector. Note that an extractor's path necessarily lies within a single patrol's sector, and so adding capture probabilities remains justified.

Other forest shapes. In the noncircular forest context, permitting extractors to traverse any length-bounded path from their starting points implies that the pristine area determined by a given patrol strategy will again be an erosion

of the forest. Computing the erosion of an arbitrary shape is computationally intensive (Soille 2004), but it is easily computable for convex polygons, which will approximate many realistic forests. In order to be practically implementable in such cases, the patrol should be symmetric around the forest area. Our algorithms above adapt easily to the setting of convex polygon forest shapes, where pristine areas are erosions, by integrating the cost of a patrol around the forest boundary. In both cases, we replace the circle circumference $2\pi(1-x)$ with the cost of the corresponding polygon circumference. For large polygons with a reasonable number of sides, the resulting error due to corners will be insignificant.

Experiments

We implemented both our algorithms, as well as the baseline solutions of homogenous and boundary patrols. We tested these algorithms on certain realistic pairs of benefit and cost functions (with forest radius 1; see four examples in Fig. 4). We now summarize our observations on these results.

In each setting (see left subfigures), we vary the patrol budget (with $E > 1$ meaning multiple patrols), computing the patrol allocation function and hence the extractor's trespass distance d_s , for each. First, the optimal algorithm indeed dominates all the others. Both our algorithms perform much better overall than the two baselines, however, up until the point at which the budget is sufficient to deter any entry into the forest, using boundary and best ring. Best ring will consider a ring at the boundary, so it cannot do worse than boundary, and so the two curves must intersect at zero. Prior to this best ring does outperform boundary. As observed above, neither homogeneous nor boundary consistently dominates the other.

We computed ring patrols for two ring widths, one very narrow ($1/10^4$) and one less so (0.1). Interestingly, neither ring size dominates the other. With a sufficiently large budget, the rings will lie on the boundary, but a wider ring will permit some nonnegligible trespass (part way across the ring itself). With smaller budgets the rings will lie in the interior of the forest. In this case, the narrow ring will spend the entire budget at one (expensive) density level, whereas the wider ring can will (more cheaply, and hence more successfully) spend some of its budget at lower-density levels.

Next (see middle subfigures), we plot the optimal $\phi_o(\cdot)$ functions under many different budgets. As can be seen, these curves sweep out different regions of the plane, depending on the $b(\cdot)$, $c(\cdot)$ pair.

Finally (see right subfigures), we illustrate the result of applying Algorithm 1 to a rectangular forest, with one sample budget (3.5, normalized to the dimensions of the forest). The patrol density is represented by the level of shading. The border of the natural core is also shown.

Acknowledgement. This research is supported by MURI grant W911NF-11-1-0332. We thank Jo Albers for introducing the problem to us. We also thank Yi Gai for many helpful discussions.

References

Albers, H. J. 2010. Spatial modeling of extraction and enforcement in developing country protected areas. *Resource*

and Energy Economics 32:165–179.

Becker, G., and Landes, W. 1974. *Essays in the Economics of Crime and Punishment*. Columbia University Press.

Boyle, D. Dec. 21, 2011. Logging in the wild west. *The Phnom Penh Post*.

Clarke, H. R.; Reed, W. J.; and Shrestha, R. M. 1993. Optimal enforcement of property rights on developing country forests subject to illegal logging. *Resource and Energy Economics* 15:271–293.

Dixon, J. A., and Sherman, P. B. 1990. *Economics of Protected Areas: A New Look at Benefits and Costs*. Washington, DC: Island Press.

Hall, J. B., and Rodgers, W. A. 1992. Buffers at the boundary. *Rural Development Forestry Network Summer* (Paper 13a).

Hofer, H.; Campbell, K. L. I.; East, M. L.; and Huish, S. A. 2000. Modeling the spatial distribution of the economic costs and benefits of illegal game meat hunting in the serengeti. *Natural Resource Modeling* 13(1):151–177.

Lober, D. J. 1992. Using forest guards to protect a biological reserve in Costa Rica: one step towards linking parks to people. *Journal of Environmental Planning and Management* 35(1):17.

MacKinnon, J.; MacKinnon, K.; Child, G.; and Thorsell, J. 1986. *Managing Protected Areas in the Tropics*. Gland, Switzerland: IUCN.

Milliman, S. R. 1986. Optimal fishery management in the presence of illegal activity. *Journal of Environmental Economics and Management* 12:363–381.

Paruchuri, P.; Pearce, J. P.; Marecki, J.; Tambe, M.; Ordóñez, F.; and Kraus, S. 2008. Playing games with security: An efficient exact algorithm for Bayesian Stackelberg games. In *AAMAS*.

Pita, J.; Tambe, M.; Kiekintveld, C.; Cullen, S.; and Steigerwald, E. 2011. Guards - game theoretic security allocation on a national scale. In *AAMAS (Industry Track)*.

Robinson, E. J. Z.; Albers, H. J.; and Williams, J. C. 2008. Spatial and temporal modelling of community non-timber forest extraction. *Journal of Environmental Economics and Management* 56:234–245.

Robinson, E. J. Z.; Albers, H. J.; and Williams, J. C. 2011. Sizing reserves within a landscape: The roles of villagers' reactions and the ecological-socioeconomic setting. *Land Economics* 87:233–249.

Robinson, E. J. Z. 2008. India's disappearing common lands: fuzzy boundaries, encroachment, and evolving property rights. *Land Economics* 84(3):409–422.

Sanchirico, J. N., and Wilen, J. E. 2001. A bioeconomic model of marine reserve creation. *Journal of Environmental Economics and Management* 42(November):257–276.

Soille, P. 2004. *Morphological Image Analysis: Principles and Applications*. Springer.

Tsai, J.; Rathi, S.; Kiekintveld, C.; Ordóñez, F.; and Tambe, M. 2009. IRIS: a tool for strategic security allocation in transportation networks. In *AAMAS (Industry Track)*.