# Designing Patrol Strategies to Maximize Pristine Forest Area

Matthew P. Johnson, Fei Fang, Milind Tambe
University of Southern California
Los Angeles, CA 90089, USA
{matthepj, feifang, tambe}@usc.edu

H.J. Albers
Oregon State University
Corvallis, OR 97331, USA
hj.albers@oregonstate.edu

## ABSTRACT

Illegal extraction of forest resources is fought, in many developing countries, by patrols that seek to deter such activity by decreasing its profitability. 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 maximize the amount of "pristine" forest area. Prior work in forest economics has posed this problem as a Stackelberg game, but efficient optimal or approximation algorithms for generating leader strategies have not previously been found. Unlike previous work on Stackelberg games in the multiagent literature, much of it motivated by counter-terrorism, here we seek to protect a continuous area, as much as possible, from extraction by an indeterminate number of followers. The continuous nature of this problem setting leads to new challenges and solutions, very different in character from in the discrete Stackelberg settings previously studied.

In this paper, we give an optimal patrol allocation algorithm and a guaranteed approximation algorithm, the latter of which is more efficient and yields simpler, more practical patrol allocations. In our experimental investigations, we find that these algorithms perform significantly better—yielding a larger pristine area—than naive patrol allocations.

## Categories and Subject Descriptors

F.2.2 [**Analysis of Algorithms and Problem Compelxity**]: Nonnumerical Algorithms and Problems—*geometrical problems and computations*

## General Terms

Algorithms, Economics, Security, Theory

## Keywords

patrol strategy, forest protection, Stackelberg game

## 1. INTRODUCTION

Illegal extraction of fuelwood or other natural resources from forests is a problem confronted by officials in many developing countries, with only partial success [8, 4, 3, 12]. 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 [7, 16]. 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.



**Figure 1: "A truck loaded with illegally cut rosewood passes through Russey Chrum Village...in the Central Cardamom Protected Forest." Photo from [2].**

We pose 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 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 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 to him as possible.[2] More precisely, given the continuous nature of this setting, we wish to *minimize* to the amount of extraction that will yield a positive net return in his cost-benefit analysis.

**Background.** Economists have studied the relationship generally between enforcement policy for protecting natural resources and the resulting incentives for neighbors of the protected area [9, 12, 15]. Our point of departure in this paper is the influential forest protection model of [1] (see also [13, 14]), in which there is a circular forest surrounded by villages (hence potential extractors); the task is to distribute the patrols' probability density across the region of interest; the objective is to minimize the distance by which the extractors will trespass into the forest and hence (since nearby villagers will extract as a function of this distance [6]) or maximize the size of the resulting *pristine* forestland.

We strengthen this model in several ways, permitting spatial variation in patrol density, multiple patrol units, and convex polygon-

---

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

shaped forests. As has been observed [1], exogenous legal restrictions on patrol strategies, such as requiring *homogenous* patrols, can degrade protection performance [8, 5]. Unlike the existing work on this model, we bring to bear algorithmic analysis on the problem. Specifically, we show that while certain such 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 and has been applied to a number of real-world security domains, including the Los Angeles International Airport [10], the Federal Air Marshals Service [19], and the Transportation Security Administration [11]. See [18] for an overview.

The problem setting we address here differs from those considered in these previous works, most crucially in that the forest protection setting 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 the 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 highly inefficient due to the geometric density.) Once we have computed a distribution over patrol locations, selecting patrols locations is straightforward. As such, our primary focus is on the choice of distribution for patrol density over the two-dimensional forest region, i.e. a probability distribution from which to select patrols.

**Contributions.** We give a full analysis of the problem of maximizing pristine forest radius. Our main contributions are efficient optimal and 1/2-approximation algorithms for this problem, 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.

## 2. PROBLEM SETTING
In this section we present the forest model of [1] 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, though we later extend to convex polygons. An extractor's action is to choose some distance $d$ to walk into the forest, extracting on the return trip. We may assume, without loss of generality, that the extractor's route goes the chosen distance $d$ towards the forest center (on a straight line), before reversing back to his starting point $P$ on the perimeter. To see this, observe that all possible paths from $P$ will sweep out a lens-like shape but, since all points on the perimeter are possible starting points, the set of all trespass paths directed towards the center sweeps out the same area. Given our objective of maximizing *pristine* forest area, this holds true even if extractors are distributed around the perimeter *nonuniformly*, as long as there is a nonzero probability of villager presence at each point on the perimeter.

Due to symmetries and the fact that extractors' decisions are unco-

ordinated, the problem is essentially one-dimensional. Extractors incur a cost and gain a benefit *if not caught*, based on an increasing marginal cost function $c(d)$ and a decreasing marginal benefit function $b(d)$. (The instantaneous or *marginal* cost and benefit functions are the derivatives of the functions specifying the cumulative costs and benefits, respectively, of walking that far into the forest.) If caught, the extractor's benefit is 0 (the extracted resources are confiscated) but the cost is unchanged (the extractor's traveled distance does not change; there is no positive punishment beyond the confiscation itself and being prevented from engaging in further extraction while leaving the forest). Since extraction can be assumed to occur only on the return trip, and given the nature of the punishment, we may restrict our attention to detection on the return trip. Thus a given patrol strategy will reduce the extractor's expected benefit for an incursion of distance $d$ from $b(d)$ to some value $b_p(d)$.
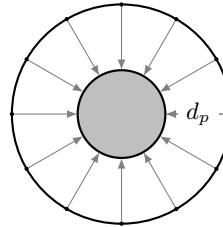
**Figure 2: The forest, with the pristine area shaded.**

For a sufficiently fast-growing cost function relative to the benefit function, there will be a "natural core" of pristine forest even with no patrolling at all [1]; that is, the optimal trespass distance will be less than 1, since the marginal cost of extraction will eventually outweigh the marginal benefit, corresponding to the point at which the curves $b(d)$ and $c(d)$ intersect (see Fig. 3). The overall result of choosing a given patrol strategy therefore is to transform the benefit curve $b(d)$ into a lower benefit curve $b_p(d)$, thus reducing the extractor's optimal incursion distance (see Fig. 3). In the language of mathematical morphology [17], the pristine forest area $P$ due to a given patrol strategy will be an *erosion* $P = F \ominus B$ of the forest $F$ by a shape $B$, where $B$ is a circle whose radius equals the trespass distance. The erosion is the locus of points reached by the center of $B$ as it moves about inside of $F$.

**Notation.** $b(x), c(x), \phi(x)$ are the *marginal* benefit, cost, and capture probability functions, respectively. $B(x), C(x), \Phi(x)$ are the corresponding *cumulative* functions. $d_p$ for $p \in \{n, o, r\}$ is the *trespass distance* under no patrols, the optimal patrol allocation, the best ring allocation, respectively. $r_p$ is the radius of the pristine forest area under some patrol $p$. (Similarly, $b_p(x), B_p(x)$.) $d_n - d_p$ is the *reduction in trespass distance* under this patrol.

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 $c < 1$) if, for every problem instance $I$, we have $ALG(I) \geq c \cdot OPT(I)$.*

The leader has a budget $E$ specifying a bound on the total detection probability mass that can be distributed across the region. The task is to choose an allocation in order to minimize the extractor's resulting optimal trespass distance $d_n$, which is equivalent to maximizing the trespass distance reduction and *implies* maximizing the pristine radius. Note that our optimal and approximation algorithms both perform binary search, and thus incur an additive error $\epsilon$.

### 2.1 Detection probability models
Let $\phi(x)$ be the detection probability density function chosen by the leader for the forest. An extractor is detected if he comes within

some distance $\Delta << 1$ 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. $\phi(x)$ 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 $\phi d$, specifically $\phi d 2\Delta/|R|$, where the total area within distance $\Delta$ of the length-$d$ walk is approximated as $d \cdot 2\Delta$. That is, probabilities are added rather than "multiplied" due to stationarity. (Here 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 $\Delta$ there. For notational convenience, we drop $\Delta$ and $|R|$ throughout the paper, assuming normalization as appropriate.

$\phi(x)$ influences the extractor's behavior in *two* ways. The rational extractor will trespass a distance into the forest that maximizes his total (or cumulative) net benefit, which is where his net marginal benefit $b(x) - c(x)$ equals zero. As the extractor moves about through a region with nonzero $\phi(x)$, his cost-benefit analysis is affected in two ways. First, the probability of reaching a given location $x$ is reduced by the cumulative probability of capture up to that point, $\Phi(x)$, and so the net marginal benefit at point $x$ is reduced from $b(x) - c(x)$ by amount $\Phi(x)b(x)$. (Recall that capture occurs on the return trip out of the forest, and so the cost $c(x)$ is paid *regardless of whether confiscation occurs*.) Second, being caught at point $x$ is $\phi(x)$ means losing the full benefit accrued so far, which further reduces the net marginal benefit at this point by amount $\phi(x)B(x)$, where $B(x) = \int_{y=0}^{x} b(y)dy$ is the cumulative benefit.

We emphasize that the extractor's strategy (trespass distance) is chosen offline (in advance), based on the expected returns of each possible strategy. Note that the extractor acquires no new information online that can affect his decision-making: the strategy consists entirely of a distance by which to *attempt* to trespass; once caught, there is no further choice.

## 3. 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 some negligible width $w$) concentric with the forest.

Boundary patrols can be superior to homogenous patrols, since homogeneous patrols waste enforcement on the natural core [1]. It is interesting to note that this is not always so. Suppose the homogenous-induced core radius is less than $1-d$, $w$ is very small, and the trip length $d$ satisfies $w < 1/2 < d \leq 1$. 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 $w \to 0$. In this case, homogeneous patrols will actually outperform boundary patrols.

Intuitively, this is because a patrol in the interior will "intersect" more trips from center to boundary than a patrol on the boundary will. Unfortunately, both boundary and homogeneous patrols can perform arbitrarily badly.

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

PROOF (SKETCH). To see this, hold the budget fixed, and consider extremely large forests and cost and benefit functions yielding an empty natural core. The relationship between the cost/benefit functions and the budget be that an optimal patrol allocation will place patrols near to the forest center, halting the extractors at some distance $r_o$ from the center, but the significant dispersions of patrols due to either boundary or homogenous allocations would mean failing to stop the extractors prior to the forest center, resulting in an approximation factor of 0. $\square$

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,[3] both concentric with the forest.

The net cumulative benefit of walking distance $x$ is $B_o(x) - C(x) = B(x) - \Phi(x)B(x) - C(x)$, where $\Phi(x)$ is the capture probability for this walk. Let $\phi(x) = d\Phi(x)/dx$ be the probability density function of the capture probability, which is proportional to patrol density. Then the probability density function corresponding to $B_o(x) - C(x)$ will be

$$d(B_o(x) - C(x))/dx = dB(x)/dx - d\Phi(x)B(x)/dx - dC(x)/dx$$
$$= (1 - \Phi(x)) \cdot b(x) - \phi(x)B(x) - c(x) \quad (1)$$

Let band $[d_o, e]$ (with $0 \leq d_o \leq e \leq d_n$) be the patrol zone chosen by Algorithm 1.

---

**Algorithm 1** Computing the optimal allocation$(b, c, E, \epsilon)$

---

1: $(d_1, d_2) \leftarrow (0, d_n)$
   *binary search:*
2: **while** $d_1 < d_2 - \epsilon/3$ or $\phi_2$ not set **do**
3:    $d \leftarrow (d_1 + d_2)/2$
4:    $.5\phi(x) \triangleq \frac{b(x)-c(x)}{B(x)} - \frac{b(x)}{B^2(x)}(B(x)-C(x)-(B(d)-C(d)))$
5:    $e \leftarrow x$ s.t. $d \leq x \leq d_n$ and $\phi(x) = 0$
6:    $cost = \int_d^e 2\pi(1-x)\phi(x)dx$
7:    $\{d_2 \leftarrow d, \quad \phi_2 \leftarrow \phi\}$ **if** $cost \leq E$ **else** $d_1 \leftarrow d$
8: **end while**
9: **return** $(d_2, \phi_2)$

---

LEMMA 1. *Without loss of generality, the optimal density $\phi(x)$ at each point $x \in [d_o, e)$ can be assumed to be the smallest possible value disincentivizing further walking from $x$, i.e., that density yielding $b_o(x) = c(x)$. Moreover, $b_o(x) < c(x)$ and $\phi(x) = 0$ for $x > e$.*

PROOF. Consider a function $\phi(\cdot)$ that successfully stops the extractor at some location $d_o$ but which violates the stated property, at

---

[3]Generalizable to other forest shapes, as discussed below.

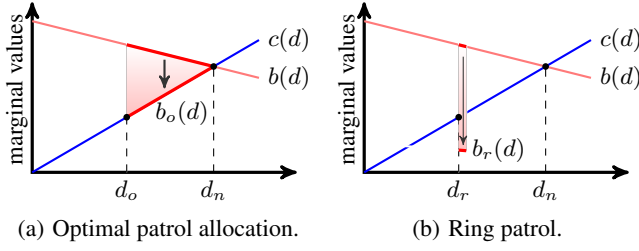(a) Optimal patrol allocation.   (b) Ring patrol.

**Figure 3: The shaded regions correspond to the reduction in marginal benefits *within the patrol zone*. Not shown are the (less dramatic) effects on $b(\cdot)$ *following* the patrol zone, due to the cumulative capture probability.**

some particular level of discretization. That is, partition the interval $[d_o, d_n]$ into $n$ equal sized subintervals, numbered $d_1, ..., d_n$. For this discretization, we write $B(d_i) = \sum_{j=1}^{i-1} b(j)$ and $\Phi(d_i) = \sum_{j=1}^{i-1} \phi(i)$ (omitting the coefficients). Let $d_i$ be the first such subinterval for which $b_o(d_i) < c(x)$, and let $d_i^+$ be shorthand for $d_i + 1$. In this case (see Eq. 1) we have $(1 - \Phi(d_i))b(d_i) - \phi(d_i)B(d_i) - c(d_i) < 0$. We correct this by subtracting a value $\delta$ from $\phi(d_i)$ to bring about equality, and adding $\delta$ to $\phi(d_i^+)$.

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

$$b(d_i^+)(1 - (\Phi(d_i^+) - \delta)) - (\phi(d_i^+) + \delta) \cdot B(d_i^+)$$
$$= b(d_i^+)(1 - \Phi(d_i^+)) + b(d_i^+)\delta - \phi(d_i^+)B(d_i^+) - \delta B(d_i^+)$$
$$= (b(d_i^+)(1 - \Phi(d_i^+)) - \phi(d_i^+)B(d_i^+)) + (b(d_i^+)\delta - \delta B(d_i^+))$$
$$= b_o(d_i^+) - \delta B(d_i)$$

Thus, since these two changes cancel out and there was no incentive for walking from $d_i$ past $d_i + 2$ 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 $(d_i, d_i^+)$, 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 total cost of the patrol. $b_o(x) < c(x)$ and $\phi(x) = 0$ for $x > e$ follows from $\phi(\cdot)$ being a band that stops the extractor at position $d_o$. $\square$

LEMMA 2. *Without loss of generality, we may assume $d_o$ kisses the outer edge of the patrol region.*

PROOF. Clearly $d_o$ will not be *prior to* the start of the patrol region. If $d_o$ lay after the beginning of the patrol region, then, by Lemma 1, the solution would have its cost only lowered by shifting the earlier patrol density past $d_o$. $\square$

Under the varying patrol density regime, the optimal patrol allocation can be computed (numerically). We remark that under the resulting patrol 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 net marginal benefit $(b(x) - c(x))$ that we need to compensate for by threat of confiscation, and yet the *magnitude* of the potential confiscation $(B(x))$ grows only larger.

THEOREM 1. *Algorithm 1 produces a near-optimal allocation (i.e., with arbitrarily small error).*

PROOF. We assume the properties stated by Lemma 1. Let $d_o$ indeed be the optimal trespass distance. Observe that for $x < d_o$, $b_o(x) = b(x)$; for $x > e$, $b_o(x)$ is determined only by $b(x)$ and the cumulative capture probability, i.e., $b_o(x) = (1 - \Phi(x)) \cdot b(x)$. $e$ is the point at which $\phi(x) = 0$ and $(1 - \Phi(x)) \cdot b(x) - c(x) = 0$. Now we compute $\phi(\cdot)$. Setting Eq. 1 to 0 yields:

$$\phi(x) = \frac{(1 - \Phi(x)) \cdot b(x) - c(x)}{B(x)} \qquad (2)$$

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

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

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

$$\int Q(x) \cdot e^{\int P(x)dx}dx = \int \frac{b(x) - c(x)}{B(x)} \cdot B(x)dx$$
$$= \int (b(x) - c(x))dx = B(x) - C(x)$$

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

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

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

$$\Phi(x) = \frac{B(x) - C(x) - (B(d_o) - C(d_o))}{B(x)}$$
$$\phi(x) = \frac{b(x) - c(x)}{B(x)} - \frac{b(x)}{B^2(x)}(B(x) - C(x) - (B(d_o) - C(d_o)))$$

Then the optimal allocation for any given budget $E$ will equal $\phi(x)$ for $x \in [d_o, d_n]$. The total cost of this is $E(d_o) = \int_{d_o}^{d*} 2\pi(1 - x)\phi(x)dx$. If $b(x)$ and $c(x)$ are polynomial functions, then $\phi(x)$ is a rational function, and so $E(d_o)$ is solvable analytically, by the method of partial fractions. In this case, we can evaluate $E(d_o)$ in constant time (for fixed $b(x)$ and $c(x)$) in a real-number computation model. Alternatively, $E(d_o)$ can be approximated within additive error $\epsilon$ in time $O(1/\epsilon)$, using standard numerical integration methods.

We can compute the smallest $d_o$ for which $E(d_o) \leq E$ by binary search. ($e$ is also found by binary search, within error $\frac{\epsilon}{3}\frac{1}{2\pi\phi(0)}$, which is a constant; recall that $\phi(x)$ is a decreasing function.) This yields a total running time of either $O(\log^2 1/\epsilon)$ or $O(1/\epsilon \log 1/\epsilon)$, depending on whether $E(d_o)$ is solved analytically or approximated. $\square$
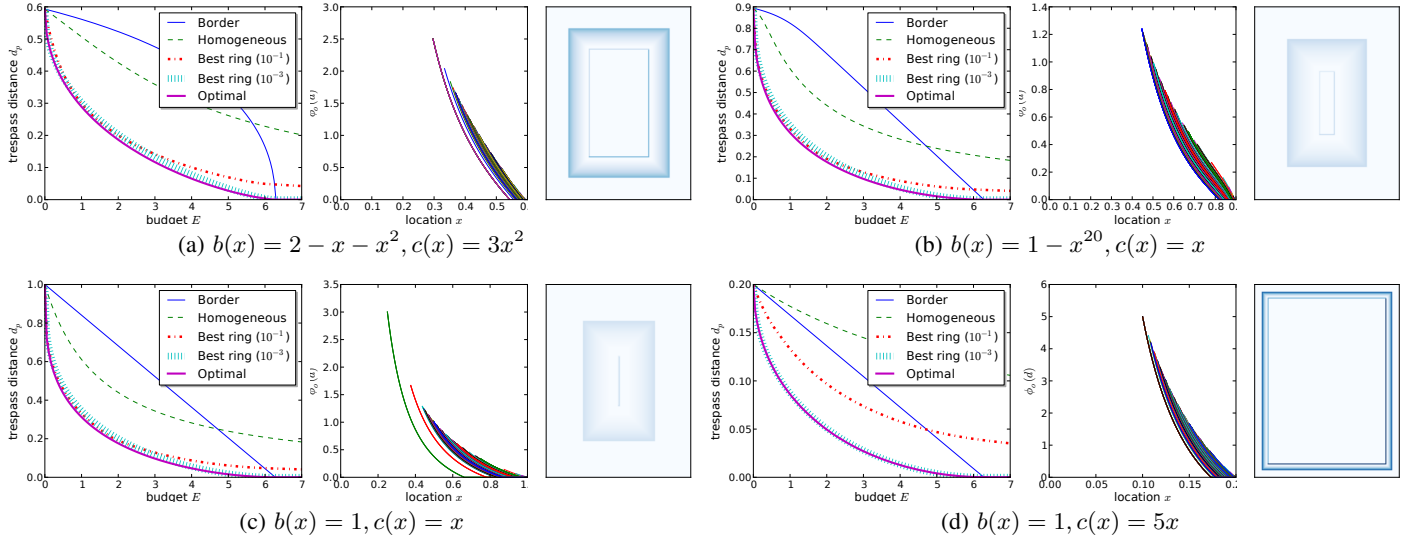
Figure 4: Patrol strategy effectiveness for sample $b(\cdot), c(\cdot)$ functions.

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

---

**Algorithm 2** Computing the best ring patrol$(b, c, E, \epsilon)$

---

1: $(d_1, d_2) \leftarrow (0, d_n)$
   *binary search:*
2: **while** $d_1 < d_2 - \epsilon$ or $\phi_2$ not set **do**
3: $\quad d \leftarrow (d_1 + d_2)/2$
4: $\quad \phi(d) \leftarrow E/(2\pi \cdot ((1 - d) - w/2) \cdot w)$
5: $\quad \Phi \leftarrow \phi \cdot w$
6: $\quad e \leftarrow x$ s.t. $(1 - \Phi)b(x) = c(x)$
7: $\quad pos \leftarrow \int_d^e (1 - \Phi)b(x) - c(x)dx$
8: $\quad neg \leftarrow \Phi \cdot B(d)$
9: $\quad \{d_2 \leftarrow d, \ \phi_2 \leftarrow \phi\}$ **if** $neg \geq pos$ **else** $d_1 \leftarrow d$
10: **end while**
11: **return** $(d_2, \phi_2)$

---

THEOREM 2. *Algorithm 2 produces a near-optimal ring patrol (i.e., within additive error at most $\epsilon$).*

PROOF. For a candidate trespass distance $d$, allocating the budget $E$ to a width-$w$ ring (where $w$ is negligible) of radius $d$ yields $\phi = E/(2\pi \cdot ((1 - d) - w/2) \cdot w)$ and $\Phi = \phi \cdot w$. In order to discourage the extractor from passing point $d$, it must be the case that the expected cost of doing so (the potential loss due to confiscated, weighted by probability: $\Phi \cdot B(d)$) exceeds the expected benefit (the remaining net benefit, weighted by probability: $(1 - \Phi) \int_d^{d_n} b(x) - c(x)dx$. We do binary search for the smallest such value $d$. □

THEOREM 3. *Algorithm 2 provides a 1/2-approximation, both in trespass distance reduction and pristine radius.*

PROOF. Let $r_n = 1 - d_n$ be the radius of the natural core. Let $r_o = 1 - d_o$ be the pristine area radius under the optimal patrol allocation $\phi_o(\cdot)$. We know that $\phi_o(\cdot)$ will be nonzero over the range $[d_o, d_n]$. Consider locations $x$ within this range. As $x$ grows from $d_o$ to $d_n$, the marginal benefit $b(x)$ falls monotonically while $c(x)$ grows, and the cumulative benefit $B(x)$ and cumulative capture probability $\Phi(x)$ both grow monotonically. Thus by Eq. 2, $\phi(x)$ falls monotonically over $[d_o, d_n]$.

Now consider the radius $r_r = (r_o + r_n)/2$ and the corresponding location $d_r = 1 - r_r$, which divides the range $[d_o, d_n]$ into two halves. Because $\phi(x)$ is monotonic decreasing, $\phi(x)$ has at least as much total mass in the first half than in the second, i.e., $\int_{d_o}^{d_r} \phi(x)dx \geq \int_{d_r}^{d_n} \phi(x)dx$. Because the total cost of patrol density $\phi(x)$ at location $x$, rotated about the entire circle, is $2\pi\phi(x)$, "flattening" $\phi(x)$ over the range $[d_o, d_n]$ (i.e., setting it equal to $\frac{1}{d_n - d_o} \int_{d_o}^{d_n} \phi(x)dx$) will only lower the total cost. (Though doing so will sacrifice the guarantee of trespass distance $d_o$.) Then "compressing" this total probability mass $\int_{d_o}^{d_n} \phi(x)dx$ from the range $[d_o, d_n]$ to the point $d_r$ will not change the cost any further, since the mean circle circumference for radii in $[r_o, r_n]$ is $2\pi(r_o + r_n)/2$, which is the same as that for radius $r_r$.

We now claim that the constructed negligible-width ring patrol at $d_r$ will deter the extractors from crossing it, by accounting for the two "halves" of $\phi_o(x)$. First, the "left" half of $\phi_o(x)$ transferred to $d_r$ will yield a cumulative detect probability of $\Phi_o(d_r)$, just as under the optimal patrol. Second, the "right" half of $\phi_o(x)$ will inflict the same total reduction in net benefits for the action of traversing $[d_r, d_n]$ as the optimal patrol does. After passing $d_r$, each additional step would provide a positive net marginal benefit, until regaining the pre-$d_r$ cumulative net benefit only at point $d_n$, after which all net marginal benefits are negative. Thus every stopping point after $d_r$ will have cumulative net benefit lower than this value immediately before $d_r$.

We have constructed a ring patrol allocation that reduces the trespass distance by at least half the optimal such value, i.e., $(r_n - r_o)/2$, yielding pristine radius $r_r = (r_o + r_n)/2 \geq r_o/2$, and so

the result follows. □

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, $\Phi_i(x)$ 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.

### 3.1 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 $E_i$ (e.g., $1/k$) 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 below, with budget $E$. Then we position patrol unit $i$ at a location within sector $i$, chosen according to the computed $\phi(\cdot)$.

**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 [17], 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.

### 4. 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. 3). We now summarize our observations on these results.

In each setting (see left subfigures), we vary the patrol budget, computing the patrol allocation function and hence the extractor's trespass distance $d_p$, 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 $(1/10^3)$ 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.

### 5. CONCLUSION

In this paper, we have presented a Stackelberg security game setting that differs significantly from those previously considered in the AI literature, which necessitates the use of very different techniques from those used in the past. At the same time, this work opens up an exciting new area of research for AI at the intersection of forest economics and game theory. Eventually, as with counterterrorism Stackelberg games studied in the literature, we aim to deploy our solutions in real-world settings. Potential sites for such deployments include Tanzania's aforementioned Kibaha Ruvu Forest Reserves and the mangrove forests of Mnazi Bay Ruvuma Estuary Marine Park.

### 6. ACKNOWLEDGEMENT

### 7. REFERENCES

[1] H. J. Albers. Spatial modeling of extraction and enforcement in developing country protected areas. *Resource and Energy Economics*, 32:165–179, 2010.

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

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

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

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

[6] H. Hofer, K. L. I. Campbell, M. L. East, and S. A. Huish. 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, 2000.

[7] D. J. Lober. 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, 1992.

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

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

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

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

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

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

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

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

[16] A. R. E. Sinclair and P. Arcese. *Serengeti II: Dynamics, Management, and Conservation of an Ecosystem*. University of Chicago Press, Chicago, 1995.

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

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

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