

Challenges in Patrolling to Maximize Pristine Forest Area (Position Paper)

Matthew P. Johnson and Fei Fang and Rong Yang and Milind Tambe

University of Southern California, Los Angeles, CA 90089, USA
{matthepj, feifang, yangrong, 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 through the forest that seek to deter such activity by decreasing its profitability. With limited resources for performing such patrols, a patrol strategy will seek to distribute the patrols throughout the forest, in space and time, in order to minimize the resulting amount of extraction that occurs or maximize the degree of forest protection, according to one of several potential metrics. We pose this problem as a Stackelberg game. We adopt and extend the simple, geometrically elegant model of (Robinson 2010). First, we study optimal allocations of patrol density under generalizations of this model, relaxing several of its assumptions. Second, we pose the problem of generating actual schedules whose site visit frequencies are consistent with the analytically computed optimal patrol densities.

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 (MacKinnon et al. 1986; Dixon and Sherman 1990; Clarke, Reed, and Shrestha 1993; Robinson 2008). To cite just two examples, Tanzania’s Kibaha Ruvu North and South 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; Sinclair and Arcese 1995). With limited resources for performing such patrols, a patrol strategy will seek to distribute the patrols throughout the forest, in space and time, in order to minimize the resulting amount of extraction that occurs or



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).

maximize the degree of forest protection, according to one of several potential metrics.

Although in some cases patrols may catch would-be extractors as they are about to perform the extraction, and therefore will directly prevent extraction from occurring, we primarily focus on the extraction that the patrols *deter*. That is, 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 *minimize* the amount of extraction that will yield a positive net return in his cost-benefit analysis. (To cite two real-life examples of this phenomenon, when the New York City Police Department issued zero-tolerance policies on “squeegeeing” and subway fare-beating, arrests for these crimes reportedly went down rather than up (Kleiman and Kilmer 2009; Kleiman 2009).) Therefore we pose the 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.

Economists have studied the relationship generally between enforcement policy for protecting natural resources and the resulting incentives for neighbors of the PA (Milliman 1986; Robinson 2008; Sanchirico and Wilen 2001). A number of models and problem formulations of forestry protection specifically have been proposed, including dynamic models (Robinson, Albers, and Williams 2008), forest shape modification to increase protection (Robinson, Albers, and Williams 2011), and combinations of forest protection with local economic activity such as bee-keeping (Albers and Robinson December 2011). Our point of departure in this paper is the simple, geometrically elegant forest protection model of (Robinson 2010), 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 (Skonhoft and Solstad 1996; Albers and Grinspoon 1997; MacDonald, Adamowicz, and Luckert 1998; Hofer et al. 2000)) maximize the size of the resulting *pristine* forestland. We probe the probabilistic assumptions of this model and consider a number number of more realistic extensions to it. One example is permitting spatial variation in patrol density. As has been observed (Robinson 2010), exogenous restrictions on patrol strategies, whether adopted for simplicity or to comply with legal restrictions, can degrade protection performance (MacKinnon et al. 1986; Hall and Rodgers 1992). We extend what is known analytically about the core model to these new settings and provide new results. We also pose the problem of computing actual patrol paths consistent with the desired densities.

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 (Paruchuri et al. 2008), the Federal Air Marshals Service (Tsai et al. 2009), and the Transportation Security Administration (Pita et al. 2011). See (Tambe 2011) for an overview.

The problem setting we address here differs from those considered in these previous works in several crucial ways. The biggest difference is that this 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 or segments to protect (e.g., modeled as nodes of a graph), whereas in the our setting ideally the entire forest area would be protected from extraction. As such, our primary focus is, at least initially, on the choice of distribution for patrol density over the two-dimensional forest region, i.e. a probability distribution from which to select patrols. (Of course, the continuous space can




Figure 2: The forest, with the pristine area shaded.

be discretized by overlaying a grid on it, but any choice of grid granularity will involve some level of approximation and error.) Second, the followers make choose a distance based on a cost-benefit analysis, considering both their expected returns and their labor costs based on distance.

The rest of the paper is organized as follows. After reviewing the model of (Robinson 2010), including several proposed patrol density strategies, we will consider several extensions and generalizations to the model, including relaxing assumptions on both the problem setting and the permitted patrol strategies. Second, we consider the problem of computing actual patrol strategies consistent with the computed patrol densities, i.e., computing (sets of) walks within the region whose frequencies of visiting locations correspond (approximately) to those locations' patrol densities. We conclude with directions for ongoing and future work.

Core models

In this section we review and extend the model of (Robinson 2010). The forest is a circular region of radius 1, with villagers uniformly distributed about its perimeter (see Fig. 2). A villager's action is to choose some distance d to walk on a line from the perimeter towards the forest center before returning his starting point, extracting on the return trip. Note that in this setting, due to symmetries and the fact that villagers' decisions are uncoordinated, the problem is essentially one-dimensional; we may assume that the extractor extracts only on the return trip from the forest. Villagers incur a cost and gain a benefit *if not caught*, based on a convex increasing cost function $C(d)$ and a concave increasing benefit function $B(d)$. If caught, the villager'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 forest³). Thus

³Observe that since extraction occurs only on the return trip and given the nature of the punishment, we may restrict our attention to detection on the return trip.




Figure 3: Shaping the extractor’s cost-benefit analysis.

a given patrol strategy will reduce the extractor’s expected benefit for an incursion of distance d from $B(d)$ to some value $B'(d)$.

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 (Robinson 2010); that is, the optimal value d^* 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 $B'(d)$, thus reducing the extractor’s optimal incursion distance (see Fig. 3).

In the simplest (linear) cost setting, the leader has a budget E specifying (perhaps after application of a suitable normalization constant) 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^* .

Detection probability models

Let Φ be the detection probability chosen by the leader, constant over some subregion R of the forest. In the model of (Robinson 2010) (a version of which we present in this paragraph, motivated slightly differently), the detection probability for a walk of distance d is equal to $d\Phi$. This setting is best understood as one operating under a time model in which 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. We assume that an extractor is detected if he comes within some distance Δ of the patrol. Then indeed the probability of detection for an extraction path of length d (when there is a single patrol unit) will be proportional to Φd , specifically $\Phi d 2\Delta / |R|$, where $|R|$ indicates the area of R and the total area within distance Δ of the length- d walk is approximated as $d \cdot 2\Delta$. Intuitively, the more steps the extractor takes undetected, the higher the probability he will be detected on the next step. (Here we assume the patrol unit is *not* visible to the extractor.) Suppose the available patrol budget de-

termines the sensing range Δ ; this is equivalent to it determining the detection probability Φ .

Alternatively, we may consider settings in which Δ is constant and the budget can be spent on multiple patrol units. If there are some number of units u , then detection probability Φ is the joint result of them together, i.e., $\Phi = 1 - (1 - \Phi_u)^u$, where Φ_u is the detection probability due to each particular unit. Then the probability of capture from a length- d extraction path becomes $1 - (1 - \Phi_u d 2\Delta / |R|)^u$.

An alternate time model is one in which the extractor moves much less quickly than the patrol units, and so from the extractor’s point of view the patrol units are newly assigned random locations (with replacement) at each moment. If we approximate a walk of length d as consisting of $n = d/\Delta$ separate “trials”, then the detection probability will be $1 - (1 - \Phi)^n$, regardless of whether there is one patrol unit or many.

Patrol allocations

Three patrol allocation strategies have previously been considered for the linear cost model:

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

It is observed in (Robinson 2010) that boundary patrols will often be superior to homogenous patrols, since homogeneous patrols waste enforcement on the natural core. It is interesting to note that this need not always be the case. Consider the variable Δ setting, and let the budget be E ; assume 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, the detection probability will be simply $E/\pi \cdot d$. With boundary patrols, however, the detection probability 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 a patrol in the interior will “intersect” more trips from boundary to center than a patrol on the boundary will.

It would be interesting to classify the situations under which one patrol allocation is superior to another, for the different probability settings. The best ring strategy can only be better than boundary patrols, since the latter is a special case of the former.

Superlinear cost

A second cost function considered by (Robinson 2010) is a superlinear cost in which achieving detection probability Φ at a point has cost proportional to Φ^γ for some $\gamma > 1$. This cost model leads to *band patrols*, in which patrol density is allocated *evenly* in a band (typically of

width greater than w above) demarcated by two concentric circles, consistent with a constraint that patrols be allocated uniformly over the patrolled region.

Absent that constraint, an optimal distribution of patrol mass can be computed, in which the level of patrol density will vary smoothly throughout the forest.

Other extensions

A number of additional potential extensions suggest themselves.

- **Asymmetry:** The core problem setting is *highly* symmetrical; the symmetry assumptions could be relaxed by degrees. We could allow the weight of villager populations to vary around the forest perimeter, for example. More dramatically, we could allow other, noncircular forest shapes.
- **Nonuniform importance:** Currently, all the forest area is equally important, both to the extractor and to the protector. Each of these assumptions could be relaxed, independently.
- **Continuity:** As discussed above, an essential feature of this problem setting is the continuity of space. Without subverting this, we could introduce obstacles to movement within the forest. With enough obstacles, however, the forest may become “sparse” enough to be representable by a graph.

Computing patrols

On whatever particular problem setting model is chosen above, once a patrol distribution is chosen, the ultimate task is to choose patrols in such a way that the frequencies with which they visit locations within the forest are consistent with the specified probability mass values of those locations. This algorithmic problem, which can be seen as sampling from a distribution or as a form of rounding, is the primary open problem to solve here.

Conclusion

In this position paper, we have presented a Stackelberg security game setting that differs significantly from those settings previously considered in the AI literature. In ongoing work, we plan to explore more realistic problem settings and tackle the problem of generating actual patrol schedules consistent with the optimal patrol probability density distributions.

Eventually, we aim to deploy the resulting patrol distributions and schedules in real-world settings. In the past, one of us (H.J. A.) has worked with the forest managers and guards at Tanzania’s aforementioned Kibaha Ruvu North and South Forest Reserves, which could provide one potential site for future research transition, as well as with the marine park managers of Mnazi Bay Ruvuma Estuary Marine Park. It may be possible to implement these models in the mangrove forests within that park, which exhibits a pattern typical among the forests prompting this work: the mangroves there are a source of wood and fuelwood/charcoal for

local people but such use of them decreases the forest’s ability to trap sediment, to clean water, and to provide a habitat to juvenile fish.

Acknowledgement

This research is supported by MURI grant W911NF-11-1- 0332.

References

- Albers, H. J., and Grinspoon, E. 1997. A comparison of the enforcement of access restrictions between xishuangbanna nature reserve (china) and khao yai national park (thailand). *Environmental Conservation* 24(4):351–362.
- Albers, H. J., and Robinson, E. J. Z. December, 2011. The trees and the bees: Using enforcement and income projects to protect forests and rural livelihoods through spatial joint production. *Agricultural and Resource Economics Review*.
- 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):151177.
- Kleiman, M., and Kilmer, B. 2009. The dynamics of deterrence. *Proceedings of the National Academy of Sciences of the United States of America* 106(34):14230–14235.
- Kleiman, M. 2009. *When Brute Force Fails: Strategic Thinking for Crime Control*. Princeton, NJ: Princeton University Press.
- 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.
- MacDonald, D. H.; Adamowicz, W.; and Luckert, M. 1998. A microeconomic analysis of choice of fuelwood collection sites in zimbabwe: valuation through behavior and caloric expenditure. In *World Congress of Environmental and Resource Economists*.
- MacKinnon, J.; MacKinnon, K.; Child, G.; and Thorsell, J. 1986. *Managing Protected Areas in the Tropics*. Gland, Switzerland: IUCN.

- Milliman, S. R. 1986. Optimal shery 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. Indias disappearing common lands: fuzzy boundaries, encroachment, and evolving property rights. *Land Economics* 84(3):409–422.
- Robinson, E. J. Z. 2010. Spatial modeling of extraction and enforcement in developing country protected areas. *Resource and Energy Economics* 32:165–179.
- 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.
- Sinclair, A. R. E., and Arcese, P. 1995. *Serengeti II: Dynamics, Management, and Conservation of an Ecosystem*. Chicago: University of Chicago Press.
- Skonhofs, A., and Solstad, J. T. 1996. Wildlife management, illegal hunting and conicts. a bioeconomic analysis. *Environment and Development Economics* 1(2):165–181.
- Tambe, M. 2011. *Security and Game Theory: Algorithms, Deployed Systems, Lessons Learned*. Cambridge University Press.
- 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)*.