

Modeling Environmental Crime in Protected Areas Using the Level Set Method

D. J. Arnold* D. Fernandez* R. Jia* C. Parkinson* D. Tonne[†] Y. Yaniv[‡]
 A. L. Bertozzi* S. J. Osher*

August 10, 2018

Abstract

National parks often serve as hotspots for environmental crime such as illegal deforestation and animal poaching. Previous attempts to model environmental crime were either discrete and network-based or required very restrictive assumptions on the geometry of the protected region and made heavy use of radial symmetry. We formulate a level set method to track criminals inside a protected region which uses real elevation data to determine speed of travel, does not require any assumptions of symmetry, and can be applied to regions of arbitrary shape. In doing so, we design a Hamilton-Jacobi equation to describe movement of criminals while also incorporating the effects of patrollers who attempt to deter the crime. We discuss the numerical schemes that we use to solve this Hamilton-Jacobi equation. Finally, we apply our method to Yosemite National Park and Kangaroo Island, Australia and design practical patrol strategies with the goal of minimizing the area that is affected by criminal activity.

1 Introduction

Environmental crime in protected national parks is a concern to authorities around the world. National parks often serve as hotspots for illegal logging and animal poaching, henceforth referred to as illegal extraction. In recent years, scientists have aided law enforcement agencies in the prevention of such crimes. One way in which they do this is by building models to describe deforestation, track animal movement, and predict adversarial behavior of criminals. Leader-Williams and Milner-Gulland [12] examined a case study of poaching in the Luangwa Valley, Zambia and concluded that increasing detection rates was a more effective deterrent to environmental crime than increasing severity of punishment. This means that models that can help improve detection rates of environmental criminals are a useful tool for patrollers of protected areas.

Authorities employ different strategies to patrol protected areas and detect illegal extractors, and research into the effectiveness of these strategies is of major interest. Researchers are challenged to find efficient and practical patrol strategies to police vast areas with limited resources. An important model of illegal extraction was given by Albers [1]. They describe illegal extraction as a continuous spatial game between patrol units and extractors, constructing a model where extractors maximize their expected profit by trading off costs of penetrating further into a protected region (increasing effort required and risk of capture) against increased benefits of extracting further from the boundary of the protected region. Albers' model was formulated as a Stackelberg game, an adversarial game with a defender (patrollers) and many attackers (extractors) with perfect information about the defender's strategy. Albers gave some qualitative results in a very simplified situation where the protected region was circular, and all quantities were radially symmetric. Johnson et al. [9] proved an additional result regarding optimal patrol strategies for this simple case.

Subsequently, models for illegal extraction in protected regions using discrete rather than continuous methods have been developed by Fang et al. [6, 7] and Kar et al. [10, 11]. These models incorporate different methods for modeling human behavior and different ways to treat the protected region. Kar et al.

*Department of Mathematics, University of California Los Angeles.

[†]Department of Mathematics, California State University Long Beach.

[‡]Department of Mathematics, University of Maryland.

[10] used repeated Stackelberg games to attempt to understand the evolution of attackers strategies, but did not consider realistic spatial domains. Fang et al. [7] developed a model including detailed terrain and spatial information by describing the protected region as a series of nodes connected by edges corresponding to natural pathways through the region (for example along rivers or walking trails). The PAWS algorithm developed by Fang et al. [7] has been deployed in Queen Elizabeth National Park (QENP) in Uganda, and in Malaysian forests. Taking another approach, Kar et al. [11] used machine learning techniques to construct a model for predicting poaching attacks, again deploying the model for a field test in QENP, Uganda.

Our goal is to generalize the model of Albers [1] so that it is applicable to realistic protected regions, can include the effects of terrain and geometry, and can be executed for real protected regions. The most important mathematical technique that enables this extension is the level set method of Osher and Sethian [15], which provides a simple way to track moving fronts through the protected region. Indeed, the level set method and Hamilton-Jacobi equations are now used extensively in modeling travel in a variety of contexts. Sethian and Vladimirsky [19, 20] discuss a level set method for optimal travel on manifolds. In application, Dubins [5], and later Takei et al. [21], modeled movement of simple cars through terrain with obstacles using methods from control theory which include Hamilton-Jacobi-Bellman equations. Tomlin has used similar methods to determine reachable sets for aircraft autolandings [2], and to model human movement in adversarial reach-avoid games [24]. Thus there is a precedent for using Hamilton-Jacobi type equations to model movement throughout domains. However, as discussed above, attempts to model environmental crime have either required overly restrictive assumptions regarding radially symmetry and geometry of the protected region, or have been discrete in nature. This work serves to fill a gap in the literature. We suggest a continuum model for environmental crime which is still able to account for realistic terrain information, and bridges the divide between environmental crime modeling and Hamilton-Jacobi formulations for modeling human movement.

This paper is structured as follows. In section 2, we describe the model of Albers [1]. In section 3, we introduce the level set method, describe the numerical methods used to solve it, and show how it can be applied to model movement through regions with terrain. In section 4, we describe our model for illegal extraction from protected regions, and provide an algorithm to calculate the regions through which extractors will pass. In section 5, we present results and discuss their implications towards both patrol strategies and our modeling approach. Finally, in section 6, we make conclusions and identify potential directions for further research.

2 The Illegal Logging Model of Albers [1]

Albers [1] presented a model for illegal extraction from protected regions that includes spatial effects. They use a Stackelberg game in a circular region between defenders and extractors. In their model, an extractor gains benefit and incurs cost based on the depth d into the region where they choose to extract. The benefit is associated with the extraction point, and is modeled as a concave increasing function $B(d)$. Cost is based on the trip distance, and is modeled as a quasiconvex increasing function $C(d)$, but no specific functions are given. Then the extractor's profit associated with extracting at depth d is

$$P(d) = B(d) - C(d). \quad (1)$$

If the region is patrolled with density $\psi(d)$, the probability that the extractor is detected is the cumulative probability $\Psi(d) = \int_0^d \psi(r)dr$ over the trip out of the protected region. If the illegal extractor is caught, any resources they've extracted are confiscated and they must leave empty-handed. The extractor's expected profit can therefore be written as

$$P(d) = (1 - \Psi(d))B(d) - C(d). \quad (2)$$

Given the extractors have perfect information of the patrol strategy, they know the function $\Psi(d)$, and hence their task is to find the optimal distance d^* to penetrate into the protected region. Simultaneously, the patrollers task is to pick the patrol strategy $\psi(d)$ that minimizes d^* . Figure 1 illustrates the model, showing the pristine region where extractors are never present and the outer region through which they pass while travelling into or out of the protected region.

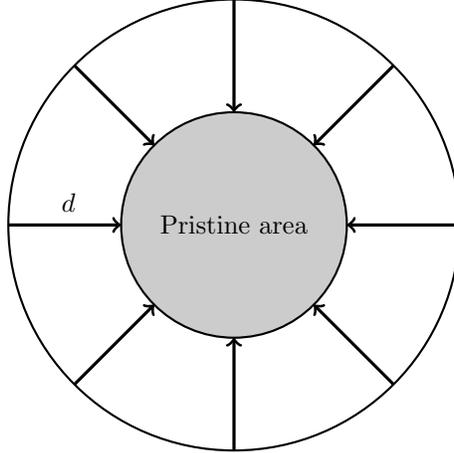


Figure 1: Illustration of the Albers [1] model, where extraction occurs at depth d in a circular region. Reproduction of Figure 2 from Johnson et al. [9].

3 The Level-Set Method

The level-set method is a powerful technique that we use to model human movement through complex domains. Originally developed by Osher and Sethian [15], the level set method models fronts which propagate with a prescribed velocity. The front is embedded as the zero level set of an auxiliary function, which is then evolved according to a partial differential equation. We describe a simple version of the method in \mathbb{R}^2 , sufficient for our needs. Suppose we are given a simple, closed curve (or a collection of non-intersecting, simple, closed curves) denoted by Γ . To evolve Γ via level set flow, we begin by finding a Lipschitz continuous function $\phi_0 : \mathbb{R}^2 \rightarrow \mathbb{R}$ such that ϕ_0 is positive inside Γ and negative outside Γ , and thus $\Gamma = \{x \in \mathbb{R}^2 : \phi_0(x) = 0\}$. Next, we evolve $\phi : \mathbb{R}^2 \times [0, \infty) \rightarrow \mathbb{R}$ using the PDE

$$\begin{aligned} \phi_t + v(x) |\nabla \phi| &= 0, \\ \phi(x, 0) &= \phi_0(x), \end{aligned} \tag{3}$$

for some non-negative velocity function v . Define $\Gamma(t) = \{x \in \mathbb{R}^2 : \phi(x, t) = 0\}$. This represents an evolution of $\Gamma = \Gamma(0)$. Indeed, we can re-write equation 3 as $\phi_t + v(x) \left(\frac{\nabla \phi}{|\nabla \phi|} \right) \cdot \nabla \phi = 0$. Note that $\nabla \phi / |\nabla \phi|$ is a unit vector which is normal to the level contour $\Gamma(t)$. Thus, locally, this equation models advection in the direction normal to $\Gamma(t)$ with velocity $v(x)$. This causes $\Gamma(t)$ to deform with normal velocity $v(x)$. With this in mind, one can view $\Gamma(t)$ as the set of points which could be reached after walking inward from the boundary for a time t with normal velocity $v(x)$. For example, if $v(x) \equiv 1$, then for $t > 0$, the curve $\Gamma(t)$ is the set of points a distance t from the original curve.

This procedure is shown in figure 2. Figure 2a shows the boundary of a region defined by $(x(\theta), y(\theta)) = (\cos(\theta), \sin(\theta) + \sin(3\theta)/2)$ and a series of contours of points accessible after traveling a certain time from the boundary. Figure 2b shows the graph of the level set function and with the same level sets overlaid on the surface. This example also displays one of the strengths of the level set method. Depending on the original curve, the the level contours may break apart into disconnected pieces or merge into into a single piece. More primitive methods for curve evolution (many of which involve tracking points along the curve) have trouble handling these changes in topology but the level set method accounts for them with no special considerations. For more information regarding level set methods, see Osher and Fedkiw [14].

4 Illegal Deforestation

Our model for the actions of illegal extractors committing crimes in protected regions is based on Albers [1], as described in section 2. They derived a game-theoretic model for illegal deforestation in an idealized

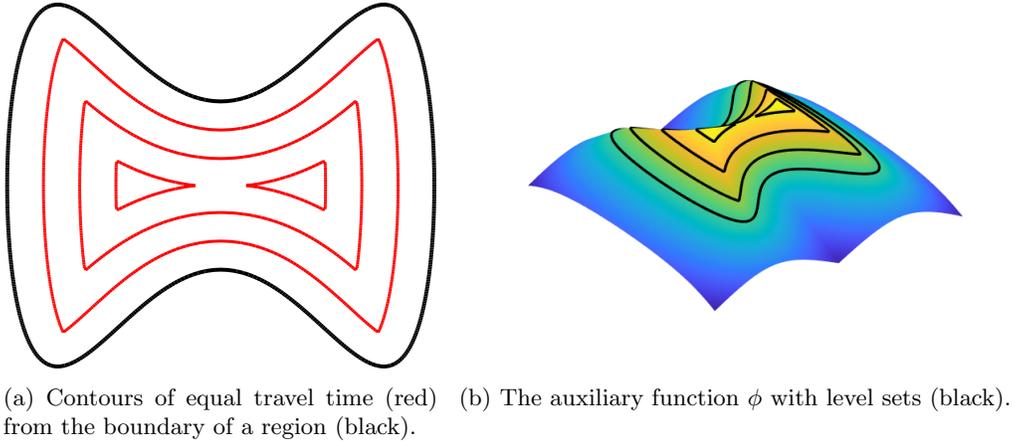


Figure 2: Using the level set method to find contours of equal travel time.

setting, assuming the protected region was circular to significantly simplify the problem, and neglecting the effects of terrain. We wish to relax these restrictions which leads to several significant difficulties.

We are interested in real national parks and hence the protected region will not be symmetric and the benefit and cost functions could be arbitrary. In Albers [1], the profit is maximized on a ring with some radius and extractors will enter from any point on the boundary of the region to travel to the maximum profit ring. Without radial symmetry however, the profit will, in general, be maximized at an isolated point rather than along a closed curve, and correspondingly there will be a single optimal route from this point to the boundary of the protected region. Hence illegal extractors only ever occupy a set of points of measure zero. In our model, we resolve this issue by assuming that extractors will tolerate extracting anywhere where the profit is close to the maximum possible profit.

The cost function of Albers [1] does not generalize easily to irregular geometries. Instead of using a predefined formula based on distance, we constructively determine the cost by considering the effects that would detract from the final profit with the help of the level set method. Given terrain and elevation information for a real protected park, the paths taken by extractors will not be simple straight lines. Therefore it is also necessary to find the paths extractors will follow to exit of the protected region. In the next section we describe an algorithm to find the pristine and non-pristine regions in a protected area with arbitrary shape, terrain, benefit distribution and patrol strategy.

4.1 Problem Description

Let Ω represent the area that needs to be protected and let $\psi(x, y)$ be the patrol density function measuring the likelihood of an extractor being caught at position (x, y) . The patrol budget E is defined as $E = \int_{\Omega} \psi(x, y) dx dy$. The budget measures how many patrolling resources can be allocated to protect the domain Ω . Effectively, the budget scales ψ . Extracting at a position (x, y) gives the extractor an amount of benefit $B(x, y)$. For example, the benefit could depend on the quantity, value, or species of trees. We make no assumptions about B except that it is known. $C(x, y)$ is the expected cost associated with extracting at position (x, y) . The cost is based on two factors, the effort involved in traveling from the extraction point to the boundary of the protected region, and the risk of being caught by patrols whilst inside the domain. The further into the protected the region the extractors penetrate, the higher the cost as they must expend more time/energy traveling, and are more likely to be captured by patrollers.

This highlights a difference between our model and Albers [1]: we include the effect of the patrol directly in the cost function, whereas they include it as a modification to the benefit. In our model, the profit an extractor expects from extracting at a point (x, y) in the protected region is simply $P(x, y) = B(x, y) - C(x, y)$. We assume the extractors accept a position to extract if the profit is within some tolerance of the maximum obtainable profit. Finally, we determine the paths that extractors will take when leaving the

protected area (where they can be captured by the patrol). We assume the extractors will always take an optimal path from the extraction point to the boundary, and can only be caught after they have finished extracting and started heading back to the boundary (since no crime has been committed before they have extracted).

4.2 Calculating the Profit Function

Since the benefit function is known in advance, it remains to calculate the expected cost function. As explained in the previous section, the cost function depends on many factors, including the optimal route to a given point. The level set method avoids most of the difficulties one might expect when calculating optimal paths from all points inside the protected region to the boundary. We use the level set method to find contours of equal cost in the protected region. The level set method implicitly uses the optimal path between the boundary and an interior point, and can be efficiently implemented.

The expected cost associated with extracting at a given point (x_0, y_0) is calculated with the following level set equation, using the boundary of the domain Ω as the zero-level set of the function $\phi(x, y, 0)$:

$$\frac{\partial \phi}{\partial t} = -\frac{1}{1/v(x, y) + \alpha\psi(x, y)B(x_0, y_0)} |\nabla \phi|. \quad (4)$$

The cost $C(x_0, y_0)$ is implicitly defined by $\phi(x_0, y_0, C(x_0, y_0)) = 0$. Intuitively, the point (x_0, y_0) lies somewhere inside the curve $\Gamma = \partial\Omega$. When we evolve the boundary using 4, the zero level-sets $\Gamma(t) = \{\phi(x, y, t) = 0\}$ represent the set of points which require equal cost to reach from the boundary. For some $C > 0$, we will have $(x_0, y_0) \in \Gamma(C)$. This C is the cost of extracting at (x_0, y_0) . This is demonstrated in figure 3.

The important term from a modeling perspective is the coefficient of $|\nabla \phi|$. The denominator of this term has two summands, $1/v(x, y)$ and $\alpha\psi(x, y)B(x_0, y_0)$. For the first summand, $v(x, y)$ is the walking speed at (x, y) which is determined by the terrain at (x, y) . It is calculated by using the steepness at (x, y) and a walking speed function adapted from [8] (discussed in section 4.2.1). Increasing v makes the level sets move faster, and hence decreases the cost to reach a given point. The second summand, $\alpha\psi(x, y)B(x_0, y_0)$, is a term that penalizes extractors for spending time in highly-patrolled areas. The patrol density function ψ measures the likelihood of an extractor getting caught at position (x, y) , and $B(x_0, y_0)$ is the benefit acquired at (x_0, y_0) . If the extractor gets caught, they lose all of the accrued benefit, so the impact of being caught should scale with the benefit gained. When $\psi(x, y)B(x_0, y_0)$ is larger, the level sets move more slowly, and the cost associated with extracting at a point increases.

To compare the summands we include a parameter α , which can be thought of as a risk aversion factor. This parameter weighs the willingness of an extractor to assume risk of being captured versus the willingness to expend time or energy. A large α signifies that the extractor will accept less risk of being captured, and will be willing to exert more effort to reach their chosen extraction point. Alternatively, a small α means the extractor is willing to accept a higher risk of being captured to decrease the physical effort required to travel in and out of the protected region. Note that this summand includes information about the extraction point, through the benefit $B(x_0, y_0)$. This means that the cost function must be calculated separately for each extraction point.

4.2.1 The walking speed equation

Tobler [23] gave a formula for walking speed based on slope. Recently, Irmischer and Clarke [8] gave an

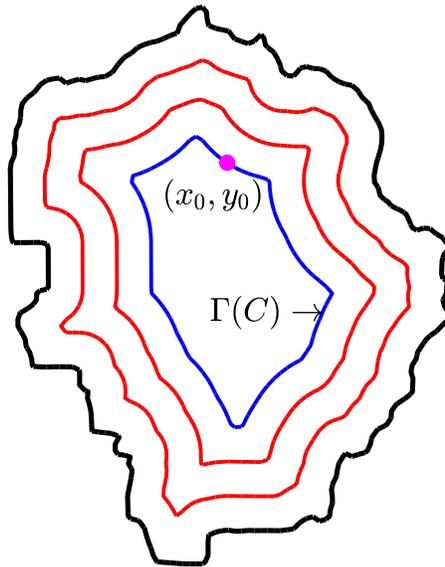


Figure 3: Evolve level sets in from the boundary until the $C > 0$ such that $(x_0, y_0) \in \Gamma(C)$.

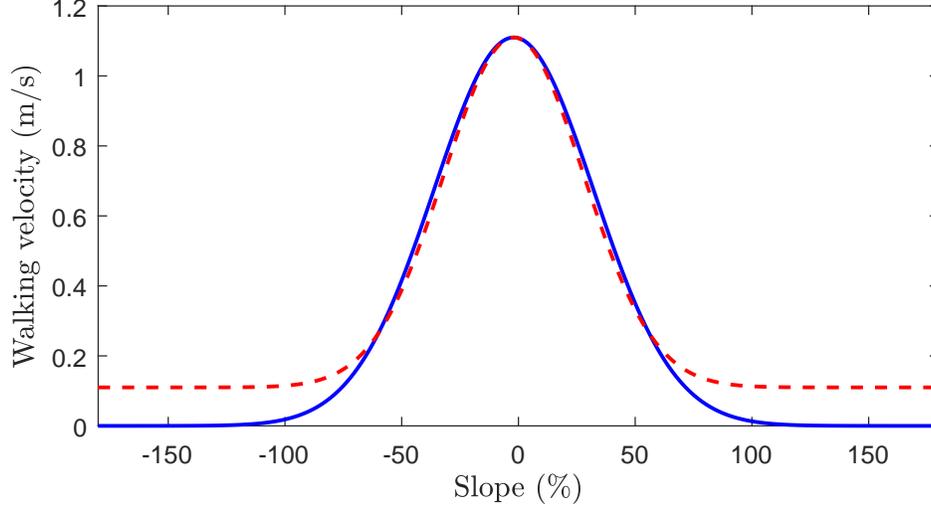


Figure 4: Our velocity function 6 (solid blue line), and the Irmischer and Clarke function 5 (dashed red line). Our velocity function decays to zero for high slopes while the function suggest by Irmischer and Clarke [8] does not.

improved formula that they claimed more accurately predicted travel times for humans walking on roads,

$$v = 0.11 + \exp\left(-\frac{(100s + 2)^2}{1800}\right), \quad (5)$$

where $100s$ is the grade in percent, and v is the corresponding speed in m/s. This formula agrees well with experimental results given in [8], but has a drawback for our model, that the speed does not vanish as the slope goes to infinity. We hence modify the given equation so that the speed drops to zero as the slope becomes more extreme and use

$$v = 1.11 \exp\left(-\frac{(100s + 2)^2}{2345}\right), \quad (6)$$

which matches the maximum speed, gives similar results for grades less than 60% and decays to zero for more extreme slopes as can be seen in figure 4.

We note that the specific walking velocity function is not crucial in our model, and other functions could be used. We also are able to incorporate areas where the walking speed is faster or slower which could model, for example, overgrown scrub or walking trails. It is also possible to set v to zero in places to model impassable obstacles.

4.3 The Expected Cost Algorithm

The expected profit function is calculated using the following algorithm.

1. Since the cost level set equation 4 depends on the particular benefit level where extraction takes place, we must solve it for all possible values of B . We pick N benefit values in $[\min_{\Omega} B, \max_{\Omega} B]$ and evaluate the cost function using each benefit value B_i , by performing steps 2-4 for each B_i .
2. For a given B_i , first find the contour(s) of points where $B(x, y) = B_i$. In figure 5a, the black contour is the boundary of the domain, and the blue contour is an equal benefit contour with the same benefit B_i .
3. Evolve the expected cost level set function with the benefit B_i using equation 4. In figure 5b, the red contours are the equal-cost level set contours. Each of the contours represents a different cost value.

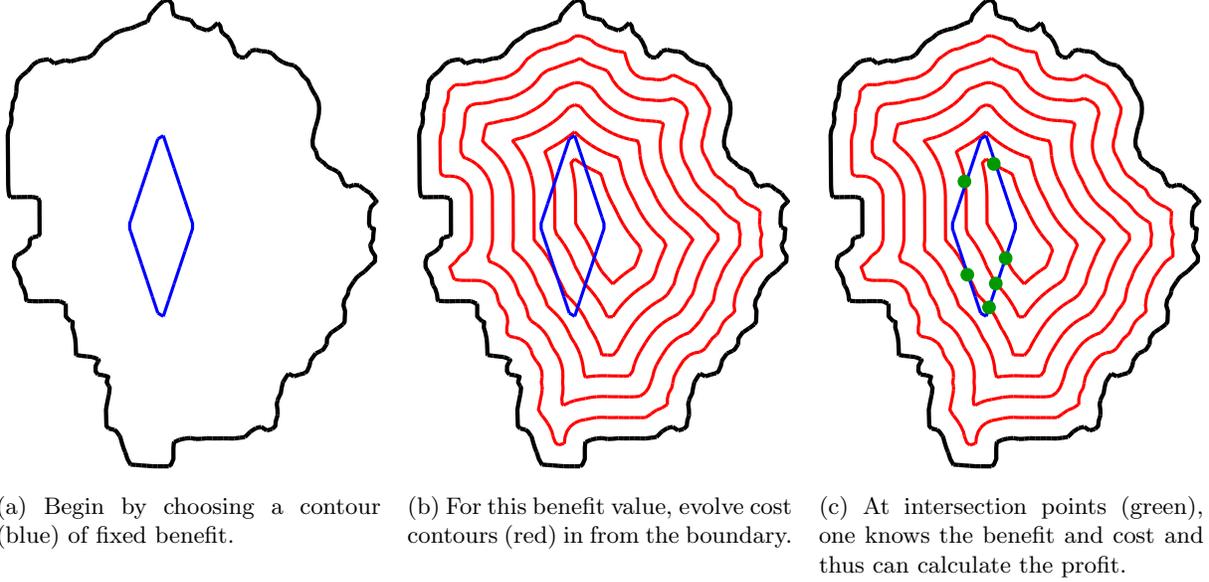


Figure 5: Determining the profit function throughout the region. These images correspond to steps 2, 3 and 4 in the expected cost algorithm (section 4.3) respectively.

4. The cost level sets only apply to the points whose benefit has been used to calculate the level sets. That is, we have calculated the cost assuming the benefit gained is B_i , so our results are only valid at the points (x, y) such that $B(x, y) = B_i$. Therefore we find the intersections between the cost level contours $\phi(x, y, C) = 0$ and the benefit contour $B(x, y) = B_i$. At these intersections (marked in green), we can calculate $C(x, y)$.
5. Repeat steps 2 to 4 for each benefit level B_i , until values of $C(x, y)$ are known throughout the region.

Once the expected cost has been found it can be interpolated, and can be used to find the expected profit $P(x, y) = B(x, y) - C(x, y)$. This expression for the profit includes the effect of the defenders.

4.4 Finding the Pristine Region

Once the expected profit function is found, the next step is to find the pristine region, the region that extractors will never occupy. In Albers [1], the pristine region was simply the part of the circular forest between its center and the ring of maximum profit. In general, however, the pristine region must be defined differently. We assume that extractors will extract anywhere where the profit is close to the maximum possible profit. Denote P_{max} as the maximum expected profit, and $(1 - \varepsilon)P_{max}$ as the lowest expected profit that the extractor will accept for some tolerance ε . We also assume that the extractors will take the optimal path away from their extraction point to the boundary of the region.

The pristine region is determined using the following algorithm.

1. Find the equal profit contour corresponding to $(1 - \varepsilon)P_{max}$. The extractors will only extract within this high-profit region.
2. Pick points inside the high-profit region uniformly at random, and calculate the optimal path from that point to the boundary of the region. The optimal path can be found easily by solving the gradient descent equation

$$\frac{d\vec{x}}{dt} = -\nabla C_{B(\vec{x}_0)}(\vec{x}), \quad (7)$$

where $\vec{x}(0) = \vec{x}_0$ and $C_{B(\vec{x}_0)}(\vec{x})$ is the cost function associated with the benefit value at the extraction point \vec{x}_0 .

3. The non-pristine region includes the high-profit region, and all points within some (small) distance of one of the optimal paths found in the previous step. The rest of the region is pristine, and is not traversed by extractors.

4.5 Metrics for Measuring Patrol Effectiveness

To measure the effectiveness of a given patrol strategy ψ , we use two metrics, although others could be used. A simple measure, as used in Albers [1], is to calculate the proportion of the protected area that is pristine, using,

$$\text{Pristine Proportion} = \frac{\int_{\Omega} \chi(x, y) dx dy}{\int_{\Omega} dx dy},$$

where $\chi(x, y)$ is an indicator function which is 1 in the pristine region and 0 in the non-pristine region.

Some parts of the protected region may be more valuable to protect than other regions, and we could instead choose to measure the proportion of the net value protected using the following integral,

$$\text{Proportion of Value Protected} = \frac{\int_{\Omega} B(x, y)\chi(x, y) dx dy}{\int_{\Omega} B(x, y) dx dy},$$

where χ is defined as previously. In place of the benefit function, other functions could be used that measure the relative importance of protecting certain areas.

4.6 Equivalence of this model and [1]

Although our model is more sophisticated than that given in [1], it does give the same profit function in some simple symmetric cases. Consider a circular forest of radius 1 with walking speed 1 and a uniform patrol strategy with patrol budget 1 (so $\psi = 1/\pi$). Further assume that the benefit associated with extracting at depth d is $B(d) = 2d$ and that the cost is $C(d) = d$ (equivalent to the travel time with walking speed 1 in our model). The capture probability associated with extracting at depth d is $\Psi(d) = \int_0^d \psi(r)dr = d/\pi$, and so the profit function based on equation 2 is

$$P(d) = \left(1 - \frac{d}{\pi}\right) 2d - d = d - \frac{2d^2}{\pi}. \quad (8)$$

We can arrive at the same result using our model. We assume that the protected region is the unit circle centered at the origin so $\phi_0(x) = 1 - |x|$. For x in the unit circle, set $B(x) = 2(1 - |x|)$; this is analogous to setting $B(d) = 2d$ in the Albers model. Further, consider constant walking velocity $v = 1$, constant patrol density $\psi = 1/\pi$ and risk factor $\alpha = 1$. Fix x_0 in the unit circle. Note that $|\nabla\phi| = 1$ almost everywhere and the normal velocity $(1 + 2(1 - |x_0|)/\pi)^{-1}$ in equation 4 is constant with respect to x . Thus the solution to equation 4 is given by

$$\phi(x, t) = 1 - |x| - \frac{t}{1 + \frac{2(1 - |x_0|)}{\pi}}. \quad (9)$$

Recall that the cost $C(x_0)$ associated with extracting at the point x_0 is given implicitly by $\phi(x_0, C(x_0)) = 0$. Solving this equation, we have

$$C(x_0) = (1 - |x_0|) \left(1 + \frac{2(1 - |x_0|)}{\pi}\right) \quad (10)$$

and the profit $P(x_0) = B(x_0) - C(x_0)$ is given by

$$P(x_0) = (1 - |x_0|) \left(1 - \frac{2(1 - |x_0|)}{\pi}\right). \quad (11)$$

Note that this profit function is radial and that for any x_0 at depth d , we have $P(x_0) = d(1 - \frac{2d}{\pi})$ exactly as in equation 8.

In general, with arbitrary benefit function $B(d)$, linear cost $C(d) = cd$, and homogenous patrol density $\psi(d) = p$, Albers gives

$$P(d) = (1 - pd)B(d) - cd, \quad (12)$$

while our proposed model gives (with speed set to $v = 1/c$ and $\alpha = 1$)

$$C(d) = d(c + pB(d)), \quad (13)$$

and so

$$P(d) = B(d) - cd - pdB(d) = (1 - pd)B(d) - cd. \quad (14)$$

The two models agree even if the patrol density is piecewise constant and radially symmetric. If the cost function and/or patrol density are more complicated, then ϕ_t will not be constant, and the PDE solution will be more complicated. Our approach to calculating the cost and benefit functions is therefore equivalent to the approach given by Albers [1] in simple scenarios but is capable of modeling much more general cases.

4.7 Numerical Implementation

We briefly discuss the numerical implementation of our model. Numerical solutions to Hamilton-Jacobi equations have been well-studied for several years. Recent approaches overcome the curse of dimensionality by using Hopf-Lax formulations and optimization [3], but since we are concerned with two spatial dimensions, grid-based finite difference methods are sufficient. Recall, a Hamilton-Jacobi equation is an equation of the form $\phi_t + H(x, \nabla\phi) = 0$. Due to the nonlinear dependence on $\nabla\phi$, we cannot use basic differencing methods since they may lead to instabilities or fail to revolve the viscosity solution to the equation [4]. Rather, in order to minimize oscillation and track the direction in which information is traveling, we replace the Hamiltonian $H(x, y, \phi_x, \phi_y)$ with a numerical Hamiltonian $\hat{H}(x, y, \phi_x^+, \phi_x^-, \phi_y^+, \phi_y^-)$ where ϕ_x^+, ϕ_x^- are the forward and backward difference approximations to ϕ_x respectively, and similarly for ϕ_y^+, ϕ_y^- . There are many different acceptable numerical Hamiltonians [16]. We use the Godunov Hamiltonian given by

$$\hat{H}(x, y, \phi_x^+, \phi_x^-, \phi_y^+, \phi_y^-) = \underset{u \in I(\phi_x^-, \phi_x^+)}{\text{ext}} \underset{v \in I(\phi_y^-, \phi_y^+)}{\text{ext}} H(x, y, u, v) \quad (15)$$

where

$$I(a, b) = [\min(a, b), \max(a, b)] \quad (16)$$

and

$$\underset{x \in I(a, b)}{\text{ext}} = \begin{cases} \min_{a \leq x \leq b} & \text{if } a \leq b, \\ \max_{b \leq x \leq a} & \text{if } a > b. \end{cases} \quad (17)$$

Performing the minimization or maximization can be difficult if the Hamiltonian is complicated, but in simple cases, they can be resolved explicitly. In our case, the Hamiltonian is $H(x, y, \phi_x, \phi_y) = \tilde{v}(x, y) |\nabla\phi|$ where the velocity function \tilde{v} is positive. When the intervals $I(\phi_x^-, \phi_x^+)$ and $I(\phi_y^-, \phi_y^+)$ do not include 0, this Hamiltonian is monotone in each of ϕ_x and ϕ_y on the rectangle $I(\phi_x^-, \phi_x^+) \times I(\phi_y^-, \phi_y^+)$ and the extrema in 15 must occur at the corners of the rectangle. If these intervals do contain 0, some adjustment is necessary, but one can still explicitly resolve the extrema and find that

$$\hat{H}(x, y, \phi_x^+, \phi_x^-, \phi_y^+, \phi_y^-) = \tilde{v}(x, y) \sqrt{\max\{(\phi_x^-)_+^2, (\phi_x^+)_-^2\} + \max\{(\phi_y^-)_+^2, (\phi_y^+)_-^2\}} \quad (18)$$

where $(A)_+ = \max(A, 0)$ and $(A)_- = \min(A, 0)$. The Godunov Hamiltonian \hat{H} gives a first-order approximation to the Hamiltonian H . Following Osher and Shu [16], we approximate the derivatives ϕ_x, ϕ_y at second order and we use second-order total variation diminishing Runge-Kutta time stepping to evolve the solution. In doing so, we have constructed a second order accurate essentially non-oscillatory scheme for 4.

There is an implementation issue which is worth mentioning. Note that the initial function $\phi(x, 0)$ can be taken to be precisely the signed distance from x to the initial contour:

$$\phi(x, 0) = \text{dist}(x, \Gamma) = \begin{cases} \inf_{y \in \Gamma} |x - y|, & x \text{ inside } \Gamma, \\ -\inf_{y \in \Gamma} |x - y|, & x \text{ outside } \Gamma. \end{cases} \quad (19)$$

This is desirable since the distance functions has gradient 1 almost everywhere and thus if we can preserve this property (that is, if we can ensure that $\phi(x, t)$ is the signed distance to $\Gamma(t)$ for all $t > 0$) we will observe the exact level set motion that we want. However as the level sets evolve, there is some distortion so that numerically for $t > 0$, we no longer have $\phi(x, t) = \text{dist}(x, \Gamma(t))$. If $|\nabla\phi|$ becomes too small, it is difficult to

resolve the level sets accurately and if $|\nabla\phi|$ becomes too large, we are required to choose exceedingly fine time discretization or else the numerical solution may develop instabilities. We can prevent these from happening by occasionally replacing ϕ with the signed distance function to $\Gamma(t)$. That is, every so often, we halt the time integration, find the current zero level contour $\Gamma(t)$, reset $\phi(x, t) = \text{dist}(x, \Gamma(t))$ and continue. This process is known as *re-distancing* and is discussed by Lee et al. [13].

5 Results

We present results for our algorithm applied to two real world locations: Yosemite National Park in California, and Kangaroo Island in South Australia. Yosemite National Park is a mountainous area with steep mountains and long valleys. Kangaroo Island has an interesting shape, with a narrow neck separating the main part of the island from a smaller part at the eastern end. For both locations we use real elevation data, sourced from the United States Geographical Survey (Yosemite National Park data) and the Foundation Spatial Data Framework (Kangaroo Island data). The elevation profiles for Yosemite National Park and Kangaroo Island are displayed in figure 6. The data was processed using QGIS [22] and imported to MATLAB using TopoToolbox [17, 18]. We apply several patrol strategies identified by Albers [1] and Johnson et al. [9] before suggesting some simple and more effective patrols that account for the geometry of the regions.

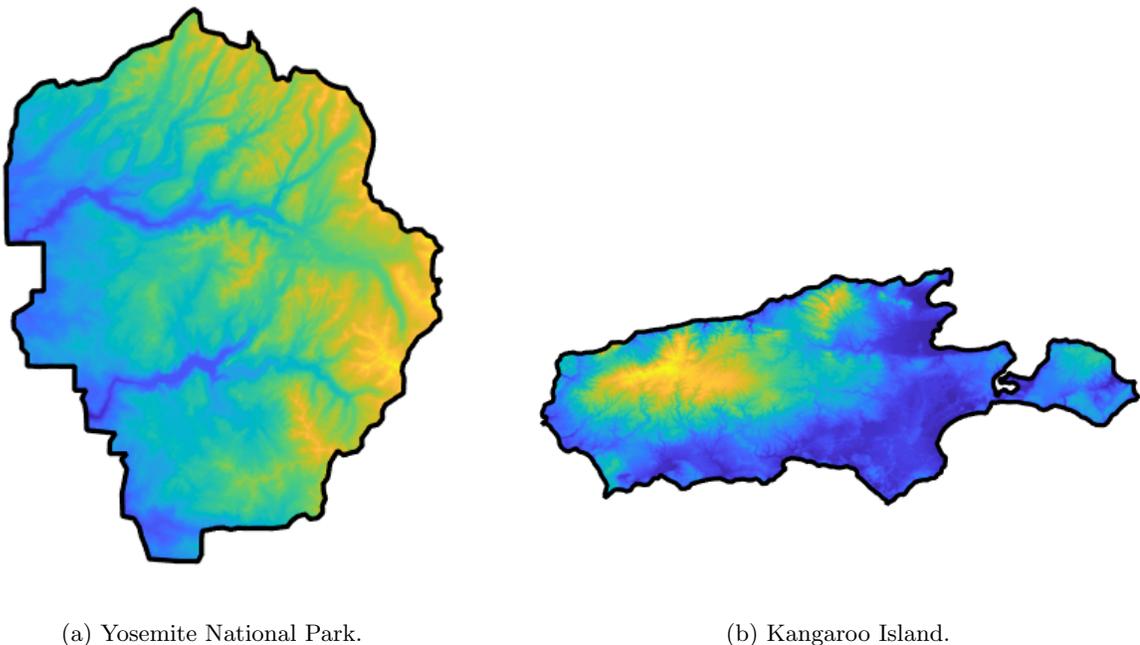
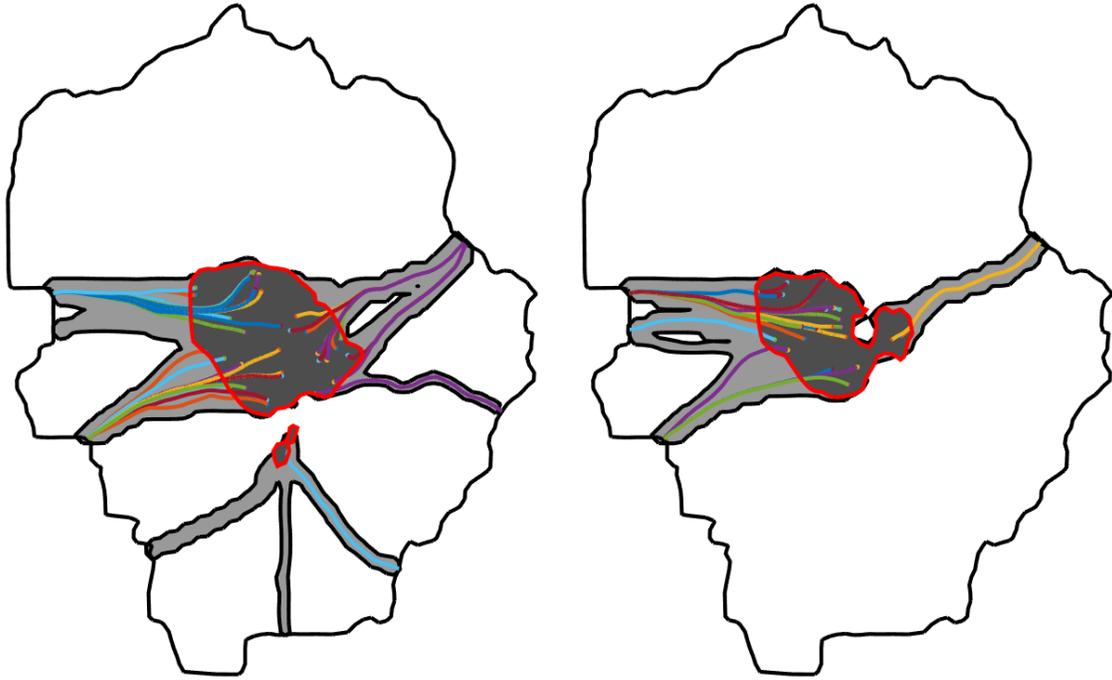


Figure 6: Elevation profiles for Yosemite National Park and Kangaroo Island. Yellow corresponds to higher elevation, blue to lower elevation (scales different in each figure).

5.1 Yosemite National Park without patrols

We first consider the case without any patrols, so the cost function depends only on the effort required to travel from the extraction point to the boundary of the park. In figure ??, we present results for two cases with different benefit functions. Both benefit functions have the same form, a quadratic increase from 0 at the boundary to a maximum value at the point furthest from the boundary, but figure ?? has maximum benefit double that as in figure ?. The lower benefit case has a larger pristine area and smaller high-profit region. In the high-benefit case, there is enough incentive for extractors to expend more effort and extract from more locations within the protected region, obtaining much higher profits (although not doubled). A selection of the optimal paths from extraction points to the boundary of the protected region are also shown. In the sections that follow, we will use $k = 8$, the high-benefit case.

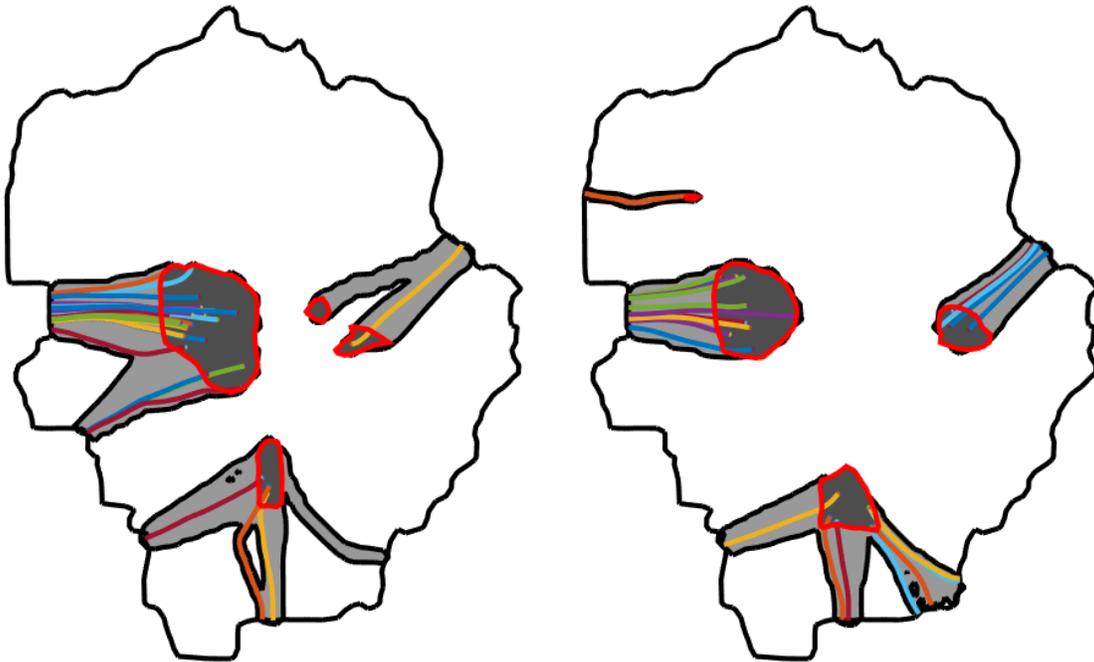


(a) $k = 4$, 24 paths out of the total 192 are shown. Maximum profit is 7.06×10^4 and pristine proportion is 0.867. (b) $k = 8$, 37 paths out of the total of 294 are shown. Maximum profit is 1.65×10^5 and pristine proportion is 0.784.

Figure 7: Two cases for Yosemite National Park with no patrols, benefit based on distance from boundary d as $B(d) = kd(2d_m - d)/d_m$ (where $d_m = \max_{\Omega} d$) and $k = 4$ in (a) and 8 in (b). The dark gray denotes the high-profit region where extraction occurs, and light gray shows the envelope of paths the extractors follow when exiting the protected region. Some of the individual paths are shown for illustrative purposes.

5.2 Homogeneous patrols

In figure ?? we consider the simplest nonzero patrol strategy, a homogeneous patrol in which the entire protected area is patrolled with equal intensity. Two cases are shown, both with the high-benefit case $k = 8$ from the previous section, but with different patrol budgets E . The patrol strategy is simply $\psi = E/A$ where E is the budget and A is the area of the protected region. We specify $\alpha = 1$ as the risk-aversion parameter, as this choice gives agreement between our model and that of Albers [1] in simple symmetric case (discussed in section 4.6). As expected, when the patrol budget increases the pristine proportion increases and the maximum profit obtained by the extractor decreases.



(a) $E = 3 \times 10^4$, 26 paths out of the total of 205 are shown. Maximum profit is 1.27×10^5 and pristine proportion is 0.809. (b) $E = 6 \times 10^4$, 23 paths out of the total of 179 are shown. Maximum profit is 9.86×10^4 and pristine proportion is 0.841.

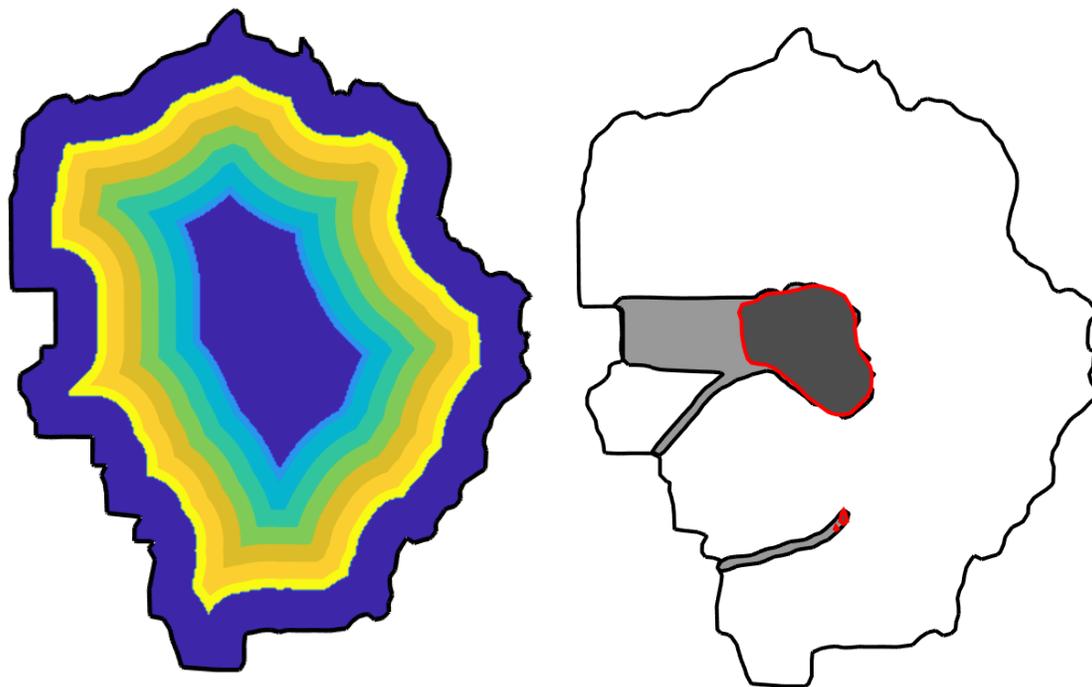
Figure 8: Two cases for Yosemite National Park with homogeneous patrols, benefit based on distance from boundary d as $B(d) = kd(2d_m - d)/d_m$ (where $d_m = \max_{\Omega} d$ and $k = 8$) and different budgets $E = 3 \times 10^4$ in (a) and $E = 6 \times 10^4$ in (b).

In figure ??, doubling the patrol budget decreases the maximum profit obtained by extractors by roughly 22%, increases the pristine proportion by roughly 4%, and increases the proportion of benefit protected by roughly 8.6%. These are clear improvements, but one perhaps would expect more change upon doubling the patrol density. Johnson et al. [9] claimed that the homogeneous patrol was not optimal, and this agrees with our intuition that some areas of the protected region will never be penetrated by extractors, so there is no need to patrol there.

5.3 Band Patrols

Johnson et al. [9] identified the band patrol as the optimal strategy in symmetric circular protected regions. In the circular case, the band patrol is a band between two distances from the boundary, $0 < d_o < d_i < d_m$, where d_m is the maximum distance from the boundary, with highest patrol density at the outer extent d_o , and decreasing density moving towards the inner-most extent of the band d_i . In the general case, we set up a band patrol by patrolling between $0.3d_m$ and $0.7d_m$ from the boundary. figure ?? shows the patrol strategy and the results of our algorithm. We did not implement the algorithm presented by Johnson et al.

[9] to find the optimal band patrol, instead testing a number of band patrols and choosing the one that gave the best results.



(a) The patrol strategy. Yellow corresponds to more intense patrolling, blue to no patrolling. (b) The maximum profit is 1.25×10^5 and the pristine proportion is 0.901.

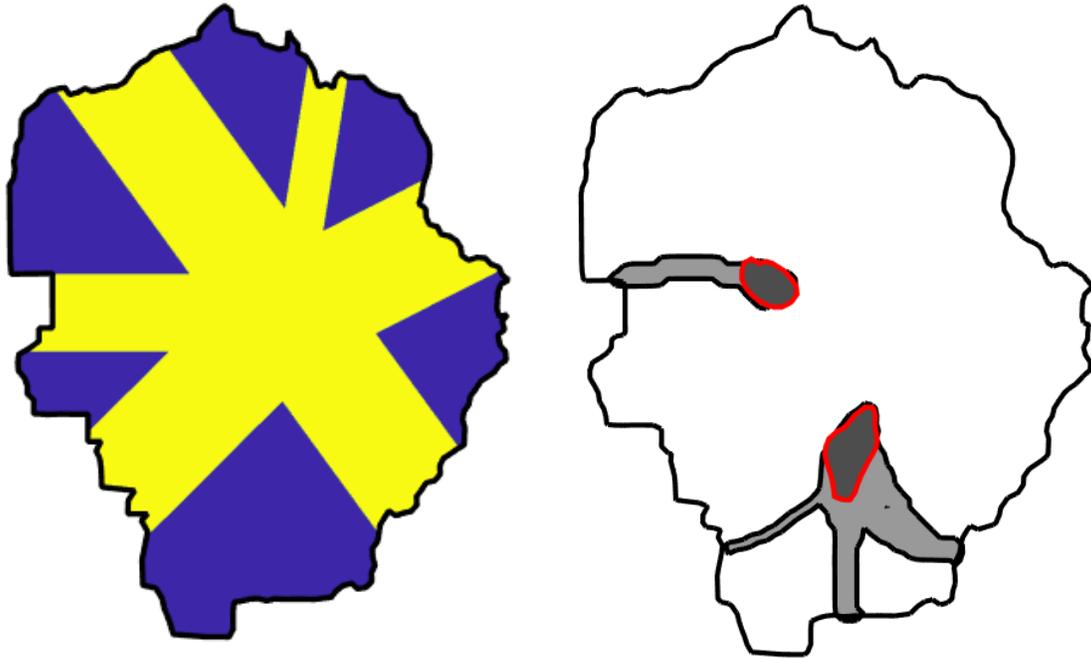
Figure 9: A band patrol in Yosemite National Park. The patrol is based on distance from boundary d and decreases linearly from $d = 0.3$ to $d = 0.7$, and is zero elsewhere. The benefit is $B(d) = kd(2d_m - d)/d_m$ (where $d_m = \max_{\Omega} d$ and $k = 8$) and the budget $E = 3 \times 10^4$.

The patrol budget in figure ?? is 3×10^4 , the same as for the homogeneous patrol in figure ??, but the outcome is significantly better for the patrollers. The homogeneous patrol had a pristine proportion 0.809 and was able to protect 76.9% of the total benefit, whereas the band patrol has a pristine proportion of 0.901 and protects 86% of the total benefit. The optimal band patrol of Johnson et al. [9] has the property that extractors do not enter the patrolled region, which ours does not. Despite this, it is still superior to the the homogenous patrol (even the homogeneous patrol with twice the budget!) and the other band patrols tested.

5.4 Asymmetric Patrols

The patrol strategies employed above were identified by Albers [1] and are radially symmetric with respect to the geometry of Yosemite National Park. That is, these strategies depended only on the depth d of the point (x, y) : $\psi(x, y) = \psi^*(d)$. Our method does not require patrols to be radially symmetric, and we here present an asymmetric patrol that outperforms the symmetric patrols discussed previously. Observing the results of the other patrols, it seems that extractors prefer to enter and leave the park at certain portions of the boundary: the portions which are most concave. This is intuitive, as entering at those areas will ensure less travel distance to reach the center. In response, we can design strategies to preferentially patrol those regions through which extractors are more likely to enter. figure ?? shows just such a patrol strategy.

Exploiting the geometry of the park, rather than patrolling in a depth-dependent way, the pristine proportion is increased to 0.92 and 90.7% of the benefit is protected. These results are better than the band or homogeneous patrols with the same budget, and could likely be improved upon even further by finding even more specialized strategies. This result shows the critical importance of the geometry of the protected



(a) The patrol strategy. Yellow corresponds to more intense patrolling, blue to no patrolling.

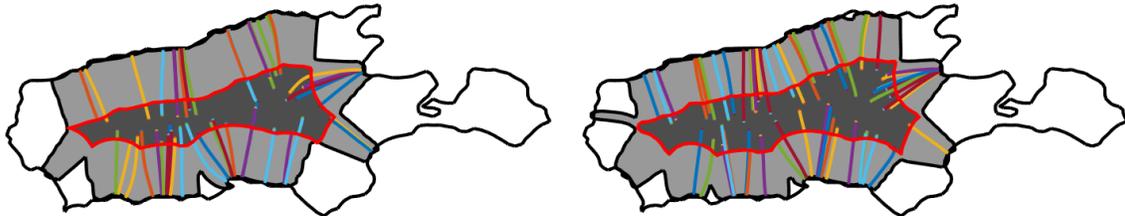
(b) The maximum profit is 1.52×10^5 and the pristine proportion is 0.920.

Figure 10: A custom patrol in Yosemite National Park which is designed to patrol the concavities in the boundary of the park. The benefit is $B(d) = kd(2d_m - d)/d_m$ (where $d_m = \max_{\Omega} d$ and $k = 8$) and the budget $E = 3 \times 10^4$.

region. To a very rough approximation Yosemite is fairly circular, but the relatively small-scale concavities are very important in understanding the behaviour of extractors.

5.5 Kangaroo Island

We now apply our method to Kangaroo Island, a small island off the southern coast of Australia. We present two patrol strategies (a homogeneous patrol and a custom designed patrol) which once again emphasize the importance of accounting for the geometry of the protected area. Kangaroo Island is fairly long and narrow with a portion at the eastern end connected to the rest of the island by a small neck.



(a) $E = 3 \times 10^4$, 44 paths out of the total of 431 are shown. Maximum profit is 1.34×10^5 and pristine proportion is 0.349. (b) $E = 6 \times 10^4$, 58 paths out of the total of 579 are shown. Maximum profit is 1.15×10^5 and pristine proportion is 0.305.

Figure 11: Two cases for Kangaroo Island with homogeneous patrols, benefit based on distance from boundary d as $B(d) = kd(2d_m - d)/d_m$ (where $d_m = \max_{\Omega} d$ and $k = 8$) and different budgets $E = 3 \times 10^4$ in (a) and $E = 6 \times 10^4$ in (b).

Figure 11 shows results of homogeneous patrols applied to Kangaroo Island. In this case, we observe that, somewhat counter-intuitively, the increased patrol budget results in a smaller pristine area proportion. The reason for this is that increasing the patrol budget decreases the profit function but causes the high profit area to spread out. That is, with increased patrol the surface profile of the profit function will look more like a plateau which has a smaller maximum than in the low budget case, but has a larger near-maximal area. Although the pristine region is smaller, the maximum profit obtained by the extractors is much lower, decreasing from 1.34×10^5 to 1.15×10^5 , so the patrol does have a significant effect on the extractors. We also observe that a large portion of the non-pristine area comprises the paths leaving the region as opposed to the high-profit region itself. This should inform our decision regarding how to more effectively patrol the region. In any radially symmetric patrol, the high-profit region will also be radially symmetric meaning that it will likely occupy the middle of the island and the paths will cover a very large area as they enter and exit from the north and south side of the island. If we can force the high profit region out of the middle of the island, then the paths will instead travel throughout the east and west portions of the island, hopefully occupying a smaller area. Another key geographic feature of the island is the peninsula at the eastern tip of the island. The peninsula is connected to the island by an isthmus thin enough that, for our purposes, the peninsula can almost be considered an independent region. Extractors will likely be uninterested by the peninsular region since it has much less depth and thus offers much less benefit than the main body of the island. Hence any patrol in that region is likely wasted effort.

With the homogeneous results in mind, figure ?? shows the results of a more effective patrol, where the middle of the island is patrolled uniformly and the eastern and western ends are not patrolled. This patrol was able to increase the pristine area proportion to 0.875 compared to the homogeneous patrol which gave 0.349 with the same budget. Again, designing this patrol required some simple observations regarding the geometry of the region, and once again shows the importance of explicitly incorporating geographical information into the model.

6 Conclusion

Modeling of environment crime in protected regions has been the subject of previous research, most of which has been focused on discrete network-based methods. Attempts at modelling environmental crime in

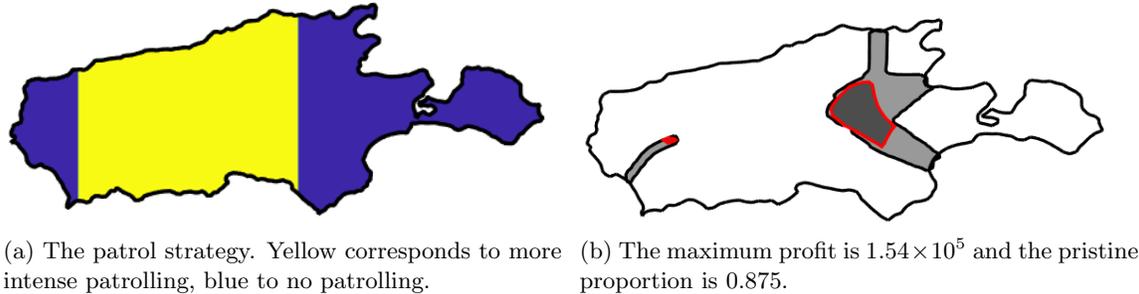


Figure 12: A custom patrol in Kangaroo Island which is designed to push the high profit region away from the center of the island. The benefit is $B(d) = kd(2d_m - d)/d_m$ (where $d_m = \max_{\Omega} d$ and $k = 8$) and the budget $E = 3 \times 10^4$.

continuous settings have required restrictive assumptions of symmetry and are only applicable to regions with very simple geometries. In this paper we have formulated a generalized version of the model of Albers [1] which can account for regions of arbitrary geometry, does not require any assumptions of symmetry and can incorporate topographical features. In doing so, we have unified environmental crime modeling and continuum, PDE models for tracking human movement.

The level set method of Osher and Sethian [15] allows us to perform calculations in regions with complicated geometry. Using the level set method, we track the movement of environmental criminals (extractors) by treating them as a front propagating from the boundary of the protected region toward its center while accounting for changes in travel velocity due to elevation. The level set method is very powerful, and is suitable for many more applications involving movement through regions with complicated geometry.

Working with arbitrary geometries increases the complexity of the model. Accordingly, we have proposed a novel method for incorporating patrol strategy into the calculation of cost to the extractors and suggested an efficient algorithm for resolving the cost function while performing relatively few level set computations. Further, we have re-defined the pristine area by considering an extraction area (an area of near-maximal profit to extractors) as well as paths that extractors will traverse when leaving the region. This is a necessary deviation from the Albers [1] model that arises due to the different topology of the point(s) that maximize the extractors' profit.

We have applied our model to two different regions: Yosemite National Park in California and Kangaroo Island in South Australia. In doing so, we tested several different patrol strategies, including some proposed by Albers [1] and Johnson et al. [9]. Then, making some basic observations regarding the geometry of each region, we designed asymmetric patrols which protected the regions more effectively than any of those previously suggested. The success of these asymmetric patrols shows the importance of explicitly treating the geometry. We expect that even more effective patrol strategies exist, and determining them is a potential avenue for further research.

The model presented here could be extended and modified in many ways. For example, the velocity function could be made to depend on ground cover (i.e. higher speeds on open plains and lower speeds in heavy scrub) or other effects such as lakes or paths. Likewise, the benefit and cost functions could be calculated differently depending on the situation. In the proposed model, the cost of travelling to the extraction point is not considered, only the trip out of the region after the crime has been committed. This follows Albers [1], who neglected the cost of the inbound trip. Adding the extra cost of the inbound trip would be a minor modification of the algorithm as presented. We have only considered benefit functions that depend on distance from the boundary, which is likely not realistic. The proposed model can be evaluated with any benefit function (continuity is not necessary), however defining a realistic and accurate benefit function for a given national park would require input from agencies involved in managing the park.

We have defined patrol strategies as density functions, but these do not correspond to patrol routes that would be followed. Determining specific patrol routes that give the desired patrol density is a task requiring significant work. Fang et al. [7] used a discrete model to find explicit patrol routes but in our continuous setting the problem is very complicated. Another opportunity for further research is studying the effects of time-dependent benefit functions. We have used a constant benefit function but in some situations this may

not be appropriate. For example, considering animal poaching, one might wish to account for the movement of animals over time. This would require significant modification to the model we have presented.

7 Acknowledgements

The authors acknowledge the financial support of the NSF (grant DMS-1737770) and the Department of Defense (approved for public release, #18-666).

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] A. Bayen, I. M. Mitchell, M. Oishi, and C. J. Tomlin. Aircraft autolander safety analysis through optimal control-based reach set computation. 30:68–77, 01 2007.
- [3] Y. T. Chow, J. Darbon, S. Osher, and W. Yin. Algorithm for overcoming the curse of dimensionality for state-dependent Hamilton-Jacobi equations. *preprint*. URL <https://arxiv.org/abs/1704.02524>.
- [4] M. G. Crandall and P.-L. Lions. Two approximations of solutions of Hamilton-Jacobi equations. *Mathematics of Computation*, 43(167):1–19, 1984.
- [5] L. E. Dubins. On curves of minimal length with a constraint on average curvature, and with prescribed initial and terminal positions and tangents. *American Journal of Mathematics*, 79(3):497–516, 1957.
- [6] F. Fang, A. X. Jiang, and M. Tambe. Optimal patrol strategy for protecting moving targets with multiple mobile resources. In *Proceedings of the 12th International Conference on Autonomous Agents and Multiagent Systems*, pages 957–964. International Foundation for Autonomous Agents and Multiagent Systems, 2013.
- [7] F. Fang, T. H. Nguyen, R. Pickles, W. Y. Lam, G. R. Clements, B. An, A. Singh, B. C. Schwedock, M. Tambe, and A. Lemieux. PAWS – a deployed game-theoretic application to combat poaching. *AI Magazine*, 38(1):23–36, March 2017.
- [8] I. J. Irmischer and K. C. Clarke. Measuring and modeling the speed of human navigation. *Cartography and Geographic Information Science*, 45(2):177–186, 2017.
- [9] M. P. Johnson, F. Fang, and M. Tambe. Patrol strategies to maximize pristine forest area. In *Proceedings of the Twenty-Sixth AAAI Conference on Artificial Intelligence*. AAAI, 2012.
- [10] D. Kar, F. Fang, F. Delle Fave, N. Sintov, and M. Tambe. A game of thrones: when human behavior models compete in repeated Stackelberg security games. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems*, pages 1381–1390. International Foundation for Autonomous Agents and Multiagent Systems, 2015.
- [11] D. Kar, B. Ford, S. Gholami, F. Fang, A. Plumbtree, M. Tambe, M. Driciru, F. Wanyama, and A. Rwtisiba. Cloudy with a chance of poaching: adversary behavior modeling and forecasting with real-world poaching data. In S. Das, E. Durfee, K. Larson, and M. Winikoff, editors, *Proceedings of 16th International Conference on Autonomous Agents and Multiagent Systems*, AAMAS ’17, pages 159–167. International Foundation for Autonomous Agents and Multiagent Systems, 2017.
- [12] N. Leader-Williams and E. J. Milner-Gulland. Policies for the enforcement of wildlife laws: the balance between detection and penalties in Luangwa Valley, Zambia. *Conservation Biology*, 7(3):611–617, 1993.
- [13] B. Lee, J. Darbon, S. Osher, and M. Kang. Revisiting the redistancing problem using the Hopf–Lax formula. *Journal of Computational Physics*, 330(1):268–281, February 2017.

- [14] S. Osher and R. P. Fedkiw. *Level set methods and dynamic implicit surfaces*, volume 153 of *Applied Mathematical Sciences*. Springer–Verlag, 2003.
- [15] S. Osher and J. Sethian. Fronts propagating with curvature dependent speed: Algorithms based on Hamilton–Jacobi formulations. *Journal of Computational Physics*, 79:12–49, 1988.
- [16] S. Osher and C. W. Shu. High order essentially non–oscillatory schemes for Hamilton–Jacobi equations. *SIAM Journal of Numerical Analysis*, 28(4):907–922, August 1991.
- [17] W. Schwanghart and N. Kuhn. Topotoolbox: a set of Matlab functions for topographic analysis. *Environmental Modelling & Software*, 25(6):770–781, 2010.
- [18] W. Schwanghart and D. Scherler. Topotoolbox 2: MATLAB-based software for topographic analysis and modeling in Earth surface sciences. *Earth Surface Dynamics*, 2(1):1–7, 2014.
- [19] J. Sethian and A. Vladimirsky. Ordered upwind methods for static hamilton–jacobi equations. *Proceedings of the National Academy of Sciences*, 98(20):11069–11074, 2001. ISSN 0027-8424.
- [20] J. Sethian and A. Vladimirsky. Ordered upwind methods for static hamilton–jacobi equations: Theory and algorithms. *SIAM Journal on Numerical Analysis*, 41(1):325–363, 2003.
- [21] R. Takei, R. Tsai, H. Shen, and Y. Landa. A practical path-planning algorithm for a simple car: a Hamilton-Jacobi approach. In *Proceedings of the 2010 American Control Conference*, pages 6175–6180, June 2010.
- [22] Quantum GIS Development Team. Quantum GIS geographic information system, 2017. URL <http://qgis.osgeo.org>.
- [23] W. Tobler. Three presentations on geographical analysis and modeling. Technical Report 93-1, National Center for Geographic Information and Analysis, February 1993.
- [24] Z. Zhou, J. Ding, H. Huang, R. Takei, and C. Tomlin. Efficient path planning algorithms in reach-avoid problems. *Automatica*, 89:28 – 36, 2018.