

CONTROL-THEORETIC MODELS OF ENVIRONMENTAL CRIME*

ELLIOT CARTEE[†] AND ALEXANDER VLADIMIRSKY[‡]

Abstract. We present two models of perpetrators' decision-making in extracting resources from a protected area. It is assumed that the authorities conduct surveillance to counter the extraction activities, and that perpetrators choose their post-extraction paths to balance the time/hardship of travel against the expected losses from a possible detection. In our first model, the authorities are assumed to use ground patrols and the protected resources are confiscated as soon as the extractor is observed with them. The perpetrators' path-planning is modeled using the optimal control of randomly terminated process. In our second model, the authorities use aerial patrols, with the apprehension of perpetrators and confiscation of resources delayed until their exit from the protected area. In this case the path-planning is based on multiobjective dynamic programming. Our efficient numerical methods are illustrated through several examples with complicated geometry and terrain of protected areas, nonuniform distribution of protected resources, and spatially nonuniform detection rates due to aerial or ground patrols.

Key words. optimal control, Hamilton–Jacobi, multiobjective path-planning, randomly terminated processes, environmental crime

AMS subject classifications. 49N90, 90C29, 35F21, 65M06, 91B76

DOI. 10.1137/19M1270483

1. Introduction. Mathematical modeling of environmental crimes in national parks has been a growing area of research. Both the authorities and nongovernmental conservation organizations strive to prevent *illegal extraction activities*, including wildlife poaching and illegal logging. In 2010, Albers [1] introduced a framework where the government chooses a patrol density (determining the probability of detection at each location) in such a way as to maximize the pristine area of the park. The perpetrators then choose where to extract resources in such a way as to maximize their expected payoff. The problem of choosing patrol densities to maximize the pristine area has also been explored in [23]. In recent years, patrol strategies based on graph-theoretic models (such as PAWS [19] and INTERCEPT [25]) have even been deployed in actual national parks.

The original Albers model [1] was posed with many restrictive/unrealistic assumptions, including the radial symmetry of the protected area, patrol strategies, and resource distribution. In a recent paper by Arnold et al [4], these restrictions were largely removed, with numerical simulations on arbitrary two-dimensional domains and the terrain directly affecting the speed of perpetrators. The key mathematical tool employed in [4] is the level set method, developed to track the propagation of interfaces by solving time-dependent Hamilton–Jacobi PDEs [29]. We begin by briefly reviewing both of these models in section 2.

In this paper, we offer a significantly different perspective on how to generalize the Albers model to arbitrary domains. As explained in section 3, our main focus is on

*Received by the editors June 25, 2019; accepted for publication (in revised form) February 19, 2020; published electronically June 9, 2020.

<https://doi.org/10.1137/19M1270483>

Funding: This work was partially supported by the National Science Foundation through grant DMS-1738010. The work of the second author was supported by the Simons Foundation Fellowship.

[†]Department of Mathematics, Cornell University, Ithaca, NY 14853 (evc34@cornell.edu).

[‡]Department of Mathematics and Center for Applied Mathematics, Cornell University, Ithaca, NY 14853 (vladimirsky@cornell.edu).

a careful treatment of the path-planning problem faced by the extractors. Capturing their rational behavior yields a more accurate prediction of affected areas and allows for a better planning of optimal patrol strategies. We show that the perpetrators optimization problem changes significantly based on the consequences of being spotted while transporting the protected resources.

In our first “Model G,” where authorities use “ground patrols” to spot the extractors, any detection leads to an immediate confiscation, with perpetrators subsequently switching to the fastest path to leave the protected area. The probability of being detected along any originally chosen path is not a sufficient description here, since it does not reflect the expected duration of that post-confiscation trip. We handle this challenge by finding the optimal controls for a suitable *randomly terminated process*. In our second “Model A,” where authorities run “aerial patrols,” it is assumed that the perpetrators remain oblivious of when a detection happens and are only apprehended later, once they reach the boundary of the protected area. As a result, they always use their preselected trajectories, which are chosen to balance the probability of detection against the time or (cumulative difficulty) of the post-extraction trip. In standard path-planning applications, the goal is usually to minimize the integral of some “running cost” along a chosen trajectory. The key difficulty of this second model is that no such running cost can be explicitly defined to recover trajectories maximizing the extractor’s expected profit. Nevertheless, in section 4 we show that this problem can be solved using the tools of *multiobjective* optimal control theory.

We discuss efficient numerical methods for the associated stationary Hamilton–Jacobi–Bellman HJB PDEs in section 5. Our algorithms are illustrated, through a number of examples in section 6, including realistic terrain maps from Yosemite National Park in California and Kangaroo Island in South Australia. We show that Models A and G often yield very different perpetrator trajectories and also differ in their predictions of the subdomain unaffected by extraction activities. We conclude by discussing possible extensions and directions for future work in section 7.

2. Prior work. The model proposed by Albers in [1] is based on a *one-period Stackelberg game* between the protected area managers (PAMs) and the extractors targeting a protected domain $\Omega \subset \mathbb{R}^2$. PAMs decide on their patrol strategy, defining the pointwise rate of extractor-detection $\psi : \Omega \rightarrow \mathbb{R}_{+,0}$ and respecting the known budget constraints for patrolling efforts,

$$(2.1) \quad \int_{\Omega} \psi^{\gamma}(\mathbf{x}) d\mathbf{x} \leq E,$$

with a constant $\gamma \geq 1$ modeling a possible increase in marginal costs of maintaining a higher detection rate.

Once ψ is chosen, it becomes known to extractors, who select which parts of Ω to target, trying to maximize their expected profit (defined below). Correspondingly, PAMs’ goal is to select ψ that will minimize the impact of the extractors’ best/rational response. This can be interpreted as either maximizing the “pristine” area unaffected by the extractors or maximizing the amount of nonextracted resources. In Albers’ original model both of these interpretations are equivalent since she assumes a homogeneous distribution of protected resources.

Even more restrictively, the domain Ω is assumed to be a unit disk, with extractors evenly distributed along its boundary. This radial symmetry dramatically simplifies the problem and allows for analytical solutions for the optimal behavior of all participants. For example, it is easy to see that only radially symmetric $\psi(\mathbf{x})$ would need to

be considered — Albers further restricts this to piecewise constant radial functions, with positive rate of detection in an annulus between radial distances d_e and d_i from $\partial\Omega$ and $\psi = 0$ everywhere else. The same radial symmetry ensures that extractors travel along straight lines from $\partial\Omega$ toward the center of Ω and their entire strategy can be encoded by the chosen depth of penetration $d \in [0, 1]$. Albers assumes that extractors travel to their chosen location unimpeded (since they have not broken any law yet), then turn around and start extracting resources as they retrace their steps to $\partial\Omega$. If they are detected by PAMs on their way back, the extracted resources are confiscated and they have to return to $\partial\Omega$ empty-handed, incurring the time/opportunity cost without any reward. Thus, each extractor chooses d to maximize their expected profit

$$(2.2) \quad P(d) = (1 - \Psi(d)) B(d) - C(d) = B(d) - (\Psi(d)B(d) + C(d)),$$

where Ψ , B , and C are respectively the cumulative probability of being detected, the value of extracted resources, and the cost of traversing the path — all corresponding to going distance d into Ω before starting to extract resources. While this is not explicitly specified in [1], C can also include the cost of getting to the point from which the extraction starts. A more subtle point is that B includes the value of *all* resources extracted along the way back. In the general setting, this continuous-along-the-path extraction would induce a competition among the extractors searching for the best paths through Ω , but the radial symmetry mitigates this effect. However, this makes it harder to justify the *one-period* nature of the model: as the extractors' actions deplete the resources in parts of the protected area, that would obviously change the future optimizations problems (faced by both PAMs and the extractors) even in a radially symmetric setting. Such dynamic consequences of “incomplete enforcement” are acknowledged in [1] and also well-known in practice; see e.g., [31, 32].

A recent paper by Arnold et al. [4] extended the above general approach to more realistic protected areas. Their model allows for general domain geometries, resource distributions, terrain effects, and PAMs' strategies. One important distinction is that, unlike Albers, they model the extraction as occurring at isolated points. An advantage of this approach is that this reduces competition among the extractors (since the paths-to- $\partial\Omega$ traversed by others don't affect the value of what you are currently trying to extract). For an extractor targeting a location $\mathbf{x} \in \Omega$, they define the expected profit as

$$(2.3) \quad P(\mathbf{x}) = B(\mathbf{x}) - \tilde{C},$$

where B is the value of resources available for extraction in a small neighborhood of \mathbf{x} , while \tilde{C} is the effective cost of a post-extraction trajectory, accounting both for the time/difficulty of travel from \mathbf{x} to $\partial\Omega$ and for the expected losses from a possible capture. A perpetrator chooses this trajectory to minimize \tilde{C} (and thus maximize P). In [4] this optimization is performed implicitly by using a level set method originally developed for tracking the propagation of interfaces [29]. The authors show that, in an Albers-style radially symmetric setting with unit walking speed of perpetrators, their approach produces the same $P(\mathbf{x})$ as predicted by formula (2.2). A comparison with (2.3) thus suggests an explicit interpretation $\tilde{C}(\mathbf{x}) = \Psi B(\mathbf{x}) + C$.

For any chosen constant $b > 0$, Arnold et al. compute $\tilde{C}(\mathbf{x}_0)$ simultaneously for all extraction positions \mathbf{x}_0 such that $B(\mathbf{x}_0) = b$. The idea is to find equal-effective-cost curves by using the level sets of an auxiliary time-dependent function $\phi(\mathbf{x}, t)$ satisfying a Hamilton–Jacobi PDE

$$(2.4) \quad \phi_t + F(\mathbf{x}, b)|\nabla\phi| = 0, \quad t > 0, \mathbf{x} \in \Omega,$$

with initial conditions $\phi(\mathbf{x}, 0) = \text{distance from } \mathbf{x} \text{ to } \partial\Omega$. Assuming that $\Gamma(t)$ is some interface starting from $\Gamma(0) = \partial\Omega$ and monotonically advancing into Ω with the normal speed F , its evolution is encoded by the zeroth level set of ϕ solving the above PDE, i.e., $\phi(\mathbf{x}, t) = 0 \iff \mathbf{x} \in \Gamma(t)$. For the purposes of this model, the authors choose their normal speed of the interface to be

$$(2.5) \quad F(\mathbf{x}, b) = \frac{1}{1/f(\mathbf{x}) + \alpha\psi(\mathbf{x})b},$$

where $f(\mathbf{x})$ is the extractor's local speed of motion defined by the terrain at \mathbf{x} , and a parameter $\alpha > 0$ encodes the degree of extractor's risk-aversion. This formula is motivated phenomenologically: $1/f(\mathbf{x})$ represents the extractor's "slowness of motion" and $\psi(\mathbf{x})b$ represents his expected losses (per unit time, while traveling through \mathbf{x}). An increase in either of these terms should make the effective cost higher. Once ϕ is computed, the authors define \tilde{C} implicitly through

$$\phi(\mathbf{x}_0, \tilde{C}(\mathbf{x}_0)) = 0 \quad \forall \mathbf{x}_0 \text{ such that } B(\mathbf{x}_0) = b,$$

with an optimal path to $\partial\Omega$ from any such \mathbf{x}_0 recovered via gradient descent in ϕ . The process is then repeated for multiple values of b , until \tilde{C} (and thus also P) are defined for all gridpoints of some grid imposed on Ω .

In another departure from Albers' original model, the authors assume that extractors will target all locations where P is within $(1 - \varepsilon)$ factor from its maximal value, i.e., the extraction happens on the set

$$\Omega_e = \left\{ \bar{\mathbf{x}} \in \Omega \mid P(\bar{\mathbf{x}}) \geq (1 - \varepsilon) \max_{\mathbf{x} \in \Omega} P(\mathbf{x}) \right\}.$$

(One drawback of this assumption is that extraction will always happen on some part of Ω , regardless of how high the detection rate $\phi(\mathbf{x})$ might be.) The authors then define the pristine area by excluding not only the above but also the union of narrow corridors around optimal trajectories from Ω_e . The latter are approximated by starting from a finite number of points selected uniformly at random in Ω_e . Since there is no extraction along the path, the motivation for this exclusion might perhaps be based on extraction-unrelated damage to the protected area. While [4] does not present a systematic method for optimizing PAMs' choice of ψ , the authors use the above interpretation to compare the effectiveness of several (ad hoc) patrol strategies staying within the specified patrol budget. Their simulations (using $\gamma = 1, \alpha = 1$, and $\varepsilon = 0.085$) are based on the terrain data from Yosemite National Park in California and Kangaroo Island in South Australia, which they have also kindly shared with us to conduct the numerical experiments presented in section 6.

We end this section by discussing the difference between path-planning to avoid detection versus capture/interception. Even if the former always leads to the latter, there is still a very subtle modeling question of whether the capture happens immediately upon detection and, if there is a significant delay, whether the extractors learn that they have been detected before the capture. These issues are not important in [1] since the extractors' paths are always radial in Albers' setting. But Arnold et al. describe path-planning that strikes a balance between the time of travel and risk of detection all the way to the boundary — with no provision for switching to

quicker/easier paths if authorities manage to intercept the extractors en route and confiscate their haul. Thus, we believe that the model in [4] is more suitable for the setting where the extractors remain oblivious of their detection (perhaps in the case of aerial surveillance) and are intercepted only upon reaching $\partial\Omega$.

3. Extractor's path-planning. This section defines the optimization problem faced by a perpetrator deciding whether to extract resources from a fixed location $\mathbf{x} \in \Omega$. As in [4], our approach is based on “pointwise” extraction decisions, but since these resources are typically continuously distributed in the protected area Ω , this requires further clarification. We will consider a fine grid \mathcal{X} of possible extraction sites, with each $\mathbf{x} \in \mathcal{X}$ representing a center of its extraction cell $N_{\mathbf{x}}$, and $B(\mathbf{x})$ representing the value of all resources that can be extracted from that cell. The same grid \mathcal{X} will also be used to solve the PDEs numerically in section 5.

We will further assume the following:

1. The extractor can enter Ω through any part of $\partial\Omega$ and has no reason to avoid surveillance from PAMs until he reaches the extraction site \mathbf{x} .
2. While traveling through Ω , the extractor may freely choose any measurable “control function” $\mathbf{a} : \mathbb{R} \rightarrow S^1$, specifying his chosen direction of motion for all times t . This defines the corresponding trajectory since his (isotropic) speed of motion $f : \Omega \rightarrow \mathbb{R}_+$ reflects the local terrain and is a part of the input data.
3. The time/difficulty cost of traveling through Ω is reflected by integrating some known running cost $K : \Omega \rightarrow \mathbb{R}_+$ along his chosen trajectory. One natural choice is $K(\mathbf{x}) = \kappa$, reflecting the monetary value the extractor associates with each hour of his time spent on a trip. (To normalize the units, all of our numerical examples use $\kappa = 1$ for all pre-extraction and post-extraction travel through Ω .) But a non-constant K can also be used to reflect the extractor's distaste for some parts of the protected area (e.g., due to vegetation).
4. Once he reaches \mathbf{x} , the extraction in $N_{\mathbf{x}}$ is instantaneous and the extractor immediately starts along his (possibly different) path to the boundary.
5. On the way back to $\partial\Omega$, he is trying to minimize both the time/cost of travel and the probability of being spotted by PAMs through aerial surveillance.
6. The extractor has full prior knowledge of PAMs' location-dependent detection rate $\psi : \Omega \rightarrow \mathbb{R}_{+,0}$.
7. Rational extractors will target all sites $\mathbf{x} \in \mathcal{X}$ for which the expected profit $P(\mathbf{x})$ of their optimal roundtrip is above some threshold level \tilde{p} . This leaves a pristine area

$$\Omega_p = \{\mathbf{x} \mid P(\mathbf{x}) \leq \tilde{p}\}.$$

(Our numerical experiments use $\tilde{p} = 0$.) Note that this is different from the approach in [4] not only in providing an absolute threshold but also because we don't exclude from Ω_p any corridors around post-extraction trajectories.

We believe that this is more consistent with the notion of localized extractions.

Computing the expected profit $P(\mathbf{x})$ requires solving two different optimization problems: (a) finding the best (PAMs-ignoring) pre-extraction trajectory from $\partial\Omega$ to \mathbf{x} and then (b) finding the best post-extraction trajectory from \mathbf{x} to $\partial\Omega$, balancing the cost of travel and risk of detection. We will start by addressing (a) in section 3.1 since it is easy to solve by standard tools of *single-criterion deterministic optimal control theory*. Below we provide only a very brief introduction, referring to standard references (e.g., [10] and [5]) for technical details.

We will then use two models to address (b) under very different assumptions about the nature of PAMs' enforcement activities and the information available to perpetrators:

Model G (ground patrols): If the extractor is detected while traveling with protected resources, he is immediately apprehended and the resources are confiscated. From there on, the perpetrator does not attempt extracting anything on this trip and simply seeks to minimize the integral of K on his way to $\partial\Omega$.

Model A (aerial patrols): The extractor has no way of learning if or when the detection occurs. But if it happens, he is later apprehended with probability one upon reaching $\partial\Omega$ and his entire haul is confiscated.

We will show that these two interpretations require fairly different mathematical tools: Model G is best recast as an optimal control of a randomly terminated process (subsection 3.2), while Model A requires the techniques from multiobjective dynamic programming (subsection 3.3). The numerical methods for both versions are covered in section 5.

3.1. Getting to the extraction point. We will use $R(\mathbf{x})$ to denote the cost of the optimal pre-extraction trajectory from $\partial\Omega$ to \mathbf{x} . But to make the discussion closer to the canonical optimal control setting, we will reverse the direction of the trajectory in (a), looking instead for a K -optimal way to reach $\partial\Omega$ from \mathbf{x} . (Based on our isotropy assumptions, this direction reversal will change neither the time of travel nor the integral of K along any considered trajectory.) The added benefit is that the same discussion will be a useful starting point for solving (b).

The extractor's time-dependent position in the domain will be specified by

$$(3.1) \quad \mathbf{y}'(t) = f(\mathbf{y}(t)) \mathbf{a}(t), \quad \mathbf{y}(0) = \mathbf{x}.$$

Once $\mathbf{a}(\cdot)$ is chosen, this defines the exit time $T_{\mathbf{x}, \mathbf{a}(\cdot)} = \min\{t \geq 0 \mid \mathbf{y}(t) \in \partial\Omega\}$ and the cumulative cost of starting at \mathbf{x} and using $\mathbf{a}(\cdot)$:

$$\mathcal{J}(\mathbf{x}, \mathbf{a}(\cdot)) = \int_0^{T_{\mathbf{x}, \mathbf{a}(\cdot)}} K(\mathbf{y}(t)) dt.$$

If the extractor only cares about selecting a path to minimize this \mathcal{J} , it is easy to accomplish by traditional tools of single-criterion dynamic programming introduced by Richard Bellman in the 1950s.

The *value function* $u(\mathbf{x})$ is defined as the minimum cost one has to pay starting from \mathbf{x} , i.e., $u(\mathbf{x}) = \inf_{\mathbf{a}(\cdot)} \mathcal{J}(\mathbf{x}, \mathbf{a}(\cdot))$. A Taylor series expansion along the optimal trajectory can be used to derive an HJB PDE that u must satisfy if it is sufficiently smooth. For our isotropic cost and dynamics, this PDE is merely an Eikonal equation

$$(3.2) \quad |\nabla u(\mathbf{x})| f(\mathbf{x}) = K(\mathbf{x})$$

solved with Dirichlet boundary conditions $u = 0$ on $\partial\Omega$. Unfortunately, in general (3.2) does not have a classical/smooth solution, and there are infinitely many weak (Lipschitz-continuous) solutions. Additional test conditions were introduced by Crandall and Lions [15] to select among them the unique *viscosity solution* coinciding with the value function of the above control problem. Several highly efficient methods for solving Eikonal PDEs numerically were developed in the past 25 years; see a brief overview at the end of section 5.

We note that taking $K = 1$ results in $J(\mathbf{x}, \mathbf{a}(\cdot)) = T_{\mathbf{x}, \mathbf{a}(\cdot)}$. So, both the distance-from-boundary $d(\mathbf{x})$ and the min-time-from-boundary $\tau(\mathbf{x})$ can also be found by solving the Eikonal equations $|\nabla d| = 1$ and $|\nabla \tau|f = 1$ with zero boundary conditions. If $K = \kappa$, then the minimum cost of reaching \mathbf{x} from $\partial\Omega$ is simply

$$R(\mathbf{x}) = \kappa\tau(\mathbf{x}).$$

The characteristic curves of HJB equations are the optimal trajectories, which in the Eikonal case coincide with the gradient lines of the viscosity solution. That is, once the value function is computed, an optimal trajectory can be found via gradient descent in u . Since this function is differentiable almost everywhere, the optimal trajectories are unique for almost all $\mathbf{x} \in \Omega$.

Remark 3.1. We note that the function \tilde{C} used in [4] can also be found directly (without using a level set formulation) by solving an Eikonal equation $|\nabla \tilde{C}|F = 1$. While the level set method is a much more general interface tracking approach, the fast Eikonal solvers are likely to be more efficient for the current application.

3.2. Returning with the loot (Model G: Ground patrols). Here we assume that a detection leads to an immediate apprehension and confiscation of extracted resources. Thus, it is natural to split the post-extraction trajectory into two parts: a pre-detection trajectory (leading all the way to $\partial\Omega$ if the detection was avoided) and a post-detection trajectory (chosen by a perpetrator to simply minimize the integral of K up to $\partial\Omega$ after the confiscation). Suppose the perpetrator starts from an extraction point \mathbf{x}_0 and manages to reach some $\mathbf{x} \in \Omega$ undetected. From there on, he uses a control $\mathbf{a}(\cdot)$, moving from $\mathbf{x} = \mathbf{y}(0)$ along a trajectory $\mathbf{y}(t)$ with a possible detection at some time T before reaching the boundary. His remaining cost is

$$\mathcal{J}(\mathbf{x}, \mathbf{a}(\cdot), T) = \begin{cases} \int_0^T K(\mathbf{y}(t)) dt + B(\mathbf{x}_0) + R(\mathbf{y}(T)) & \text{if } T < T_{\mathbf{x}, \mathbf{a}(\cdot)} \\ & \text{(i.e., if detected),} \\ \int_0^{T_{\mathbf{x}, \mathbf{a}(\cdot)}} K(\mathbf{y}(t)) dt & \text{otherwise.} \end{cases}$$

Of course, the time of capture T is random and its distribution depends on the chosen pre-detection trajectory. So, the perpetrator's value function is defined to minimize the expected remaining cost (from \mathbf{x} to $\partial\Omega$):

$$\bar{u}(\mathbf{x}) = \inf_{\mathbf{a}(\cdot)} \mathbb{E}_T [\mathcal{J}(\mathbf{x}, \mathbf{a}(\cdot), T)].$$

We note that it also depends on the value of extracted resources, $b = B(\mathbf{x}_0)$, but we leave this implicit to simplify the notation.

To obtain the Bellman optimality condition, we assume that the detection will not happen for the next s seconds but might occur after that period with probability $s\psi(\mathbf{x})$. This yields

$$\bar{u}(\mathbf{x}) = \inf_{\mathbf{a}(\cdot)} \left\{ \int_0^s K(\mathbf{y}(t)) dt + s\psi(\mathbf{x}) [b + R(\mathbf{y}(s))] + (1 - s\psi(\mathbf{x})) \bar{u}(\mathbf{y}(s)) \right\} + o(s).$$

Assuming that \bar{u} is sufficiently smooth, a Taylor expansion of the above yields a Hamilton–Jacobi PDE,

$$(3.3) \quad \begin{aligned} |\nabla \bar{u}(\mathbf{x})|f(\mathbf{x}) &= K(\mathbf{x}) + \psi(\mathbf{x})(b + R(\mathbf{x}) - \bar{u}(\mathbf{x})) && \text{on } \Omega; \\ \bar{u}(\mathbf{x}) &= 0 && \text{on } \partial\Omega. \end{aligned}$$

Since this PDE generally does not have a classical/smooth solution, that derivation is only formal, but viscosity solution theory [5, 15] allows us to pick the weak solution coinciding with the value function. The above derivation is similar to what is used in optimal control of *randomly terminated process*. Such processes arise in production/maintenance planning [9], economic growth and global climate change modeling [22], multigenerational games [21], and optimizing the routing of emergency response vehicles [3]. In the current context, “termination” is either “detection/capture” (since we already know the cost of optimal actions after that point) or “reaching $\partial\Omega$ undetected.” The latter possibility yields Dirichlet boundary conditions in (3.3) and avoids the issues related to “free boundary” and quasi-variational inequalities considered in [3].

Once \bar{u} is found, the optimal pre-detection trajectories can be found by gradient descent. The extractor’s expected profit can also be computed as

$$P^g(\mathbf{x}_0) = b - \bar{u}(\mathbf{x}_0) - R(\mathbf{x}_0) \quad \forall \mathbf{x}_0 \in X \text{ such that } B(\mathbf{x}_0) = b,$$

where $R(\mathbf{x}_0)$ accounts for the cost of a pre-extraction trajectory.

We note that it is easy to find an upper bound on \bar{u} by committing to use some $\mathbf{a}(\cdot)$ up until $\partial\Omega$ even if the detection occurs earlier. If that $\mathbf{a}(\cdot)$ is selected to minimize the integral of K , this yields $\bar{u}(\mathbf{x}_0) \leq b + R(\mathbf{x}_0)$, but a sharper bound is obtained by choosing a control optimal for Model A considered below. This is why $P^g \geq P^a$ and $\Omega_p^g \subset \Omega_p^a$.

3.3. Returning with the loot (Model A: Aerial patrols). On the way back to $\partial\Omega$, the extractor needs to balance the cumulative cost of the path with expected losses resulting from a possible detection by PAMs. This is even more important under the current model, where the extractor does not learn whether he was detected until reaching $\partial\Omega$. Here we formally derive the probability that an extractor is not detected while traversing a path $\mathbf{y} : [0, T] \rightarrow \Omega$. Suppose the probability of detection at a location \mathbf{x} over a small time interval Δt is $\psi(\mathbf{y})\Delta t$. Then the probability \mathcal{P}_N of making it through N such consecutive intervals of length $\Delta t = T/N$ without ever being detected is

$$\mathcal{P}_N = (1 - \psi(\mathbf{y}_1)\Delta t)(1 - \psi(\mathbf{y}_2)\Delta t) \cdots (1 - \psi(\mathbf{y}_N)\Delta t) = \prod_{i=1}^N (1 - \psi(\mathbf{y}_i)\Delta t).$$

Taking the logarithm of both sides,

$$\log(\mathcal{P}_N) = \sum_{i=1}^N \log(1 - \psi(\mathbf{y}_i)\Delta t) = - \sum_{i=1}^N \left(\psi(\mathbf{y}_i)\Delta t + \mathcal{O}(\Delta t^2) \right).$$

Taking the limit as $N \rightarrow \infty$, the probability of not being detected along this path until at least the time T is

$$(3.4) \quad \mathcal{P} = e^{- \int_0^T \psi(\mathbf{y}(t))dt}.$$

Our emphasis on careful modeling of the cumulative probability of detection $\Psi = 1 - \mathcal{P}$ is perhaps the main distinction from earlier models. Both [1] and [4]

implicitly assume $\Psi \approx \int_0^T \psi(\mathbf{y}(t))dt$. As a first-order approximation of a concave function ($1 - e^{-\xi} = \xi + O(\xi^2)$), this overestimates Ψ for every trajectory. It also yields a simplified optimization problem, but with significant limitations described in Remark 3.2. Instead, we find the extractor's expected profit by using the accurate \mathcal{P} to optimize his post-extraction trajectory:

$$(3.5) \quad \begin{aligned} \mathbf{y}'(t) &= f(\mathbf{y})\mathbf{a}(t), \quad \mathbf{y}(0) = \mathbf{x}, \\ T_{\mathbf{x}, \mathbf{a}(\cdot)} &= \min\{t \geq 0 \mid \mathbf{y}(t) \in \partial\Omega\}, \\ \mathcal{J}_1(\mathbf{x}, \mathbf{a}(\cdot)) &= \int_0^{T_{\mathbf{x}, \mathbf{a}(\cdot)}} \psi(\mathbf{y}(t))dt, \\ \mathcal{J}_2(\mathbf{x}, \mathbf{a}(\cdot)) &= \int_0^{T_{\mathbf{x}, \mathbf{a}(\cdot)}} K(\mathbf{y}(t))dt, \\ P^a(\mathbf{x}) &= \sup_{\mathbf{a}(\cdot) \in \mathcal{A}} \left\{ B(\mathbf{x})e^{-\mathcal{J}_1(\mathbf{x}, \mathbf{a}(\cdot))} - \mathcal{J}_2(\mathbf{x}, \mathbf{a}(\cdot)) \right\} - R(\mathbf{x}), \end{aligned}$$

where \mathcal{A} is the set of *admissible* controls, i.e., the set of measurable functions from \mathbb{R} to S^1 , and $R(\mathbf{x})$ is the optimal cost of a pre-extraction path defined in subsection 3.1.

Remark 3.2. If $\mathcal{J}_1(\mathbf{x}_0, \mathbf{a}(\cdot)) \ll 1$ for all \mathbf{x}_0 and all controls $\mathbf{a}(\cdot)$ close to optimizing (3.5), a linear approximation of the detection probability $\Psi \approx \mathcal{J}_1$ is reasonable and leads to a significantly simpler optimization problem similar to those considered in [1] and [4]. In that case,

$$(3.6) \quad P^a(\mathbf{x}) \approx B(\mathbf{x}) - \inf_{\mathbf{a}(\cdot) \in \mathcal{A}} \left\{ B(\mathbf{x})\mathcal{J}_1(\mathbf{x}, \mathbf{a}(\cdot)) + \mathcal{J}_2(\mathbf{x}, \mathbf{a}(\cdot)) \right\} - R(\mathbf{x}).$$

Since the linearized approach exaggerates Ψ , we know that this estimate is actually a lower bound on $P^a(\mathbf{x})$. The minimizing control in (3.6) can be found by solving the Eikonal PDE

$$(3.7) \quad |\nabla u|f(\mathbf{x}) = B(\mathbf{x}_0)\psi(\mathbf{x}) + K(\mathbf{x})$$

with zero boundary conditions. If one prefers the level set formulation, this is equivalent to using the interface speed of

$$F(\mathbf{x}, b) = f(\mathbf{x}) / (K(\mathbf{x}) + b\psi(\mathbf{x})),$$

which is similar to but not quite the same as the formula (2.5) used in [4] — even in the case of risk neutral extractors (i.e., $\alpha = 1$) concerned with their time along the trajectory (i.e., $K(\mathbf{x}) = 1$).

Unfortunately, \mathcal{J}_1 can be arbitrarily large, which makes this “ Ψ -linearization” unsuitable for most realistic situations. (Indeed, in our numerical examples of section 6, $\mathcal{J}_1 > 1$ is often observed along many optimal paths.) This is the key reason why the more complicated multiobjective path-planning methods considered below are actually needed.

Remark 3.3. The nonlinear dependence on \mathcal{J}_1 is the reason why P^a cannot be computed by standard tools of single-objective optimal control on Ω . But this actually can be accomplished if we are willing to increase the dimension of our planning space. We note that \mathcal{P} satisfies an ODE $\mathcal{P}'(t) = -\psi(\mathbf{y}(t))\mathcal{P}(t)$ along the chosen trajectory, with $\mathcal{P}(0) = 1$, $\mathbf{y}(0) = \mathbf{x}_0$. One could define a new value function $w(\mathbf{x}, p)$ as the minimum expected cost of the remaining path to $\partial\Omega$ if the extractor has already

moved from \mathbf{x}_0 to \mathbf{x} and the probability that he has not been detected so far is p . A formal control-theoretic argument shows that w should satisfy the following Hamilton–Jacobi PDE on $\Omega \times (0, 1]$:

$$(3.8) \quad \psi(\mathbf{x})p \frac{\partial w}{\partial p} + |\nabla_{\mathbf{x}} w|f(\mathbf{x}) = K(\mathbf{x})$$

with boundary conditions $w(\mathbf{x}, p) = (1 - p)b$ on $\partial\Omega \times (0, 1]$. Once w is computed, the expected profit can be found as $P^a(\mathbf{x}_0) = b - w(\mathbf{x}_0, 1) - P_0(\mathbf{x}_0)$ for every \mathbf{x}_0 such that $B(\mathbf{x}_0) = b$. The need to solve this PDE on a higher dimensional domain for many b values would make this approach computationally expensive, with additional numerical difficulties due to that fact that the coefficient $(\psi(\mathbf{x})p)$ can be arbitrarily small or even zero. This is why we opt instead to use multiobjective dynamic programming on the original planning space Ω .

4. Multiobjective approach (Model A).

For convenience, we will define

$$\mathcal{F}(\mathcal{J}_1, \mathcal{J}_2) = Be^{-\mathcal{J}_1} - \mathcal{J}_2$$

and, in a slight abuse of notation, drop the first argument from both \mathcal{J}_1 and \mathcal{J}_2 (since \mathbf{x} is fixed throughout this section) and also drop the argument from $\mathbf{a}(\cdot)$.

Ideally, one would prefer to find a control minimizing both \mathcal{J}_1 and \mathcal{J}_2 simultaneously. Somewhat surprisingly this is actually possible when K is constant and PAMs use “ τ -banded” patrol densities, i.e., when ψ is really a function of the min-time-from- $\partial\Omega$ $\tau(\mathbf{x})$ defined in section 3. It is not hard to show that in this case the trajectories recovered from the Eikonal (3.2) will also minimize the probability of detection. When f is also constant, such trajectories will be simply straight lines (from the extraction point \mathbf{x} to the closest point on $\partial\Omega$), which will also minimize the probability of detection for all “distance-to- $\partial\Omega$ -banded” patrol strategies $\psi(d)$ (including the homogeneous $\psi(\mathbf{x}) = \psi_0$) considered in both [1] and [4]. This is also the setting where there is no difference between the predictions based on Model A and Model G — even if the loot is confiscated, there is no incentive to switch to another trajectory.

For the above scenarios the extractor’s path-planning is fairly trivial. But in the general case, these two minimization criteria are in conflict, and one has to switch to a notion of *Pareto optimality*.

4.1. Pareto front. We will say that a control \mathbf{a}_1 *dominates* another control \mathbf{a}_2 if $\mathcal{J}_1(\mathbf{a}_1) \leq \mathcal{J}_1(\mathbf{a}_2)$ and $\mathcal{J}_2(\mathbf{a}_1) \leq \mathcal{J}_2(\mathbf{a}_2)$ with at least one of those inequalities strict. All \mathbf{a} ’s not dominated by any other control are called *Pareto optimal*. They correspond to trajectories that cannot be improved with respect to both criteria simultaneously. It will be useful to consider the *Pareto front*

$$PF = \{(\mathcal{J}_1(\mathbf{a}), \mathcal{J}_2(\mathbf{a})) \mid \mathbf{a} \text{ is Pareto optimal}\}.$$

In addition to control functions, the extractor could in principle also consider mixed/probabilistic strategies, which specify a probability distribution over a set of available controls. A strategy selecting any control with probability one is usually called pure or deterministic. Since \mathcal{F} is decreasing in both arguments, it follows that $\mathcal{F}(\mathcal{J}_1(\mathbf{a}_1), \mathcal{J}_2(\mathbf{a}_1)) > \mathcal{F}(\mathcal{J}_1(\mathbf{a}_2), \mathcal{J}_2(\mathbf{a}_2))$ for any control \mathbf{a}_1 dominating \mathbf{a}_2 . From this we see that any optimizer of (3.5) over the set of pure strategies must belong to PF. We also note that the extractor has no incentive to consider mixed strategies, since the expected payoff for using strategy \mathbf{a}_1 with probability θ and \mathbf{a}_2 with probability $1 - \theta$ is

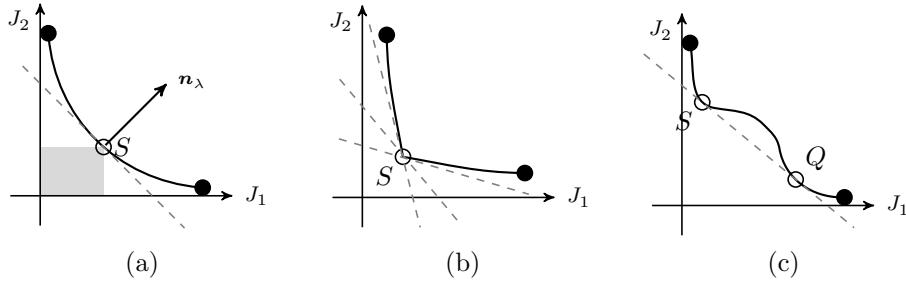


FIG. 1. Common Pareto front scenarios. (a) Convex smooth PF with a point S corresponding to some specific λ . The dashed line perpendicular to $\mathbf{n}_\lambda = (\lambda, 1 - \lambda)$ is tangent to PF at S . If any path yielded a (J_1, J_2) point below this line, the path corresponding to S would not be λ -optimal. Since the gray rectangle (comprising of points dominating S) lies below that line, λ -optimality implies Pareto optimality for all $\lambda \in (0, 1)$. (b) Convex nonsmooth PF with a “kink” at the point S makes the corresponding path λ -optimal for a range of λ ’s, with a different “support hyperplane” corresponding to each of them. (c) Nonconvex smooth PF. Points S and Q correspond to two different λ -optimal paths. The portion of PF between S and Q cannot be found by scalarization.

$$\theta \mathcal{F}(\mathcal{J}_1(\mathbf{a}_1), \mathcal{J}_2(\mathbf{a}_1)) + (1 - \theta) \mathcal{F}(\mathcal{J}_1(\mathbf{a}_2), \mathcal{J}_2(\mathbf{a}_2)) \leq \max_{i=1,2} \mathcal{F}(\mathcal{J}_1(\mathbf{a}_i), \mathcal{J}_2(\mathbf{a}_i)).$$

So in practice, we will compute $P^a(\mathbf{x})$ by maximizing $\mathcal{F}(\mathcal{J}_1, \mathcal{J}_2)$ over the $(\mathcal{J}_1, \mathcal{J}_2) \in PF$.

There are a number of techniques for finding Pareto optimal controls including scalarization on Ω [28] and other methods relying on augmented PDEs on an expanded state space [17, 26]. To scalarize, one chooses some $\lambda \in [0, 1]$ and optimizes a convex combination $\mathcal{J}^\lambda = \lambda \mathcal{J}_1 + (1 - \lambda) \mathcal{J}_2$. The resulting λ -optimal trajectory is also Pareto optimal (see Figure 1), and the procedure is applied repeatedly for different λ values. Unfortunately, such scalarization only finds points on the convex part of the Pareto front [16]. This is a serious disadvantage since nonconvexities of PF are both common and important in many control applications including robotic navigation [14, 26]. In other cases, such as our recent work on surveillance-evasion games [11, 20], special features of the problem might make scalarization sufficient. Below we explain why this is also the case for the post-extraction path-planning.

OBSERVATION. *Any global maximizer of (3.5) will correspond to a point on the convex part of PF.*

Proof. Suppose a control \mathbf{a} corresponds to a point $(\mathcal{J}_1(\mathbf{a}), \mathcal{J}_2(\mathbf{a}))$ on a nonconvex section of PF. Then there exist controls $\mathbf{a}_1, \mathbf{a}_2$ and a constant $0 \leq \theta \leq 1$ such that

$$\begin{aligned} \theta \mathcal{J}_1(\mathbf{a}_1) + (1 - \theta) \mathcal{J}_1(\mathbf{a}_2) &\leq \mathcal{J}_1(\mathbf{a}), \\ \theta \mathcal{J}_2(\mathbf{a}_1) + (1 - \theta) \mathcal{J}_2(\mathbf{a}_2) &\leq \mathcal{J}_2(\mathbf{a}) \end{aligned}$$

with at least one of the inequalities strict. Since \mathcal{F} is decreasing and convex in both arguments,

$$\begin{aligned} \mathcal{F}(\mathcal{J}_1(\mathbf{a}), \mathcal{J}_2(\mathbf{a})) &< \mathcal{F}\left(\theta \mathcal{J}_1(\mathbf{a}_1) + (1 - \theta) \mathcal{J}_1(\mathbf{a}_2), \theta \mathcal{J}_2(\mathbf{a}_1) + (1 - \theta) \mathcal{J}_2(\mathbf{a}_2)\right) \\ &\leq \theta \mathcal{F}\left(\mathcal{J}_1(\mathbf{a}_1), \mathcal{J}_2(\mathbf{a}_1)\right) + (1 - \theta) \mathcal{F}\left(\mathcal{J}_1(\mathbf{a}_2), \mathcal{J}_2(\mathbf{a}_2)\right) \\ &\leq \max_{i=1,2} \{\mathcal{F}(\mathcal{J}_1(\mathbf{a}_i), \mathcal{J}_2(\mathbf{a}_i))\}. \end{aligned}$$

Thus, such \mathbf{a} cannot be a maximizer in (3.5). \square

4.2. Scalarized Eikonal. For any $\lambda \in [0, 1]$, we define the scalarized running cost K^λ as

$$(4.1) \quad K^\lambda(\mathbf{x}) = \lambda\psi(\mathbf{x}) + (1 - \lambda)K(\mathbf{x}).$$

A standard optimal control argument [5] shows that the value function

$$(4.2) \quad u^\lambda(\mathbf{x}) = \inf_{\mathbf{a}(\cdot) \in \mathcal{A}} \left\{ \lambda\mathcal{J}_1(\mathbf{x}, \mathbf{a}(\cdot)) + (1 - \lambda)\mathcal{J}_2(\mathbf{x}, \mathbf{a}(\cdot)) \right\}$$

is the unique viscosity solution [15] of the scalarized Eikonal equation:

$$(4.3) \quad \begin{aligned} f(\mathbf{x}) |\nabla u^\lambda(\mathbf{x})| &= K^\lambda(\mathbf{x}), & \mathbf{x} \in \Omega, \\ u^\lambda(\mathbf{x}) &= 0, & \mathbf{x} \in \partial\Omega. \end{aligned}$$

For every specific λ , we can recover the corresponding λ -optimal paths by gradient descent in u^λ . We note that taking $\lambda = 0$ corresponds to detection-indifferent optimization, yielding the PDE (3.2), which we already needed to solve to find R along with the optimal pre-extraction path from $\partial\Omega$ to each site \mathbf{x}_0 . The case $\lambda = 1$ yields the PDE for recovering the time-indifferent least-detectable paths.

Remark 4.1. We note that a proper choice of λ can also represent the PDE model based on the “ $\Psi \approx \mathcal{J}_1$ ” approximation discussed in Remark 3.2. In particular, for $\lambda_\sharp = B(\mathbf{x}_0)/(B(\mathbf{x}_0) + 1)$, (4.2) can be rewritten as

$$(B(\mathbf{x}_0) + 1) f(\mathbf{x}) |\nabla u^{\lambda_\sharp}(\mathbf{x})| = B(\mathbf{x}_0)\psi(\mathbf{x}) + K(\mathbf{x}).$$

As a result, $(B(\mathbf{x}_0) + 1) u^{\lambda_\sharp}(\mathbf{x})$ is the viscosity solution of (3.7). When $\mathcal{J}_1 \ll 1$, this yields “nearly optimal” post-extraction paths. Based on (3.6), we also know that

$$(4.4) \quad P^a(\mathbf{x}_0) \geq P_\sharp(\mathbf{x}_0) = B(\mathbf{x}_0) - (B(\mathbf{x}_0) + 1) u^{\lambda_\sharp}(\mathbf{x}_0) - R(\mathbf{x}_0).$$

4.3. Auxiliary PDEs. However, we are ultimately interested in the expected payoff (3.5), not the value function $u^\lambda(\mathbf{x})$. Under suitable regularity assumptions on the speed f and running cost K , it can be shown that an optimal control $\mathbf{a}(\cdot)$ exists for every $\mathbf{x} \in \Omega$, and the infimum in formula (4.2) can be replaced with a minimum [5]. We define the set of λ -optimal controls $\mathcal{A}_\mathbf{x}^\lambda$ starting from $\mathbf{x} \in \Omega$ as

$$(4.5) \quad \mathcal{A}_\mathbf{x}^\lambda = \arg \min_{\mathbf{a}(\cdot) \in \mathcal{A}} \left\{ \lambda\mathcal{J}_1(\mathbf{x}, \mathbf{a}(\cdot)) + (1 - \lambda)\mathcal{J}_2(\mathbf{x}, \mathbf{a}(\cdot)) \right\}.$$

We note that the optimal control $\mathbf{a}(\cdot)$ will be unique at every point \mathbf{x} where $u^\lambda(\mathbf{x})$ is differentiable. Given our conditions on f and K , $u^\lambda(\mathbf{x})$ is Lipschitz-continuous, and so $\mathcal{A}_\mathbf{x}^\lambda$ will be a singleton for almost every $\mathbf{x} \in \Omega$.

We will now define the λ -optimal-restricted value functions v_1^λ and v_2^λ as

$$(4.6) \quad v_1^\lambda(\mathbf{x}) = \inf_{\mathbf{a}(\cdot) \in \mathcal{A}_\mathbf{x}^\lambda} \{ \mathcal{J}_1(\mathbf{x}, \mathbf{a}(\cdot)) \},$$

$$(4.7) \quad v_2^\lambda(\mathbf{x}) = \inf_{\mathbf{a}(\cdot) \in \mathcal{A}_\mathbf{x}^\lambda} \{ \mathcal{J}_2(\mathbf{x}, \mathbf{a}(\cdot)) \}.$$

Given the solution $u^\lambda(\mathbf{x})$ of (4.3), we can recover v_1^λ and v_2^λ by solving the following system of auxiliary linear PDEs introduced in [28]:

$$(4.8) \quad \begin{aligned} \nabla v_1^\lambda(\mathbf{x}) \cdot \nabla u^\lambda(\mathbf{x}) &= \frac{\psi(\mathbf{x}) K^\lambda(\mathbf{x})}{f(\mathbf{x})^2}, \\ \nabla v_2^\lambda(\mathbf{x}) \cdot \nabla u^\lambda(\mathbf{x}) &= \frac{K(\mathbf{x}) K^\lambda(\mathbf{x})}{f(\mathbf{x})^2} \end{aligned}$$

with $v_1^\lambda = v_2^\lambda = 0$ on $\partial\Omega$.

Since we have shown in section 4.1 that the global optimizer of P lies on the convex part of the $(\mathcal{J}_1, \mathcal{J}_2)$ Pareto front, it can be recovered through scalarization. More precisely, the optimal expected payoff satisfies

$$P^a(\mathbf{x}) = \sup_{\lambda \in [0,1]} \left\{ B(\mathbf{x}) e^{-v_1^\lambda(\mathbf{x})} - v_2^\lambda(\mathbf{x}) \right\} - R(\mathbf{x}),$$

where the optimal λ will be usually \mathbf{x} -dependent.

5. Numerical methods. Since the domain Ω is meant to represent some protected area such as a national park, it will in general be irregularly shaped. One way of handling this would be to discretize Ω directly through the use of $\partial\Omega$ -conforming triangulated meshes. Efficient methods for numerically solving (4.3) on such meshes are already known [7, 30, 35]. For this paper, we will instead use a Cartesian grid \mathcal{X} to discretize a rectangular region in \mathbb{R}^2 containing Ω . Without loss of generality, suppose that Ω is contained within a rectangle $[0, X_{max}] \times [0, Y_{max}]$. Our grid spacing will be $\Delta x = X_{max}/N_x$ in the x -direction and $\Delta y = Y_{max}/N_y$ in the y -direction. The coordinates of our gridpoints will be $x_i = i\Delta x$ for $i = 0, 1, \dots, N_x$ and $y_j = j\Delta y$ for $j = 0, 1, \dots, N_y$. For Model A, we will also discretize λ as $\lambda_k = k/N_\lambda$ for $k = 0, 1, \dots, N_\lambda$. For Model G, we will also discretize b as $b_m = \hat{B} + m(\check{B} - \hat{B})/N_b$, with $\hat{B} = \min_{\mathbf{x}} B(\mathbf{x})$, $\check{B} = \max_{\mathbf{x}} B(\mathbf{x})$, and $m = 0, 1, \dots, N_b$.

5.1. Scalarized Eikonal equation (Model A). We will approximate our value function with

$$U_{i,j}^k \approx u^{\lambda_k}(x_i, y_j)$$

and our λ -optimal-restricted value functions with

$$V_{i,j}^{k,1} \approx v_1^{\lambda_k}(x_i, y_j), \quad V_{i,j}^{k,2} \approx v_2^{\lambda_k}(x_i, y_j).$$

For each λ_k value, we will need to solve the scalarized Eikonal equation (4.3) and the coupled system of auxiliary linear PDEs (4.8) simultaneously. To approximate the gradients of the value functions, we will make heavy use of the following first-order difference operators:

$$\begin{aligned} D_{i,j}^{+x}[W] &= \frac{W_{i+1,j} - W_{i,j}}{\Delta x}, & D_{i,j}^{-x}[W] &= \frac{W_{i,j} - W_{i-1,j}}{\Delta x}, \\ D_{i,j}^{+y}[W] &= \frac{W_{i,j+1} - W_{i,j}}{\Delta y}, & D_{i,j}^{-y}[W] &= \frac{W_{i,j} - W_{i,j-1}}{\Delta y}, \end{aligned}$$

where W could represent any of the value function approximations U^k , $V^{k,1}$, or $V^{k,2}$.

Since the scalarized Eikonal equation (4.3) and the auxiliary PDEs (4.8) are solved along the same characteristics, we will make sure that our approximations of the gradients are always using the same stencils. The upwind difference operator \mathcal{D}_{ij}^x in the x -direction will be

$$\mathcal{D}_{i,j}^x[W] = \begin{cases} D_{i,j}^{+x}[W] & \text{if } D_{i,j}^{+x}[U^k] \leq \min(-D_{i,j}^{-x}[U^k], 0); \\ D_{i,j}^{-x}[W] & \text{if } -D_{i,j}^{-x}[U^k] < \min(D_{i,j}^{+x}[U^k], 0); \\ 0 & \text{otherwise;} \end{cases}$$

with $\mathcal{D}_{i,j}^y$ defined accordingly for the y -direction. Note that the stencil being used will always be determined by U^k regardless of whether \mathcal{D} is being applied to U^k , $V^{k,1}$, or $V^{k,2}$.

Using these first-order difference operators, we can write the first-order upwind discretization of our system of PDEs as

$$(5.1) \quad \begin{aligned} f_{i,j} \left[\left(\mathcal{D}_{i,j}^x[U^k] \right)^2 + \left(\mathcal{D}_{i,j}^y[U^k] \right)^2 \right]^{1/2} &= K_{i,j}^{\lambda_k}, \\ \mathcal{D}_{i,j}^x[U^k] \cdot \mathcal{D}_{i,j}^x[V^{k,1}] + \mathcal{D}_{i,j}^y[U^k] \cdot \mathcal{D}_{i,j}^y[V^{k,1}] &= \frac{\psi_{i,j} K_{i,j}^{\lambda_k}}{f_{i,j}^2}, \\ \mathcal{D}_{i,j}^x[U^k] \cdot \mathcal{D}_{i,j}^x[V^{k,2}] + \mathcal{D}_{i,j}^y[U^k] \cdot \mathcal{D}_{i,j}^y[V^{k,2}] &= \frac{K_{i,j} K_{i,j}^{\lambda_k}}{f_{i,j}^2} \end{aligned}$$

with the boundary conditions

$$U_{i,j}^k = V_{i,j}^{k,1} = V_{i,j}^{k,2} = 0 \quad \text{for } (x_i, y_j) \notin \Omega.$$

Since this is a monotone and consistent finite difference approximation of (4.3), U^k will converge under grid refinement to the viscosity solution u^{λ_k} [6]. But on any fixed grid, finding U^k values is not straightforward since the above discretized system is coupled and nonlinear. On a grid with M gridpoints, straightforward Gauss–Jacobi iterations result in a high computational cost of $O(M^2)$, and Gauss–Seidel iterations with any fixed ordering of gridpoints are similarly inefficient.

A significant speed-up is gained by exploiting the causal properties of upwind difference operators. Based on the definition of \mathcal{D}^x and \mathcal{D}^y , each $U_{i,j}^k$ depends only on the *smaller* neighboring U^k values. This *monotone causality* allows for the use of noniterative Dijkstra-like “fast marching” methods [33, 34, 38], which compute U^k in $O(M \log M)$ operations. They mimic the logic of a classical shortest path algorithm on graphs [18] and effectively decouple this discretized system by determining the correct ordering of gridpoints/equations on the fly.

Another popular approach for efficient Eikonal solvers is to use Gauss–Seidel iterations, but alternating through a list of direction-aligned grid orderings (from SW, SE, NE, and NW). Such “fast sweeping” techniques [8, 37, 39] aim to guess the direction of characteristics for a part of the domain, recovering many correct U^k values in each iteration. These methods converge in $O(\omega M)$ operations, but the factor ω is not known a priori and depends on the number of times the characteristics change direction from quadrant to quadrant. Both classes of methods have their preferred sets of problems. For example, fast sweeping is clearly better when characteristics are largely straight lines, while fast marching is typically preferable on domains with complex geometry or with rapidly changing PDE coefficients. More recent hybrid (two-scale) methods aim to combine the best features of marching and sweeping. We refer to [12, 13] for a detailed description, benchmarking, and comparison with other sequential and parallel Eikonal solvers.

While our own implementation is based on fast marching, other methods mentioned above can be similarly adapted to the current application. (Since the directions of derivative approximations for $V^{k,1}$ and $V^{k,2}$ are determined by U^k values, these two grid functions are also easily computed simultaneously with U^k .) Which method would be more efficient will depend on the geometry of Ω and the properties of ψ, K , and f .

Remark 5.1. The above grid functions are needed to predict the decisions of a rational extractor. For any single site $\mathbf{x} = (x_i, y_j)$, whether it will be extracted under Model A can be answered as soon as we find any k such that

$$B_{i,j} e^{-V_{i,j}^{k,1}} - V_{i,j}^{k,2} - R_{i,j} > \tilde{p}.$$

The exact $P^a(x_i, y_j)$ could be computed by any continuous optimization in λ . But to approximate P^a on the entire spatial grid \mathcal{X} , it is more efficient to compute $V^{k,1}$ and $V^{k,2}$ for all $k = 0, 1, \dots, N_\lambda$ and then use a λ -grid search:

$$(5.2) \quad P_{i,j}^a \approx \max_k \left\{ B_{i,j} e^{-V_{i,j}^{k,1}} - V_{i,j}^{k,2} \right\} - R_{i,j}.$$

5.2. Randomly terminated Eikonal equation (Model G). We will define \bar{U} as a grid function approximating the value function \bar{u} of section 3.2. The upwind difference operators can be similarly used to approximate the gradient and then used to obtain a consistent, monotone discretization of PDE (3.3):

$$(5.3) \quad \begin{aligned} & \left[\min \left(D_{i,j}^{+x} [\bar{U}^m], -D_{i,j}^{-x} [\bar{U}^m], 0 \right) \right]^2 + \left[\min \left(D_{i,j}^{+y} [\bar{U}^m], -D_{i,j}^{-y} [\bar{U}^m], 0 \right) \right]^2 \\ & = \left[K_{i,j} + \psi_{ij} \left(b_m + R_{i,j} - \bar{U}_{i,j}^m \right) \right]^2 / f_{i,j}^2, \end{aligned}$$

with the boundary conditions $\bar{U}_{i,j}^m = 0$ for $(x_i, y_j) \notin \Omega$.

This is essentially the same discretization previously used in [3], where it was shown that its causal properties allow for a noniterative Dijkstra-like algorithm with $O(M \log M)$ complexity. We refer readers to [3, Appendix B] for a quadrant-by-quadrant interpretation of (5.3) and a detailed discussion of efficient implementation. This approach is arguably even better suited in the current context: the “free boundary” present in [3] made it necessary to use a special initialization, but our Dirichlet boundary conditions make it possible to use the original fast marching.

Remark 5.2. For any single site $\mathbf{x} = (x_i, y_j)$, the expected profit under Model G could be computed by using the exact $B_{i,j}$ instead of b_m in (5.3). But to approximate P^g on the entire spatial grid \mathcal{X} , it is more efficient to compute \bar{U}^m for all $m = 0, 1, \dots, N_b$. If m is such that $B_{i,j} \in (b_m, b_{m+1}]$, we can define

$$P_{i,j}^{g-} = B_{i,j} - \bar{U}_{i,j}^m - R_{i,j} \quad \text{and} \quad P_{i,j}^{g+} = B_{i,j} - \bar{U}_{i,j}^{m+1} - R_{i,j}.$$

Either one of these can be used as a good approximation since $P_{i,j}^g \in [P_{i,j}^{g-}, P_{i,j}^{g+}]$ and $P_{i,j}^{g+} - P_{i,j}^{g-} \leq (\check{B} - \hat{B})/N_b$. We know that (x_i, y_j) will be extracted if $P_{i,j}^{g-} > \tilde{p}$ and will not be extracted if $P_{i,j}^{g+} \leq \tilde{p}$. In our numerical tests we use the \tilde{p} level set of $\frac{1}{2}(P^{g-} + P^{g+})$ to approximate $\partial\Omega_p^g$. A more conservative discrete approximation of that boundary could be also obtained from the set $\{(x_i, y_j) \mid P_{i,j}^{g-} \leq \tilde{p} \leq P_{i,j}^{g+}\}$.

6. Examples. In all of our numerical examples, we use a homogeneous running cost $K(\mathbf{x}) = \kappa = 1$. As a result, the pre-extraction cost $R(\mathbf{x})$ coincides with $\tau(\mathbf{x})$, the minimum time required starting from $\partial\Omega$ to reach the extraction point \mathbf{x} . Similarly, $\mathcal{J}_2(\mathbf{x}, \mathbf{a}(\cdot))$ represents the time to $\partial\Omega$ starting from \mathbf{x} and using the post-extraction control $\mathbf{a}(\cdot)$. In Examples 1–5 the perpetrators can move with speed $f = 1$ (and thus, $\tau(\mathbf{x}) = d(\mathbf{x})$), while in Examples 6 and 7, their speed is terrain dependent.

To save space, we primarily focus on the expected profit P^a corresponding to Model A (the aerial patrols). Our figure captions report the maximum attainable expected profit $\bar{P}^a = \max_{\mathbf{x} \in \Omega} P^a(\mathbf{x})$, the proportion of protected area $A_p^a = |\Omega_p^a|/|\Omega|$, and the proportion of protected value $V_p^a = \int_{\Omega_p^a} B(\mathbf{x})d\mathbf{x} / \int_{\Omega} B(\mathbf{x})d\mathbf{x}$. Implementing the same patrol density ψ in both Model A and Model G (the ground patrols) might be hard in practice, but for the sake of quantitative comparison, we assume that this can be done. Since in Model G the costs resulting from a detection are typically lower, this results in $P^g \geq P^a$ and $\Omega_p^g \subset \Omega_p^a$. We demonstrate this in Examples 5–7, where we compare $\partial\Omega_p^a$ with $\partial\Omega_p^g$ (using $N_\lambda = N_b$) and provide the related statistics.

In most cases we consider ψ to be a part of the problem statement, but in Examples 3 and 4 we also perform a straightforward optimization of ψ to maximize A_p^a while respecting the specified patrol density budget E , using (2.1) with $\gamma = 1$.

Our numerical approach is quite efficient. Even though the implementation is not highly optimized, it was easy to perform all simulations on a midrange laptop computer (MacBook Pro with a 2GHz Intel Core i5 processor and 8GB of RAM). For a fixed ψ and any specific λ , the execution times varied from 120 milliseconds on smaller grids (Examples 3–4) to 50 seconds on larger grids (Examples 6–7). The overall time depended on the desired λ -resolution.

6.1. Example 1: A circle. We start with a radially symmetric example close to the original setting of Albers [1]. The domain Ω is a disk of diameter 1, and we use a constant benefit function $B = 2$. But we assume a more general pointwise detection rate based on the distance d from the boundary:

$$(6.1) \quad \psi(d) = \frac{\mu}{50(d - 0.3)^2 + 0.5}, \quad d \in [0, 0.5],$$

where μ is a constant chosen to enforce the budget constraint $\int_{\Omega} \psi(\mathbf{x})d\mathbf{x} = E$ with $E = 2.5$. A contour plot of ψ can be found in Figure 2(b). The expected profit P^a is shown in Figure 2(c), with the dashed line indicating the boundary of the pristine region Ω_p^a . The dotted line indicates the zero level set of P_{\sharp}^a , the approximation to P^a discussed in Remark 4.1.

We use a uniform rectangular grid with dimensions $N_x = 501$, $N_y = 501$. Since f is uniformly one, with ψ and K depending only on the distance from the boundary, this example falls into the category of “ τ -banded” scenarios discussed in section 4. The optimal trajectories are simply straight lines to the closest point on the boundary. Since the same trajectory minimizes both \mathcal{J}_1 and \mathcal{J}_2 , the Pareto front consists of a single point, which can be recovered using any λ with a single PDE solve.

6.2. Example 2: Banded patrol density on a square. We use a sum of two Gaussians to specify the benefit function

$$B(x, y) = e^{-10((x-0.25)^2 + (y-0.5)^2)} + e^{-10((x-0.75)^2 + (y-0.5)^2)}$$

on a unit square Ω ; see Figure 3(a). We assume the same “banded” detection rate defined in (6.1), but with μ selected based on $E = 2$. A contour plot of ψ is shown

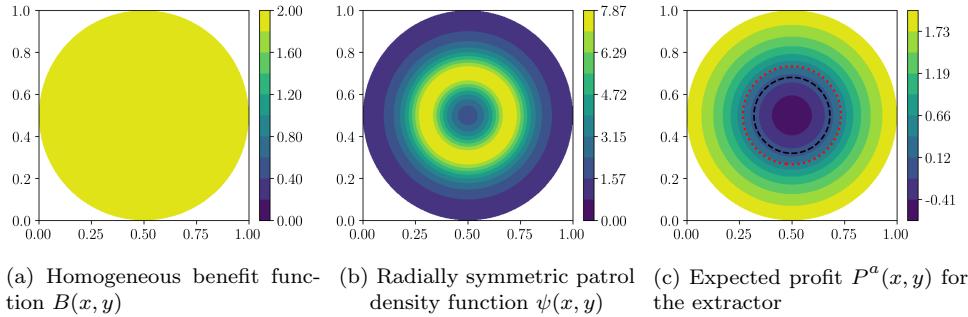


FIG. 2. A disk-like domain. The maximum profit $\bar{P}^a = 2$ is attained at $\partial\Omega$, the pristine area proportion is $A_p^a = 13.02\%$ and, since $B(\mathbf{x})$ is constant, the same $V_p^a = 13.02\%$ of the value is protected. If P_\sharp was used instead of P^a , it would overestimate the pristine proportion and value protected as $A_p^a = V_p^a = 21.77\%$.

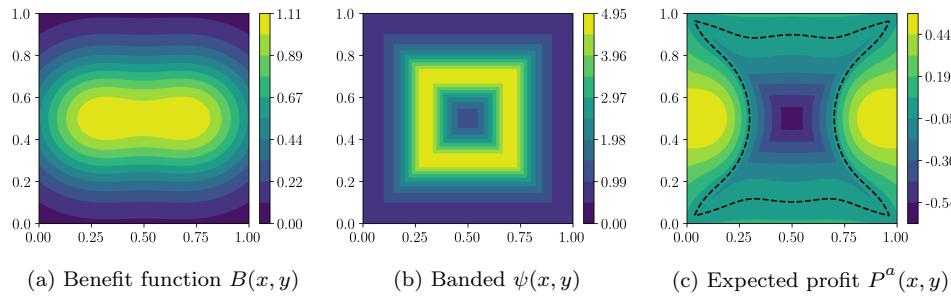


FIG. 3. Banded patrol density on a square. $\bar{P}^a = 0.56$, $A_p^a = 45.67\%$, and $V_p^a = 52.99\%$.

in Figure 3(b). The expected profit P^a can be found in Figure 3(c), with the dashed line showing $\partial\Omega_p^a$.

This example was computed on a uniform rectangular grid with dimensions $N_x = 501$, $N_y = 501$. As in Example 1, this is another case of a “ τ -banded” patrol density, so optimal trajectories are straight lines to the closest point on $\partial\Omega$, and any single u^λ is sufficient to compute $P^a(\mathbf{x})$. (Note that this property holds regardless of a particular benefit function B . It also implies that both Models A and G will yield exactly the same predictions.)

6.3. Example 3: Placing a patrol station. As noted in [1], in realistic situations PAMs face additional constraints and the detection rate typically decreases with the distance to patrol stations (in our case — the locations where the surveillance drones are based). We will suppose that the detection rate $\psi(x, y)$ is now a single Gaussian “centered” at a patrol station $(\bar{x}, \bar{y}) \in \Omega$:

$$\psi_{\bar{x}, \bar{y}}(x, y) = \mu e^{-30((x-\bar{x})^2 + (y-\bar{y})^2)}$$

with μ selected based on $E = 2$. Using the same domain, speed, and benefit functions defined in Example 2, we will now search for the patrol station location that maximizes the pristine area proportion A_p^a . In principle, this could be done with any black-box derivative-free optimization method on Ω , but for the sake of simplicity we have

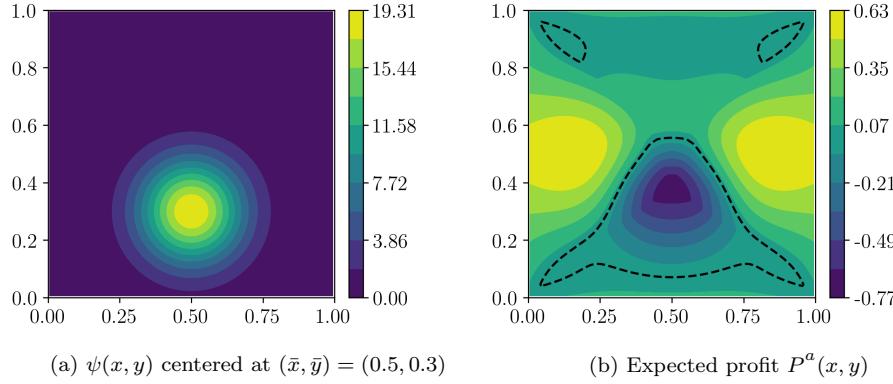


FIG. 4. *Optimal patrol station position. $\bar{P}^a = 0.63$, $A_p^a = 23.53\%$, and $V_p^a = 24.54\%$*

implemented a grid search on a coarse (11×11) grid. For each candidate location, the pristine region Ω_p is computed on a uniform rectangular grid with $N_x = 201$, $N_y = 201$, and $N_\lambda = 101$.

Due to the symmetry, there are actually two A_p^a -maximizing locations: $(\bar{x}, \bar{y}) = (0.5, 0.3)$ and $(\bar{x}, \bar{y}) = (0.5, 0.7)$. Figure 4(a) shows ψ centered at the former. The corresponding expected profit P^a can be found in Figure 4(b), with the dashed line showing $\partial\Omega_p^a$.

6.4. Example 4: Allocating surveillance resources. Given the nonunique maxima (\bar{x}, \bar{y}) in Example 3, we now suppose that PAMs decide to build a separate drone station at each of these two locations. The next question is how to share the surveillance resources between them. Assuming that the overall number of drones is fixed and that the detection rate generated by each station is proportional to the number of drones stationed there, any allocation corresponds to choosing $\mathbf{w} = (w_1, w_2)$ with $w_1, w_2 \geq 0$, $w_1 + w_2 = 1$, and defining

$$\begin{aligned}\psi_{\mathbf{w}}(x, y) &= \mu_{\mathbf{w}}(w_1 G_1(x, y) + w_2 G_2(x, y)), \\ G_1(x, y) &= e^{-30((x-0.5)^2 + (y-0.3)^2)}, \\ G_2(x, y) &= e^{-30((x-0.5)^2 + (y-0.7)^2)},\end{aligned}$$

with $\mu_{\mathbf{w}}$ chosen to ensure $\int_{\Omega} \psi_{\mathbf{w}}(x, y) dx dy = E = 2$. For each \mathbf{w} , we compute the A_p^a on a uniform rectangular grid with $N_x = 201$, $N_y = 201$, and $N_\lambda = 101$. A grid search over 101 values of w_1 is used to find the A_p^a -maximizing allocation. As in the previous example, the answer is not unique: both $\mathbf{w} = (0.43, 0.57)$ and $\mathbf{w} = (0.57, 0.43)$ are optimal.

For the first of these, Figure 5(a) shows a contour plot of $\psi_{\mathbf{w}}$ along with four different Pareto optimal paths starting from an extraction point $\mathbf{x}_0 = (0.555, 0.315)$. The path maximizing the expected profit P^a is shown in magenta, the path minimizing the probability of detection Ψ is shown in green, and the path minimizing the time of travel is shown in black. The red path corresponds to the “nearly optimal” $\lambda_{\sharp} = \frac{B}{B+1}$ of the linearized- Ψ problem, as discussed in Remark 4.1. The corresponding expected profit P^a can be found in Figure 5(b), with the dashed line showing $\partial\Omega_p^a$.

In Figure 6(a), we consider the λ -optimization problem (5.2) for the same extraction point \mathbf{x}_0 in detail. Figure 6(a) shows the λ_k -optimal payoffs $B(x, y)e^{-V_{i,j}^{k,1}} -$

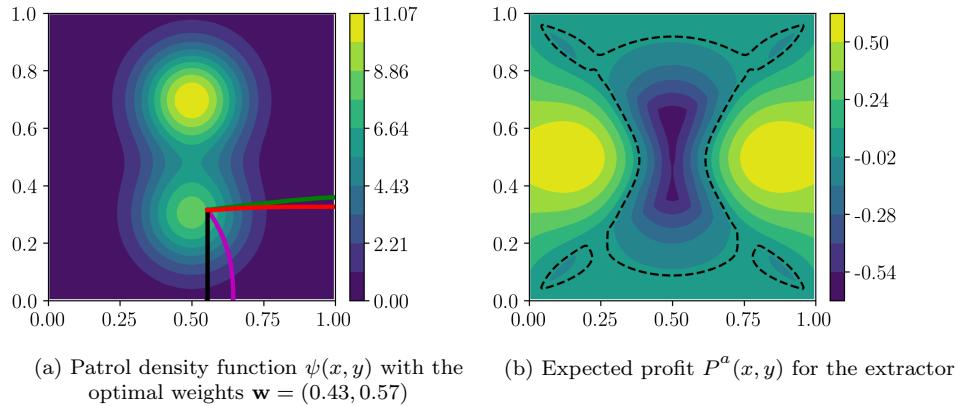


FIG. 5. Allocating surveillance resources. $\bar{P}^a = 0.63$, A_p^a has increased to 35.26%, and V_p^a has increased to 38.48%. We note that from the PAMs' perspective, the banded patrol from Example 2 still performs better while using the same budget ($E = 2$).

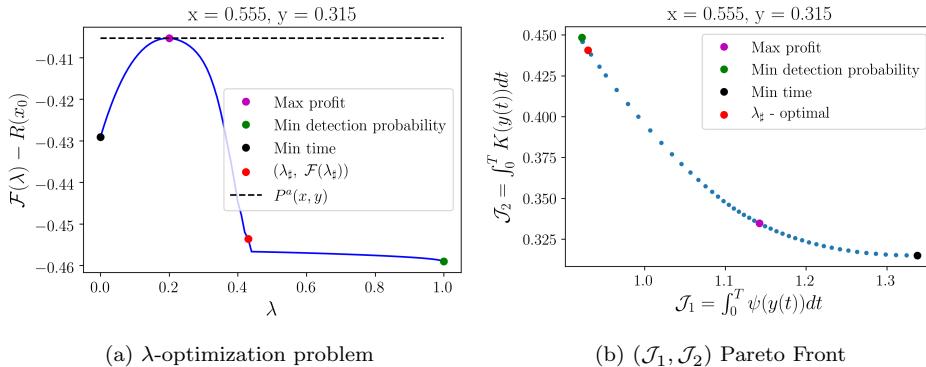


FIG. 6. Lambda optimization and Pareto front at $\mathbf{x}_0 = (0.555, 0.315)$ in Example 4.

$V_{i,j}^{k,2} - R(x, y)$ in blue, with the maximum profit $P^a(x, y)$ as a black dashed line. The choice of path based on λ_{\sharp} is clearly far from optimal, which is not surprising since \mathcal{J}_1 is not small. The linearized- Ψ approach significantly overestimates the probability of detection along all trajectories. In fact, if we were to base the estimated “profit” on $u^{\lambda_{\sharp}}$, the formula (4.4) would yield an even more pessimistic $P_{\sharp}(\mathbf{x}_0) = -0.699$.

Figure 6(b) shows the $(\mathcal{J}_1, \mathcal{J}_2)$ Pareto front for the same \mathbf{x}_0 . The color of points in 6(a) and 6(b) matches the color of the corresponding paths in 5(a).

6.5. Example 5: Exploiting patrol gaps. The main purpose of this example is to highlight the differences in post-extraction trajectories between Models A and G. The domain Ω is again the unit square $[0, 1] \times [0, 1]$, with a uniform speed of $f = 1$. The benefit function B is

$$B(x, y) = 3 + 7.5 \exp \left[-10 \left((x - 0.7)^2 + (y - 0.5)^2 \right) \right]$$

and can be seen in Figure 7(a). The patrol density $\psi(x, y)$ is a sum of eight Gaussians with different means and standard deviations and can be seen in Figure 7(b).

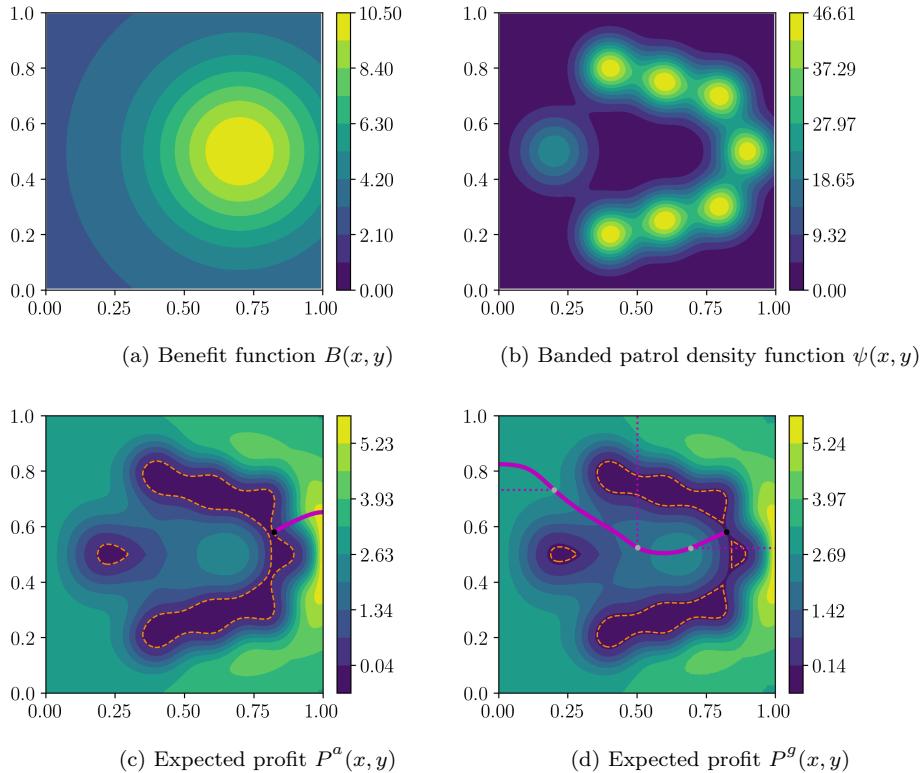


FIG. 7. Differences between path-planning in Model A ($\bar{P}^a = 5.87$, $A_p^a = 11.06\%$, $V_p^a = 14.66\%$) and Model G ($\bar{P}^g = 5.88$, $A_p^g = 8.63\%$, $V_p^g = 11.02\%$).

In Figure 7(c) we see the expected profit P^a from Model A computed on a uniform rectangular grid with $N_x = N_y = 501$, and $N_\lambda = 401$. The optimal post-extraction trajectory corresponding to $\hat{\mathbf{x}} = (0.824, 0.58)$ is shown in magenta, and $\partial\Omega_p^a$ is shown in orange.

In Figure 7(d) we see the expected profit P^g from Model G computed on the same grid, and $\partial\Omega_p^g$ is shown in orange. While P^g is only slightly larger than P^a , a significant difference can be seen in the post-extraction trajectories. The solid magenta line represents the optimal post-extraction (pre-detection) trajectory starting from $\hat{\mathbf{x}}$. The dotted magenta lines show possible post-detection trajectories from three different positions. Since $f = 1$, every post-detection trajectory is just a straight line to the closest point on $\partial\Omega$.

6.6. Example 6: Yosemite National Park. Here we use the domain and terrain data from Yosemite National Park. To facilitate the comparison, we use the same B , ψ , and f defined in [4].

The benefit function shown in Figure 8(a) is $B(d) = 8d(2d_m - d)/(d_m)$, where $d_m = \max_{\mathbf{x} \in \Omega} d(\mathbf{x})$. The banded patrol density shown in Figure 8(b) is defined by

$$\psi(d) = \begin{cases} \frac{\mu(0.7d_m - d)}{d_m}, & 0.3d_m < d < 0.7d_m, \\ 0, & \text{otherwise,} \end{cases}$$

where μ is selected based on $E = 3 \times 10^4$.

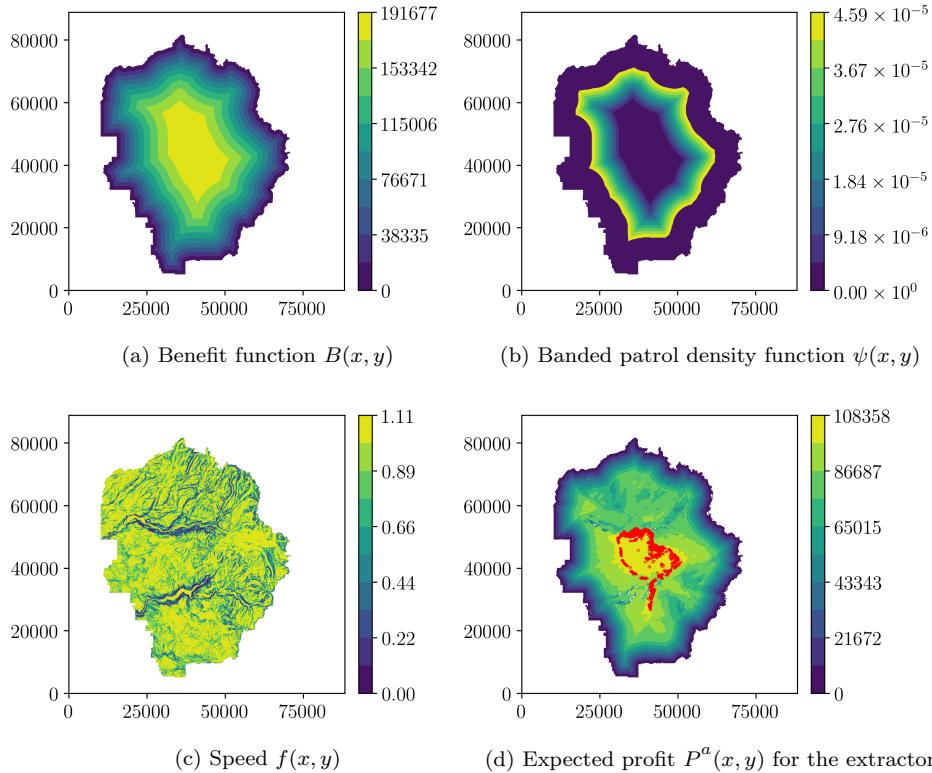


FIG. 8. *Yosemite National Park with $E = 3 \times 10^4$, corresponding to Figure 9b in [4]. $\bar{P}^a = 108358$, $A_p^a = 1.90\%$, and $V_p^a = 1.86\%$. Dashed red line indicates $\partial\Omega_e$.*

The extractors' isotropic speed of motion is shown in Figure 8(c) and depends on the slope of the elevation map $z(\mathbf{x})$. Namely, $s = |\nabla z(\mathbf{x})|$ is the grade in the direction of $\nabla z(\mathbf{x})$ and

$$(6.2) \quad f(s) = 1.11 \exp\left(-\frac{(100s + 2)^2}{2345}\right).$$

This example was computed on a uniform rectangular grid with $N_x = 5406$, $N_y = 4325$, and $N_\lambda = 21$. The expected profit P^a can be seen in Figure 8(d). The red dashed line corresponds to the boundary of the “high value region”

$$(6.3) \quad \Omega_e = \left\{ \bar{\mathbf{x}} \in \Omega \mid P^a(\bar{\mathbf{x}}) + R(\bar{\mathbf{x}}) \geq (1 - \varepsilon) \max_{\mathbf{x} \in \Omega} (P^a(\mathbf{x}) + R(\mathbf{x})) \right\},$$

with $\varepsilon = 0.85$. This definition is chosen for the sake of direct comparison with Figure 9 in [4] since that paper does not consider the cost of pre-extraction trajectories and their definition of pristine region is also different.

Under our interpretation of the extraction criteria, the extractors will target more than 98% of the park, and the patrol budget E is clearly insufficient. We repeat this experiment with a higher $E = 2 \times 10^5$; see Figure 9.

This example presents a nontrivial Pareto front since the detection rate ψ is d -banded (rather than τ -banded) and the time-optimal paths are generally not detection-minimizing. As a result, the predicted profit in Model A is different from Model G,

but the differences are not very large here, e.g., $A_p^a - A_p^g \approx 2.16\%$ even with the higher patrol budget used in Figure 9.

It is also worth noting that the definition (6.2) results in very small f (and occasionally even zero up to machine precision) on all steep parts of Ω . As a result, the cost of targeting those points becomes extremely high even if everything around them is extracted. This explains the “noisy” look of $\partial\Omega_e$ in Figure 8(d) and $\partial\Omega_p$ in Figure 8(d). Such unextractable points are not identified in Figure 9b of [4] since they are masked by the b -interpolation procedure used in that paper.

6.7. Example 7: Kangaroo Island. We use the domain and terrain data corresponding to Kangaroo Island with the same B, ψ , and f considered in [4].

The functions B and f are defined exactly as in Example 6; see Figures 10(a) and 10(c). The detection rate ψ is a positive constant from $x = 30,000$ to $x = 90,000$ and zero elsewhere, scaled to match the budget $E = 3 \times 10^4$; see Figure 10(b).

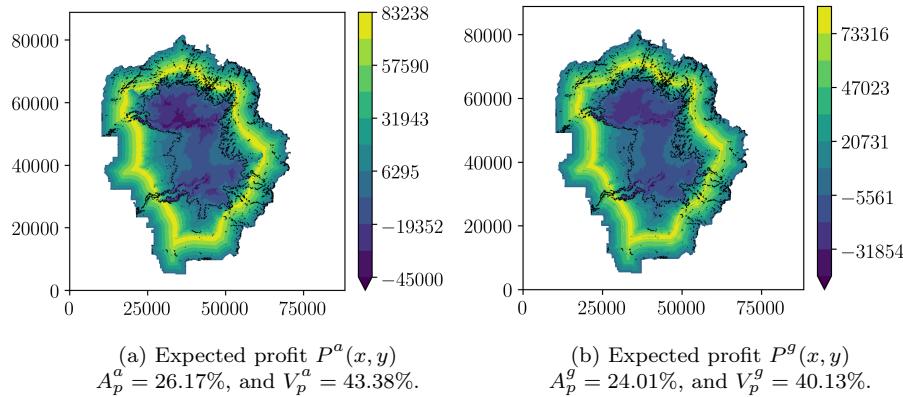


FIG. 9. Yosemite National Park with $E = 2 \times 10^5$. Black line indicates $\partial\Omega_p$.

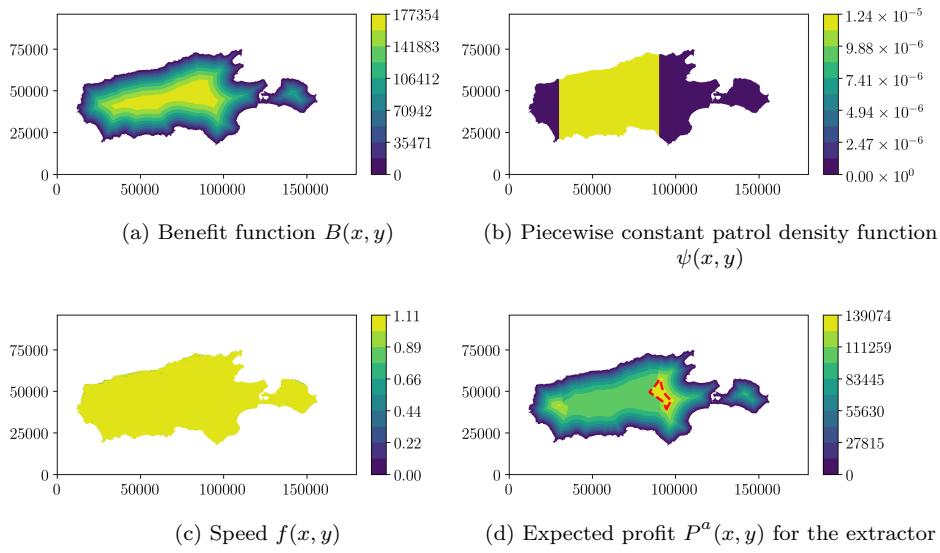


FIG. 10. Kangaroo Island with $E = 3 \times 10^4$, corresponding to Figure 9b in [4]. $\bar{P}^a = 139074$, $A_p^a = 0.16\%$, and $V_p^a = 0.0042\%$. Dashed red line indicates $\partial\Omega_e$.

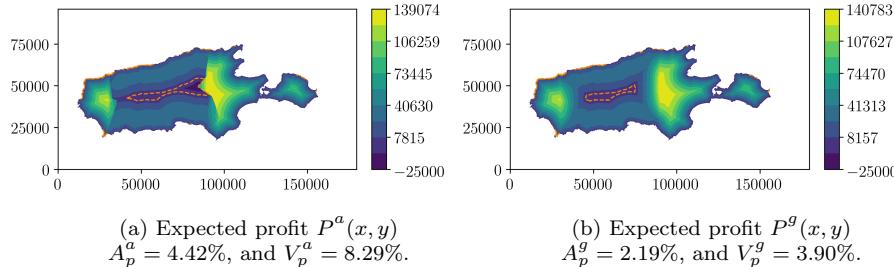


FIG. 11. Kangaroo Island with $E = 2.5 \times 10^5$. Dashed orange line indicates $\partial\Omega_p$.

This example was computed on a uniform rectangular grid with $N_x = 7159$, $N_y = 3111$, and $N_\lambda = 21$. The expected profit P^a can be seen in Figure 10(d). The red dashed line corresponds to the boundary of the “high value region” defined in (6.3) and can be compared with Figure 12b in [4]. Again, almost the entire Ω is targeted by the extractors. Since the terrain is mostly flat, this remains true even with a much higher budget $E = 2.5 \times 10^5$; see Figure 11(a).

In Figure 11(b) we examine Model G with the same parameters. Since a detection is much less costly in Model G, the extractors are even more willing to enter the patrolled area, and the pristine area (shown in orange) is about two times smaller than in Figure 11(a).

7. Conclusions. We have presented two control-theoretic models of illegal resource extraction counteracting the detection efforts of PAMs. The perpetrators are assumed to choose potential extraction sites based on their expected profit. They take into account the distribution of resources, the geometry and terrain of the protected area Ω , and the localized detection rate $\psi : \Omega \rightarrow \mathbb{R}_{+,0}$ due to PAMs’ patrol activities. The cost of extracting at $\mathbf{x} \in \Omega$ is based on a roundtrip $\partial\Omega \rightarrow \mathbf{x} \rightarrow \partial\Omega$. The two portions of the trip usually follow different paths. The $\partial\Omega \rightarrow \mathbf{x}$ path is selected to minimize the time and hardship of travel, while on the way back ($\mathbf{x} \rightarrow \partial\Omega$) these are balanced against the expected losses resulting from a possible detection by PAMs.

A detection is always followed by confiscation of resources carried by the extractor. But the details of how and when a confiscation happens significantly influence the perpetrators’ decision-making process. Our Model G assumes that the extractors are spotted by ground patrols with an immediate capture/confiscation after their detection. It is further assumed that after the confiscation the extractors switch to the quickest path to leave the protected area. We find their best pre-detection trajectories by recasting this as an optimal control of a suitable randomly terminated process. Our Model A assumes that the extractors are spotted by aerial patrols but remain oblivious when this happens, continuing along their originally chosen trajectory until they are captured at $\partial\Omega$. The key challenge of this application is the impossibility of writing a single running cost to be integrated along the post-extraction path unless we resort to a linearizing approximation for the probability of detection (see Remark 3.2) or agree to increase the dimension of our planning space (see Remark 3.3). We show that the former often leads to very inaccurate predictions while the latter leads to expensive numerical methods. Instead, we opt to use the techniques of *multiobjective* dynamic programming on the original domain Ω , resulting in an accurate and efficient numerical implementation. Our method is illustrated through a range of examples, with two of them based on real terrain maps from Yosemite National Park in California and Kangaroo Island in South Australia. In each case, we

compute the perpetrators' expected profit from all possible extraction sites as well as the "pristine" region $\Omega_p \subset \Omega$ not affected by extraction. We also show that Model G generally results in higher expected profits, smaller Ω_p , and different post-extraction trajectories.

Many of our modeling assumptions were made primarily for the simplicity of exposition and could be easily relaxed in the future. For example, it would not be hard to incorporate extractor's preferences for starting and terminal positions on $\partial\Omega$. One could also easily model the time needed to extract at each site and the resulting increase in the risk of detection. It would be similarly trivial to include additional fines in case of detection/capture. Following [4], we have assumed that the extractors' speed of motion f is isotropic. That is, if $z(\mathbf{x})$ is the elevation map, the speed f was based on the grade in the direction of $\nabla z(\mathbf{x})$ rather than on the grade in the chosen direction of motion \mathbf{a} . In reality, a pedestrian's speed should depend on both, particularly for moderate $|\nabla z|$. Any such model would replace Eikonal PDEs with anisotropic HJB equations. Efficient numerical methods for the latter have also been developed in both fast marching [2, 27, 36] and fast sweeping [24, 37] frameworks.

While our current implementation is sequential, we note that it would be easy to build a parallel algorithm with excellent scalability for computing $P(\mathbf{x})$: for each specific λ_k or b_m , the value function could be computed by a separate CPU core. The resulting speed-up will be very useful in optimizing the choice of detection rate ψ to maximize the area of Ω_p while staying within PAMs' enforcement/surveillance budget. In the current paper we perform some of this by straightforward grid searches, but a more general ψ -optimization will require higher computational efficiency. The problem of optimally allocating surveillance resources across several patrol stations (e.g., Example 4) will become particularly expensive as the number of stations grows. We believe that an efficient solution will be based on generalizing our recent work on surveillance-evasion games [11, 20].

We conclude by listing two significant challenges not tackled in the current paper. First, a "single period" assumption for this Stackelberg game is rather limiting. As extractors continue degrading the protected area Ω , the distribution of resources changes and PAMs should accordingly reallocate the surveillance resources, adjusting ψ to increase the pristine region Ω_p . Second, in practice the detection rate ψ is not chosen directly but results from implementing specific surveillance strategies (e.g., choices of particular patrol trajectories and patrol frequencies). It would be far more realistic and more difficult to optimize those strategies directly.

Acknowledgments. We are very grateful to all authors of [4] for answering our questions about their model and for sharing the terrain data used in their numerical experiments. The current paper was written during AV's sabbatical visit to ORFE/Princeton, and he would like to thank ORFE for its hospitality.

REFERENCES

- [1] H. J. ALBERS, *Spatial modeling of extraction and enforcement in developing country protected areas*, Resource Energy Econom., 32 (2010), pp. 165–179.
- [2] K. ALTON AND I. M. MITCHELL, *An ordered upwind method with precomputed stencil and monotone node acceptance for solving static convex Hamilton-Jacobi equations*, J. Sci. Comput., 51 (2012), pp. 313–348.
- [3] J. ANDREWS AND A. VLADIMIRSKY, *Deterministic control of randomly-terminated processes*, Interfaces Free Bound., 16 (2014), pp. 1–40, <https://doi.org/10.4171/IFB/312>.
- [4] D. J. ARNOLD, D. FERNANDEZ, R. JIA, C. PARKINSON, D. TONNE, Y. YANIV, A. L. BERTOZZI, AND S. J. OSHER, *Modeling environmental crime in protected areas using the level set*

method, SIAM J. Appl. Math., 79 (2019), pp. 802–821, <https://pubs.siam.org/doi/abs/10.1137/18M1205339>.

[5] M. BARDI AND I. CAPUZZO-DOLCETTA, *Optimal Control and Viscosity Solutions of Hamilton-Jacobi-Bellman Equations*, Springer, New York, 2008.

[6] G. BARLES AND P. E. SOUGANIDIS, *Convergence of approximation schemes for fully nonlinear second order equations*, Asymptot. Anal., 4 (1991), pp. 271–283.

[7] T. J. BARTH AND J. A. SETHIAN, *Numerical schemes for the Hamilton–Jacobi and level set equations on triangulated domains*, J. Comput. Phys., 145 (1998), pp. 1–40.

[8] M. BOUÉ AND P. DUPUIS, *Markov chain approximations for deterministic control problems with affine dynamics and quadratic cost in the control*, SIAM J. Numer. Anal., 36 (1998), pp. 667–695.

[9] E. K. BOUKAS, A. HAURIE, AND P. MICHEL, *An optimal control problem with a random stopping time*, J. Optim. Theory Appl., 64 (1990), pp. 471–480, <https://doi.org/10.1007/BF00939419>.

[10] A. BRESSAN AND B. PICCOLI, *Introduction to the Mathematical Theory of Control*, Vol. 1, American Institute of Mathematical Sciences, Springfield, MO, 2007.

[11] E. CARTEE, L. LAI, Q. SONG, AND A. VLADIMIRSKY, *Time-dependent surveillance-Evasion games*, 2019 IEEE 58th Conference on Decision and Control (CDC), Nice, France, 2019, pp. 7128–7133, <https://doi.org/10.1109/CDC40024.2019.9029329>.

[12] A. CHACON AND A. VLADIMIRSKY, *Fast two-scale methods for eikonal equations*, SIAM J. Sci. Comput., 34 (2012), pp. A547–A578.

[13] A. CHACON AND A. VLADIMIRSKY, *A parallel two-scale method for eikonal equations*, SIAM J. Sci. Comput., 37 (2015), pp. A156–A180.

[14] Z. CLAWSON, X. DING, B. ENGLOT, T. A. FREWEN, W. M. SISSON, AND A. VLADIMIRSKY, *A Bi-Criteria Path Planning Algorithm for Robotics Applications*, <https://arxiv.org/abs/1511.01166>, 2015.

[15] M. G. CRANDALL AND P.-L. LIONS, *Viscosity solutions of Hamilton–Jacobi equations*, Trans. Amer. Math. Soc., 277 (1983), pp. 1–42.

[16] I. DAS AND J. E. DENNIS, *A closer look at drawbacks of minimizing weighted sums of objectives for Pareto set generation in multicriteria optimization problems*, Struct. Optim., 14 (1997), pp. 63–69.

[17] A. DESILLES AND H. ZIDANI, *Pareto front characterization for multi-objective optimal control problems using Hamilton–Jacobi approach*, SIAM J. Control Optim., 57 (2019), pp. 3884–3910, <https://pubs.siam.org/doi/abs/10.1137/18M1176993>.

[18] E. W. DIJKSTRA, *A note on two problems in connexion with graphs*, Numer. Math., 1 (1959), pp. 269–271.

[19] 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 (2017), pp. 23–36.

[20] M. GILLES AND A. VLADIMIRSKY, *Evasive path planning under surveillance uncertainty*, Dyn. Games Appl., 10 (2020), pp. 391–416, <https://doi.org/10.1007/s13235-019-00327-x>.

[21] A. HAURIE, *A multigenerational game model to analyze sustainable development*, Ann. Oper. Res., 137 (2005), pp. 369–386, <https://doi.org/10.1007/s10479-005-2267-2>.

[22] A. HAURIE AND F. MORESINO, *A stochastic control model of economic growth with environmental disaster prevention*, Automatica, 42 (2006), pp. 1417–1428.

[23] M. P. JOHNSON, F. FANG, AND M. TAMBE, *Patrol strategies to maximize pristine forest area*, in Proceedings of the 26th AAAI Conference on Artificial Intelligence, 2012.

[24] C. KAO, S. OSHER, AND Y. TSAI, *Fast sweeping methods for static Hamilton–Jacobi equations*, SIAM J. Numer. Anal., 42 (2005), pp. 2612–2632, <https://doi.org/10.1137/S0036142902419600>.

[25] D. KAR, B. FORD, S. GHOLAMI, F. FANG, A. PLUMPTRE, M. TAMBE, M. DRICIRU, F. WANYAMA, A. RWETSIBA, M. NSUBAGA, ET AL., *Cloudy with a chance of poaching: Adversary behavior modeling and forecasting with real-world poaching data*, in Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems, International Foundation for Autonomous Agents and Multiagent Systems, 2017, pp. 159–167.

[26] A. KUMAR AND A. VLADIMIRSKY, *An efficient method for multiobjective optimal control and optimal control subject to integral constraints*, J. Comput. Math., (2010), pp. 517–551.

[27] J.-M. MIREBEAU, *Efficient fast marching with Finsler metrics*, Numer. Math., 126 (2014), pp. 515–557.

[28] I. M. MITCHELL AND S. SASTRY, *Continuous path planning with multiple constraints*, in Proceedings of the 42nd IEEE International Conference on Decision and Control, Vol. 5, IEEE, 2003, pp. 5502–5507.

- [29] S. OSHER AND J. A. SETHIAN, *Fronts propagating with curvature-dependent speed: Algorithms based on Hamilton-Jacobi formulations*, J. Comput. Phys., 79 (1988), pp. 12–49, [https://doi.org/10.1016/0021-9991\(88\)90002-2](https://doi.org/10.1016/0021-9991(88)90002-2).
- [30] J. QIAN, Y.-T. ZHANG, AND H.-K. ZHAO, *Fast sweeping methods for eikonal equations on triangular meshes*, SIAM J. Numer. Anal., 45 (2007), pp. 83–107.
- [31] E. J. ROBINSON, H. J. ALBERS, AND J. C. WILLIAMS, *Spatial and temporal modeling of community non-timber forest extraction*, J. Environ. Econom. Manag., 56 (2008), pp. 234–245, <https://doi.org/10.1016/j.jeem.2008.04.002>.
- [32] E. J. Z. ROBINSON, *India's disappearing common lands: Fuzzy boundaries, encroachment, and evolving property rights*, Land Econom., 84 (2008), pp. 409–422, <http://www.jstor.org/stable/27647834>.
- [33] J. A. SETHIAN, *A fast marching level set method for monotonically advancing fronts*, Proc. Natl. Acad. Sci., 93 (1996), pp. 1591–1595.
- [34] J. A. SETHIAN, *Fast marching methods*, SIAM Rev., 41 (1999), pp. 199–235.
- [35] J. A. SETHIAN AND A. VLADIMIRSKY, *Fast methods for the Eikonal and related Hamilton–Jacobi equations on unstructured meshes*, Proc. Natl. Acad. Sci., 97 (2000), pp. 5699–5703.
- [36] J. A. SETHIAN AND A. VLADIMIRSKY, *Ordered upwind methods for static Hamilton–Jacobi equations: Theory and algorithms*, SIAM J. Numer. Anal., 41 (2003), pp. 325–363.
- [37] Y.-H. R. TSAI, L.-T. CHENG, S. OSHER, AND H.-K. ZHAO, *Fast sweeping algorithms for a class of Hamilton–Jacobi equations*, SIAM J. Numer. Anal., 41 (2003), pp. 673–694.
- [38] J. N. TSITSIKLIS, *Efficient algorithms for globally optimal trajectories*, IEEE Trans. Automat. Control, 40 (1995), pp. 1528–1538.
- [39] H. ZHAO, *A fast sweeping method for eikonal equations*, Math. Comp., 74 (2005), pp. 603–627.