

Optimal Average Disk-Inspection via Fermat’s Principle

Konstantinos Georgiou   

Toronto Metropolitan University, Canada

Abstract

This work resolves the optimal average-case cost of the Disk-Inspection problem, a variant of Bellman’s 1955 lost-in-a-forest problem. In Disk-Inspection, a mobile agent starts at the center of a unit disk and follows a trajectory that inspects perimeter points whenever the disk does not obstruct visibility. The worst-case cost was solved optimally in 1957 by Isbell [31], but the average-case version remained open, with heuristic upper bounds proposed by Gluss [29] in 1961 and improved only recently in [16].

Our approach applies Fermat’s Principle of Least Time from optics to the discretization framework of [16], showing that optimal solutions are captured by a one-parameter family of recurrences independent of the discretization size. In the continuum limit these recurrences give rise to a single-parameter optimal control problem, whose trajectories coincide with limiting solutions of the original Disk-Inspection problem. A crucial step is proving that the optimal initial condition generates a trajectory that avoids the unit disk, thereby validating the optics formulation and reducing the many-variable optimization to a rigorous one-parameter problem. In particular, this disproves Gluss’s conjecture [29] that optimal trajectories must touch the disk.

Our analysis determines the exact optimal average-case inspection cost, equal to 3.549259... and certified to at least six digits of accuracy.

2012 ACM Subject Classification Theory of computation → Computational geometry; Theory of computation → Theory and algorithms for application domains

Keywords and phrases Inspection, Disk, Average-Case Performance

Digital Object Identifier 10.4230/LIPIcs.STACS.2026.44

Related Version *Full Version:* <https://arxiv.org/abs/2509.06334> [26]

Funding Research supported in part by an NSERC Discovery Grant and by the Toronto Metropolitan University Faculty of Science Dean’s Research Fund.

Acknowledgements We thank Derek Muller and his team at *Veritasium* for an expository discussion of Fermat’s principle of least time, in particular the episode “*The Closest We’ve Come to a Theory of Everything*”, without which the perspective adopted in this work would likely not have emerged. Special thanks also to Caleb Jones for early discussions of the project.

1 Introduction

A hiker (mobile agent) stands in a forest, knowing only that the boundary is a straight infinite line at distance one but not its orientation. The task is to design a trajectory that guarantees reaching the boundary. Performance is evaluated in the worst-case (adversarial orientation) or in the average-case (orientation drawn uniformly at random).

The above setting is Bellman’s 1955 *lost-in-a-forest* problem specialized to a halfspace with known distance, commonly called the *shoreline problem with known distance*. Isbell [31] solved the worst-case version optimally in 1957. The average-case variant was first examined by Gluss [29], who proposed heuristic search strategies to the problem. He further conjectured that optimal solutions, as in the worst-case setting, must touch the unit disk. More recently, the same problem was studied again from the equivalent perspective of *Disk-Inspection*, leading



© Konstantinos Georgiou;

licensed under Creative Commons License CC-BY 4.0

43rd International Symposium on Theoretical Aspects of Computer Science (STACS 2026).

Editors: Meena Mahajan, Florin Manea, Annabelle McIver, and Nguyễn Kim Thăng

Article No. 44; pp. 44:1–44:19



Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



to a high-dimensional nonconvex NonLinear Program (NLP) [16] modelling a discretized version of the problem, which achieved an average-cost upper bound of 3.5509015. However, that approach could not certify global optimality of the NLP solutions (which would imply that the reported bound is close to the true optimum), the derived solutions did not scale computationally (which would theoretically give better upper bounds), and it offered no structural characterization of optimal trajectories in the continuous setting. Moreover, no nontrivial lower bound for the average-case cost was known, leaving it entirely unclear how close the reported bound was to the truth.

In this work, we resolve the Average-Case Disk-Inspection problem. We prove that the exact optimum is

$$3.54925958\dots,$$

accurate to at least six digits, and that it is realized by trajectories determined through the solutions of an ODE system. Our key technical contribution is to reformulate the discretized problem, via Fermat's Principle of Least Time, as an optics problem. Assuming an optimal trajectory does not touch the disk, we show that optimal solutions to the discrete inspection problem admit a recurrence structure that, in the continuum limit, yields a single-parameter optimal control problem. In this optimization problem, feasible solutions are trajectories of an ODE system determined by one initial condition, and optimizing over this parameter gives a rigorously certifiable optimum. A crucial part of the analysis is to prove that the non-touching condition is in fact satisfied, so the reduction to the optics problem is valid. Overall, this transforms the previous nonconvex many-variable optimization of [16] with no guarantees into a tractable one-parameter problem, completing the proof of optimality and establishing that an optimal trajectory avoids the unit disk, contrary to Gluss's conjecture [29], thereby settling the optimal solution to the average-case problem definitively.

This paper is an extended abstract; several proofs and technical details are omitted due to space limitations. The full version of this work appears in [26].

Related Work

In the mid-50s, Bellman [9] introduced what is nowadays called *lost-in-a-forest-problem*, one of the earliest formal questions in search theory. The input is a region R (the forest) containing a point P (starting position of a searcher/mobile agent) chosen at random, and the objective is to design a trajectory that minimizes the expected time to reach the boundary, starting from P . This formulation became a cornerstone for later work on search under uncertainty, and still many variations that have been proposed through the decades remain open.

Surveys of the lost-in-a-forest problem are given in [10, 22]. Work on specific domains includes convex polygons [28], strips [22] and general regions analyzed through competitive ratios [34]. When region R is a half-space, and the starting position of the agent is fixed deterministically, the problem is known as the *shoreline problem* and was first studied for a single agent in [7] and later extended to several agents in [5, 6, 8, 32]. For one agent, the logarithmic spiral is the best strategy known, with ratio 13.81 [7, 23]. The strongest lower bounds known are 6.3972 in general [6] and 12.5385 for cyclic searches [36]. For two agents, a double logarithmic spiral achieves ratio 5.2644 [6]. The optimal solutions to the one and two agent search problems are still open. For $n \geq 3$, trajectories along rays achieve at most $1/\cos(\pi/n)$ [6], and lower bounds matching these values were proven for $n \geq 4$ and for $n = 3$ in [1] and [18], respectively.

The variation to lost-in-a-forest pertaining to our results is the *shoreline problem with known distance* where the searcher knows her distance to the boundary of the forest. For minimizing the worst-case cost, the problem was solved optimally by Isbell [31] for one agent, and by Dobrev et al. in [18] for two agents, while [16] recently extended the results to the case where the hidden shorelines are tangent to a contiguous portion of the disk, rather than the entire disk.

The average-case version of the shoreline problem with known distance was first studied by Gluss [29], who proposed two heuristic strategies and carried out a rigorous expected-cost analysis for each. He further conjectured that an optimal trajectory, as in the worst-case setting, must touch the disk. The numerical values reported in [29] were miscalculated, and as explained in the full version [15] of the conference paper [16], the best heuristic bound due to Gluss is $3.63489\dots$. The first systematic approach to the average-case problem was given in the conference version [16] and its full version [15]. Their method introduced a discretization framework and formulated a nonconvex NonLinear Program (NLP) whose feasible solutions correspond to valid inspection trajectories, with the objective exactly capturing their average cost. Solving this NLP for large discretization parameter k produced a sequence of feasible trajectories and rigorous upper bounds, resulting in a reported value of $3.5509015\dots$. However, because the NLP is nonconvex, the obtained solutions could not be certified as globally optimal, and its computational hardness restricted k to moderate values, preventing sharper approximations. No lower bounds for the average-case cost had been studied or reported, so it remained unclear how close the best previously reported bound was to the true optimum.

Beyond shoreline search and lost-in-a-forest formulations, related problems in search theory have been studied extensively. The field itself is well established, with books and surveys providing systematic accounts of its models and applications [2, 3, 4, 17, 25]. Within point search in the plane, Langetepe [35] proved that the logarithmic spiral is optimal for a single robot, while [27] recently analyzed the case of multi-speed agents. Other studies considered search for a circle in a plane [30], multi-agent and grid searches with memory constraints [20, 21, 24, 37], and more recent work investigated searches in geometric terrains [12, 38] and the role of information in determining search cost [39, 40], to name only a few.

2 Inspection Problems and New Contributions

In this section we introduce the inspection problems studied in this work and state our main result, before outlining the technical framework that supports it.

2.1 Problem Definition and the Main Result

We begin by defining inspective trajectories and formalizing both the worst-case and average-case versions of Disk-Inspection, leading up to the main theorem.

► **Definition 1** (Inspective Trajectory). *A continuous and piecewise differentiable path in \mathbb{R}^2 (i.e., a curve with two endpoints) is called an inspective trajectory if its convex hull contains the unit disk centered at one of its endpoints.*

The requirement that the curve segment be piecewise differentiable stems from the fact that the problem we study concerns arc lengths. This leads to the following problem, also known as the shoreline problem with known distance.

44:4 Optimal Average Disk-Inspection via Fermat's Principle

► **Problem 2** (Worst-Case Disk-Inspection). Find an inspective trajectory of minimum length.

The Worst-Case Disk-Inspection problem was solved in [31], with optimal cost $1 + \sqrt{3} + 7\pi/6 \approx 6.39724$. Inspective trajectories admit two equivalent descriptions, see [16]. First, they *inspect* every point of the unit disk in the following sense. For any perimeter point P there exists a point A on the trajectory such that $\|\lambda A + (1 - \lambda)P\| \geq 1$ for all $\lambda \in [0, 1]$ (that is, A sees P without obstruction by the disk). Second, the trajectories intersect every line tangent to the unit disk (i.e. they discover eventually all shorelines). In what follows we use the inspection perspective, which will also define the average-case problem.

Given an inspective trajectory and a point P on the unit disk, define the inspection time I_P as the length of the trajectory from the origin to the first point that inspects P (equivalently, the time for a unit-speed inspector starting at the origin to see P from outside the disk). The Worst-Case Disk-Inspection, i.e. Problem 2, asks for the inspective trajectory that minimizes $\sup_P I_P$, where the supremum is over all perimeter points P . The average-case version replaces the supremum by expectation.

► **Problem 3** (Average-Case Disk-Inspection – ADI). Find an inspective trajectory that minimizes $\mathbb{E}_P[I_P]$, where P is chosen uniformly on the perimeter of the disk.

The cost of ADI is defined with respect to an inspector starting at the disk's center and following an inspective trajectory. One may also allow randomized algorithms that draw a curve from a distribution of inspective trajectories. After the algorithm is fixed, the adversary selects a point P , and the algorithm's performance is the expected inspection time at P . If the algorithm first applies a uniform random rotation, then all perimeter points are symmetric, and the adversary's choice of P has no effect. In this case the cost of a distribution equals the expected inspection time of a randomly chosen point, which is minimized by the deterministic curve achieving the smallest such expectation. Thus the deterministic ADI problem already captures the randomized model.

The average-case problem was first studied by Gluss [29], who proposed two heuristic strategies and conjectured that optimal trajectories touch the disk. More recently, a discretization-based nonlinear programming framework was developed [15, 16], yielding the first systematic upper bounds, with best reported value 3.5509015. However, the nonconvexity of this approach precluded global optimality guarantees, and no lower bounds were known, so the exact optimum remained unresolved. We resolve this problem with the following result, establishing the optimal value of ADI with high numerical accuracy. In particular, we prove that the previously reported upper bound in [16] was indeed very close to optimal, and we obtain this conclusion through a new continuous framework that replaces the earlier discrete, upper-bound approach. The technical contributions underlying this framework are presented in Section 2.2. We now state our main result.

► **Theorem 4.**

$$|C_{\text{opt}} - 3.549259| < 10^{-7},$$

where C_{opt} is the cost for an optimal solution to ADI.

All numerical errors are controlled as in Section 5.3, yielding at least six correct decimal digits, with a discussion on numerical robustness in Section 5.4. The trajectory certifying Theorem 4 is shown in Figure 1 and quantified in the next section.

2.2 New Technical Contributions

Our second contribution is a continuous framework that replaces the discrete approach of [16]. In this framework, candidate inspection trajectories are described, partially, by trajectories generated from an ordinary differential equation (ODE) system. The system is governed by a single real parameter, which turns ADI into a one-parameter optimal control problem. We now introduce the ODE system and the associated trajectories, which form the basis of the technical result leading to Theorem 4.

► **Definition 5** (ODE system $\text{SYS}(\tau_0)$ for (ψ, τ)). For a parameter $\tau_0 \in \mathbb{R}_{\geq 0}$, let $\psi, \tau : [0, 1] \rightarrow \mathbb{R}$ be the unique solution to

$$\begin{aligned} \psi'(x) &= -2\pi + \frac{\cot \psi(x)}{x}, & \psi(0) &= \frac{\pi}{2}, \\ \frac{1}{2\pi} \tau'(x) &= \tau(x) \cot \psi(x) - 1, & \tau(0) &= \tau_0. \end{aligned}$$

It can be shown that $\text{SYS}(\tau_0)$ is well defined near $x = 0$ and extends uniquely to $[0, 1]$. We later solve this system numerically to identify the trajectories defined below.

► **Definition 6** (Curve \mathcal{T}^{τ_0}). Let (ψ, τ) be the solution to $\text{SYS}(\tau_0)$. The associated curve is

$$\mathcal{T}^{\tau_0} : [0, 1] \rightarrow \mathbb{R}^2, \quad \mathcal{T}^{\tau_0}(x) = (\mathcal{T}_1^{\tau_0}(x), \mathcal{T}_2^{\tau_0}(x)),$$

where

$$\begin{aligned} \mathcal{T}_1^{\tau_0}(x) &= \cos(2\pi x) - \tau(x) \sin(2\pi x), \\ \mathcal{T}_2^{\tau_0}(x) &= -\sin(2\pi x) - \tau(x) \cos(2\pi x). \end{aligned}$$

We refer to \mathcal{T}^{τ_0} , the solution to ODE $\text{SYS}(\tau_0)$, as a curve. For suitable $\tau_0 > 0$, this curve corresponds to a portion of the optimal inspective trajectory certifying Theorem 4. Not every choice of τ_0 yields an inspective trajectory, which is why in the following definition we distinguish some τ_0 that result to curves that do not intersect the unit disk.


► **Definition 7.** $\tau_0 \in \mathbb{R}_{\geq 0}$ is called inspection-feasible for $\text{SYS}(\tau_0)$ if:

- for all $x \in [0, 1]$, we have $\|\mathcal{T}^{\tau_0}(x)\| > 1$, and
- there exists $\xi \in (0, 1]$ with $\mathcal{T}_1^{\tau_0}(\xi) = 1$.

The set of all inspection-feasible values is denoted by \mathcal{I} . For each $\tau_0 \in \mathcal{I}$, we denote by $\xi(\tau_0)$ the smallest value of $\xi > 0$ that satisfies the second condition, and call it the deployment parameter of τ_0 .¹

Following the standard shoreline framework (see, e.g., [16], building on [29]), we restrict attention to ADI trajectories that start with a so-called *deployment phase*. This is defined as a line segment connecting the center of the disk to some point A_∞ on the line $x = 1$, where the slope θ of segment OA_∞ will be referred to as the *deployment angle*, and hence $A_\infty = (1, \tan \theta)$ with $\theta \in [0, \pi/2)$. The reader may consult Figure 1 which shows the optimal inspective trajectory certifying Theorem 4. The initial deployment phase is followed by the so-called *inspection phase* shown as the green curve whose other endpoint is A_0 . In the discrete setting (Figure 3), the endpoint is denoted A_k , where k is proportional to a discretization parameter that tends to infinity. It is therefore natural to denote the limiting endpoint by A_∞ .

¹ The terminology reflects that ξ determines the initial deployment angle $\theta = (1 - \xi)\pi$ of the deployment phase, see Figure 1.



■ **Figure 1** The optimal inspective trajectory certifying Theorem 4. The trajectory starts with the deployment phase OA_∞ , and is followed by the green curve (inspection phase) \mathcal{T}^{τ_0} , solution to $\text{SYS}(\tau_0)$, for $\tau_0 \approx 1.64697686$, which is inspection-feasible with deployment parameter $\xi(\tau_0) \approx 0.81190987$. The inspection phase is therefore the curve $\mathcal{T}^{\tau_0}(x)$ for $x \in [0, \xi(\tau_0)]$, with endpoints $A_0 = \mathcal{T}^{\tau_0}(0)$ and $A_\infty = \mathcal{T}^{\tau_0}(\xi(\tau_0))$. The dotted blue line marks $x = 1$, corresponding to the feasibility condition $\mathcal{T}_1^{\tau_0}(\xi) = 1$. The deployment angle is $\theta = (1 - \xi(\tau_0))\pi$, with labelled points $P_0 = (1, 0)$, $A_\infty = (1, \tan(\theta))$, and $A_0 = (1, -\tau_0)$.

We are now ready to state the main technical result, which expresses the cost of ADI as a function of the curve \mathcal{T}^{τ_0} and of (ψ, τ) , the solution to $\text{SYS}(\tau_0)$. For ease of reference, and inspired by standard terminology in control theory, we call the resulting optimization problem a *Single-Parameter Optimal Control Problem* (SPOCP) where each feasible solution is the trajectory of an ODE system determined by a single initial condition serving as the decision variable. In our setting the initial parameter is τ_0 , and the corresponding feasible trajectories are generated by $\text{SYS}(\tau_0)$. We refer to this problem as $\text{SPOCP}(\tau_0)$, and is formally stated in Theorem 8 below.

► **Theorem 8** ($\text{SPOCP}(\tau_0)$ Formulation). *If the inspection phase of an optimal trajectory for ADI does not touch the unit disk, then the inspection phase is the subcurve $(\mathcal{T}^{\tau_0}(x))_{0 \leq x \leq \xi}$ for some inspection-feasible τ_0 , where (ψ, τ) solves $\text{SYS}(\tau_0)$, and $\xi = \xi(\tau_0) > 1/2$ is the deployment parameter of τ_0 . Together with the deployment segment $O\mathcal{T}^{\tau_0}(\xi)$, the optimal cost to ADI is*

$$\frac{1}{2\pi} \log\left(\frac{1 + \sin(\xi\pi)}{1 - \sin(\xi\pi)}\right) + \frac{\xi}{\cos((1 - \xi)\pi)} + 2\pi \int_0^\xi \frac{x \cdot \tau(x)}{\sin(\psi(x))} dx, \quad (1)$$

minimized over all inspection-feasible $\tau_0 \in \mathcal{I}$, where $\theta = (1 - \xi)\pi$ is the corresponding deployment angle.

We prove Theorem 8 in Section 5.1. In Section 5.3 we show that the inspection phase of the optimal inspective trajectories do not touch the unit disk. This allows us to invoke Theorem 8 and optimize (1), thereby solving the underlying one-parameter optimal control problem and establishing Theorem 4.

3 Background Machinery

We begin by reviewing the main tools introduced in [16], which we follow throughout Section 3 apart from minor reformulations. These ideas play only a preliminary role in our development, since they provide the initial reduction and notation on which our main arguments build. That earlier work proposed a discretized version of ADI together with a Nonlinear Programming (NLP) formulation whose solutions yield upper bounds. In Section 3.1 we introduce a related intermediate “partial” problem that simplifies the reduction from the continuous setting. Section 3.2 then describes the discretized problem and concludes with a brief sketch of the NLP formulation, which, while not directly relevant here, was previously used to derive upper bounds to ADI.

3.1 Reduction to the Partial Average-Case Disk-Inspection Problem

It is convenient to host our arguments on the Cartesian plane and require that the inspecting curves have one endpoint at the origin, where the disk to be inspected is also centered. More specifically, we parameterize the unit disk perimeter by points

$$P_\phi := (\cos(\phi), \sin(\phi)), \quad (2)$$

where $\phi \in [0, 2\pi)$. Therefore, the subject inspective trajectories are functions $T : [0, 1] \rightarrow \mathbb{R}^2$, with $T(0) = (0, 0)$. Towards defining the discretized ADI, we first need to introduce a useful *partial inspection* variant to the problem.

► **Definition 9** (*θ -Inspective Curves*). *Let $\theta \in [0, \pi/2)$. A continuous and piece-wise differentiable curve segment in \mathbb{R}^2 (i.e. a curve with two endpoints) is called θ -inspective if its convex hull contains a unit disk centered at a point which is $1/\cos(\theta)$ away from one of the curve’s endpoints.*

The motivation for introducing θ -inspective curves is that inspection may begin not from the disk center but from a point located at distance $1/\cos(\theta)$ from it. We reserve the term *trajectory* for full solutions to ADI, while *curve* refers to this partial variant. Placing the disk at the origin, we may assume that the starting point is $(1, \tan(\theta))$, which already inspects all boundary points P_ϕ with $\phi \in [0, 2\theta]$ at zero cost. An example is illustrated in Figure 1, where the θ -inspective curve has endpoints $A_\infty = (1, \tan(\theta))$ and A_0 on the line $x = 1$.

► **Problem 10** (*θ -Average-Disk-Inspection – θ -ADI*). Given $\theta \in [0, \pi/2)$, find a θ -inspective curve that minimizes $\mathbb{E}_\phi[I_{P_\phi}]$, where the expectation is over $\phi \in [2\theta, 2\pi]$ chosen uniformly at random.

Any θ -inspective curve can be extended to an inspective trajectory by appending the line segment from the origin to $(1, \tan(\theta))$ (previously referred to as the deployment phase). One of the results in [16] was the following explicit relation between the two problems.

► **Theorem 11.** *If θ -ADI admits a solution of average cost $s = s(\theta)$, then ADI admits a solution of average cost*

$$B_\theta(s) := \frac{1}{2\pi} \log \left(\frac{1 + \sin(\theta)}{1 - \sin(\theta)} \right) + \left(1 - \frac{\theta}{\pi} \right) \left(\frac{1}{\cos(\theta)} + s \right).$$

In this expression, the logarithmic term accounts for the average cost of inspecting P_ϕ with $\phi \in [0, 2\theta]$ during the initial deployment segment, while the remainder reflects the

contribution of the θ -inspective curve. Thus, for each fixed θ , the best partial cost $s_0(\theta)$ yields a full trajectory of cost $B_\theta(s_0)$. It follows that the optimal solution to ADI is given by

$$\inf_{\theta \in [0, \pi/2)} B_\theta(s_0). \quad (3)$$

The task is therefore reduced to selecting the deployment angle θ and designing the corresponding θ -inspective curve of minimal average cost $s_0(\theta)$. The starting point for our work is the formulation that lead to the upper bound on $B_\theta(s_0)$ established in [16].

Theorem 11 is the continuous analogue of Lemma 3 in [16], which is stated for a discrete inspection setting. The proof there is geometric and does not rely on discreteness: the total average cost splits into two independent contributions, namely (i) the average cost of inspecting the initial arc of length 2θ during the deployment phase, and (ii) the average cost of inspecting the remaining portion of the perimeter, regardless of whether it consists of finitely many or infinitely many points.

Every feasible ADI trajectory admits such a decomposition. Following the classical shoreline analysis of Isbell and Gluss, one may restrict attention to trajectories whose rim begins with a straight deployment segment from the origin to a point $(1, \tan \theta)$ for some $\theta \in [0, \pi/2)$, and is angularly monotone thereafter. For a fixed θ , the contribution of the continuation to the total ADI cost depends only on the average cost of the corresponding θ -ADI solution. If the continuation were not optimal for that θ -ADI instance, replacing it by a better θ -ADI curve would strictly improve the total ADI cost. Hence, the continuation of any optimal ADI solution must itself be an optimal solution of the associated θ -ADI problem.

3.2 Disk-Inspection via Discretization (and Nonlinear Programming)

We now introduce a discretized version of the partial inspection problem and show how it can be modeled as a nonconvex Nonlinear Program (NLP) with $\Theta(k)$ variables, yielding a $(1 + 1/k)$ -approximation to the continuous θ -inspection problem.

Fix $\theta \in [0, \pi/2)$ and $k \geq 5$. Define

$$\phi_i := 2\pi - (\pi - \theta) \frac{2i}{k}, \quad i = 0, \dots, k, \quad (4)$$

and let P_{ϕ_i} denote the corresponding $k + 1$ equidistant points on the arc of the unit circle of length $2\pi - 2\theta$, see Figure 3. For convenience, we will abbreviate P_{ϕ_i} by P_i when the index meaning is clear from context.

We assume $k \geq 5$ so that the angular spacing between consecutive discretization angles is strictly less than $\pi/2$; this condition is required to ensure that feasibility of the discretized formulation implies feasibility of the corresponding continuous inspection trajectory, as already used in [16].

► **Definition 12** ((θ, k) -Inspective Curves). *A continuous piecewise-differentiable curve segment in \mathbb{R}^2 is called (θ, k) -inspective if its convex hull contains all points P_0, \dots, P_k of a unit disk centered at a point $1/\cos(\theta)$ away from one endpoint of the curve.*

As $k \rightarrow \infty$, (θ, k) -inspective curves approximate θ -inspective curves. In particular, scaling a (θ, k) -inspective curve by a factor $1 + O(1/k)$ produces a θ -inspective curve, which is one interpretation of the upper bounds derived in [16].

► **Problem 13** ((θ, k) -Average-Disk-Inspection – (θ, k) -ADI). Given $\theta \in [0, \pi/2)$ and $k \geq 5$, find a (θ, k) -inspective curve that minimizes $\mathbb{E}_i[I_{P_i}]$, where i is chosen uniformly at random from $\{0, \dots, k\}$.

Let now $L_i(t)$ be the tangent line at P_i , parameterized by

$$L_i(t) := \begin{pmatrix} \cos(\phi_i) \\ \sin(\phi_i) \end{pmatrix} + t \begin{pmatrix} \sin(\phi_i) \\ -\cos(\phi_i) \end{pmatrix}, \quad t \geq 0. \quad (5)$$

Each line $L_i(t)$ passes through P_i at $t = 0$ and is tangent to the unit disk. For parameters $t_i \geq 0$, define $A_i := L_i(t_i)$. In particular, we fix $t_k = \tan(\theta)$ so that $A_k = (1, \tan(\theta))$. A candidate (θ, k) -inspective curve is then the polygonal chain

$$A_k \rightarrow A_{k-1} \rightarrow \cdots \rightarrow A_0,$$

see Figure 3.

The next lemma expresses the average inspection cost in terms of the parameters t_i .

► **Lemma 14.** *Let $\theta \in [0, \pi/2)$, $k \geq 5$, and $t = (t_0, \dots, t_k) \in \mathbb{R}_{\geq 0}^{k+1}$ with $t_k = \tan(\theta)$. Define $A_i = L_i(t_i)$. Then the polygonal chain $A_k A_{k-1} \cdots A_0$ is (θ, k) -inspective with average inspection cost*

$$C_{\theta, k}(t) := \frac{1}{k+1} \sum_{i=1}^k i \|A_i - A_{i-1}\|. \quad (6)$$

Proof. Points $A_i \in L_i$ inspect P_i , so the polygonal chain is (θ, k) -inspective. Each segment $A_{i-1}A_i$ contributes to the inspection time of exactly i points, and averaging over $k+1$ points gives (6). ◀

It was shown in [16] that if $t_i = \Omega(1/k)$ for all i , then the polygonal chain induces a θ -inspective curve for the corresponding continuous instance. Under this feasibility requirement, the discrete average $C_{\theta, k}(t)$ is a Riemann-sum approximation of the continuous average inspection cost associated with the induced θ -inspective curve, since inspection times vary continuously with the endpoint angle at scale $O(1/k)$. In particular, optimizing $C_{\theta, k}(t)$ over feasible discretizations yields asymptotically tight upper bounds on the continuous θ -ADI cost, as used in [16].

The expression $C_{\theta, k}(t)$ is non-convex in the variables t and θ . Under the feasibility constraints $t_i = \Omega(1/k)$, it serves as the objective of a nonlinear program characterizing the cost of discretized solutions to ADI. In what follows, we show how this program can be reduced to a single-parameter optimization problem and analyze its limiting behavior as $k \rightarrow \infty$, resulting in the promised SPOCP.


4 Fermat's Principle Solves (θ, k) -ADI

In this section we characterize any optimal trajectory to (θ, k) -ADI and its cost via a recursion based on Fermat's Principle. This marks the beginning of our new contributions.

4.1 The Principle of Least Time

We begin by introducing Snell's Law, along with terminology that will be used in subsequent sections. The exposition starts with the *Principle of Least Time*, also known as *Fermat's Principle*, which postulates that the trajectory of a light ray between two given points is the one that minimizes travel time. The principle has been confirmed through experimental observation and explains the rules of refraction in ray optics.

44:10 Optimal Average Disk-Inspection via Fermat's Principle



■ **Figure 2** Light refraction across two media M_1 and M_2 , separated by interface line ℓ with normal η . A ray from A_1 in M_1 crosses ℓ at L and continues in M_2 towards A_2 , forming angles α_1, α_2 with η . The same trajectory also applies in reverse, illustrating Snell's Law (Theorem 15).

Consider two media with constant speeds s_1, s_2 , separated by a line ℓ with normal η ; see Figure 2. Assume light has constant speed in each medium, hence it travels along a straight ray within each. A ray from A_1 to A_2 refracts at $L \in \ell$, forming angles α_1, α_2 with normal η and obeying Snell's Law below. For simplicity, we refer to phase velocity as speed.

► **Theorem 15 (Snell's Law).** *If the speed of light in M_1, M_2 is s_1, s_2 , respectively, then the incidence angle α_1 and refraction angle α_2 satisfy*

$$\frac{\sin(\alpha_1)}{\sin(\alpha_2)} = \frac{s_1}{s_2}.$$

Although Snell's Law is an experimental law of optics, it can be derived rigorously from Fermat's Principle. We present the formal claim next.

► **Lemma 16.** *Let M_1, M_2 be two media with constant speeds s_1, s_2 . Among all continuous paths connecting $A_1 \in M_1$ to $A_2 \in M_2$, the unique trajectory minimizing travel time is piecewise-linear path made of two straight segments A_1L^*, L^*A_2 , where $L^* \in \ell$ is chosen so that refraction at L^* satisfies Snell's Law.*

4.2 Optimal Solution to (θ, k) -ADI via Recursion

In this section we compute the optimal solution to (θ, k) -ADI and its cost using recursion derived from the optics principles of the previous section.

Fix $\theta \in [0, \pi/2)$ and $k \in \mathbb{Z}$, see Figure 3. A solution to (θ, k) -ADI is determined by values $t_i \geq 0$ that specify points $A_i = L_i(t_i)$ on tangent halflines $L_i(t)$, $i = 0, \dots, k$, which in turn inspect perimeter points P_i . For given t_i we define the following counterclockwise angles:

$$\begin{aligned} x_i &:= \text{angle formed by } A_iA_{i-1} \text{ and } A_iP_i, \quad i = 1, \dots, k, \\ y_i &:= \text{angle formed by } A_iA_{i+1} \text{ and } L_i(t), \quad t \geq t_i, \quad i = 0, \dots, k-1, \end{aligned}$$


and distances

$$d_i := \|A_i - A_{i-1}\|, \quad i = 1, \dots, k.$$

Let also denote

$$\alpha := \frac{2(\pi - \theta)}{k},$$

that is, α is the angular distance between two consecutive points P_i, P_{i+1} on the perimeter.



■ **Figure 3** Geometric setup of the discrete trajectory for (θ, k) -ADI. Points P_i on the perimeter, tangent lines L_i , and trajectory points A_i define angles x_i, y_i and distances d_i . The regions M_i represent optical media used in the refraction based interpretation leading to the recursions of Lemma 17.

► **Lemma 17.** *Given $\theta \in [0, \pi/2)$ and $k \geq 5$, suppose an optimal trajectory to (θ, k) -ADI is identified by $A_i = L_i(t_i)$ with $t_i \geq \tan(\frac{\alpha}{2})$ (so it does not intersect the unit disk).² Then the trajectory and its cost are characterized by the recursions*

$$x_i = y_{i-1} - \alpha, \tag{7}$$

$$y_i = \arccos\left(\frac{i}{i+1} \cos(y_{i-1} - \alpha)\right), \tag{8}$$

$$t_i = \left(t_{i-1} - \tan\left(\frac{\alpha}{2}\right)\right) \frac{\sin(y_{i-1})}{\sin(x_i)} - \tan\left(\frac{\alpha}{2}\right), \tag{9}$$

$$d_i = \left(t_{i-1} - \tan\left(\frac{\alpha}{2}\right)\right) \frac{\sin(\alpha)}{\sin(x_i)}, \tag{10}$$

for $i = 1, \dots, k$. The initial conditions are $y_0 = \pi/2$ and $t_k = \tan(\theta)$.

5 Single-Parameter Reformulation and Continuum Limit

The backbone of the argument relies on (3), which states that the optimal cost to ADI can be computed as $\inf_{\theta \in [0, \pi/2)} B_\theta(s)$, where $s = s(\theta)$ is the optimal solution to θ -ADI. We show that $B_\theta(s)$ can, under suitable conditions, be minimized as a SPOCP in parameter τ_0 . Thus

² The condition $t_i \geq \tan(\frac{\alpha}{2})$ is the standard geometric feasibility requirement in the (θ, k) -inspection framework of [16], ensuring that the discrete trajectory does not intersect the unit disk. We note that in the parameter ranges relevant to our later analysis and numerical evaluations (with k large), the optimized values of t_i are observed to be bounded well away from 0, so this condition is comfortably satisfied in practice.

our first task is to verify that these conditions hold for the deployment parameter θ and the corresponding τ_0 that generate solutions to $\text{SYS}(\tau_0)$, and then to make the relation between θ and τ_0 explicit.

Section 5.1 formulates the minimization of $B_\theta(s(\theta))$ as a SPOCP(τ_0). Section 5.2 restricts the range of deployment angles θ so that the conditions apply. Section 5.3 relates θ and τ_0 and solves the resulting SPOCP(τ_0) numerically. Section 5.4 discusses the robustness of these numerical results.

5.1 A Single Parameter Optimal Control Problem – Proof of Theorem 8

We now prove Theorem 8. The starting point is the recurrence of Lemma 17, which describes the optimal inspective curve as an optics trajectory with refraction angles independent of $t_k = \tan(\theta)$. The deployment angle θ still determines the trajectory and its cost, and fixes the endpoint $A_0 = L_0(t_0) = (1, -t_0)$. Thus the trajectory connecting two points of the line $x = 1$ in the first and fourth quadrants (see Figure 3) may be viewed either as a ray starting from A_k in the first quadrant or from A_0 in the fourth.

For technical reasons, rather than parameterizing the recurrence by $t_k = \tan(\theta)$, we work with $t_0 = \tau_0$, restricted to the feasible range of Definition 7. The challenge is then to determine the point where the trajectory intersects the line $x = 1$ again, and hence the corresponding deployment angle θ . In Lemma 17 the angular step is $\alpha = 2(\pi - \theta)/k$, still vanishing as $k \rightarrow \infty$ but expressed in terms of θ . To eliminate this dependence we place n equispaced points P_0, \dots, P_n on the unit circle, with step $\alpha = 2\pi/n$, and recover θ as a function of τ_0 .

This setup leads to the following continuum limit as $n \rightarrow \infty$, namely the ODE system $\text{SYS}(\tau_0)$ of Definition 5. The corresponding trajectory \mathcal{T} begins at $(1, -\tau_0)$, remains outside the unit disk, and intersects the line $x = 1$ in the first quadrant at some point $\mathcal{T}(\xi)$, where ξ is the deployment parameter. We now state a sequence of technical results establishing this limit.

First, we show that if the discrete recursion is extended to continuous functions by connecting consecutive values linearly, the resulting piecewise-linear functions converge to the solution of the ODE system $\text{SYS}(\tau_0)$ of Definition 5.

► **Lemma 18.** *Let $(y_i)_{i=0}^n$ and $(t_i)_{i=0}^n$ be defined by (7)–(9) with $\alpha = 2\pi/n$, $y_0 = \pi/2$, and $t_0 = \tau_0$. Define piecewise-linear interpolants $\psi_n, \tau_n : [0, 1] \rightarrow \mathbb{R}$ by $\psi_n(i/n) = y_i$ and $\tau_n(i/n) = t_i$, extended linearly on each $[i/n, (i+1)/n]$. Then (ψ_n, τ_n) converges uniformly on compact subsets of $(0, 1]$ to (ψ, τ) , the unique solution to the ODE system $\text{SYS}(\tau_0)$ with $\psi(0) = \pi/2$ and $\tau(0) = \tau_0$.*

We are now ready to provide the limiting behavior of the inspecting curve, giving rise to \mathcal{T}^{τ_0} as in Definition 6. In the remainder of the section, for notational convenience, we drop the superscript and write simply \mathcal{T} . The next lemma shows that the polygonal trajectories converge indeed to \mathcal{T} .

► **Lemma 19.** *Let ϕ_i and $L_i(t)$ be defined as in (4) and (5), and let $A_i = L_i(t_i)$, where (t_i) is defined by (9) with $t_0 = \tau_0$. Assume that τ_0 is inspection-feasible (Definition 7), and let $\xi = \xi(\tau_0)$ be its deployment parameter, that is, the smallest $\xi \in (0, 1]$ such that $\mathcal{T}_1(\xi) = 1$ for $\mathcal{T} = \mathcal{T}^{\tau_0}$ of Definition 6. Let $\tilde{\mathcal{T}}_n$ be the polygonal path obtained by linear interpolation through the points $(A_i)_{i=0}^n$ with parameter values $h_i = i/n$.*

Then $\tilde{\mathcal{T}}_n$ converges to \mathcal{T} on $[0, \xi]$, in the sense that for every $0 < \delta \leq \xi' < \xi$,

$$\sup_{x \in [\delta, \xi']} \|\tilde{\mathcal{T}}_n(x) - \mathcal{T}(x)\| \rightarrow 0 \quad \text{as } n \rightarrow \infty.$$

Moreover, $\mathcal{T}(0) = A_0$.

Since \mathcal{T} arises as the continuum limit of the discrete trajectories, the endpoint A_k of the polygonal path (see Figure 3) corresponds to an index $k = \Theta(n)$. As $n \rightarrow \infty$ this endpoint is denoted A_∞ in Figure 1. We can now justify the definition of the deployment parameter.

► **Lemma 20.** *Let τ_0 be inspection-feasible to the ODE system of Definition 5. Then its deployment parameter $\xi = \xi(\tau_0)$ satisfies $\mathcal{T}_2(\xi) = \tan(\theta)$, where $\theta = (1 - \xi)\pi$.*

From the progress above, and starting with inspection-feasible τ_0 , we compute θ -inspective curve to θ -ADI, where $\theta = \theta(\xi)$, and $\xi = \xi(\tau_0)$. The next lemma uses again the continuum construction to derive the cost of \mathcal{T} to θ -ADI, as a function of τ_0 , and in terms of the solution to ODE system $\text{SYS}(\tau_0)$.

► **Lemma 21.** *Let τ_0 be inspection-feasible with deployment parameter $\xi = \xi(\tau_0)$. Then \mathcal{T} is a feasible solution to θ -ADI, where $\theta = (1 - \xi)\pi$, and the average cost equals*

$$\frac{2\pi}{\xi} \int_0^\xi \frac{x\tau(x)}{\sin(\psi(x))} dx.$$

We can now conclude the proof of Theorem 8. By (3) and the discussion of Section 3.1, the optimal inspective curve has cost $\inf_{\theta \in [0, \pi/2)} B_\theta(s_0)$, where $s_0 = s_0(\theta)$ is the cost to some θ -ADI instance. Since the optimal curve does not touch the disk, the deployment angle is $\theta = (1 - \xi)\pi$ for some inspection-feasible τ_0 . Lemma 21 gives s_0 as a function of ξ , which can then be substituted into Theorem 11 to yield the expression of Theorem 8.

Finally, since $\theta \in [0, \pi/2)$ we have $\xi > 1/2$, ensuring that expression (1) is well defined. To conclude, it remains to show that the expression attains a minimum, which follows from continuity together with the Extreme Value Theorem.

► **Lemma 22.** *Fix an inspection-feasible initial value τ_0 with deployment parameter $\xi \in (1/2, 1]$. Define*

$$I(\xi) := 2\pi \int_0^\xi \frac{x\tau(x)}{\sin(\psi(x))} dx.$$

Then $I(\xi)$ is well defined and belongs to $C^1([0, \xi])$, with

$$I'(\xi) = 2\pi\xi \frac{\tau(\xi)}{\sin(\psi(\xi))}.$$

5.2 Bounds on the Optimal Deployment Angle

In light of Theorem 8 proved in the previous section, the natural approach to finding the optimal solution to ADI is to solve the underlying single-parameter (here τ_0) optimal control problem with objective (1). Put differently, we are back to determining the optimal solution as $\inf_{\theta \in [0, \pi/2)} B_\theta(s_0)$, where $s_0 = s_0(\theta)$ is the optimal solution to θ -ADI, see (3). However, not all deployment angles θ give rise to optimal trajectories that avoid intersecting the disk, which is a premise of Theorem 8.

For this reason, we restrict the range $[0, \pi/2)$ of deployment angles. Small values of θ are excluded because the corresponding optimal trajectories touch the disk. Large values of θ are excluded for a different reason: as $\theta \rightarrow \pi/2$, the deployment point moves arbitrarily far from the origin, and we show in the next section that such angles cannot be optimal. Consequently, these angles need not be considered in the optimization. We then search numerically for the minimizing deployment angle over the remaining range. Restricting to this range also avoids numerical instability in the associated SPOCP, which would otherwise arise from the large scales involved. The numerical procedures used in this restricted regime are justified in Section 5.4.

To summarize, in this section we show a refinement of (3), as follows.

44:14 Optimal Average Disk-Inspection via Fermat's Principle

► **Lemma 23.** *The optimal solution to ADI is given by*

$$\inf_{\theta \in [0.52, 1.148]} B_\theta(s_0).$$

We start by showing that the deployment angle cannot be too large.

► **Lemma 24.** $\theta_0 \leq 1.148$ *for the optimal deployment angle θ_0 to ADI.*

Then, we need a lower bound to the optimal solution to θ -ADI.

► **Lemma 25.** *Fix $\theta \in [0, \pi/2)$ and $k \geq 5$. Then, the optimal solution to θ -ADI is at least the minimum of*

$$\sum_{i=1}^k \frac{i-1}{k} \|A_{i-1}A_i\|$$

over all points A_i that form (θ, k) -inspective curves.

We are ready to show that the deployment angle cannot be too small.

► **Lemma 26.** $\theta_0 \geq 0.52$ *for the optimal deployment angle θ_0 to ADI.*

Note that Lemma 23 is directly implied by Lemmata 24 and 26.

5.3 The Optimal Solution to ADI – Proof of Theorem 4

In this section, we show that the premise of Theorem 8 is satisfied, namely that the optimal inspective curve does not touch the unit disk. Consequently, this allows us to optimize expression (1) and thus conclude with the proof of Theorem 4. We rely on the theoretical foundations established previously, and, as is necessary, we make use of further numerical calculations implemented in the `Julia` programming language [11]. Many of the arguments below are based on numerical comparisons. In Section 5.4 we provide the justifications that these computations are robust and consistent with the accuracy promised in the main theorem.

The next lemma analyzes inspection trajectories in a carefully chosen regime of initial values τ_0 to the ODE system $\text{SYS}(\tau_0)$.

► **Lemma 27.** *Every $\tau_0 \in [1.64697, 1.6525]$ is inspection feasible, and the corresponding deployment parameters $\xi = \xi(\tau_0)$ determine deployment angles $\theta = (1 - \xi)\pi$ whose range covers interval $[0.52, 1.148]$.*

Proof. We solve the ODE system \mathcal{T}^{τ_0} of Definition 5 on a grid of 2000 sample points for initial conditions $\tau_0 \in [1.64697, 1.6525]$. The resulting values are summarized in Figure 4. Recall that the initial condition $\tau(0) = \tau_0$ determines ξ , θ and the entire curve $\mathcal{T}(\cdot)$.

Figure a reports $\xi = \xi(\tau_0)$, the parameter with $\mathcal{T}(0) = A_0$ and $\mathcal{T}(\xi) = A_\infty$. In Figure b, we compute the minimum of $\tau(x)$ over $x \in [0, \xi]$, which is bounded below by 0.2. Since $\|\mathcal{T}(x)\| = \sqrt{1 + \tau(x)^2}$, the distance of $\mathcal{T}(x)$ from the unit circle is $\sqrt{1 + \tau(x)^2} - 1$. Therefore a uniform bound $\tau(x) \geq 0.2$ implies a radial clearance of at least


$$\sqrt{1 + 0.2^2} - 1 \approx 0.01980198 \dots$$

This proves inspection feasibility for every τ_0 in the stated range. Accuracy of the ODE integration and stability checks are deferred to Section 5.4.

Finally, Figure c shows the corresponding deployment angles $\theta = (1 - \xi)\pi$. The endpoints satisfy

$$\theta(1.64697) \approx 0.501177\dots \quad \text{and} \quad \theta(1.6525) \approx 1.1600947\dots,$$

and the image of $[1.64697, 1.6525]$ under $\theta(\cdot)$ contains $[0.52, 1.148]$, as indicated by the horizontal reference lines in the figure. ◀



(a) Plot of the deployment parameter $\xi = \xi(\tau_0)$ against τ_0 . (b) Plot of the minimum value of $\tau(x)$, where $x \in [0, \xi]$, against τ_0 . (c) Plot of the deployment angle $\theta = \theta(\tau_0)$ against τ_0 .

■ **Figure 4** Plots of parameters of trajectory \mathcal{T} as obtained by the solution to the ODE system of Definition 5 for initial conditions $\tau_0 \in [1.64697, 1.6525]$.

We are now ready to conclude with the proof of Theorem 4.

Proof of Theorem 4. By Lemma 23, the optimal cost to ADI is

$$\inf_{\theta \in [0.52, 1.148]} B_\theta(s(\theta)).$$

By Lemma 27, all $\tau_0 \in [1.64697, 1.6525]$ are inspection feasible, and the corresponding deployment angles cover $[0.52, 1.148]$. Thus the inspective curves do not touch the disk, and Theorem 8 applies.


Therefore, in order to determine the optimal solution it remains to minimize the cost expression of Theorem 8 over the admissible trajectories. The ODE formulation shows that each trajectory is uniquely determined by the initial condition $\tau(0) = \tau_0$, so that the cost becomes a function of this single parameter. Hence the problem reduces to SPOCP(τ_0), namely the problem of minimizing (1) over $\tau_0 \in [1.64697, 1.6525]$.

Since the admissible family is one-dimensional, we evaluate this function numerically and refine the search interval around the minimizer. Figure 5 summarizes this coarse-to-fine evaluation over increasingly refined intervals of 2000 grid points each. This refinement brackets the minimizer and provides the required numerical precision.

The minimum is sandwiched between 3.5492598 and 3.54925986. For

$$\tau_0 = 1.6469768608776936$$

(exact value) we obtain $\xi = 0.8119098734258519\dots$, $\theta = 0.5909025598581181\dots$, and a corresponding inspective curve with minimum distance at least $0.0302318\dots$ to the disk boundary (corresponding to $\tau(\chi) \approx 0.24774522\dots$ for some χ), and reported cost approximately $3.5492595860809693\dots$. Numerical accuracy of this computation is discussed in Section 5.4. ◀



■ **Figure 5** Plot of the cost to ADI as given by Theorem 8 against various initial conditions $\tau(0) = \tau_0$ shown on the horizontal axis. The vertical green dotted line shows $\tau_0 = 1.64697686$. In Figure b, the dotted red horizontal line corresponds to value 3.5492595. In Figure c, the dotted red horizontal lines correspond to values 3.549259, 3.549260.

5.4 Numerical Methods and Accuracy Guarantees

We used the Julia programming language [11] for all computations.

For the convex nonlinear program that lower bounds $s(\theta)$ in Lemma 26, we model it with JuMP [19] and solve it using the interior point method Ipopt [14, 42]. The program is convex since the feasible set $\{t_i \geq 0\}$ is convex and the objective is a nonnegative conical combination of norms of affine maps. Hence any KKT point is globally optimal. In practice, the solver returns primal and dual feasible solutions with residuals below 10^{-14} and objective values stable to at least 10^{-9} . Because convexity guarantees global optimality, these certificates validate the solutions and support the digits reported in our lower bounds.

For the ODE system of Definition 5, used in Section 5.3 to establish feasibility of trajectories and to evaluate the cost functional, we proceed as follows. The equation for ψ has a singularity at $x = 0$, so we begin the integration at $x_0 = 10^{-6}$ using the asymptotic expansion $\psi(x_0) = \pi/2 - \pi x_0 + (\pi^2/2)x_0^2$, whose truncation error is $O(x_0^3) \approx 10^{-18}$ and therefore negligible compared with the solver tolerances. We integrate ψ on $[x_0, 1]$ with the adaptive stiff solver Rodas5() from the DifferentialEquations.jl library [41], setting absolute and relative tolerances to 10^{-12} , and then integrate τ on the same interval using the dense output of ψ in the right-hand side.

The deployment parameter ξ is obtained in two stages. A uniform grid of 10,000 points provides a coarse ξ_{approx} as the last feasible point according to a geometric predicate that tolerates floating-point noise at the level 10^{-12} . This value is refined by bisection to absolute tolerance 10^{-8} . To further confirm stability, we repeat the bisection at tolerance $5 \cdot 10^{-10}$ and recompute the objective, reporting the discrepancy between the two runs, which is consistently negligible. In addition, the minimum of $\tau(\cdot)$ on $[x_0, \xi_{\text{approx}}]$ is computed by Brent's method [13] with tolerance 10^{-10} to certify that all trajectories remain uniformly outside the disk.

The integral of Lemma 21 is evaluated using the adaptive Gauss-Kronrod quadrature routine quadgk from QuadGK.jl [33], with relative tolerance 10^{-12} and absolute tolerance 10^{-14} . Substituting (ξ, θ) into the expression of Theorem 8 then yields the values reported in Section 5.3.

Each source of numerical error is explicitly controlled. The asymptotic initialization at $x_0 = 10^{-6}$ contributes error below 10^{-18} . The ODE solver controls local error to 10^{-12} . The quadrature routine bounds the relative error for computing integral of Lemma 21 to 10^{-12} . The binary search tolerance 10^{-8} induces uncertainty on θ of at most $\pi \cdot 10^{-8}$, with an even tighter self-check available. Finally, the Brent minimization shows that $\tau(x) \geq 0.2$

throughout, implying $\|\mathcal{T}(x)\| \geq 1.0198$, so the curves remain at least 10^{-2} outside the disk. These guarantees confirm that the cost values reported in Theorem 4 are reliable to at least six decimal digits.

6 Discussion

Bellman introduced the famous lost-in-a-forest problem and proposed several variants almost seventy years ago [9]. In this work we resolve one of these variants, the Average-Case Disk-Inspection problem. The line of inquiry began with the heuristics of Gluss [29] in the 1960s and continued through the discretization framework developed recently in [15, 16]. Our analysis not only establishes the exact optimum with certified numerical accuracy, but also reveals the structural nature of optimal trajectories. We show that they arise from a reformulation of the problem as an optics model based on Fermat's Principle of Least Time, which leads to a single-parameter ODE system. Crucially, the resulting optimal trajectories avoid the unit disk, contrary to the conjecture of Gluss.

Beyond closing this specific problem, the methods introduced here suggest a possible direction for approaching other geometric search questions. The reformulation of a many-variable nonconvex program into a single-parameter optimal control problem shows how optics-inspired principles can reduce complexity and expose structure that is otherwise hidden. While our techniques were developed for the disk-inspection setting, the interplay between discrete recursions and continuum limits may find use in related problems where discretization has been the standard tool but has remained difficult to analyze or scale.

References

- 1 S. Acharjee, K. Georgiou, S. Kundu, and A. Srinivasan. Lower bounds for shoreline searching with 2 or more robots. In *23rd OPODIS*, volume 153 of *LIPIcs*, pages 26:1–26:11. Schloss Dagstuhl – LZI, 2019. doi:10.4230/LIPIcs.OPODIS.2019.26.
- 2 R. Ahlswede and I. Wegener. *Search problems*. John Wiley & Sons, Inc., 1987.
- 3 S. Alpern, R. Fokkink, L. Gasieniec, R. Lindelauf, and V. S. Subrahmanian. *Search theory*. Springer, 2013.
- 4 S. Alpern and S. Gal. *The theory of search games and rendezvous*, volume 55. Springer Science & Business Media, 2006.
- 5 R. Baeza-Yates. Searching: an algorithmic tour. *Encyclopedia of Computer Science and Technology*, 37:331–359, 1997.
- 6 R. Baeza-Yates and R. Schott. Parallel searching in the plane. *Computational Geometry*, 5(3):143–154, 1995. doi:10.1016/0925-7721(95)00003-R.
- 7 R. A. Baeza-Yates, J. C. Culberson, and G. J. E. Rawlins. Searching with uncertainty. In *Scandinavian Workshop on Algorithm Theory*, pages 176–189. Springer, 1988.
- 8 R. A. Baeza-Yates, J. C. Culberson, and G. J. E. Rawlins. Searching in the plane. *Information and computation*, 106(2):234–252, 1993. doi:10.1006/INCO.1993.1054.
- 9 R. Bellman. Minimization problem. *Bull. Amer. Math. Soc.*, 62(3):270, 1956.
- 10 G. Berzsenyi. Lost in a forest (a problem area initiated by the late Richard E. Bellman). *Quantum (November/December, 1995)*, 41, 1995.
- 11 J. Bezanson, A. Edelman, S. Karpinski, and V. B. Shah. Julia: A fresh approach to numerical computing. *SIAM review*, 59(1):65–98, 2017. doi:10.1137/141000671.
- 12 S. Bouchard, Y. Dieudonné, A. Pelc, and F. Petit. Deterministic treasure hunt in the plane with angular hints. In *29th International Symposium on Algorithms and Computation, ISAAC 2018*, volume 123, pages 48:1–48:13. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2018. doi:10.4230/LIPIcs.ISAAC.2018.48.
- 13 R. P. Brent. *Algorithms for minimization without derivatives*. Courier Corporation, 2013.

- 14 COIN-OR. Ipopt: Interior point optimizer. <https://github.com/coin-or/Ipopt>. Accessed: 2024-06-19.
- 15 J. Conley and K. Georgiou. Disk and partial disk inspection: Worst- to average-case and Pareto upper bounds. *arXiv preprint arXiv:2411.15391*, 2025.
- 16 J. Conley and K. Georgiou. Multi-agent disk inspection. In Ulrich Schmid and Roman Kuznets, editors, *Structural Information and Communication Complexity*, pages 262–280, Cham, 2025. Springer Nature Switzerland.
- 17 J. Czyzowicz, K. Georgiou, and E. Kranakis. Group search and evacuation. In P. Flocchini, G. Prencipe, and N. Santoro, editors, *Distributed Computing by Mobile Entities; Current Research in Moving and Computing*, chapter 14, pages 335–370. Springer, 2019. doi:10.1007/978-3-030-11072-7_14.
- 18 S. Dobrev, R. Kráľovič, and D. Pardubská. Improved lower bounds for shoreline search. In *International Colloquium on Structural Information and Communication Complexity*, pages 80–90. Springer, 2020.
- 19 I. Dunning, J. Huchette, and M. Lubin. Jump: A modeling language for mathematical optimization. *SIAM review*, 59(2):295–320, 2017. doi:10.1137/15M1020575.
- 20 Y. Emek, T. Langner, D. Stolz, J. Uitto, and R. Wattenhofer. How many ants does it take to find the food? *Theoretical Computer Science*, 608:255–267, 2015. doi:10.1016/J.TCS.2015.05.054.
- 21 Y. Emek, T. Langner, J. Uitto, and R. Wattenhofer. Solving the ants problem with asynchronous finite state machines. In *Proceedings of International Colloquium on Automata, Languages, and Programming (ICALP), LNCS 8573*, pages 471–482, 2014.
- 22 S. R. Finch and J. E. Wetzel. Lost in a forest. *The American Mathematical Monthly*, 111(8):645–654, 2004. URL: <http://www.jstor.org/stable/4145038>.
- 23 S. R. Finch and L.-Y. Zhu. Searching for a shoreline. *arXiv preprint math/0501123*, 2005.
- 24 G. M. Fricke, J. P. Hecker, A. D. Griego, L. T. Tran, and M. E. Moses. A distributed deterministic spiral search algorithm for swarms. In *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 4430–4436. IEEE, 2016.
- 25 S. Gal. Search games. *Wiley Encyclopedia of Operations Research and Management Science*, 2010.
- 26 K. Georgiou. Optimal average disk-inspection via fermat's principle. *arXiv:2509.06334 preprint*, 2025.
- 27 K. Georgiou, C. Jones, and M. Madej. Spirals and beyond: Competitive plane search with multi-speed agents. arXiv:2508.10793, 2025. Accepted to LATIN 2026.
- 28 P. Gibbs. Bellman's escape problem for convex polygons, 2016.
- 29 B. Gluss. An alternative solution to the “lost at sea” problem. *Naval Research Logistics Quarterly*, 8(1):117–122, 1961.
- 30 B. Gluss. The minimax path in a search for a circle in a plane. *Naval Research Logistics Quarterly*, 8(4):357–360, 1961.
- 31 J. R. Isbell. An optimal search pattern. *Naval Research Logistics Quarterly*, 4(4):357–359, 1957.
- 32 A. Jež and J. Łopuszański. On the two-dimensional cow search problem. *Information Processing Letters*, 109(11):543–547, 2009. doi:10.1016/J.IPL.2009.01.020.
- 33 Steven G. Johnson. QuadGK.jl: Gauss–Kronrod integration in Julia. <https://github.com/JuliaMath/QuadGK.jl>, 2013.
- 34 D. Kübel and E. Langetepe. On the approximation of shortest escape paths. *Computational Geometry*, 93:101709, 2021. doi:10.1016/J.COMGEO.2020.101709.
- 35 E. Langetepe. On the optimality of spiral search. In *Proceedings of the twenty-first annual ACM-SIAM symposium on Discrete Algorithms*, pages 1–12. SIAM, 2010.
- 36 E. Langetepe. Searching for an axis-parallel shoreline. *Theoretical Computer Science*, 447:85–99, 2012. doi:10.1016/J.TCS.2011.12.069.

- 37 T. Langner, B. Keller, J. Uitto, and R. Wattenhofer. Overcoming obstacles with ants. In E. Anceaume, C. Cachin, and M. G. Potop-Butucaru, editors, *International Conference on Principles of Distributed Systems (OPODIS)*, volume 46 of *LIPICs*, pages 9:1–9:17. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2015. doi:10.4230/LIPICs.OPODIS.2015.9.
- 38 A. Pelc. Reaching a target in the plane with no information. *Information Processing Letters*, 140:13–17, 2018. doi:10.1016/J.IPL.2018.04.006.
- 39 A. Pelc and R. N. Yadav. Information complexity of treasure hunt in geometric terrains. *arXiv preprint arXiv:1811.06823*, 2018.
- 40 A. Pelc and R. N. Yadav. Cost vs. information tradeoffs for treasure hunt in the plane. *arXiv preprint arXiv:1902.06090*, 2019.
- 41 C. Rackauckas and Q. Nie. Differentialequations.jl—a performant and feature-rich ecosystem for solving differential equations in julia. *Journal of open research software*, 5(1):15–15, 2017.
- 42 A. Wächter and L. T. Biegler. On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming. *Mathematical programming*, 106(1):25–57, 2006. doi:10.1007/S10107-004-0559-Y.