

Worst-Case Optimal Exploration of Terrains with Obstacles[☆]

Jurek Czyzowicz^{a,1}, David Ilcinkas^{b,3}, Arnaud Labourel^{c,2,3,*}, Andrzej Pelc^{a,4}

^a *Département d'informatique, Université du Québec en Outaouais, Gatineau, Québec J8X 3X7, Canada.*

^b *LaBRI, CNRS & Université de Bordeaux, 33405 Talence, France.*

^c *LIF, Aix-Marseille Université, 13288 Marseille, France.*

Abstract

A mobile robot represented by a point moving in the plane has to explore an unknown flat terrain with impassable obstacles. Both the terrain and the obstacles are modeled as arbitrary polygons. We consider two scenarios: the *unlimited vision*, when the robot situated at a point p of the terrain explores (sees) all points q of the terrain for which the segment pq belongs to the terrain, and the *limited vision*, when we require additionally that the distance between p and q is at most 1. All points of the terrain (except obstacles) have to be explored and the performance of an exploration algorithm, called its complexity, is measured by the length of the trajectory of the robot.

For unlimited vision we show an exploration algorithm with complexity $O(P + D\sqrt{k})$, where P is the total perimeter of the terrain (including perimeters of obstacles), D is the diameter of the convex hull of the terrain, and k is the number of obstacles. We do not assume knowledge of these parameters. We also prove a matching lower bound showing that the above complexity is optimal, even if the terrain is known to the robot. For limited vision we show exploration algorithms with complexity $O(P + A + \sqrt{Ak})$, where A is the area of the terrain (excluding obstacles). Our algorithms work either for arbitrary terrains (if one of the parameters A or k is known) or for c -fat terrains, where c is any constant (unknown to the robot) and no additional knowledge is assumed. (A terrain

[☆]A preliminary version of this paper appeared in Proc. 12th Scandinavian Symposium and Workshops on Algorithm Theory (SWAT 2010), LNCS 6139.

*Corresponding author

Email addresses: jurek@uqo.ca (Jurek Czyzowicz), david.ilcinkas@labri.fr (David Ilcinkas), arnaud.labourel@lif.univ-mrs.fr (Arnaud Labourel), pelc@uqo.ca (Andrzej Pelc)

¹Partially supported by NSERC discovery grant.

²This work was done during this author's stay at the Université du Québec en Outaouais as a postdoctoral fellow.

³Partially supported by the ANR project ALADDIN, the INRIA project CEPAGE and by a France-Israel cooperation grant (Multi-Computing project).

⁴Partially supported by NSERC discovery grant and by the Research Chair in Distributed Computing at the Université du Québec en Outaouais.

\mathcal{T} with obstacles is c -fat if $R/r \leq c$, where R is the radius of the smallest disc containing \mathcal{T} and r is the radius of the largest disc contained in \mathcal{T} .) We also prove a matching lower bound $\Omega(P + A + \sqrt{Ak})$ on the complexity of exploration for limited vision, even if the terrain is known to the robot.

Keywords: mobile robot, exploration, polygon, obstacle.

1. Introduction

The background and the problem. Exploring unknown terrains by mobile robots has important applications when the environment is dangerous or of difficult access for humans. Such is the situation when operating in nuclear plants or cleaning toxic wastes, as well as in the case of underwater or extra-terrestrial operations. In many cases a robot must inspect an unknown terrain and come back to its starting point. Due to energy and cost saving requirements, the length of the robot's trajectory should be minimized.

We model the exploration problem as follows. The terrain is represented by an arbitrary polygon \mathcal{P}_0 with pairwise disjoint polygonal obstacles $\mathcal{P}_1, \dots, \mathcal{P}_k$, included in \mathcal{P}_0 , i.e., the terrain is $\mathcal{T} = \mathcal{P}_0 \setminus (\mathcal{P}_1 \cup \dots \cup \mathcal{P}_k)$. We assume that boundaries of all polygons \mathcal{P}_i belong to the terrain. The robot is modeled as a point moving along a polygonal line inside the terrain. It should be noted that the restriction to polygons is only to simplify the description, and all our results hold in the more general case where polygons are replaced by bounded subsets of the plane homeotopic with a disc (i.e., connected and without holes) and regular enough to have well-defined area and boundary length. Every point of the trajectory of the robot is called *visited*. We consider two scenarios: the *unlimited vision*, when the robot visiting a point p of the terrain \mathcal{T} *explores* (sees) all points q for which the segment pq is entirely contained in \mathcal{T} , and the *limited vision*, when we require additionally that the distance between p and q is at most 1. In both cases the task is to explore all points of the terrain \mathcal{T} . The cost of an exploration algorithm is the length of the trajectory of the robot, which should be as small as possible. The *complexity* of an algorithm is the order of magnitude of its cost. We assume that the robot does not know the terrain before starting the exploration, but it has unbounded memory and can record the portion of the terrain seen so far and the already visited portion of its trajectory.

Our results. For unlimited vision we show an exploration algorithm with complexity $O(P + D\sqrt{k})$, where P is the total perimeter of the terrain (including perimeters of obstacles), D is the diameter of the convex hull of the terrain, and k is the number of obstacles. We do not assume knowledge of these parameters. We also prove a matching lower bound for exploration of some terrains (even if the terrain is known to the robot), showing that the above complexity is worst-case optimal.

For limited vision we show exploration algorithms with complexity $O(P +$

$A + \sqrt{Ak}$), where A is the area of the terrain¹. Our algorithms work either for arbitrary terrains, if one of the parameters A or k is known, or for c -fat terrains, where c is any constant larger than 1 (unknown to the robot) and no additional knowledge is assumed. (A terrain \mathcal{T} is c -fat if $R/r \leq c$, where R is the radius of the smallest disc containing \mathcal{T} and r is the radius of the largest disc contained in \mathcal{T} .) We also prove a matching lower bound $\Omega(P + A + \sqrt{Ak})$ on the complexity of exploration, even if the terrain is known to the robot.

The main open problem resulting from our research is whether exploration with asymptotically optimal cost $O(P + A + \sqrt{Ak})$ can be performed in arbitrary terrains without *any* a priori knowledge. Another interesting open problem is whether such worst-case performance can be obtained by an $O(k)$ -competitive algorithm. (Our algorithms are a priori not competitive.)

Related work. Exploration of unknown environments by mobile robots was extensively studied for both the unlimited and limited vision. Most of the research in this domain concerns the competitive framework, where the trajectory of the robot not knowing the environment is compared to that of the optimal exploration algorithm having full knowledge.

One of the most important works for unlimited vision is [1]. The authors gave a 2-competitive algorithm for rectilinear polygon exploration without obstacles. The case of non-rectilinear polygons (without obstacles) was also studied in [2, 3, 4] and competitive algorithms were given.

For polygonal environments with an arbitrary number of polygonal obstacles, it was shown that no competitive strategy exists, even if all obstacles are parallelograms [1]. Later, this result was improved in [5] by providing a lower bound in $\Omega(\sqrt{k})$ for the competitive ratio of any on-line algorithm exploring a polygon with k obstacles. This bound remains true even for rectangular obstacles. On the other hand, there exists an algorithm with competitive ratio in $O(k)$ [2, 3]. Moreover, for particular shapes of obstacles (convex and with bounded aspect ratio) the optimal competitive ratio $\Theta(\sqrt{k})$ has been proven in [3].

Exploration of polygons by a robot with limited vision has been studied, e.g., in [6, 7, 8, 9, 10, 11]. In [6] the authors described an on-line algorithm with competitive ratio $1 + 3(\Pi S/A)$, where Π is a quantity depending on the perimeter of the polygon, S is the area seen by the robot, and A is the area of the polygon. The exploration in [6, 7] fails on a certain type of polygons, such as those with narrow corridors. In [8], the authors consider exploration in discrete steps. The robot can only explore the environment when it is motionless, and the cost of the exploration algorithm is measured by the number of stops during

¹Since parameters D, P, A are positive reals that may be arbitrarily small, it is important to stress that complexity $O(P + A + \sqrt{Ak})$ means that the trajectory of the robot is at most $c(P + A + \sqrt{Ak})$, for some constant c and *sufficiently large* values of P and A . Similarly for $O(P + D\sqrt{k})$. This permits to include, e.g., additive constants in the complexity, in spite of arbitrarily small parameter values.

the exploration. In [9, 10], the complexity of exploration is measured by the trajectory length, but only terrains composed of identical squares are considered. In [11] Ntafos studied off-line exploration of the boundary of a terrain with limited vision.

An experimental approach was used in [12] to show the performance of a greedy heuristic for exploration in which the robot always moves to the frontier between explored and unexplored area. Practical exploration of the environment by an actual robot was studied, e.g., in [13, 14]. In [14], a technique is described to deal with obstacles that are not in the plane of the sensor. In [13] landmarks are used during exploration to construct the skeleton of the environment.

Navigation is a closely related task which consists in finding a path between two given points in a terrain with unknown obstacles. Navigation in a $n \times n$ square containing rectangular obstacles aligned with sides of the square was considered in [15, 16, 17, 18]. It was shown in [15] that the navigation from a corner to the center of a room can be performed with a competitive ratio $O(\log n)$, only using tactile information (i.e., the robot modeled as a point sees an obstacle only when it touches it). No deterministic algorithm can achieve better competitive ratio, even with unlimited vision [15]. For navigation between any pair of points, there is a deterministic algorithm achieving a competitive ratio of $O(\sqrt{n})$ [17], and no deterministic algorithm can achieve a better competitive ratio [18]. However, there is a randomized approach performing navigation with a competitive ratio of $O(n^{\frac{4}{9}} \log n)$ [16]. Navigation with little information was considered in [19]. In this model, the robot cannot perform localization nor measure any distances or angles. Nevertheless, the robot is able to learn the critical information contained in the classical shortest-path roadmap and perform locally optimal navigation.

2. Unlimited vision

Let S be a smallest square in which the terrain \mathcal{T} is included. Our algorithm constructs a *quadtrees decomposition* of S . A quadtree is a rooted tree with each non-terminal node having four children. Each node of the quadtree corresponds to a square. The children of any non-terminal node v correspond to four identical squares obtained by partitioning the square of v using its horizontal and vertical axes of symmetry. This implies that the squares of the terminal nodes form a partition of the root². More precisely,

1. $\{S\}$ is a quadtree decomposition of S .
2. If $\{S_1, S_2, \dots, S_j\}$ is a quadtree decomposition of S , then $\{S_1, S_2, \dots, S_{i-1}, S_{i_1}, S_{i_2}, S_{i_3}, S_{i_4}, S_{i+1}, \dots, S_j\}$, where $S_{i_1}, S_{i_2}, S_{i_3}, S_{i_4}$ form a partition of S_i using its vertical and horizontal axes of symmetry, is a quadtree decomposition of S .

²In order to have an exact partition we assume that each square of the quadtree partition contains its East and South edges but not its West and North edges.

The trajectory of the robot exploring \mathcal{T} will be composed of parts which will follow the boundaries of \mathcal{P}_i , for $0 \leq i \leq k$, and of straight-line segments, called *approaching segments*, joining the boundaries of \mathcal{P}_i , $0 \leq i \leq k$. Obviously, the endpoints of an approaching segment must be mutually visible. The quadtree decomposition will be dynamically constructed in a top-down manner during the exploration of \mathcal{T} . At each moment of the exploration we consider the set \mathcal{Q}_S of all squares of the current quadtree and the set $\mathcal{Q}_\mathcal{T}$ of squares being the terminal nodes of the current quadtree. We will also construct dynamically a bijection $f : \{\mathcal{P}_0, \mathcal{P}_1, \dots, \mathcal{P}_k\} \longrightarrow \mathcal{Q}_S \setminus \mathcal{Q}_\mathcal{T}$.

When a robot moves along the boundary of some polygon \mathcal{P}_i , it may be in one of the two possible modes: the *recognition mode* – when it goes around the entire boundary of a polygon without any deviation, or in the *exploration mode* – when, while moving around the boundary, it tries to detect (and approach) new obstacles. When the decision to approach a new obstacle is made at some point r of the boundary of \mathcal{P}_i , the robot moves along an approaching segment to reach the obstacle, processes it by a recursive call, and (usually much later), returning from the recursive call, it moves again along this segment in the opposite direction in order to return to point r and to continue the exploration of \mathcal{P}_i . However, some newly detected obstacles may not be immediately approached. Namely, we say that, when the robot is at position $r \in \mathcal{P}_i$, an obstacle \mathcal{P}_j is *approachable*, if there exists a point $q \in \mathcal{P}_j$, belonging to a square $S_t \in \mathcal{Q}_\mathcal{T}$ of diameter $D(S_t)$ such that $|rq| \leq 2D(S_t)$. It is important to state that if exactly one obstacle becomes approachable at moment t , then it is approached immediately and if more than one obstacle become approachable at a moment t , then one of them (chosen arbitrarily) is approached immediately, and the others are approached later, possibly from different points of the trajectory. Each time a new obstacle is visited by the robot (i.e., all the points of its boundary are visited in the recognition mode), the terminal square of the current quadtree containing the first visited point of the new obstacle is partitioned. This square is then associated to this obstacle by function f . The trajectory of the robot is composed of three types of sections: *recognition sections*, *exploration sections* and *approaching sections*. The boundary of each polygon will be traversed twice: first time contiguously during a recognition section, and second time through exploration sections, which may be interrupted several times in order to approach and visit newly detected obstacles. We say that an obstacle is *completely explored*, if each point on the boundary of this obstacle has been traversed by an exploration section. We will prove that the sum of the lengths of the approaching sections is $O(D\sqrt{k})$.

Algorithm ExpTrav (polygon R , starting point r^* on the boundary of R)

- 1 Make a recognition traversal of the boundary of R .
- 2 Partition square $S_t \in \mathcal{Q}_\mathcal{T}$ containing r^* into four identical squares.
- 3 $f(R) := S_t$
- 4 **repeat**
- 5 Traverse the boundary of R until, for the current position r , there exists a visible point q of a new obstacle Q belonging to square $S_m \in \mathcal{Q}_\mathcal{T}$,

such that $|rq| \leq 2D(S_m)$.

- 6 Traverse the segment rq .
- 7 ExpTrav(Q, q)
- 8 Traverse the segment qr .
- 9 **until** R is completely explored.

Before the initial call of ExpTrav, the robot reaches a position r_0 at the boundary of the polygon \mathcal{P}_0 . This is done as follows. At its initial position v , the robot chooses an arbitrary half-line α which it follows as far as possible. When it hits the boundary of a polygon \mathcal{P} , it traverses the entire boundary of \mathcal{P} . Then it computes the point u which is the farthest point from v in $\mathcal{P} \cap \alpha$. It goes around \mathcal{P} until reaching u again and progresses on α , if possible. If this is impossible, the robot recognizes that it went around the boundary of \mathcal{P}_0 and it is positioned on this boundary. It initialises the quadtree decomposition to a smallest square S containing \mathcal{P}_0 . This square is of size $O(D(\mathcal{P}_0))$. The length of the above walk is less than $3P$.

Fig. 1 shows an example of execution of Algorithm ExpTrav. The robot starts at point q_1 . It reaches the boundary of \mathcal{P}_0 at point q_2 . It makes a recognition traversal of \mathcal{P}_0 , then it sees the obstacle \mathcal{P}_1 . It goes to \mathcal{P}_1 (reaching point q_3), partitions the corresponding square and makes a recognition traversal of \mathcal{P}_1 . Then it sees obstacle \mathcal{P}_2 , goes to it through point q_4 and partitions the corresponding square. Then the robot makes a recognition traversal of \mathcal{P}_2 followed by another traversal in exploration mode. Since it sees no approachable obstacles, it comes back to q_3 . From q_3 , it goes to the approachable obstacle \mathcal{P}_3 through q_5 (partitioning the corresponding square). Then the robot makes a recognition traversal of \mathcal{P}_3 followed by an exploration traversal. Finally, it returns to q_3 , and then to q_2 to finish the exploration traversal of \mathcal{P}_0 .

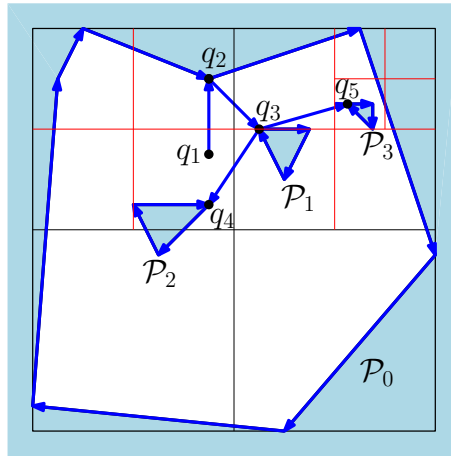


Figure 1: An execution of Algorithm ExpTrav.

Lemma 2.1. *Algorithm ExpTrav visits all boundary points of all obstacles of the terrain \mathcal{T} .*

Proof: Note that Algorithm ExpTrav always terminates. Indeed, since there is a finite number of obstacles, there is a finite number of calls of ExpTrav and steps 5-8 are executed a finite number of times. Moreover, since each obstacle has a finite boundary, step 1 and the repeat loop are always completed.

Consider the quadtree decomposition \mathcal{Q} of S arising at the completion of the algorithm. Suppose, for contradiction, that there exists a point on the boundary of an obstacle \mathcal{P}_i which was never visited. Let p be a point among all unvisited boundary points for which its terminal square S_j belonging to the quadtree \mathcal{Q} has the smallest possible diameter. Observe that S_j cannot be the root of the quadtree decomposition, since the root is always partitioned (line 2 of Algorithm ExpTrav). Consider the square S_m , the parent of S_j in \mathcal{Q} . S_m was partitioned in step 2 of some call of ExpTrav as a result of detecting some obstacle \mathcal{P}_l intersecting S_m . Consider the segment qp , where $q \in S_m$ belongs to the boundary of \mathcal{P}_l . Since both points p and q belong to the boundary of \mathcal{T} and q was visited while p was not, there exists a pair of points q' and p' ($p' = p$ and $q' = q$ may hold), both belonging to the boundary of \mathcal{T} and to the segment qp , such that q' was visited, p' was not and p' is visible from q' (see Fig. 2). Such a pair exists because at the end of the exploration the boundary of each polygon is either entirely visited or not at all. Consider the quadtree at the moment t when the robot visited point q' , and its terminal square S_i containing point p' . Clearly, $D(S_i) \geq D(S_j)$, because S_j is a square with the smallest diameter containing unvisited boundary points. Hence $|q'p'| \leq |qp| \leq D(S_m) = 2D(S_j) \leq 2D(S_i)$ and p' was approachable from q' at time t , a contradiction.

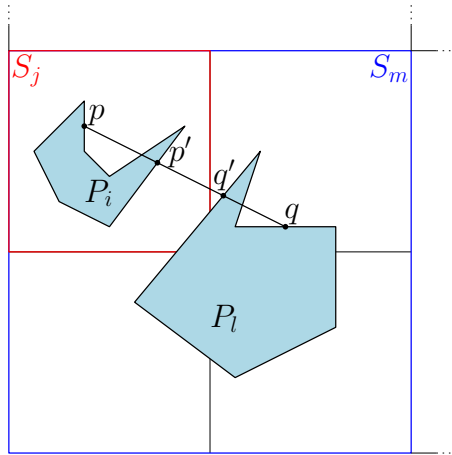


Figure 2: The existence of a pair $p'q'$.

□

Lemma 2.2. *Function f is a bijection from $\{\mathcal{P}_0, \mathcal{P}_1, \dots, \mathcal{P}_k\}$ to $\mathcal{Q}_S \setminus \mathcal{Q}_T$, where \mathcal{Q}_S and \mathcal{Q}_T correspond to the final quadtree decomposition produced by Algorithm ExpTrav.*

Proof: When ExpTrav is called for the first time, the robot is on the boundary of \mathcal{P}_0 and the quadtree has exactly one non-terminal node – its root S , and $f(\mathcal{P}_0) = S$. By induction, each time a new obstacle Q is approached in step 6 of a call of ExpTrav, $f(Q)$ is set to some terminal node S_t intersecting Q and S_t becomes a nonterminal node of the quadtree in step 2. Hence each square corresponding to a non-terminal node of the quadtree is an image of a different polygon \mathcal{P}_i , $0 \leq i \leq k$. \square

Lemma 2.3. *For any quadtree T , rooted at a square of diameter D and having x non-terminal nodes, the sum $\sigma(T)$ of diameters of these nodes is at most $2D\sqrt{x}$.*

Proof: The diameter of a square at depth t of the quadtree is $\frac{D}{2^t}$. We prove first that among all quadtrees with x non-terminal nodes, $\sigma(T)$ is maximized for the quadtree having all terminal nodes of at most two consecutive depths. Suppose, to the contrary, that there exists a quadtree T maximizing $\sigma(T)$ having terminal nodes of two different depths t_1 and t_2 with $t_1 < t_2 - 1$. Let p be a terminal node of T of depth t_1 and q be a non-terminal node at depth $t_2 - 1$. Let T' be the tree obtained from T by detaching from T node $p \in T$ and the subtree of T rooted at q and exchanging their places. T' is again a quadtree with x non-terminal nodes, having one less non-terminal node of depth $t_2 - 1$ and one extra non-terminal node of depth t_1 . Hence $\sigma(T') \geq \sigma(T) - \frac{D}{2^{t_2-1}} + \frac{D}{2^{t_1}} > \sigma(T)$, which contradicts the maximality of $\sigma(T)$.

Therefore it is sufficient to consider only quadtrees having terminal nodes of at most two consecutive depths t and $t + 1$. Suppose that there are y terminal nodes of depth $t + 1$, $0 \leq y \leq 4^{t+1}$. Then the number x of non-terminal nodes equals $\frac{4^t - 1}{3} + \frac{y}{4}$ and

$$\begin{aligned} \sigma(T) &= \frac{yD}{4 \cdot 2^t} + \sum_{i=0}^{t-1} 4^i \frac{D}{2^i} \\ &= D \left(2^t - 1 + \frac{y}{2^{t+2}} \right) \end{aligned}$$

We need to prove that $\sigma(T) \leq 2D\sqrt{x}$, i.e. that

$$D \left(2^t - 1 + \frac{y}{2^{t+2}} \right) \leq 2D \sqrt{\frac{4^t - 1}{3} + \frac{y}{4}}.$$

Hence it is sufficient to show that

$$\left(2^t - 1 + \frac{y}{2^{t+2}} \right)^2 \leq 4 \left(\frac{4^t - 1}{3} + \frac{y}{4} \right).$$

i.e., that

$$0 \leq y \left(1 + \frac{1}{2^{t+1}} - \frac{1}{2} - \frac{y}{4^{t+2}} \right) + \left(4 \cdot \frac{4^t - 1}{3} - 2^{2t} - 1 + 2^{t+1} \right).$$

The second term is clearly non-negative for $t \geq 0$ and the first term is also positive, since $0 \leq y \leq 4^{t+1}$. We conclude that $\sigma(T) \leq 2D\sqrt{x}$. \square

Theorem 2.1. *Algorithm ExpTrav explores the terrain \mathcal{T} of perimeter P and convex hull diameter D with k obstacles in time $5P + 12D\sqrt{k}$.*

Proof: Take an arbitrary point p inside \mathcal{T} and a ray outgoing from p in an arbitrary direction. This ray reaches the boundary of \mathcal{T} at some point q . Since, by Lemma 2.1 point q was visited by the robot, p was visible from q during the robot's traversal, and hence p was explored.

To prove the complexity of the algorithm, observe that the robot traverses twice the boundary of each polygon of \mathcal{T} , once during its recognition in step 1 and the second time during the iterations of step 5. Hence the sum of lengths of the recognition and exploration sections is $2P$. The only other portions of the trajectory are produced in steps 6 and 8, when the obstacles are approached and returned from. According to the condition from step 5, an approaching segment is traversed in step 6 only if its length is shorter than twice the diameter of the associated square. If $k = 0$ then the sum of lengths of all approaching segments is 0, due to the fact that exploration starts at the external boundary of the terrain. In this case the length of the trajectory is at most $5P$ (since the length of at most $3P$ was traversed before the initial call). Hence we may assume that $k > 0$. By Lemma 2.2 each obstacle and \mathcal{P}_0 are associated with a different non-terminal node of the quadtree, and the number x of non-terminal nodes of the quadtree equals $k + 1$. Hence the sum of lengths of all approaching segments is at most $2\sigma(T)$. By Lemma 2.3 we have $\sigma(T) \leq 2D\sqrt{x} = 2D\sqrt{k+1}$, hence the sum of lengths of approaching segments is at most $2\sigma(T) \leq 4D\sqrt{k+1} \leq 4D\sqrt{2k} \leq 6D\sqrt{k}$. Each segment is traversed twice, so the total length of this part of the trajectory is at most $12D\sqrt{k}$. It follows that the total length of the trajectory is at most $5P + 12D\sqrt{k}$. \square

Theorem 2.2. *Any algorithm for a robot with unlimited visibility, exploring polygonal terrains with k obstacles, having total perimeter P and the convex hull diameter D , produces trajectories in $\Omega(P + D\sqrt{k})$ in some terrains, even if the terrain is known to the robot.*

Proof: In order to prove the lower bound, we show two families of terrains: one for which $P \in \Theta(D)$ (P cannot be smaller), D and k are unbounded, and the exploration cost is $\Omega(D\sqrt{k})$, and the other in which P is unbounded, D is arbitrarily small, $k = 0$, and still the exploration cost is $\Omega(P)$.

Consider the terrain from Fig. 3(a) where k identical tiny obstacles are distributed evenly at the $\sqrt{k} \times \sqrt{k}$ grid positions inside a square of diameter D . The distance between obstacles is at least $\frac{D\sqrt{2}}{2(\sqrt{k+1})} - \epsilon$ where $\epsilon > 0$ may be as small as necessary by choosing obstacles sufficiently small. The obstacles are such that to explore the small area inside the convex hull of the obstacle the robot must enter this convex hull. Since each such area must be explored, the trajectory of the robot must be of size at least $(k-1) \left(\frac{D\sqrt{2}}{2(\sqrt{k+1})} - \epsilon \right)$, which is clearly in $\Omega(D\sqrt{k})$. Note that the perimeter P is in $\Theta(D)$.




Figure 3: Lower bound for unlimited visibility

The terrain from Fig. 3(b) is a polygon of arbitrarily small diameter (without obstacles), whose exploration requires a trajectory of size $\Omega(P)$, where P is unbounded (as the number of “corridors” can be unbounded). Indeed, each “corridor” must be traversed almost completely to explore points at its end. The two families of polygons from Fig. 3 lead to the $\Omega(P + D\sqrt{k})$ lower bound. \square

3. Limited vision

In this section we assume that the vision of the robot has range 1. The following algorithm is at the root of all our positive results on exploration with limited vision. The idea of the algorithm is to partition the environment into small parts called *cells* (of diameter at most 1) and to visit them using a depth-first traversal. The local exploration of cells can be performed using Algorithm `ExpTrav`, since the vision inside each cell is not limited by the range 1 of the vision of the robot. The main novelty of our exploration algorithm is that the robot completely explores *any* terrain. This should be contrasted with previous algorithms with limited visibility, e.g. [6, 7, 9, 10] in which only a particular class of terrains with obstacles is explored, e.g., terrains without narrow corridors or terrains composed of complete identical squares. This can be done at cost $O(A)$. Our lower bound shows that exploration complexity of arbitrary terrains depends on the perimeter and the number of obstacles as well. The complete exploration of arbitrary terrains achieved by our algorithm significantly complicates both the exploration process and its analysis.

Our algorithms `LETA` and `LETk`, and the tourist algorithm described in [3] share a similar approach to exploration, i.e., using several square decompositions of the terrain with different side lengths to figure out the characteristics of the terrain and achieve efficient exploration. However, our algorithms differ from the tourist algorithm in two important ways : (1) the exploration within each square is done with an optimal algorithm (Section 2) instead of a greedy one, and (2) the limited visibility forces an upper bound on the side of the square (significantly complicating `LETk`). More important, due to the numerous

differences between our model and the one of [3], the analyses of the complexities of the algorithms are unrelated.

Algorithm `LimExpTrav` (*LET*, for short)

INPUT: A point s inside the terrain \mathcal{T} and a positive real $F \leq \sqrt{2}/2$.

OUTPUT: An exploration trajectory of \mathcal{T} , starting and ending at s .

Tile the area with squares of side F , such that s is on the boundary of a square. The connected regions obtained as intersections of \mathcal{T} with each tile are called *cells*. For each tile S , maintain a quadtree decomposition Q_S initially set to $\{S\}$. Then, arbitrarily choose one of the cells containing s to be the starting cell C and call `ExpCell`(C, s).

Procedure `ExpCell`(current cell C , starting point $r^* \in C$)

- 1 Record C as visited
- 2 `ExpTrav`(C, r^*) using the quadtree decomposition Q_S ; S is the tile s.t. $C \subseteq S$
- 3 **repeat**
- 4 Traverse the boundary of C until the current position r belongs to an unvisited cell U
- 5 `ExpCell`(U, r)
 (if r is in several unvisited cells, choose arbitrarily one to be processed)
- 6 **until** the boundary of C is completely traversed

It is worth to note that, at the beginning of the exploration of the first cell belonging to a tile S , the quadtree of this tile is set to a single node. However, at the beginning of explorations of subsequent cells belonging to S , the quadtree of S may be different. So the top-down construction of this quadtree may be spread over the exploration of many cells which will be visited at different points in time.

Consider a tile T and a cell $C \subseteq T$. Let A_C be the area of C and B_C the length of its boundary. Let P_C be the length of the part of the boundary of C included in the boundary of the terrain \mathcal{T} , and let R_C be the length of the remaining part of the boundary of C , i.e., $R_C = B_C - P_C$.

Lemma 3.1. *The inequality $R_C \leq 8(A_C/F + P_C)$ holds for any cell C .*

Proof: We consider three cases:

Case 1: $P_C < F/2$ and $A_C < F^2/2$

In this case, we will show that $R_C \leq 3P_C$. We call a *borderline* a maximum connected part of the boundary of \mathcal{T} inside the tile T delimiting the cell C . Let $\mathcal{L} = \{L_1, \dots, L_l\}$ be the set of borderlines of C . There are two types of borderlines: the *linking* borderlines that link two points of the boundary of T and the *closed* borderlines that are closed polygonal lines inside S . A borderline L separates the tile S into two connected regions, the *inside* region denoted by I_L , i.e., the region containing C , and the *outside* region, denoted by O_L . If the area of I_L is smaller than that of O_L , then L is a *small-inside* borderline, otherwise L is a *large-inside* borderline. We denote by M_L the region among I_L and O_L which has the smaller area.

Notice that the two endpoints of a *linking* borderline can either be on the same side of S or on two adjacent sides. Indeed, any borderline linking two points on opposite sides of S would have a length at least F , a contradiction with the inequality $P_C < F/2$. If L is a linking borderline with both endpoints x and y on the same side of S , then the length of segment xy is smaller than $|L|$. Hence, the perimeter of M_L is smaller than $2 \cdot |L|$. If L is a linking borderline with endpoints x and y on two sides that intersect at a vertex v , then the lengths of segments vx and vy are both less than $|L|$. Therefore the perimeter of M_L is smaller than $3 \cdot |L|$. If L is a closed borderline, then the perimeter of M_L is exactly L . Hence, the perimeter of M_L is always less than $3 \cdot |L|$. Moreover, the area of M_L is less than $(3|L|)^2/4\pi$ by the isoperimetric inequality [20].

Claim 1. At least one borderline in \mathcal{L} is a small-inside borderline.

On the contrary, suppose that for all $i = 1, \dots, l$, the borderline L_i is a large-inside borderline. Then, since $M_{L_i} = O_{L_i}$, we have $C = S \setminus \bigcup_{i=1}^l M_{L_i}$. It follows that:

$$\begin{aligned} A_C &= \text{Area}(S) - \sum_{i=1}^l \text{Area}(M_{L_i}) \\ &\geq F^2 - \sum_{i=1}^l \frac{(3|L_i|)^2}{4\pi}, \quad \text{since } (3|L_i|)^2/4\pi \geq \text{Area}(M_{L_i}) \text{ for all } i = 1, \dots, l \\ &\geq F^2 - \frac{9P_C^2}{4\pi}, \quad \text{since } P_C^2 = \left(\sum_{i=1}^l |L_i| \right)^2 \geq \sum_{i=1}^l |L_i|^2 \\ &> \frac{F^2}{2}, \quad \text{since } F/2 > P_C. \end{aligned}$$

We obtain $A_C > F^2/2$, a contradiction. This ends the proof of Claim 1.

Let $L \in \mathcal{L}$ be a small-inside borderline. We have $C \subseteq M_L$ and thus $R_C < 3|L| \leq 3P_C$.

Case 2: $P_C < F/2$ and $A_C \geq F^2/2$

We have:

$$R_C \leq 4F \leq \frac{8}{F}A_C, \quad \text{since } F \leq 2A_C/F.$$

Case 3: $P_C \geq F/2$

We have:

$$R_C \leq 4F \leq 8P_C, \quad \text{since } F \leq 2P_C.$$

In all cases, we have $R_C \leq 8(A_C/F + P_C)$. \square

The following is the key lemma for all upper bounds proved in this section. Let $\mathcal{S} = \{T_1, T_2, \dots, T_n\}$ be the set of tiles with non-empty intersection with \mathcal{T} and $\mathcal{C} = \{C_1, C_2, \dots, C_m\}$ be the set of cells that are intersections of tiles from \mathcal{S} with \mathcal{T} . For each $T \in \mathcal{S}$, let k_T be the number of obstacles of \mathcal{T} entirely contained in T .

Lemma 3.2. *For any $F \leq \sqrt{2}/2$, Algorithm LET explores the terrain \mathcal{T} of area A and perimeter P , using a trajectory of length $O(P + A/F + F \sum_{i=1}^n \sqrt{k_{T_i}})$.*

Proof: First, we show that Algorithm LET explores the terrain \mathcal{T} . Consider the graph G whose vertex set is \mathcal{C} and edges are the pairs $\{C, C'\}$ such that C and C' have a common point at their boundaries. The graph G is connected, since \mathcal{T} is connected. Note that for any cell C and point r on the boundary of C , ExpTrav on C and r and thus ExpCell on C and r starts and ends on r . Therefore, Algorithm LET performs a depth first traversal of graph G , since during the execution of $\text{ExpCell}(C, \dots)$, procedure $\text{ExpCell}(U, \dots)$ is called for each unvisited cell U adjacent to C . Hence, $\text{ExpCell}(C, \dots)$ is called for each cell $C \in \mathcal{C}$, since G is connected. During the execution of $\text{ExpCell}(C, r)$, C is completely explored by $\text{ExpTrav}(C, r)$ by the same argument as in the proof of Lemma 2.1, since the convex hull diameter of C is less than one.

It remains to show that the length of the LET trajectory is $O(P + A/F + F \sum_{i=1}^n \sqrt{k_{T_i}})$. For each $j = 1, \dots, m$, the part of the LET trajectory inside the cell C_j is produced by the execution of $\text{ExpCell}(C_j, \dots)$. In step 2 of $\text{ExpCell}(C_j, \dots)$, the robot executes ExpTrav with $D = \sqrt{2}F$ and $P = P_{C_j} + R_{C_j}$. The sum of lengths of recognition and exploration sections of the trajectory in C_j is at most $2(P_{C_j} + R_{C_j})$. The sum of lengths of approaching sections of the trajectory in T_i is at most $6\sqrt{2}F\sqrt{k_{T_i}}$ and each approaching section is traversed twice (cf. proof of Theorem 2.1). In step 3 of $\text{ExpCell}(C_j, \dots)$, the robot only makes the tour of the cell C_j , hence the distance traveled by the robot is at most $P_{C_j} + R_{C_j}$. It follows that:

$$\begin{aligned} |LET| &\leq 3 \sum_{j=1}^m (P_{C_j} + R_{C_j}) + 12\sqrt{2}F \sum_{i=1}^n \sqrt{k_{T_i}} \\ &\leq 3 \sum_{j=1}^m ((1+8)P_{C_j} + cA_{C_j}/F) + 12\sqrt{2}F \sum_{i=1}^n \sqrt{k_{T_i}} \quad \text{by Lemma 3.1} \\ &\leq 27P + 24A/F + 12\sqrt{2}F \sum_{i=1}^n \sqrt{k_{T_i}}. \end{aligned}$$

□

In view of Lemma 3.2, exploration of a particular class of terrains can be done at a cost which will be later proved optimal.

Theorem 3.1. *Let $c > 1$ be any constant. Exploration of a c -fat terrain of area A , perimeter P and with k obstacles can be performed using a trajectory of length $O(P + A + c\sqrt{Ak})$ (without any a priori knowledge).*

Proof: The robot executes Algorithm LET with $F = \sqrt{2}/2$. By Lemma 3.2, the total cost is $O(P + A + \sum_{i=1}^n \sqrt{k_{T_i}})$. Recall that n is the number of tiles that have non-empty intersection with the terrain. We have $\sum_{i=1}^n \sqrt{k_{T_i}} \leq$

$\sum_{i=1}^n \sqrt{\frac{k}{n}} = \sqrt{nk}$. By definition of a c -fat terrain, there is a disk D_1 of radius r included in the terrain and a disk D_2 of radius R that contains the terrain, such that $\frac{R}{r} \leq c$. There are $\Theta(r^2)$ tiles entirely included in D_1 and hence in the terrain. So, we have $A = \Omega(r^2)$. $\Theta(R^2)$ tiles are sufficient to cover D_2 and hence the terrain. So $n = O(R^2)$. Hence, we obtain $n = O(c^2A)$ in view of $R \leq cr$. We have $O(P + A + \sum_{i=1}^n \sqrt{kT_i}) = O(P + A + \sqrt{nk}) = O(P + A + c\sqrt{Ak})$. \square

Consider any terrain \mathcal{T} of area A , perimeter P and with k obstacles. We now turn attention to the exploration problem if some knowledge about the terrain is available a priori. Notice that if A and k are known before the exploration, Lemma 3.2 implies that Algorithm *LET* executed for $F = \min\{\sqrt{A/k}, \sqrt{2}/2\}$ explores *any* terrain at cost $O(A+P+\sqrt{Ak})$. (Indeed, if $F = \sqrt{A/k}$ then $A/F = \sqrt{Ak}$ and $kF = \sqrt{Ak}$, while $F = \sqrt{2}/2$ implies $A/F = \Theta(A)$ and $kF = O(A)$.) This cost will be later proved worst-case optimal. It turns out that a much more subtle use of Algorithm *LET* can guarantee the same complexity assuming only knowledge of A or k . We present two different algorithms depending on which value, A or k , is known to the robot. Both algorithms rely on the same idea. The robot executes Algorithm *LET* with some initial value of F until either the terrain is completely explored, or a certain termination condition, depending on the algorithm, is satisfied. This execution constitutes the first stage of the two algorithms. If exploration was interrupted because of the termination condition, then the robot proceeds to a new stage by executing Algorithm *LET* with a new value of F . Values of F decrease in the first algorithm and increase in the second one. The exploration terminates at the stage when the terrain becomes completely explored, while the termination condition is never satisfied. In each stage the robot is oblivious of the previous stages, except for the computation of the new value of F that depends on the previous stage. This means that in each stage exploration is done “from scratch”, without recording what was explored in previous stages. In order to test the termination condition in a given stage, the robot maintains the following three values: the sum A^* of areas of explored cells, updated after the execution of *ExpTrav* in each cell; the length P^* of all the boundaries traversed by the robot, continuously updated when the robot moves along a boundary for the first time (i.e., in the recognition mode); and the number k^* of obstacles approached by the robot, updated when an obstacle is approached. The values of A^* , P^* and k^* at the end of the i -th stage are denoted by A_i , P_i and k_i , respectively. Let F_i be the value of F used by Algorithm *LET* in the i -th stage. Now, we are ready to describe the termination conditions and the values F_i in both algorithms.

Algorithm LET_A , for A known before exploration

The value of F used in Algorithm *LET* for the first stage is $F_1 = \sqrt{2}/2$. The value of F for subsequent stages is given by $F_{i+1} = \frac{A}{k_i F_i}$. The termination condition is $\{k^* F_i \geq 2A/F_i \text{ and } k^* F_i \geq P^* + 1\}$.

Algorithm LET_k, for k known before exploration

The value of F used in Algorithm *LET* for the first stage is $F_1 = \frac{1}{k+\sqrt{2}}$.
 The value of F for subsequent stages is given by $F_{i+1} = \min \left\{ \frac{A_i}{kF_i}, \frac{\sqrt{2}}{2} \right\}$. The
 termination condition is $\{A^*/F_i \geq 2kF_i \text{ and } A^*/F_i \geq P^* + 1 \text{ and } F_i < \sqrt{2}/2\}$.

Let \mathcal{C}_t be the set of cells recorded as visited by Algorithm *LET* at time moment t , and let \mathcal{O}_t be the set of obstacles approached by the robot until time t . For each $C \in \mathcal{C}_t$, let E_C be the length of the intersection of the exterior boundary of cell C with the boundary of the terrain. For each $O \in \mathcal{O}_t$, let $|O|$ be the perimeter of obstacle O and let $k_t = |\mathcal{O}_t|$. For any set of obstacles \mathcal{O} , let $\|\mathcal{O}\|$ be the sum of perimeters of obstacles in \mathcal{O} , i.e., $\|\mathcal{O}\| = \sum_{D \in \mathcal{O}} |D|$. The following proposition is proved similarly as Lemma 3.2.

Proposition 3.1. *There is a positive constant d such that the trajectory of the robot until any time t , during the execution of Algorithm *LET*, is at most $d(\sum_{C \in \mathcal{C}_t} (E_C + A_C/F) + (k_t + 1) \cdot F + \|\mathcal{O}_t\|)$.*

Proof: Let t be any moment in time during the execution of Algorithm *LET*. Let U be the current cell at moment t , which means that the instance of procedure *ExpCell* executed at moment t was called with U as its first parameter. Let \mathcal{O}'_t be the set of obstacles in \mathcal{O}_t that are not included in U and let $\mathcal{C}'_t = \mathcal{C}_t \setminus \{U\}$. All the obstacles included in $C \in \mathcal{C}'_t$ were approached and visited by the robot. Hence, we have $\sum_{C \in \mathcal{C}'_t} P_C = \sum_{C \in \mathcal{C}'_t} E_C + \|\mathcal{O}'_t\|$. By the same arguments as in the proof of Lemma 3.2, it follows that there is a constant d' such that the total length of the trajectory of the robot in the cells in \mathcal{C}'_t , for any time moment t , is at most $d'(\sum_{C \in \mathcal{C}'_t} (P_C + A_C/F) + F|\mathcal{O}'_t|) = d'(\sum_{C \in \mathcal{C}'_t} (E_C + A_C/F) + \|\mathcal{O}'_t\| + F|\mathcal{O}'_t|)$. The length of the trajectory of the robot in U is at most $3G_U + 2\|\mathcal{O}_t \setminus \mathcal{O}'_t\| + |\mathcal{O}_t \setminus \mathcal{O}'_t| \cdot F$, where G_U is the length of the exterior boundary of cell U . Hence, there is a constant d such that the length of the trajectory of the robot until any time t is at most $d(\sum_{C \in \mathcal{C}_t} (E_C + A_C/F) + (k_t + 1) \cdot F + \|\mathcal{O}_t\|)$, since $G_U \leq 4F + E_U$. \square

The following lemma establishes the complexity of exploration if either the area of the terrain or the number of obstacles is known a priori.

Lemma 3.3. *Algorithm LET_A (resp. LET_k) explores a terrain \mathcal{T} of area A , perimeter P and with k obstacles, using a trajectory of length $O(P + A + \sqrt{Ak})$, if A (resp. k) is known before exploration.*

Proof:

Part 1: complexity of Algorithm LET_A

First, we show that the algorithm eventually terminates and completely explores \mathcal{T} . If there is no obstacles in \mathcal{T} ($k = 0$), then the termination condition is never satisfied. In this case, the total cost of the solution is $O(P + A)$ by Lemma 3.2, since all the terrain is explored during the first stage of the algorithm (with $F = \sqrt{2}/2$). Hence it remains to show the complexity of the algorithm

for $k > 0$. Observe that for each $i > 1$, $F_i = \frac{A}{k_{i-1}F_{i-1}} \leq \frac{F_{i-1}}{2}$, since $k_{i-1}F_{i-1} \geq 2\frac{A}{F_{i-1}}$. Since $F_1 = \sqrt{2}/2$, for each $i \geq 1$, we have $F_i \leq \frac{\sqrt{2}}{2^i}$. The algorithm eventually terminates, since there exists $m \in \mathbb{N}$ such that $kF_m < 2A/F_1 < 2A/F_m$ and the termination condition is never satisfied in this case. In the last stage, Algorithm *LET* performs complete exploration of the terrain by Lemma 3.2, since the value of F used for exploration is at most $\sqrt{2}/2$.

Let D_i be the distance traveled by the robot during the i -th stage and let n be the number of stages. By Proposition 3.1, if stage i ends at moment t_i then $D_i \leq d(\sum_{C \in \mathcal{C}_{t_i}} (E_C + A_C/F_i) + (|\mathcal{O}_{t_i}| + 1) \cdot F_i + \|\mathcal{O}_{t_i}\|)$ for each $i \geq 1$. Since the algorithm can only interrupt stage $i < n$ when approaching an obstacle, we have $P_i = \sum_{C \in \mathcal{C}_{t_i}} E_C + \|\mathcal{O}_{t_i}\|$. We obtain that $D_i \leq d(P_i + A/F_i + (k_i + 1)F_i)$ for each $i \geq 1$.

If $n = 1$, then the termination condition is never satisfied and $kF_1 < \max\{2A/F_1, P + 1\}$. The total cost is at most $d(P + A/F_1 + (k + 1)F_1) = O(P + A)$, since $F_1 = \sqrt{2}/2$. On the other hand, if $n > 1$, then for each i such that $1 \leq i < n$, we have $D_i \leq d(P_i + A/F_i + (k_i + 1)F_i) \leq 3dk_iF_i$. Indeed, we have $k_iF_i \geq A/F_i$ and $k_iF_i \geq P_i + F_i$, since the termination condition is satisfied at the end of the i -th stage. Moreover, we have $\frac{1}{2}k_iF_i \geq A/F_i = k_{i-1}F_{i-1}$, for all $1 < i < n$. It follows that $\sum_{i=1}^{n-1} D_i \leq 6dk_{n-1}F_{n-1}$. In order to show that the total cost is $O(P + A + \sqrt{Ak})$, it is sufficient to show that $P + A/F_n + kF_n = O(P + A + \sqrt{Ak})$, since $k_{n-1}F_{n-1} = A/F_n$ and $D_n \leq d(P_n + A/F_n + (k + 1)F_n)$.

Take the moment t_{n-1} when the $(n - 1)$ -th stage was interrupted, i.e., when both inequalities of the termination condition started to be satisfied. Consider the inequality which was not satisfied just before moment t_{n-1} . If this is the first of the two inequalities, then at time t_{n-1} the Algorithm *LET_A* must have increased the value of k^* , hence just before moment t_{n-1} we had $(k_{n-1} - 1)F_{n-1} < 2AF_{n-1}$. Similarly, if the second inequality was not satisfied just before moment t_{n-1} then we had $(k_{n-1} - 1)F_{n-1} < P_{n-1} + 1$.

Case 1: $(k_{n-1} - 1)F_{n-1} < 2A/F_{n-1}$

We have:

$$\begin{aligned} k_{n-1}F_{n-1} &\leq 2\frac{A}{F_{n-1}} + F_{n-1} \\ \frac{A}{F_n} &\leq 2\frac{A}{F_{n-1}} + 1, \quad \text{since } F_{n-1} = \frac{A}{k_{n-1}F_n} \text{ and } F_{n-1} \leq 1 \\ \frac{A}{F_n} &\leq \sqrt{2}\sqrt{Ak_{n-1}} + 1, \quad \text{since } k_{n-1}F_{n-1} \geq 2\frac{A}{F_{n-1}} \\ &\quad \text{and thus } \left(\frac{A}{F_{n-1}}\right)^2 \leq \frac{1}{2}k_{n-1}F_{n-1}\frac{A}{F_{n-1}} \\ \frac{A}{F_n} &\leq 2\sqrt{Ak} + 1, \quad \text{since } k_{n-1} \leq k \end{aligned}$$

Since the termination condition in the last stage is not satisfied, we have

$kF_n \leq \max\{2A/F_n, P+1\}$. We obtain $P + A/F_n + kF_n = O(P + A/F_n)$. Hence, we have $P + A/F_n + kF_n = O(P + \sqrt{Ak})$.

Case 2: $(k_{n-1} - 1)F_{n-1} < P_{n-1} + 1$

We have:

$$\begin{aligned} k_{n-1}F_{n-1} &\leq P_{n-1} + F_{n-1} + 1 \\ \frac{A}{F_n} &\leq P_{n-1} + F_{n-1} + 1, \quad \text{since } k_{n-1}F_{n-1} = \frac{A}{F_n} \\ \frac{A}{F_n} &\leq P + 2, \quad \text{since } F_{n-1} \leq 1 \text{ and } P_{n-1} \leq P \end{aligned}$$

Since the termination condition in the last stage is not satisfied, we have $kF_n \leq \max\{2A/F_n, P+1\}$, as before. We obtain $P + A/F_n + kF_n = O(P + A/F_n) = O(P)$.

Part 2: complexity of Algorithm LET_k

The proof is similar to that concerning Algorithm LET_A. The main difference follows from the additional clause $F_i < \sqrt{2}/2$ in the termination condition. This clause was not necessary in Algorithm LET_A because, as opposed to the current case, sides of tiles were decreasing. First, we show that the algorithm eventually terminates and completely explores \mathcal{T} . Observe that for each $i > 1$, we have $F_i = \min\{A_{i-1}/kF_{i-1}, \sqrt{2}/2\} \geq \min\{2F_{i-1}, \sqrt{2}/2\}$, since $A_{i-1}/F_{i-1} \geq 2kF_{i-1}$. Hence, the algorithm eventually terminates. Indeed, even if the first two inequalities remain true, the third must become false at some point. Notice that $F_1 \leq \sqrt{2}/2$, since $k \geq 0$, and for all $i > 1$ we have $F_i \leq \sqrt{2}/2$. In the last stage, Algorithm *LET* performs complete exploration of the terrain by Lemma 3.2, since the value of F used for the exploration is at most $\sqrt{2}/2$.

Let D_i be the distance traveled by the robot during the i -th stage and let n be the number of stages. By Proposition 3.1, if stage i ends at moment t_i then $D_i \leq d(\sum_{C \in \mathcal{C}_{t_i}} (E_C + A_C/F_i) + (|\mathcal{O}_{t_i}| + 1) \cdot F_i + \|\mathcal{O}_{t_i}\|)$ for each $i \geq 1$. Since the algorithm can only stop when completing the exploration of a cell, we have $P_i = \sum_{C \in \mathcal{C}_{t_i}} E_C + \|\mathcal{O}_{t_i}\|$ and $A_i = \sum_{C \in \mathcal{C}_{t_i}} A_C$. We obtain that $D_i \leq d(P_i + A_i/F_i + (k+1)F_i)$ for each $i \geq 1$.

If $n = 1$, then the termination condition is never satisfied and either $A/F_1 \leq \max\{2kF_1, P+1\}$ or $F_1 = \sqrt{2}/2$. In the first case, the total cost is at most $d(P + A/F_1 + (k+1)F_1) = O(P)$, since $kF_1 \leq 1$. In the second case, we have $k = 0$ and the total cost is at most $d(P + A/F_1 + (k+1)F_1) = O(P + A)$. On the other hand, if $n > 1$ then for each i such that $1 \leq i < n$, we have $D_i \leq d(P_i + A_i/F_i + (k+1)F_i) \leq 3dA_i/F_i$. Indeed, we have $\frac{A_i}{F_i} \geq kF_i$ and $\frac{A_i}{F_i} \geq P_i + F_i$, since the termination condition is satisfied at the end of the i -th stage. Moreover, we have $\frac{A_i}{2F_i} \geq kF_i = \frac{A_{i-1}}{F_{i-1}}$ for all $1 < i < n$. It follows that $\sum_{i=1}^{n-1} D_i \leq 6d\frac{A_{n-1}}{F_{n-1}}$.

Notice that the third inequality of the termination condition is always satisfied during the $(n-1)$ -stage. Take the moment t_{n-1} when the $(n-1)$ -th stage was interrupted, i.e., when the two first inequalities of the termination condition started to be satisfied. Consider the inequality which was not satisfied just before moment t_{n-1} . If this is the first of the two inequalities, then at time t_{n-1} the Algorithm LET_k must have increased the value of A^* by at most F_{n-1}^2 , since each cell is included in a square of size F_{n-1} . Hence just before moment t_{n-1} , we had $\frac{A_{n-1}-F_{n-1}^2}{F_{n-1}} < 2kF_{n-1}$. Similarly, if the second inequality was not satisfied just before moment t_{n-1} then we had $\frac{A_{n-1}-F_{n-1}^2}{F_{n-1}} < P_{n-1} + 1$.

Case 1: $\frac{A_{n-1}-F_{n-1}^2}{F_{n-1}} < 2kF_{n-1}$

Notice that $k \geq 1$, since otherwise $F_1 = \sqrt{2}/2$ and the termination condition would never be satisfied in the first stage. We have:

$$\begin{aligned} \frac{A_{n-1}}{F_{n-1}} &\leq 2kF_{n-1} + F_{n-1} \\ kF_n &\leq 2kF_{n-1} + 1, \quad \text{since } F_n \leq \frac{A_{n-1}}{kF_{n-1}} \text{ and } F_{n-1} \leq 1 \\ kF_n &\leq \sqrt{2}\sqrt{A_{n-1}k} + 1, \quad \text{since } \frac{A_{n-1}}{F_{n-1}} \geq 2kF_{n-1} \\ &\quad \text{and thus } (kF_{n-1})^2 \leq \frac{1}{2} \frac{A_{n-1}}{F_{n-1}} kF_{n-1} \\ kF_n &\leq 2\sqrt{Ak} + 1, \quad \text{since } A_{n-1} \leq A \end{aligned}$$

Since the termination condition in the last stage is not satisfied, we have either $A/F_n \leq \max\{2kF_n, P+1\}$ and thus $A/F_n = O(P+kF_n)$, or $F_n = \sqrt{2}/2$ and thus $A/F_n = O(A)$. We obtain that $D_n \leq d(P+A/F_n+(k+1)F_n) = O(P+\sqrt{Ak}+A)$. From the previous serie of inequalities, we also have $\sum_{i=1}^{n-1} D_i \leq 6d(A_{n-1}/F_{n-1}) = O(\sqrt{Ak})$. Hence, the total cost is $O(P+\sqrt{Ak}+A)$.

Case 2: $\frac{A_{n-1}-F_{n-1}^2}{F_{n-1}} < P_{n-1} + 1$

We have:

$$\begin{aligned} \frac{A_{n-1}}{F_{n-1}} &\leq P_{n-1} + F_{n-1} + 1 \\ kF_n &\leq P_{n-1} + F_{n-1} + 1, \quad \text{since } kF_n \leq \frac{A_{n-1}}{F_{n-1}} \\ kF_n &\leq P+2, \quad \text{since } F_{n-1} \leq 1 \text{ and } P_{n-1} \leq P \end{aligned}$$

As before, we have $A/F_n = O(P+kF_n)$ or $A/F_n = O(A)$. We obtain that $D_n \leq d(P+A/F_n+(k+1)F_n) = O(P+A)$. We have $\sum_{i=1}^{n-1} D_i \leq 6d(A_{n-1}/F_{n-1}) = O(P)$. Hence, the total cost is $O(P+A)$. \square

The following theorem shows that the lengths of trajectories in Lemma 3.3 and in Theorem 3.1 are asymptotically worst-case optimal.

Theorem 3.2. *There are some terrains in which any algorithm for a robot with limited visibility, exploring polygonal terrains of area A , perimeter P and with k obstacles, produces trajectories of length $\Omega(P + A + \sqrt{Ak})$, even if the terrain is known to the robot.*

Proof: In order to prove our lower bound we present three families of terrains. A terrain in the first family (cf. Fig. 3(a)) is a square with identical obstacles of diameter ϵ located on a $\sqrt{k} \times \sqrt{k}$ grid. The side of the square is $\sqrt{A+x}$, where x is the negligible total area of all k obstacles and the total perimeter of all obstacles is 1. By the same arguments as in the proof of Theorem 2.2, any exploration trajectory must be of length at least $(k-1) \left(\frac{\sqrt{A+x}}{\sqrt{k+1}} - \epsilon \right)$, which is in $\Omega(\sqrt{Ak})$. At the same time we have $P = \Theta(\sqrt{A})$ (P cannot be smaller). A terrain in the second family (cf. Fig. 3(b)) is a polygon of arbitrarily small area (without obstacles), whose exploration requires a trajectory of size $\Omega(P)$. A terrain in the third family is the empty square of side \sqrt{A} . Now we have $P = \Theta(\sqrt{A})$ and $k = 0$. Finally, when the robot traverses a distance d , it explores a new area of at most $2d$. So, the robot has to travel a distance $\Omega(A)$ to explore such a terrain. These three families of terrains justify our lower bound. \square

The examples from the above proof can be adjusted so that all terrains are c -fat, for any constant $c > 1$. Lemma 3.3 and Theorem 3.2 imply

Theorem 3.3. *Consider terrains of area A , perimeter P and with k obstacles. If either A or k is known before the exploration, then the exploration of any such terrain can be performed using a trajectory of length $O(P + A + \sqrt{Ak})$, which is asymptotically worst-case optimal.*

Notice that in order to explore a terrain at cost $O(P + A + \sqrt{Ak})$, it is enough to know the parameter A or k up to a multiplicative constant, rather than the exact value. This can be proved by a careful modification of the proof of Lemma 3.3. For the sake of clarity, we stated and proved the weaker version of Lemma 3.3, with knowledge of the exact value.

Suppose now that no a priori knowledge of any parameters of the terrain is available. We iterate Algorithm LET_A or LET_k for A (resp. k) equal $1, 2, 4, 8, \dots$ interrupting the iteration and doubling the parameter as soon as the explored area (resp. the number of obstacles seen) exceeds the current parameter value. The algorithm stops when the entire terrain is explored (which happens at the first probe exceeding the actual unknown value of A , resp. k). We get an exploration algorithm using a trajectory of length $O((P + A + \sqrt{Ak}) \log A)$, resp. $O((P + A + \sqrt{Ak}) \log k)$. By interleaving the two procedures we get the minimum of the two costs. Thus we have the following corollary.

Corollary 3.1. *Consider terrains of area A , perimeter P and with k obstacles. Exploration of any such terrain can be performed without any a priori knowledge at cost differing from the worst-case optimal cost with full knowledge only by a factor $O(\min\{\log A, \log k\})$.*

- [1] X. Deng, T. Kameda, C. Papadimitriou, How to learn an unknown environment. i: the rectilinear case, *Journal of the ACM (JACM)* 45 (2) (1998) 215–245.
- [2] X. Deng, T. Kameda, C. Papadimitriou, How to learn an unknown environment, in: *Proceedings of the 32nd annual symposium on Foundations of computer science*, IEEE Computer Society, 1991, pp. 298–303.
- [3] B. Kalyanasundaram, K. Pruhs, A competitive analysis of algorithms for searching unknown scenes, *Computational Geometry* 3 (3) (1993) 139–155.
- [4] F. Hoffmann, C. Icking, R. Klein, K. Kriegel, The polygon exploration problem, *SIAM Journal on Computing* 31 (2) (2001) 577–600.
- [5] S. Albers, K. Kursawe, S. Schuierer, Exploring unknown environments with obstacles, *Algorithmica* 32 (2002) 123–143.
- [6] Y. Gabriely, E. Rimon, Spanning-tree based coverage of continuous areas by a mobile robot, *Annals of Mathematics and Artificial Intelligence* 31 (1) (2001) 77–98.
- [7] Y. Gabriely, E. Rimon, Competitive on-line coverage of grid environments by a mobile robot, *Computational Geometry* 24 (3) (2003) 197–224.
- [8] S. K. Ghosh, J. W. Burdick, A. Bhattacharya, S. Sarkar, Online algorithms with discrete visibility-exploring unknown polygonal environments, *Robotics & Automation Magazine*, IEEE 15 (2) (2008) 67–76.
- [9] C. Icking, T. Kamphans, R. Klein, E. Langetepe, Exploring an unknown cellular environment, in: *In Abstracts of the 16th European Workshop on Computational Geometry*, 2000, pp. 140–143.
- [10] A. Kolenderska, A. Kosowski, M. Małafiejski, P. Żyliński, An improved strategy for exploring a grid polygon, in: *Proceedings of the 16th international conference on Structural Information and Communication Complexity*, Springer-Verlag, 2009, pp. 222–236.
- [11] S. Ntafos, Watchman routes under limited visibility, *Computational Geometry: Theory and Applications* 1 (3) (1992) 149–170.
- [12] T. Bandyopadhyay, Z. Liu, M. H. Ang, W. K. G. Seah, Visibility-based exploration in unknown environment containing structured obstacles, in: *Advanced Robotics, 2005. ICAR'05. Proceedings., 12th International Conference on*, IEEE, 2005, pp. 484–491.

- [13] N. Cuperlier, M. Quoy, P. Gaussier, Navigation and planning in an unknown environment using vision and a cognitive map, in: European Robotics Symposium 2006, Springer, 2006, pp. 129–142.
- [14] R. Sim, J. J. Little, Autonomous vision-based exploration and mapping using hybrid maps and rao-blackwellised particle filters, in: International Conference on Intelligent Robots and Systems, IEEE, 2006, pp. 2082–2089.
- [15] E. Bareli, P. Berman, A. Fiat, P. Yan, Online navigation in a room, *Journal of Algorithms* 17 (3) (1994) 319–341.
- [16] P. Berman, A. Blum, A. Fiat, H. Karloff, A. Rosén, M. Saks, Randomized robot navigation algorithms, in: Proceedings of the seventh annual ACM-SIAM symposium on Discrete algorithms, Society for Industrial and Applied Mathematics, 1996, pp. 75–84.
- [17] A. Blum, P. Raghavan, B. Schieber, Navigating in unfamiliar geometric terrain, *SIAM Journal on Computing* 26 (1) (1997) 110–137.
- [18] C. H. Papadimitriou, M. Yannakakis, Shortest paths without a map, *Theoretical Computer Science* 84 (1) (1991) 127–150.
- [19] B. Tovar, R. Murrieta-Cid, S. M. LaValle, Distance-optimal navigation in an unknown environment without sensing distances, *IEEE Transactions on Robotics* 23 (3) (2007) 506–518.
- [20] R. Osserman, The isoperimetric inequality, *Bulletin (New Series) of the American Mathematical Society* 84 (6) (1978) 1182–1238.