

Realizing Metric Spaces with Convex Obstacles

Sándor Kisfaludi-Bak  

Aalto University, Finland

Leonidas Theocharous  

University of Ottawa, Canada

Abstract

The presence of obstacles has a significant impact on distance computation, motion-planning, and visibility. These problems have been studied extensively in the planar setting, while our understanding of these problems in 3- and higher-dimensional spaces is still rudimentary. In this paper, we study the impact of different types of obstacles on the induced geodesic metric in 3-dimensional Euclidean space. We say that a finite metric space (X, dist_X) is *approximately realizable* by a collection $\mathcal{T}$ of obstacles in $\mathbb{R}^3$ if for any $\varepsilon > 0$ it can be embedded into $(\mathbb{R}^3 \setminus \bigcup_{T \in \mathcal{T}} T, \text{dist}_{\mathcal{T}})$ with worst-case multiplicative distortion $1 + \varepsilon$, where $\text{dist}_{\mathcal{T}}$ denotes the geodesic distance in the free space induced by $\mathcal{T}$. We focus on three key geometric properties of obstacles –convexity, disjointness, and fatness– and examine how dropping each one of them affects the existence of such embeddings.

Our main result concerns dropping the fatness property: we demonstrate that any finite metric space is realizable with $1 + \varepsilon$ worst-case multiplicative distortion using a collection of convex and pairwise disjoint obstacles in $\mathbb{R}^3$, even if the obstacles are congruent and equilateral triangles. Based on the same construction, we can also show that if we require fatness but drop any of the other two properties instead, then we can still approximately realize any finite metric space.

Our results have important implications on the approximability of TSP WITH OBSTACLES, a natural variant of TSP introduced recently by Alkema *et al.* (ESA 2022). Specifically, we use the recent results of Banerjee *et al.* on TSP in doubling spaces (FOCS 2024) and of Chew *et al.* on distances among obstacles (Inf. Process. Lett. 2002) to show that TSP WITH OBSTACLES admits a PTAS if the obstacles are convex, fat, and pairwise disjoint. If any of these three properties is dropped, then our results, combined with the APX-hardness of METRIC TSP, demonstrate that TSP WITH OBSTACLES is APX-hard.

2012 ACM Subject Classification Theory of computation → Computational geometry

Keywords and phrases traveling salesman, geodesic distance

Digital Object Identifier 10.4230/LIPIcs.ISAAC.2025.46

Related Version *Full Version*: <https://arxiv.org/abs/2509.14709> [26]

Funding *Sándor Kisfaludi-Bak*: This work was supported by the Research Council of Finland, Grant 363444.

1 Introduction

Understanding navigation [3, 10, 20, 36, 37], distance [5, 23, 28, 38, 40, 41] and visibility [12, 16, 27, 29, 33, 35] problems among obstacles has been one of the well-motivated research directions in computational geometry. Most of the work has focused on the planar setting, where the metric could be the geodesic distance inside a simple polygon, or more generally in a polygon with holes in the Euclidean plane. In the plane one can benefit from the fact that paths will often cross, whereas in higher dimensions paths will typically stay disjoint. The additional degree of freedom in 3- and higher-dimensional space makes the fundamental distance and navigation problems among obstacles harder, and they have not been explored as thoroughly.

Fortunately, we can still get a useful metric structure for certain obstacles in d -dimensional Euclidean space (henceforth denoted by $\mathbb{R}^d$). We say that a set $\mathcal{T}$ of obstacles is α -fat if for each obstacle object $T \in \mathcal{T}$ we have $\frac{r_{in}(T)}{r_{circum}(T)} \geq \alpha$, where $r_{in}(T)$ denotes the radius of the



© Sándor Kisfaludi-Bak and Leonidas Theocharous;
licensed under Creative Commons License CC-BY 4.0

36th International Symposium on Algorithms and Computation (ISAAC 2025).

Editors: Ho-Lin Chen, Wing-Kai Hon, and Meng-Tsung Tsai; Article No. 46; pp. 46:1–46:21

Leibniz International Proceedings in Informatics



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

maximum inscribed ball of T , and $r_{\text{circum}}(T)$ denotes the radius of the minimum enclosing ball of T . We will think of α as a universal constant and talk about collections of fat objects. For a set $\mathcal{T}$ of obstacles let $\mathcal{F}_{\mathcal{T}} = \mathbb{R}^d \setminus \bigcup_{T \in \mathcal{T}} T$ denote the *free space* outside the obstacles. We denote by $\text{dist}_{\mathcal{T}}(u, v)$ the geodesic distance between $u, v \in \mathcal{F}_{\mathcal{T}}$, which is the infimum of the length of piecewise linear curves in $\mathcal{F}_{\mathcal{T}}$ containing u and v . Chew et al. [11] established that for any set of convex, disjoint and fat obstacles in $\mathbb{R}^d$ one can navigate around the obstacles with constant overhead compared to the Euclidean distance. Their theorem implies that $(\mathcal{F}_{\mathcal{T}}, \text{dist}_{\mathcal{T}})$ has bounded *doubling dimension*, that is, for any $r > 0$ and any $x \in \mathcal{F}_{\mathcal{T}}$ the geodesic ball of center x and radius r can be covered by at most 2^δ geodesic balls of radius $r/2$, for some constant δ . The minimum number δ satisfying the above condition is the doubling dimension. In Section 2 we establish an explicit bound on the doubling dimension among pairwise disjoint convex α -fat obstacles in $\mathbb{R}^d$. Metrics of bounded doubling dimension are well-studied generalizations of Euclidean spaces, and retain many useful properties of $\mathbb{R}^d$.

A natural question is whether the conditions of fatness, convexity, and pairwise disjointness are all needed for the doubling dimension bound. More generally, one wonders what finite metric spaces can be realized if the condition of fatness, disjointness or convexity is dropped. We say that a metric space (X, dist_X) is *realizable* with a set $\mathcal{T}$ of obstacles in $\mathbb{R}^d$ if one can embed (X, dist_X) isometrically into $(\mathcal{F}_{\mathcal{T}}, \text{dist}_{\mathcal{T}})$. An *approximate realization* allows for an embedding with a worst-case multiplicative distortion of $(1 + \varepsilon)$.

The starting point of our investigation is the following question. Let $\mathfrak{P}$ denote the disjunction of some subset of properties among (i) convex (ii) pairwise disjoint (iii) fat.

► **Question 1.** *Is it possible to (approximately) realize the finite metric space (X, dist_X) in $\mathbb{R}^3$ with obstacles of property $\mathfrak{P}$?*

It is also worth exploring how various types of obstacles impact the algorithms and complexity of some classic distance problems. The study of the shortest path problem among a set of polygonal obstacles $\mathcal{P}$ goes back at least 50 years, to the work of Wangdahl et al. [41] in 1974. Their approach to compute a shortest path between two points s, t (as well as many later approaches; see for instance [5, 38]) relied on running Dijkstra’s algorithm on the visibility graph of the vertices of $\mathcal{P}$ together with s and t . Any such method has an $\Omega(n^2)$ lower bound (where n is the number of vertices in $\mathcal{P}$), due to the size of the visibility graph. To achieve subquadratic running times, later algorithms shifted to the continuous Dijkstra paradigm [23, 28]. Finally, Wang presented an optimal $O(n \log n)$ -time, $O(n)$ -space algorithm [40]. Data structures for two-point shortest path queries have also been studied and, recently, De Berg et al. presented an $O(n^{10+\varepsilon})$ -space data structure with $O(\log n)$ query time [14]. The shortest path problem among obstacles can be generalized to account for movement through regions with different weights, a variation known as the *weighted region problem*. The general problem admits a $(1 + \varepsilon)$ -approximation algorithm [30], while cases with restricted region shapes or weights have also been considered [9, 17]. While these are all positive results, computing shortest paths among 3-dimensional obstacles is a much harder problem. In fact, it is NP-hard [31]. On the other hand, the problem admits an FPTAS [13, 21].

It is also natural to consider classic geometric optimization problems in the presence of obstacles. Recently, Alkema et al. [4] introduced a natural generalization of EUCLIDEAN TSP called TSP WITH OBSTACLES. In TSP WITH OBSTACLES, we are given a set of n points to visit and a collection $\mathcal{T}$ of obstacles of total complexity m to avoid. Our goal is to find the shortest tour between a set of sites in the plane or some higher-dimensional space, while avoiding a given set of obstacles. In [4] the authors studied a variant of TSP WITH OBSTACLES called

TSP IN A SIMPLE POLYGON, where the sites lie in a simple polygon P , that is, the obstacle that the salesman has to avoid is the complement of P . They gave an exact algorithm for this problem, running in $2^{O(\sqrt{n} \log n)} + \text{poly}(n, m)$ time.

In this paper, we will demonstrate the algorithmic impact of various types of obstacles via studying approximation algorithms and APX-hardness for TSP WITH OBSTACLES. In particular, we consider the problem in $\mathbb{R}^3$ for different kinds of obstacles. Again, let $\mathfrak{P}$ denote the disjunction of some subset of properties among (i) convex (ii) pairwise disjoint (iii) fat.

► **Question 2.** *Is there a PTAS for TSP WITH OBSTACLES in $\mathbb{R}^3$ if the obstacles satisfy property $\mathfrak{P}$?*

1.1 Our contribution

Our main contribution is the following theorem, which shows that all finite metric spaces are approximately realizable with pairwise disjoint convex obstacles in $\mathbb{R}^3$.

► **Theorem 3.** *Let $\varepsilon \in (0, 1)$ and let (X, dist_X) be a metric space of size $|X| = n$ and spread $\Phi := \max_{a,b,c,d \in X} \left(\frac{\text{dist}_X(a,b)}{\text{dist}_X(c,d)} \right)$. Then there exists a collection $\mathcal{T}$ of $O(n^{17} \Phi^5 / \varepsilon^5)$ pairwise disjoint congruent equilateral triangular obstacles and an injection $f : X \rightarrow \mathbb{R}^3$ such that*

$$\text{dist}_X(a, b) \leq \text{dist}_{\mathcal{T}}(f(a), f(b)) \leq (1 + \varepsilon) \cdot \text{dist}_X(a, b).$$

for all $a, b \in X$. The set of obstacles can be constructed in $\text{poly}(n, \Phi, 1/\varepsilon)$ time.

Let $\mathfrak{P}$ denote the disjunction of some subset of properties among (i) convex (ii) pairwise disjoint (iii) fat. The theorem directly answers Question 1 if the property $\mathfrak{P}$ does not require that objects are fat. We can use some ideas from this construction to give a complete answer to Question 1 in case of approximate realizations: we can approximately realize any metric space with fat pairwise disjoint obstacles (that are non necessarily convex) and with convex fat obstacles that are not necessarily disjoint, see Theorem 15.


We note that our construction is quite generic and we get the same result with other convex shapes that are constant-fat in $\mathbb{R}^2$, such as with unit disks or squares instead of congruent equilateral triangles. Using Theorem 3 and Theorem 15 we can use TSP lower bounds designed for general metric spaces to answer Question 2.

► **Corollary 4.** *Let $\mathcal{T}$ be a collection of obstacles in $\mathbb{R}^3$ with property $\mathfrak{P}$ of total complexity $m = \text{poly}(|X|)$. Then TSP WITH OBSTACLES has a PTAS if $\mathfrak{P} = \text{“convex and pairwise disjoint and fat”}$, and it is APX-hard for all other $\mathfrak{P}$.*

The PTAS is based on the fact that for $\mathfrak{P} = \text{“convex and pairwise disjoint and fat”}$ the free space outside the obstacles has bounded doubling dimension, see Section 2. We can also efficiently compute approximate shortest paths using Har-Peled’s algorithm [21], and running any of the known PTASes for TSP in doubling metrics [6, 8] we get the desired algorithm. The lower bound is based on realizing the metric space from the lower bound of Karpinski, Lampis and Schmied [25] using obstacles. The construction in [25] has polynomial spread, which allows us to apply Theorem 3. See Section 5 for a proof of Corollary 4.

1.2 Related work

Motion planning is a central problem in robotics, where the goal is to compute a collision-free path for a set of moving objects (robots) in a space with obstacles. The complexity of the problem depends significantly on the underlying environment and the degrees of freedom



■ **Figure 1** (i) The construction of Dumitrescu and Tóth [15] for a wall in $\mathbb{R}^2$ consisting of convex obstacles. The shortest path from a to b must zigzag between the layers of segments. (ii) The view of a flat wall from $z = \infty$, with 4 shifts (iii) The flat wall construction viewed from a generic point. The planes containing the four types of squares have slightly differing z -coordinates, and they are sandwiched between the planes $z = 0$ and $z = 0.05$.

of the robots. The pioneering works of Schwartz and Sharir [36, 37], and later Canny [10], introduced the first general techniques for solving this problem in the plane. These algorithms have a polynomial dependency on n , the combinatorial complexity of the obstacles, but an exponential dependency on k , the degrees of freedom of the robots. For a detailed discussion of these foundational works, refer to [19, 20]. To highlight the problem's inherent difficulty, note that even in the restricted case of rectangular robots navigating a rectangular environment, the problem is PSPACE-hard [24]. Recently, Agarwal [3] *et al.* presented an FPTAS with a running time of $O(n^2 \varepsilon^{-O(1)} \log n)$ for the case of two square robots moving in a polygonal environment. This is the first known PTAS for two robots in such settings.


Focusing now on geodesic distance, many well-studied problems in the Euclidean setting have been explored within polygonal domains under this metric. Notable examples include geodesic convex hulls of points in polygonal domains [7, 39], the 1- and 2-center problems [32, 34], and geodesic spanners [1].

Finally, Abam and Seraji [2] constructed an $8\sqrt{3}$ -spanner of size $O(n \log^3 n)$ amid axis-parallel boxes in $\mathbb{R}^3$, which is the only known geodesic spanner result in 3-dimensional space with obstacles.

1.3 Overview of the proof of Theorem 3

Our approach to proving Theorem 3 is inspired by the work of Dumitrescu and Tóth [15], who studied a similar problem in the planar setting. They denote by $\varrho(P)$ the ratio between the geodesic and the Euclidean diameter of a polygon P , and show that over all convex polygons with h convex holes, the supremum of $\varrho(P)$ is between $\Omega(h^{1/3})$ and $O(h^{1/2})$. Crucially, their construction includes building a “wall” of polynomially many short segments that lengthen the geodesic diameter by a large factor, see Figure 1(i) for an illustration.

Our strategy to embed a metric space is as follows. We will define a custom closed surface $\mathcal{S}$ embedded in $\mathbb{R}^3$, and map each point $x_i \in X$ to some point $f(x_i)$ inside. We then place a large collection of obstacles near the surface $\mathcal{S}$ (within distance $1/2$) that ensures that any geodesic that would connect the inside and outside of $\mathcal{S}$ must be very long, longer than the maximum distance in (X, dist_X) . This means that the minimum distance between a point $f(x_i)$ and $f(x_j)$ must be realized by a geodesic that stays within $\mathcal{S}$. By setting up $\mathcal{S}$ correctly, we are able to give an upper and lower bound on the length of any geodesic connecting $f(x_i)$ and $f(x_j)$ so that we can ensure $\text{dist}_X(x_i, x_j) \leq \text{dist}_{\mathcal{T}}(f(x_i), f(x_j)) \leq (1 + \varepsilon) \cdot \text{dist}_X(x_i, x_j)$.




■ **Figure 2** The surface $\mathcal{S}$ for a 3-point metric space. The shortest geodesic from $f(x_i)$ to $f(x_j)$ is forced to go through some tubes, consisting of two vertical and a horizontal cylinder. The length of such a path can be adjusted by changing the length of the vertical cylinders.

Construction of the surface $\mathcal{S}$. Before discussing our obstacle placement, we first describe the construction of $\mathcal{S}$ and the embedding f . (Note that the main text will define the surface $\mathcal{S}$ last.) First, without loss of generality we assume that the minimum distance in (X, dist_X) is at least some large polynomial in $n := |X|$, as we can use a scaling on the final construction to get the desired distance. The function f simply assigns the points x_i to equally spaced points along the x -axis. The surface $\mathcal{S}$ will contain a cube $\mathcal{R}_i$ centered at each point $f(x_i)$.

For each pair i, j we realize the distance of $\text{dist}_X(x_i, x_j)$ via a sequence of three tubes (cylinders), where the first and third cylinder are vertical (parallel to the z -axis) and the second is parallel to the x -axis, see Figure 2. The tubes connect a hole in the top of the cube $\mathcal{R}_i$ to a hole on the cube $\mathcal{R}_j$, and in order to realize the turns, we use two small cubes with holes that connect the first and second as well as the second and third tube in the sequence. Notice that the length of the second horizontal tube is determined by $f(x_i)$ and $f(x_j)$, thus in order to set the length of the cube sequence, we can vary the length of the two vertical tubes. We place these tubes at different depths (y -coordinates) so that they remain disjoint from each other. One can see that this leads to a good approximate realization, however, this surface needs to be smoother (differentiable) in order to ease the placement of obstacles. Thus $\mathcal{S}$ is defined using differentiable surface patches: we need a rounded cube and a differentiable joint that allows the connection of a flat cube face with a circular cylinder boundary.

A flat wall example. Let us consider an example of a simple placement of convex objects that increases the geodesic distance between the half-space $z < 0$ and the half-space $z > 1$ to significantly more than 1. We call this construction a *flat wall*. In this construction we use squares rather than triangles. Let S_1 be the set of axis-parallel squares of side length 0.99 centered at the points $(a, b, 0.01)$ for $a, b \in \mathbb{Z}$ in the $z = 0.01$ plane. Repeat the same construction 3 more times at heights $z = 0.02, 0.03, 0.04$ by translating S_1 with the vectors $(1/2, 0, 0.01), (1/2, 1/2, 0.02), (0, 1/2, 0.03)$, to get the sets S_2, S_3 , and S_4 , respectively. See Figure 1(ii) and (iii). Consider a geodesic connecting the planes $z = 0$ and $z = 0.05$: its intersection point p with $z = 0$ will be within ℓ_∞ -distance 0.25 to the center of some square in one of the four shifts, and thus it will need to have length at least 0.24 to reach the boundary of this square before it can get to $z = 0.05$. Using further copies of these squares shifted to



■ **Figure 3** (i) A joint patch with its corresponding net. (ii) A quarter-cylinder patch with its corresponding net. (iii) A spherical triangle patch with its corresponding net.

different height ranges can increase the lower bound of 0.24 as necessary. By making more layers that are packed more densely we can achieve any separation between the halfspaces $z < 0$ and $z > 1$.

While the above construction is quite intuitive, we need a more careful approach to make a similar separation near an embedded closed surface. One can see that the flat construction cannot be adapted to a surface with sharp edges, so we will need to use a smoother closed surface. Even then, we will not be able to place layers of convex objects arbitrarily close together, as convex 2-dimensional shapes will not stay inside a curved surface. The challenge is to keep our objects in different layers disjoint, while creating a strong separation.

Curved wall construction via placing triangle layers. The surface $\mathcal{S}$ that we define has the crucial property that it has offset surfaces (often called parallel surfaces) that are well-defined and differentiable for any distance offset $\delta \in [-1/2, 1/2]$ from $\mathcal{S}$. The offset surface at distance δ is denoted by $\mathcal{S}(\delta)$. Since $\mathcal{S}(\delta)$ is a closed surface, we are able to talk about the inside and outside of $\mathcal{S}(\delta)$, and we can also define the region between two offsets $\mathcal{S}(\delta)$ and $\mathcal{S}(\delta')$ when $\delta < \delta'$. In our construction, the offset δ will have a role similar to the z -coordinate in the flat wall construction above: we will place triangles tangent to $\mathcal{S}(\delta)$ at some given collection of points. Using a large collection of different offsets, we can ensure the required separation between the inside of $\mathcal{S}(-1/2)$ and the outside of $\mathcal{S}(1/2)$.

We need a handful of geometric properties from $\mathcal{S}(\delta)$ to ensure that we can place our obstacles near it. The most important property is to limit their curvature¹ in some sense. We require that at any point p on $\mathcal{S}(\delta)$ there exist two balls of radius at least $1/2$ whose unique intersection point with $\mathcal{S}(\delta)$ and with each other is p . This helps to ensure that a small 2-dimensional convex shape of diameter $\mu \ll 1$ tangent to $\mathcal{S}(\delta)$ at p stays within distance $O(\mu^2)$ of $\mathcal{S}(\delta)$. In particular, a small triangle will stay between the surfaces $\mathcal{S}(\delta - \mu^2)$ and $\mathcal{S}(\delta + \mu^2)$, and thus triangles of the same size tangent to $\mathcal{S}(\delta - 2\mu^2)$ and $\mathcal{S}(\delta + 2\mu^2)$ will remain disjoint from this triangle.

Finally, we must place the triangles in neighboring layers with “shifts” to ensure that a constant number of consecutive triangular layers of side-length μ triangles creates a separation of at least $\Omega(\mu)$. Here translation is not an option, so we define an (a, b) -net, a weaker version of an ε -net (in the metric space sense). For two points $u, v \in \mathbb{R}^3$ we will denote by uv the

¹ While our results could be presented with a more general differential geometric toolbox, the need to construct these (a, b) -nets prompted us to use specific surface patches that are easy to handle algorithmically. One could adapt our techniques to more general smooth embedded surfaces. As a trade-off our surfaces are not smooth (not even twice differentiable) at the shared boundary of our surface patches. We have also used the Euclidean distance function rather than measuring distances within these surfaces to make our construction more elementary and explicit.

line segment between u and v , and by $|uv|$ the length of uv . An (a, b) -net of a surface $\mathcal{S}$ embedded in $\mathbb{R}^3$ is a point set $N \subset \mathcal{S}$ where the pairwise distance of points in N is at least a , and for any point $q \in \mathcal{S}$ there is a $p \in N$ such that $|pq| \leq b$.

We can show that the offsets of the net points provide a weaker but still useful net in the offset surfaces: a “shared” $(\Omega(\zeta), O(\zeta))$ -net can be constructed for the surfaces $\mathcal{S}(\delta)$. Here, “shared” means that we first find a $(\zeta, O(\zeta))$ -net on $\mathcal{S}$; sliding every net point along its surface normal gives, on each surface $\mathcal{S}(\delta)$, an $(\Omega(\zeta), O(\zeta))$ -net. These net points are then partitioned into constantly many classes, where the distances within points in each class C are large enough so that equilateral triangles of side length $\mu = \Theta(\zeta)$ tangent to $\mathcal{S}(\delta)$ at each point of C will remain disjoint; i.e., these triangles would correspond to a single square set S_1 from our flat wall example. Using all of these partition classes in different offsets $\mathcal{S}(\delta + i\zeta)$ ensures a separation of $\Omega(\zeta)$ for geodesics passing between these layers, as seen with the four shifts of squares in our flat wall example. In our construction we must set ζ small enough to be able to accommodate several iterations of the above construction between $\mathcal{S}(-1/2)$ and $\mathcal{S}(1/2)$, but large enough to get the desired length lower bound for any geodesic going from the inside of $\mathcal{S}(-1/2)$ to the outside of $\mathcal{S}(1/2)$. This concludes the generic construction of the curved wall, as well as the overview of the ideas behind Theorem 3.

2 Doubling dimension among convex fat disjoint obstacles

We denote by $B(p, r)$ the Euclidean ball of radius r centered at $p \in \mathbb{R}^d$, and by $B_{\mathcal{T}}(p, r)$ the geodesic ball of radius r centered at $p \in \mathcal{F}_{\mathcal{T}}$.

We recall the following key result by Chew et al. [11] on convex fat objects. Based on this, we will subsequently show that the free space induced by a set of convex, fat and disjoint obstacles in $\mathbb{R}^d$ has bounded doubling dimension, under the shortest-path metric.

► **Lemma 5.** *Let $\mathcal{T}$ be a set of convex, α -fat and disjoint obstacles in $\mathbb{R}^d$. Then, for any $u, v \in \mathcal{F}_{\mathcal{T}}$, we have that $\text{dist}_{\mathcal{T}}(u, v) \leq \beta \cdot |uv|$, where β is a constant depending on d and α defined as follows:*

$$\beta = \begin{cases} 1 + \frac{4}{\pi\alpha}, & \text{for } d = 2, \\ 1 + \frac{8d^d}{\alpha}, & \text{for } d \geq 3. \end{cases}$$

Since a set $\mathcal{T}$ of convex, fat and disjoint obstacles distorts distances by a constant factor, it follows that the doubling dimension of $(\mathcal{F}_{\mathcal{T}}, \text{dist}_{\mathcal{T}})$ will be bounded; we provide a proof for completeness.

► **Lemma 6.** *Let $\mathcal{T}$ be a set of convex, α -fat and disjoint obstacles in $\mathbb{R}^d$. Then the metric space $(\mathcal{F}_{\mathcal{T}}, \text{dist}_{\mathcal{T}})$ has doubling dimension $O(d(d \log d + \log(1/\alpha)))$.*

Proof. Let $B_{\mathcal{T}}(p, r)$ be a geodesic ball centered at $p \in \mathcal{F}_{\mathcal{T}}$. Note that $B_{\mathcal{T}}(p, r) \subseteq B(p, r)$. We will construct a covering of $B_{\mathcal{T}}(p, r)$ by geodesic balls of radius $r/2$. Towards that, we cover $B(p, r)$ by c_d Euclidean balls of radius $r/2$, where $\log(c_d) = \Theta(d)$ is the doubling dimension of $\mathbb{R}^d$ [22, Ch. 10]. Repeating this process recursively $\log(2\beta) + 1$ times, the ball $B(p, r)$ is covered by $c_d^{\log(2\beta)+1}$ Euclidean balls of radius $r/(2^{\log(2\beta)+1}) = r/(4\beta)$. Let $\mathcal{B}$ denote the resulting collection of Euclidean balls. Since $B_{\mathcal{T}}(p, r) \subseteq B(p, r)$, this collection also covers $B_{\mathcal{T}}(p, r)$.

We transform $\mathcal{B}$ into a collection $\mathcal{B}_{\mathcal{T}}$ of geodesic balls in $(\mathcal{F}_{\mathcal{T}}, \text{dist}_{\mathcal{T}})$ of radius $r/2$, such that the union of these geodesic balls covers $B_{\mathcal{T}}(p, r)$. For each Euclidean ball $B(q, r/(4\beta)) \in \mathcal{B}$, we handle the following cases:

1. If $B(q, r/(4\beta)) \cap \mathcal{F}_{\mathcal{T}} = \emptyset$, then we ignore $B(q, r/(4\beta))$ in the cover.
2. If $B(q, r/(4\beta)) \cap \mathcal{F}_{\mathcal{T}} \neq \emptyset$, let $o \in B(q, r/(4\beta)) \cap \mathcal{F}_{\mathcal{T}}$. Include the geodesic ball $B_{\mathcal{T}}(o, r/2)$ in $\mathcal{B}_{\mathcal{T}}$. To verify coverage, note that if $s \in B(q, r/(4\beta)) \cap \mathcal{F}_{\mathcal{T}}$, then:

$$|os| \leq \frac{r}{2\beta},$$

By Lemma 5, the geodesic distance between o and s satisfies:

$$\text{dist}_{\mathcal{T}}(o, s) \leq \beta \cdot |os| \leq \beta \cdot \frac{r}{2\beta} = \frac{r}{2}.$$

Hence, $s \in B_{\mathcal{T}}(o, r/2)$, and $B(q, r/(4\beta)) \cap \mathcal{F}_{\mathcal{T}} \subseteq B_{\mathcal{T}}(o, r/2)$.

Each ball in $\mathcal{B}_{\mathcal{T}}$ has radius $r/2$, and the union of these geodesic balls covers $B_{\mathcal{T}}(p, r)$. The total number of geodesic balls in $\mathcal{B}_{\mathcal{T}}$ is at most $c_d^{\log(2\beta)+1}$. Therefore, the doubling dimension of $(\mathcal{F}_{\mathcal{T}}, \text{dist}_{\mathcal{T}})$ is given by

$$\log(c_d^{\log(2\beta)+1}) = (\log(2\beta) + 1) \cdot \log(c_d) = O(d(d \log d + \log(1/\alpha))). \quad \blacktriangleleft$$

3 Realizing metric spaces with convex disjoint obstacles

This section contains the main result of our paper. We show that by using only convex and disjoint obstacles in $\mathbb{R}^3$, it is possible to realize any metric space approximately.

3.1 Surface patches

We will consider surfaces embedded in $\mathbb{R}^3$ that are differentiable (i.e., have well-defined normals at each point) described by differentiable surface patches. More precisely, our surface $\mathcal{S}$ will consist of the following types of surface patches:

- Axis-parallel squares of side length at least 1 with 0 or more disjoint circular holes of radius 2, such that each hole is at distance at least 2 from the square, and the hole centers have pairwise distance at least 6. We will refer to such a patch as a *square patch*.
- Sections of a sphere of radius 1 cut out by axis-parallel planes through the center of the sphere. These sections correspond to octants of spheres and will be referred to as *spherical triangle patches*.
- Cylinders of axis length at least 1 and radius 1, or their sections of angle $\pi/2$ around their axes, where the cylinder axis is parallel to some coordinate axis. We will refer to such a patch as a *cylinder patch* and a *quarter-cylinder patch*, respectively.
- *Joints*, which are defined as a surface of revolution obtained by rotating the quarter circular arc $\{(x, y) \mid x^2 + y^2 = 1, x \geq 0, y \geq 0\}$ around the axis $x = 2$. We only allow rotations where the axis of revolution is parallel to some coordinate axis. We will refer to such a patch as a *joint patch*.

Notice in particular that a joint patch can be used to connect a circular hole of a square patch to a cylinder patch, and the quarter cylinder and spherical triangle patches can be used to “round” an axis parallel cube, i.e., to get the boundary of $C \oplus B_1$, where C is an axis-parallel cube and B_1 is a ball of radius 1, and $\oplus$ denotes the Minkowski sum.

We say that a surface $\mathcal{S}$ is a *patchwork* if it is a closed differentiable surface that is the union of some patches of the above types.

Let $\mathcal{S}$ be a patchwork, and let $\mathbf{n} : \mathcal{S} \rightarrow \mathbb{S}^2$ denote its normal bundle. For a fixed $\delta \in (-1/2, 1/2)$ we define the *offset surface of $\mathcal{S}$ at distance δ* , denoted by $\mathcal{S}(\delta)$, as $\mathcal{S}(\delta) := \{p + \delta \mathbf{n}(p) \mid p \in \mathcal{S}\}$. Because of the restriction of $|\delta| < 1/2$ we have that $\mathcal{S}(\delta)$ is also a differentiable surface, and in fact it consists of patches of the following types, which we will call δ -patches:

Let o denote the center of B_1 and choose q_3 in the extension of q_1q towards q_1 , such that $|q_3q| = \frac{1}{2}$. Then in the right triangle $\triangle oq_1q_3$ we have: $|pq|^2 + (\frac{1}{2} - |q_1q|)^2 = (\frac{1}{2})^2$, which gives: $|q_1q| = \frac{1}{2} - \sqrt{(\frac{1}{2})^2 - |pq|^2}$. (The other root of the equation would lead to $|q_1q| > 1/2$, but this is not possible as q_1q is shorter than the radius of B_1 .) Thus:

$$|q_1q| = \frac{1}{2} \cdot \left(1 - \sqrt{1 - (2|pq|)^2}\right) \leq \frac{1}{2} \cdot (1 - (1 - (2|pq|)^2)) = 2|pq|^2 \leq 2\varepsilon^2$$

where in the second inequality above, we have used that $\sqrt{1 - x^2} \geq 1 - x^2$, for $x \in [0, 1]$. $\blacktriangleleft$

3.2 Nets, offsets, triangle layers

For the remainder of this section let $\mathcal{S}$ be a 1-patchwork surface, and let $|\delta| < 1/2$. Due to space limit, we provide proofs for the more interesting claims. For the rest of the proofs, refer to the full version of our paper [26]. We start with two technical lemmas that construct nets.

► **Lemma 8.** *For any $\zeta \leq 1/8$, we can compute a $(\zeta, 8\zeta)$ -net of size $O(\text{area}(\mathcal{S})/\zeta^2)$ on $\mathcal{S}$ in $O(\text{area}(\mathcal{S})/\zeta^2)$ time.*

► **Lemma 9.** *Let $\zeta \leq 1/8$ and let $N \subset \mathcal{S}$ be the $(\zeta, 8\zeta)$ -net of $\mathcal{S}$ constructed above. Then, for any $\delta \in (-1/2, 1/2)$, we have that $N_\delta := \{p + \delta \mathbf{n}(p) | p \in N\}$ is a $(\zeta/2, 12\zeta)$ -net of $\mathcal{S}(\delta)$.*

► **Lemma 10.** *Let N_δ be a $(\zeta/2, 12\zeta)$ -net of $\mathcal{S}(\delta)$ for each $|\delta| < 1/2$ constructed above. Then for any fixed δ , any $1 \leq t < 1/(4\zeta)$ and any $p \in N_\delta$ we have that the ball $B(p, t\zeta)$ contains at most $O(t^2)$ points of N_δ . Moreover, for any t we can partition N_δ into $O(t^2)$ point sets $N_{\delta,i}$ such that the pairwise distance of points within $N_{\delta,i}$ is at least $t\zeta$.*

Proof. First note that $t \cdot \zeta < 1/4$. Moreover, any cylinder patch of $\mathcal{S}(\delta)$, has distance at least $1/2$ from any square patch of $\mathcal{S}(\delta)$. Therefore, $B(p, t \cdot \zeta)$ cannot intersect both a square and a cylinder patch at the same time. Recall that any two holes on a square patch have pairwise distance at least 6 and are also at distance at least 2 from the boundary of the square patch. Therefore, $B(p, t \cdot \zeta)$ cannot intersect both a quarter-cylinder patch and a joint or a cylinder patch at the same time. The highest number of distinct surface patches $B(p, t \cdot \zeta)$ can intersect is 7, which may happen when it intersects a spherical triangle patch and the three quarter-cylinder patches around it and the three square patches around it. Since it can intersect at most a constant number of distinct patches, it suffices to show that it can contain at most $O(t^2)$ points of N_δ from each different type of surface patch.

To that end, let $\mathcal{P}$ denote a patch which is not a square patch and not a spherical triangle patch. For the remaining types of patches, the constructed net is taken as a point set of the form $I(A, B)$, where elements in A and B are quarter-circular arcs, line segments or circles. In all cases, we have argued already that any two elements in A have distance at least $\zeta/2$ (since now N_δ is a $(\zeta/2, 12\zeta)$ -net) and any two elements in B have also distance at least $\zeta/2$. Therefore, from any point $p \in \mathcal{P}$ and within distance $t \cdot \zeta$, we can reach at most $2t$ elements of A and $2t$ elements of B . Therefore we can reach at most $4t^2$ net points of $\mathcal{P}$. On a spherical triangle patch, the arcs of A again have pairwise distance at least $\zeta/2$, so we can reach at most $O(t)$ arcs, and on each arc we can reach at most $O(t)$ points as on each quarter-circular arc the distance between consecutive points is at least $\zeta/2$.

For a square patch, the process of construction was slightly different due to the existence of the holes. Using the same argumentation as above, we can deduce that $B(p, t \cdot \zeta)$ can contain at most $4t^2$ net points from the initially placed grid. Moreover, $B(p, t \cdot \zeta)$ can intersect

at most one hole. Note that a single hole contains $O(1/\zeta)$ net points and since $t < 1/(2\zeta)$, it contains at most $O(t)$ net points. Therefore, $B(p, t \cdot \zeta)$ can contain at most $5t^2$ net points of a square patch.

Since we argued that $B(p, t \cdot \zeta)$ can intersect at most 7 distinct surface patches, we have that it can contain at most $7 \cdot 5t^2 = 35t^2$ points of N_δ .

For the second part of the lemma create $b = \lfloor 35t^2 + 1 \rfloor$ partition classes $N_{\delta,1}, N_{\delta,2}, \dots, N_{\delta,b}$ and greedily place points in these classes: we place each point $p \in N_\delta$ in the available class of smallest index (that is, a class of minimum index not already containing any point q such that $|pq| < t \cdot \zeta$). Note that in this way, there will always be a class which is available, since for any point there are at most $35t^2$ points that cannot be in the same class with it. ◀

► **Lemma 11.** *There exists a constant ν such that for any $\zeta < 1/(2\nu)$ and $-1/2 < \delta < 1/2 - \nu\zeta^2$ there is collection $\mathcal{T}$ of $O(\text{area}(\mathcal{S})/\zeta^2)$ congruent equilateral triangles of side length $\Theta(\zeta)$ that satisfy the following properties:*

- (i) *each triangle in $\mathcal{T}$ is located between $\mathcal{S}(\delta)$ and $\mathcal{S}(\delta + \nu\zeta^2)$*
- (ii) *the triangles in $\mathcal{T}$ are pairwise disjoint, and*
- (iii) *the geodesic distance from any point of $\mathcal{S}(\delta)$ to any point of $\mathcal{S}(\delta + \nu\zeta^2)$ is at least 4ζ .*


Proof. By applying Lemma 8 and Lemma 9 for we obtain a $(\zeta/2, 12\zeta)$ -net N_δ of $\mathcal{S}(\delta)$ of size $O(\text{area}(\mathcal{S})/\zeta^2)$. We set $t = 48(1 + \sqrt{3})$, and choose ν large enough such that $\zeta < 1/(4t)$. Then, by Lemma 10, we can partition N_δ into $b = \lfloor 35t^2 + 1 \rfloor$ classes $(N_{\delta,i})_{i=1}^b$, where points within each class $N_{\delta,i}$ have pairwise distance at least $t \cdot \zeta = 48(1 + \sqrt{3})\zeta$. We will now place each class $N_{\delta,i}$ on a different offset $\mathcal{S}(\delta_i)$, defined as follows. We start by setting $\nu = 4b + 1$ and define $\delta_i := \delta + 4i \cdot \zeta^2$, for $i = 1, 2, \dots, b$. From now on, we consider the points of each class $N_{\delta,i}$ on the surface $\mathcal{S}(\delta_i)$ defined as above.

The above definition ensures $\delta_i \leq \delta + 4b\zeta^2 < \delta + \nu\zeta^2 < 1/2$. Thus, the offsets $\mathcal{S}(\delta_i)$ are well-defined. Furthermore, we also have the following:

$$\text{For } 1 \leq i \leq b, \text{ we have that } \delta_i - 2\zeta^2 < \delta_i < \delta_i + 2\zeta^2 < \delta_{i+1}. \quad (1)$$

For $1 \leq i \leq b$ and $p \in N_{\delta,i}$, let T_p denote the equilateral triangle of side length $48\sqrt{3}\zeta$ centered at p , which lies on the tangent plane H_p at p . Let $\mathcal{T}_i = \{T_p : p \in N_{\delta,i}\}$ and $\mathcal{T} = \bigcup_{i=1}^b \mathcal{T}_i$. We now argue that $\mathcal{T}$ satisfies the three properties of the Lemma:

- (i) This holds trivially since $\delta + 4\zeta^2 \leq \delta_i \leq \delta + 4b\zeta^2 < \delta + \nu\zeta^2$.
- (ii) Lemma 7 together with (1), ensure that any two triangles in $\mathcal{T}$ that belong to different offsets are disjoint. Next, we argue that any two triangles $T_p, T_q \in \mathcal{T}_i$ are disjoint. From Lemma 10 we know that $|pq| \geq 48(1 + \sqrt{3})\zeta > 96\zeta$. Since each of our triangles has circumradius 48ζ , it follows that T_p, T_q cannot intersect.
- (iii) Refer to Section 3.2 for an illustration. Let $p \in \mathcal{S}(\delta)$ and $q \in \mathcal{S}(\delta + \nu\zeta^2)$ and consider a shortest geodesic path $\pi(p, q)$. We want to show that $\text{dist}_{\mathcal{T}}(p, q) \geq c \cdot \zeta$ for some constant c . Let u be a point in $\pi(p, q) \cap \mathcal{S}(\delta_1)$ and consider the segment s between the points $u - \mathbf{n}(u)$ and $u + \mathbf{n}(u)$. We consider the following two cases:
 - *Case I: There exists $v \in \pi(p, q)$, such that $\text{dist}_{\mathbb{R}^3}(v, s) \geq 4\zeta$.*
Then π has the desired length.
 - *Case II: $\pi(p, q)$ is contained in the cylinder with axis s and radius 4ζ*
Let $\mathcal{C}(ab, x)$ denote the cylinder of axis ab and radius x . We will show that there exists $T \in \mathcal{T}$ which cuts the cylinder $\mathcal{C}(s, 4\zeta)$ in such a way that the top and bottom disk of the cylinder are completely contained in different sides of the cut. Since $\pi(p, q)$ connects these disks and it is disjoint from all triangles $T \in \mathcal{T}$, this will give a contradiction.



■ **Figure 5** Illustration for the proof of Lemma 11 (side view). The geodesic $\pi(p, q)$ is shown in grey. Note that the cylinder around s is cut by one of the red triangles, therefore $\pi(p, q)$ needs to exit the cylinder.

By the properties of N_δ , we know that there exists a surface $\mathcal{S}(\delta_i)$, such that $s \cap \mathcal{S}(\delta_i)$ is within distance 12ζ from a point $w \in N_{\delta_i}$. We will argue that T_w is the required triangle. Towards that, let $s_i = s \cap \mathcal{S}(\delta_i)$ and recall that T_w has inradius 24ζ and denote by D the incircle of T_w . Observe that $\mathcal{C}(s_{i-1}s_{i+1}, 4\zeta) \subset B(w, 24\zeta)$. By Lemma 7, we know that T_w is between $\mathcal{S}(\delta_{i-1})$ and $\mathcal{S}(\delta_{i+1})$. Since D splits $B(w, 24\zeta)$ in two pieces where s_{i-1} and s_{i+1} are in different parts, we get that D has to cut $\mathcal{C}(s_{i-1}s_{i+1}, 4\zeta)$. We conclude that indeed T_w cuts $\mathcal{C}(s, 4\zeta)$ in the described manner, which is a contradiction. ◀


By iterating the construction of Lemma 11 we get our curved wall separation theorem.

► **Theorem 12** (Polynomial separation). *Let $\mathcal{S}$ be a 1-patchwork surface. Then for any $\sigma > 1$ there is a collection of $O(\sigma^4 \text{area}(\mathcal{S}))$ pairwise disjoint congruent equilateral triangle obstacles between $\mathcal{S}(-1/2)$ and $\mathcal{S}(1/2)$, such that the geodesic distance from any point of $\mathcal{S}(-1/2)$ to any point of $\mathcal{S}(1/2)$ is at least σ .*

3.3 The construction and the proof of Theorem 3

Recall that (X, dist_X) is the metric space whose distances we want to realize, with $|X| = n$ and spread $\Phi = \max_{a,b,c,d \in X} \left(\frac{\text{dist}_X(a,b)}{\text{dist}_X(c,d)} \right)$, and $\varepsilon > 0$ denotes an accuracy parameter. We can assume that the minimum distance of any two points in X is $41n^3/\varepsilon$, as in general we can scale the whole construction so that the minimum geodesic distance exactly matches the minimum distance in X . We construct a 1-patchwork surface $\mathcal{S}$ together with an injection $f : X \rightarrow \mathbb{R}^3$ as follows.

For each point $x_i \in X$, $1 \leq i \leq n$, let $f(x_i) = (40n^2 \cdot i, 0, -10n^2)$. Let $\mathcal{C}_i$ denote the cube centered at $f(x_i)$ of side-length $20n^2 - 2$. We define the rounded cube $\mathcal{R}_i = \mathcal{C}_i \oplus B_1$, where B_1 denotes the ball of radius 1. Note that each rounded cube $\mathcal{R}_i$ has size $20n^2$, where we define the size of a rounded cube to be the side-length of its bounding box. Then the top square patch of each $\mathcal{R}_i$ lies on the plane $z = 0$, has side length $20n^2 - 2$ and is centered at the point $(40n^2 \cdot i, 0, 0)$. For each pair (i, j) where $i < j$, we model the distance between x_i and x_j as follows. We start by placing two holes $H_{i,j}, H_{j,i}$ on the top square patch of $\mathcal{R}_i$ and $\mathcal{R}_j$ respectively, such that:



■ **Figure 6** A tube consisting of three cylinders and a turn that connects the rounded cubes $\mathcal{R}_i$ and $\mathcal{R}_j$. We will set $\ell_{\text{vert}}(i, j)$ to ensure that the length of the tube corresponds to the distance $\text{dist}_X(x_i, x_j)$.

- $H_{i,j}$ has radius 2 and center $(40n^2 \cdot i, -10n^2 + 10(n-1) \cdot i + 10j - 10, 0)$, and
- $H_{j,i}$ has radius 2 and center $(40n^2 \cdot j, -10n^2 + 10(n-1) \cdot i + 10j - 10, 0)$.

Observe that the centers of any two holes $H_{i,j}, H_{j,i}$ defined this way have the same y -coordinates. Moreover, the y -coordinates of the centers of all holes are in the range $[-10n^2 + 10n, -10n]$. Since the y -coordinates of the top square patches are in the range $[-10n^2 + 1, 10n^2 - 1]$, this ensures that:

1. holes are well-defined,
2. each hole has distance at least 2 from the boundary of its square patch, and
3. for each $1 \leq i, j, k \leq n$, where i, j, k are pairwise different, the centers of $H_{i,j}$ and $H_{i,k}$ have distance at least 10. As a result, $H_{i,j}$ and $H_{i,k}$ have distance at least 6, as required in the definition of a square patch.

It will be useful for the rest of the description, to define a *cylinder-joint patch of length ℓ* as a cylinder patch of length ℓ , with two joint patches attached to its bases. Note that a cylinder-joint patch can be used to connect two holes that are opposite to each other.

Next, for each $1 \leq i < j \leq n$, we connect $H_{i,j}$ and $H_{j,i}$ as follows. We attach to both $H_{i,j}$ and $H_{j,i}$ a cylinder-joint patch of length $\ell_{\text{vert}}(i, j)$, which we denote by $\mathcal{K}_{i,j}$ and $\mathcal{K}_{j,i}$ respectively. The length $\ell_{\text{vert}}(i, j)$ will depend on ε and its exact value will be set in the proof of Theorem 3. We then attach $\mathcal{K}_{i,j}$ to a rounded cube $\mathcal{C}_{i,j}$ of size 8, through a circular hole of radius 2 centered at the center of the bottom square patch of $\mathcal{C}_{i,j}$. Note that this hole has distance at least 2 from the boundary of the square patch. Similarly, we attach $\mathcal{K}_{j,i}$ to a rounded cube $\mathcal{C}_{j,i}$ of size 8, through a circular hole of radius 2 centered at the center of the bottom square patch of the cube. We finally connect $\mathcal{C}_{i,j}$ and $\mathcal{C}_{j,i}$ through a horizontal cylinder-joint patch of length $\ell_{\text{hor}}(i, j) = 40n^2 \cdot (j - i)$, denoted by $\mathcal{J}_{i,j}$. This concludes the construction of $\mathcal{S}$. Refer to Figure 6 for an illustration of the connection between $f(x_i), f(x_j)$. After constructing $\mathcal{S}$ we apply Theorem 12 for $\sigma := \max_{x_i, x_j} \text{dist}_X(x_i, x_j) = O(\Phi \cdot n^3/\varepsilon)$, to achieve a separation of $\max_{x_i, x_j} \text{dist}_X(x_i, x_j)$ between $\mathcal{S}(-1/2)$ and $\mathcal{S}(1/2)$. We let $\mathcal{T}$ denote the resulting set of triangle obstacles.

Next, we sketch the proof of Theorem 3. Refer to the full version [26] for a detailed proof.

Proof sketch. We need to upper and lower bound the length of a geodesic from $f(x_i)$ to $f(x_j)$. The choice of σ ensures that we only need to consider geodesics within $\mathcal{S}(1/2)$. We set $\ell_{\text{vert}}(i, j)$ so that any geodesic that passes through the three-cylinder-tube connecting

$\mathcal{R}_i$ and $\mathcal{R}_j$ has length at least $\text{dist}_X(x_i, x_j)$. Now a geodesic passing through multiple such tubes will have a length that is at least the sum of the corresponding distances, thus by the triangle inequality in X any such path will also have length at least $\text{dist}_X(x_i, x_j)$.

To upper bound the geodesic distance of $f(x_i)$ to $f(x_j)$, we observe that there is a direct path in $\mathcal{S}(-1/2)$ connecting $f(x_i)$ and $f(x_j)$ through the designated three-cylinder tube. This will contain some additional short segments compared to our lower bound, namely, the segments connecting the endpoints of the geodesic ($f(x_i)$ and $f(x_j)$) to the entrance and exit of the tube. The total length of these segments is shown to be less than $\varepsilon \text{dist}_X(x_i, x_j)$. ◀

We note the following observation from within the proof of Theorem 3 for later use.

► **Observation 13.** *There is a path connecting $f(x_i)$ and $f(x_j)$ in the inside of $\mathcal{S}(-1/2)$ that has length at most $(1 + \varepsilon) \cdot \text{dist}_X(x_i, x_j)$.*

► **Remark 14 (Exact realizations).** It is particularly challenging to realize a metric *exactly* with convex obstacles. While our construction can be modified to ensure the exact length of the shortest path through the corresponding tube has the same exact length as the original distance in (X, dist_X) , this is not sufficient.

The problem can be seen by considering the metric induced by a star on 4 vertices. If a, b, c are the leaves and x is the center of the star, then the shortest $f(a) \rightarrow f(b)$ path will be realized by going through the $f(a) \rightarrow f(x)$ tube and the $f(x) \rightarrow f(b)$ tube, foregoing the small detour to $f(x)$ within its rounded cube. This path will be strictly shorter than $\text{dist}_{\mathcal{T}}(f(a), f(x)) + \text{dist}_{\mathcal{T}}(f(x), f(b))$. While a realization exists for star graph metrics using relatively open triangles, this does not generalize. Can we for example realize (X, dist_X) exactly with disjoint convex obstacles assuming the image of dist_X is $\{1, 2\}$?

4 Realizing metrics with fat obstacles

We tweak our construction to prove the following theorem and answer Question 1 for convex fat non-disjoint obstacles and for fat, disjoint but non-convex obstacles.

► **Theorem 15.** *Let $\varepsilon \in (0, 1)$ and let (X, dist_X) be a metric space of size $|X| = n$ and spread $\Phi := \max_{a,b,c,d \in X} (\frac{\text{dist}_X(a,b)}{\text{dist}_X(c,d)})$. Then there exists a collection $\mathcal{T}$ of obstacles for each of the following restrictions:*

- (i) $\mathcal{T}$ consists of $O(n^{17}\Phi^5/\varepsilon^5)$ congruent regular simplices (that are not necessarily disjoint).
- (ii) $\mathcal{T}$ consists of $O(n^5\Phi/\varepsilon)$ disjoint, similarly-sized α -fat objects that are homeomorphic to a ball. The objects have total complexity $O(n^{17}\Phi^5/\varepsilon^5)$, where $\alpha > 0$ is a constant.


In both cases, there is an injection $f : X \rightarrow \mathbb{R}^3$ such that

$$\text{dist}_X(a, b) \leq \text{dist}_{\mathcal{T}}(f(a), f(b)) \leq (1 + \varepsilon) \cdot \text{dist}_X(a, b).$$

for all $a, b \in X$. The set of obstacles can be constructed in $\text{poly}(n, \Phi, 1/\varepsilon)$ time.

Proof. (i) We modify the set of triangular obstacles in our basic construction as follows. Recall from the proof of Lemma 11 that each triangle $T \in \mathcal{T}$ is tangent at the center of T to some $\mathcal{S}(\delta)$ at a point $p \in \mathcal{S}(\delta)$, for some $-1/2 < \delta < 1/2$. We turn T into a regular tetrahedron T' where one face of T' is T , and the vertex of T' outside T is placed along the normal ray from p , that is, the tetrahedron will point to the outside of $\mathcal{S}(\delta)$. The obstacles in the new collection obtained this way, denoted by $\mathcal{T}'$, are convex and $1/3$ -fat.

We can now prove Theorem 3 for $\mathcal{T}'$ instead of $\mathcal{T}$. For the lower bound, it suffices to observe that $\mathcal{F}_{\mathcal{T}'} \subset \mathcal{F}_{\mathcal{T}}$. This implies that the set of geodesics in $\mathcal{F}_{\mathcal{T}'}$ is a subset of the geodesics in $\mathcal{F}_{\mathcal{T}}$, and thus for any two points $x_i, x_j \in X$ we have $\text{dist}_{\mathcal{T}'}(f(x_i), f(x_j)) \geq$



■ **Figure 7** (i) Side view of the region $\mathcal{U}$ (shaded in grey). (ii) The region $U_p = C_p \cap \mathcal{U}$ (in red), together with the ball $B_p \subset U_p$ (in blue), indicating that U_p is fat.

$\text{dist}_{\mathcal{T}}(f(x_i), f(x_j)) \geq \text{dist}_X(x_i, x_j)$. On the other hand, the inside of $\mathcal{S}(-1/2)$ is in $\mathcal{F}_{\mathcal{T}'}$, due to the way we constructed the tetrahedra (pointing “out” of the offset surface they are tangent to). Observation 13 now implies that $\text{dist}_{\mathcal{T}'}(f(x_i), f(x_j)) \leq (1 + \varepsilon)\text{dist}_X(x_i, x_j)$.


(ii) Let $\mathcal{T}$ be the triangle collection constructed by Section 3.3, and let $\mathcal{U}$ be the union of all offsets of $\mathcal{S}$, that is, $\mathcal{U} := \bigcup_{\delta \in [-1/2, 1/2]} \mathcal{S}(\delta)$. Refer to Figure 7 for an illustration. Let N be a $(1/8, 1)$ -net of $\mathcal{S}$ obtained from Lemma 9 with $\zeta = 1/8$. Consequently, $|N| = \text{area}(\mathcal{S}) = O(n^5 \Phi/\varepsilon)$. Compute the *Voronoi diagram* of N in $\mathbb{R}^3$, that is, a partition of $\mathbb{R}^3$ into cells according to the closest point of N . The Voronoi cells can be computed in $\text{poly}(|N|)$ time [18, Ch. 27]. For the cell C_p of $p \in N$ let $U_p := C_p \cap \mathcal{U}$; note that each U_p can be represented explicitly in $\text{poly}(|N|) = \text{poly}(|\mathcal{T}|)$ space. Observe that U_p has diameter at most $\text{diam}(U_p) \leq 3/2$: indeed, any point in $\mathcal{U}$ is within distance at most $1/2$ from $\mathcal{S}$ (along the normal direction or its inverse), and any point of $\mathcal{S}$ is within distance at most 1 from some point of N by the net property of N .

We claim that each set U_p contains a ball B_p of radius $1/16$. Clearly the Voronoi cell C_p contains this ball by the net property; it remains to show that B_p is contained in $\mathcal{U}$. Consider the line ℓ normal to $\mathcal{S}$ at p , and let $p(\delta)$ be the point of this line on $\mathcal{S}(\delta)$ where $\delta \in [-1/16, 1/16]$. Notice that a point $q \in B_p$ the nearest point of ℓ is some point $p(\delta)$, and the tangent plane H_δ of $\mathcal{S}(\delta)$ at $p(\delta)$ contains q . Since $|q, p(\delta)| \leq 1/16$, we have by Lemma 7 that q is between $\mathcal{S}(\delta - 2/16^2)$ and $\mathcal{S}(\delta + 2/16^2)$. Since $\delta < 1/16$, this implies that $q \in \mathcal{U}$, and concludes the proof that $B_p \subset U_p$.

Consider now the obstacle set $\mathcal{T}$, and notice that its triangles have positive pairwise distance from each other and also from $\mathcal{S}(-1/2)$ and from $\mathcal{S}(1/2)$. Let $\mu > 0$ be some small number such that the Euclidean μ -neighborhoods of the triangles of $\mathcal{T}$ remain pairwise disjoint, and they remain subsets of $\mathcal{U}$. For a triangle $T \in \mathcal{T}$ let $x(T)$ denote the center of T and let T^μ denote the set of points q in $\mathbb{R}^3$ such that $\text{dist}_{\mathbb{R}^3}(q, T) < \mu$ (i.e., T^μ is an open set). Here $\text{dist}_{\mathbb{R}^3}(q, T)$ denotes the minimum Euclidean distance between q and a point of T .

Let $U_p^{-\mu}$ denote the set of points q with $\text{dist}_{\mathbb{R}^3}(q, \mathbb{R}^3 \setminus U_p) \geq \mu$, that is, $U_p^{-\mu}$ is a closed subset of U_p where a small neighborhood of the boundary of U_p has been removed. Consequently, the sets $U_p^{-\mu}$ are pairwise disjoint.

We group $\mathcal{T}$ according to the location of their triangle centers. Let $\mathcal{T}_p$ denote the set of triangles $T \in \mathcal{T}$ where $x(T) \in U_p$; if $x(T)$ is on the shared boundary of several sets $(U_p)_{p \in N}$, then we assign T to one of the corresponding classes arbitrarily. As a result $\{\mathcal{T}_p\}_{p \in N}$ is a partition of $\mathcal{T}$. For each point $p \in N$ we define the following obstacle W_p by adding all triangles of $\mathcal{T}_p$ to $U_p^{-\mu}$, but subtracting T^μ for all $T \in \mathcal{T} \setminus \mathcal{T}_p$:



■ **Figure 8** (i) By shrinking U_p , we obtain the region $U_p^{-\mu}$. Square points correspond to centers of triangles in $\mathcal{T}$. Four triangles are drawn, again in side view. For illustration purposes, various objects have not been drawn to scale. (ii) The final obstacle W_p , obtained by attaching to $U_p^{-\mu}$ the two red triangles, and removing a small neighborhood of the other two.

$$W_p := \left(U_p^{-\mu} \cup \left(\bigcup_{T \in \mathcal{T}_p} T \right) \right) \setminus \left(\bigcup_{T \in \mathcal{T} \setminus \mathcal{T}_p} T^\mu \right).$$

Refer to Figure 8 for an illustration of an obstacle W_p . The sets W_p are computed in polynomial time, and they are closed sets that are pairwise disjoint. Moreover, each set W_p has diameter at most $\text{diam}(W_p) \leq \text{diam}(U_p) + s_T < 2$ where $s_T < 1/100$ is the side length of the triangles in T . Recall that each set U_p contains a ball of radius $1/16$ centered at p . It follows that W_p contains the ball of radius $1/16 - \mu - s_T > 1/32$ centered at p . We conclude that the sets W_p are $1/64$ -fat.

The proof that the obstacle set $\{W_p\}_{p \in N}$ approximately realizes (X, dist_X) can be argued as in part (i). The number of obstacles is $|N| = O(n^5 \Phi/\varepsilon)$. ◀

We note that both the size of our obstacle collection and the construction time could be further improved; here, however, we did not attempt to optimize for these aspects.

5 Algorithms and lower bounds for TSP WITH OBSTACLES

This section is devoted to proving Corollary 4. In order to give an algorithm, we will first need to (approximately) compute the distances between our input points. Let $\mathcal{T}$ denote a collection of pairwise-disjoint polyhedral obstacles in $\mathbb{R}^3$ of total complexity m . Har-Peled [21] showed how to construct an ε -approximate shortest path map from any source point $s \in \mathcal{F}_{\mathcal{T}}$, in time roughly $O(m^4/\varepsilon^6)$. From his result we can deduce the following:

► **Theorem 16** (Har-Peled [21]). *Let $\mathcal{T}$ be a collection of pairwise-disjoint polyhedral obstacles in $\mathbb{R}^3$ of total complexity m and let $S \subset \mathcal{F}_{\mathcal{T}}$, be a set of $|S| = n$ points. Then, for any $0 < \varepsilon < 1$, we can compute a $(1 + \varepsilon)$ -approximation for the all-pairs-shortest-paths problem on S in time:*

$$O\left(n \frac{m^4}{\varepsilon^2} \left(\frac{\beta(m)}{\varepsilon^4} \log \frac{m}{\varepsilon} + \log(m\rho) \log(m \log \rho) \right) \log \frac{1}{\varepsilon}\right)$$

where ρ is the ratio of the longest edge in $\mathcal{T}$ with the distance of the closest pair of points in S , and $\beta(m) = \alpha(m)^{O(\alpha(m))^{O(1)}}$, and $\alpha(m)$ is the inverse of the Ackermann function.

Note that Har-Peled's result is based on constructing a graph whose vertices correspond to points and whose edges correspond to pairwise visible point pairs, where edge lengths are high-precision estimates of the distance of the visible point pair. Har-Peled's result computes the true shortest paths in this graph. Consequently, the resulting approximate distance function $\text{dist}_{\text{approx}}$ on S forms a metric space. One can also show using arguments similar to Lemma 6 that for any $\varepsilon \in (0, 1)$ the doubling dimension of $(S, \text{dist}_{\text{approx}})$ is at most constant times the doubling dimension of $(S, \text{dist}_{\mathcal{T}})$.

We now mention the result by Banerjee, Bartal, Gottlieb and Hovav, on TSP in doubling spaces [6].

► **Theorem 17** (Corollary 35 in [6]). *Let (X, dist_X) be a metric space with $|X| = n$ points and of bounded doubling dimension δ . A $(1 + \varepsilon)$ -approximation to the optimal tour of X can be computed by a randomized algorithm in $2^{(\delta/\varepsilon)^{O(\delta)}} n + (1/\varepsilon)^{O(\delta)} n \log n$ time.*

By combining Theorem 17 and Theorem 16, we will obtain a PTAS for the case $\mathfrak{P} = \text{convex} \wedge \text{pairwise disjoint} \wedge \text{fat}$. For the other cases, the APX-hardness is based on the inapproximability result by Karpinski, Lampis and Schmied on METRIC TSP [25]. They showed an inapproximability bound of $123/122$. By inspection of the construction given in [25], we have the following crucial observation:

► **Observation 18.** *It is NP-hard to approximate METRIC TSP within a ratio of $123/122$, even when the underlying metric space has polynomial spread.*

Proof. This follows by inspection of the construction by Karpinski, Lampis and Schmied (Section 5.1 in [25]), which is based on a the shortest path metric of a connected edge-weighted graph. The weights of all edges used are small constants. Namely, the smallest edge has weight 0.5, and the longest edge has length 2. Therefore, the graph has diameter at most $2(n - 1)$, and the construction has a spread at most $\frac{2(n-1)}{0.5} = 4(n - 1) = O(n)$. ◀

Now we have all the necessary elements to prove Corollary 4.

► **Corollary 4.** *Let $\mathcal{T}$ be a collection of obstacles in $\mathbb{R}^3$ with property $\mathfrak{P}$ of total complexity $m = \text{poly}(|X|)$. Then TSP WITH OBSTACLES has a PTAS if $\mathfrak{P} = \text{"convex and pairwise disjoint and fat"}$, and it is APX-hard for all other $\mathfrak{P}$.*

Proof of Corollary 4. Let $S \subset \mathcal{F}_{\mathcal{T}}$ denote a set of n points, whose optimal tour we want to compute. We have the following cases:

■ **The obstacles in $\mathcal{T}$ are convex, α -fat and pairwise disjoint.**

We show that then TSP WITH OBSTACLES admits a PTAS. Towards that, let $\varepsilon \in (0, 1)$. By Theorem 16, we can approximate $\text{dist}_{\mathcal{T}}$ up to a factor of $1 + \varepsilon/3$. Let $\text{dist}_{\text{approx}}$ denote the approximate metric. The doubling dimension of $(S, \text{dist}_{\text{approx}})$ is at most a constant factor larger than that of $(S, \text{dist}_{\mathcal{T}})$. Therefore, by Lemma 6, it is at most $O(1)$. By Theorem 17, we can get a $(1 + \varepsilon/3)$ -approximation for TSP WITH OBSTACLES on $(S, \text{dist}_{\text{approx}})$, with running time $2^{1/\varepsilon^{O(1)}} n + 1/\varepsilon^{O(1)} n \log n$. Since $(1 + \varepsilon/3)^2 < 1 + \varepsilon$, the result is a $(1 + \varepsilon)$ -approximation for TSP WITH OBSTACLES.

■ **The obstacles in $\mathcal{T}$ are convex and pairwise disjoint but not fat, or fat and convex but non-disjoint, or fat and disjoint but non-convex.**

We will show that in this case the problem is APX-hard, with an inapproximability bound of $\frac{123}{122} - \varepsilon$, for any $\varepsilon > 0$. For a set $S \subset \mathcal{F}_{\mathcal{T}}$, let $\text{OPT}_{\mathcal{T}}(S)$ denote the length of an optimal tour through S . Assume that we could, in polynomial time, compute a tour of length $T_{\text{approx}}(S)$ such that $T_{\text{approx}}(S) \leq (\frac{123}{122} - \varepsilon) \text{OPT}_{\mathcal{T}}(S)$, for a fixed $\varepsilon > 0$.

Let (X, dist_X) denote an instance of METRIC TSP, where $|X| = n$ and X has spread $\text{poly}(n)$. From Observation 18, we have a $123/122$ inapproximability bound for computing the optimal tour on X , denoted by $\text{OPT}(X)$. From Theorem 3 or Theorem 15, we compute a set of obstacles $\mathcal{T}$ in $\mathbb{R}^3$ and an injection $f : X \rightarrow \mathbb{R}^3$, such that:

$$\text{OPT}_{\mathcal{T}}(f(X)) \leq (1 + \varepsilon') \text{OPT}(X),$$

where we set ε' to satisfy $\frac{123}{122(1+\varepsilon')} = \frac{123}{122} - \varepsilon$.

By our assumption, we can compute a value $T_{\text{approx}}(f(X))$ such that:

$$T_{\text{approx}}(f(X)) \leq \left(\frac{123}{122} - \varepsilon \right) \text{OPT}_{\mathcal{T}}(f(X))$$

Therefore, $T_{\text{approx}}(f(X)) \leq \frac{123}{122} \text{OPT}(X)$, which contradicts the APX-hardness of METRIC TSP. $\blacktriangleleft$

6 Conclusion

In this paper we have shown that all metric spaces are approximately realizable in $\mathbb{R}^3$ using convex disjoint obstacles, or using convex fat but non-disjoint obstacles, or using fat disjoint obstacles that are non-convex. Our convex construction is easily generalizable to several other obstacle types, however, it always requires objects whose minimum bounding box has two long edges (polynomially long in $n, 1/\varepsilon$ and the spread) and one short edge of length 1.

Some notable obstacle types where the possible realizable metrics remain unexplored include disjoint axis-parallel box obstacles and disjoint $1 \times 1 \times n$ boxes of arbitrary orientation. We suspect that obstacles of the former type cannot realize all metric spaces, and obstacles of the latter type can only realize spaces of small doubling dimension. It also remains open whether exact realization of any metric space is possible with convex obstacles, even if all distances in the metric space are either 1 or 2.

Finally, can we generalize the planar studies of distances among weighted regions [9, 17, 30] to $\mathbb{R}^3$? What metrics can be realized if we are allowed to pass through the objects of $\mathcal{T}$, but we need to pay some custom penalty (weight) for doing so?

References

- 1 Mohammad Ali Abam, Marjan Adeli, Hamid Homapour, and Pooya Zafar Asadollahpoor. Geometric Spanners for Points Inside a Polygonal Domain. In Lars Arge and János Pach, editors, *31st International Symposium on Computational Geometry (SoCG 2015)*, volume 34 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 186–197, Dagstuhl, Germany, 2015. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.SOCG.2015.186.
- 2 Mohammad Ali Abam and Mohammad Javad Rezaei Seraji. Geodesic spanners for points in $\mathbb{R}^3$ amid axis-parallel boxes. *Information Processing Letters*, 166:106063, 2021. doi:10.1016/j.ipl.2020.106063.
- 3 Pankaj K. Agarwal, Dan Halperin, Micha Sharir, and Alex Steiger. Near-optimal min-sum motion planning for two square robots in a polygonal environment. In *Proceedings of the 2024 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 4942–4962, 2024. doi:10.1137/1.9781611977912.176.
- 4 Henk Alkema, Mark de Berg, Morteza Monemizadeh, and Leonidas Theocharous. TSP in a Simple Polygon. In Shiri Chechik, Gonzalo Navarro, Eva Rotenberg, and Grzegorz Herman, editors, *30th Annual European Symposium on Algorithms (ESA 2022)*, volume 244 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 5:1–5:14, Dagstuhl, Germany, 2022. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.ESA.2022.5.

- 5 Takao Asano, Tetsuo Asano, Leonidas Guibas, John Hershberger, and Hiroshi Imai. Visibility of disjoint polygons. *Algorithmica*, 1(1):49–63, November 1986. doi:10.1007/BF01840436.
- 6 Sandip Banerjee, Yair Bartal, Lee-Ad Gottlieb, and Alon Hovav. Novel Properties of Hierarchical Probabilistic Partitions and Their Algorithmic Applications. In *2024 IEEE 65th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 1724–1767, Los Alamitos, CA, USA, 2024. IEEE Computer Society. doi:10.1109/FOCS61266.2024.00107.
- 7 Luis Barba, Michael Hoffmann, Matias Korman, and Alexander Pilz. Convex Hulls in Polygonal Domains. In David Eppstein, editor, *16th Scandinavian Symposium and Workshops on Algorithm Theory (SWAT 2018)*, volume 101 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 8:1–8:13, Dagstuhl, Germany, 2018. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.SWAT.2018.8.
- 8 Yair Bartal, Lee-Ad Gottlieb, and Robert Krauthgamer. The traveling salesman problem: Low-dimensionality implies a polynomial time approximation scheme. *SIAM J. Comput.*, 45(4):1563–1581, 2016. doi:10.1137/130913328.
- 9 Prosenjit Bose, Guillermo Esteban, and Anil Maheshwari. Weighted shortest path in equilateral triangular meshes. In Yeganeh Bahoo and Konstantinos Georgiou, editors, *Proceedings of the 34th Canadian Conference on Computational Geometry, CCCG 2022, Toronto Metropolitan University, Toronto, Ontario, Canada, August 25-27, 2022*, pages 60–67, 2022.
- 10 John F. Canny. *The complexity of robot motion planning*. MIT Press, Cambridge, MA, USA, 1988.
- 11 L.Paul Chew, Haggai David, Matthew J. Katz, and Klara Kedem. Walking around fat obstacles. *Information Processing Letters*, 83(3):135–140, 2002. doi:10.1016/S0020-0190(01)00321-0.
- 12 V. Chvátal. A combinatorial theorem in plane geometry. *Journal of Combinatorial Theory, Series B*, 18(1):39–41, February 1975. doi:10.1016/0095-8956(75)90061-1.
- 13 Kenneth L. Clarkson. Approximation algorithms for shortest path motion planning (extended abstract). In Alfred V. Aho, editor, *Proceedings of the 19th Annual ACM Symposium on Theory of Computing, 1987, New York, New York, USA*, pages 56–65. ACM, 1987. doi:10.1145/28395.28402.
- 14 Sarita de Berg, Tillmann Miltzow, and Frank Staals. Towards Space Efficient Two-Point Shortest Path Queries in a Polygonal Domain. In Wolfgang Mulzer and Jeff M. Phillips, editors, *40th International Symposium on Computational Geometry (SoCG 2024)*, volume 293 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 17:1–17:16, Dagstuhl, Germany, 2024. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.SoCG.2024.17.
- 15 Adrian Dumitrescu and Csaba D. Tóth. Maximal distortion of geodesic diameters in polygonal domains. In *Combinatorial Algorithms - 34th International Workshop, IWOCA 2023, Tainan, Taiwan, June 7-10, 2023, Proceedings*, volume 13889 of *Lecture Notes in Computer Science*, pages 197–208. Springer, 2023. doi:10.1007/978-3-031-34347-6_17.
- 16 Sándor P. Fekete, Stephan Friedrichs, Michael Hemmer, Joseph S. B. Mitchell, and Christiane Schmidt. On the chromatic art gallery problem. In *Canadian Conference on Computational Geometry (CCCG)*, 2014. URL: <http://www.cccg.ca/proceedings/2014/papers/paper11.pdf>.
- 17 L. Gewali, A. Meng, J. S. Mitchell, and S. Ntafos. Path planning in 0/1/ weighted regions with applications. In *Proceedings of the Fourth Annual Symposium on Computational Geometry, SCG '88*, pages 266–278, New York, NY, USA, 1988. Association for Computing Machinery. doi:10.1145/73393.73421.
- 18 Jacob E. Goodman and Joseph O'Rourke, editors. *Handbook of discrete and computational geometry*. CRC Press Series on Discrete Mathematics and its Applications. CRC Press, Boca Raton, FL, 1997.
- 19 D. Halperin, L. Kavraki, and K. Solovey. *Robotics*, chapter 51, pages 1343–1376. Chapman & Hall/CRC, 3rd edition, 2018.
- 20 D. Halperin, M. Sharir, and O. Salzman. *Algorithmic Motion Planning*, chapter 50, pages 1311–1342. Chapman & Hall/CRC, 3rd edition, 2018. doi:10.1201/9781315119601.

- 21 Sarel Har-Peled. Constructing approximate shortest path maps in three dimensions. *SIAM J. Comput.*, 28(4):1182–1197, 1999. doi:10.1137/S0097539797325223.
- 22 Juha Heinonen. *Lectures on Analysis on Metric Spaces*, volume 1 of *Universitext*. Springer New York, NY, 1 edition, 2001. doi:10.1007/978-1-4613-0131-8.
- 23 John Hershberger and Subhash Suri. An optimal algorithm for Euclidean shortest paths in the plane. *SIAM Journal on Computing*, 28(6):2215–2256, 1999. doi:10.1137/S0097539795289604.
- 24 J. E. Hopcroft, J. T. Schwartz, and M. Sharir. On the complexity of motion planning for multiple independent objects; PSPACE-hardness of the “warehouseman’s problem”. *International Journal of Robotics Research*, 3(4):76–88, 1984.
- 25 Marek Karpinski, Michael Lampis, and Richard Schmied. New inapproximability bounds for TSP. *Journal of Computer and System Sciences*, 81(8):1665–1677, 2015. doi:10.1016/j.jcss.2015.06.003.
- 26 Sándor Kisfaludi-Bak and Leonidas Theodorarous. Realizing metric spaces with convex obstacles. arXiv:2509.14709, 2025. doi:10.48550/arXiv.2509.14709.
- 27 Anna Lubiw and Anurag Murty Naredla. The Visibility Center of a Simple Polygon. In Petra Mutzel, Rasmus Pagh, and Grzegorz Herman, editors, *29th Annual European Symposium on Algorithms (ESA 2021)*, volume 204 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 65:1–65:14, Dagstuhl, Germany, 2021. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.ESA.2021.65.
- 28 Joseph S. B. Mitchell. Shortest paths among obstacles in the plane. In *Proceedings of the Ninth Annual Symposium on Computational Geometry*, SCG ’93, pages 308–317, New York, NY, USA, 1993. Association for Computing Machinery. doi:10.1145/160985.161156.
- 29 Joseph S. B. Mitchell. Approximating watchman routes. In *Proceedings of the 2013 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 844–855, 2013. doi:10.1137/1.9781611973105.60.
- 30 Joseph S. B. Mitchell and Christos H. Papadimitriou. The weighted region problem: finding shortest paths through a weighted planar subdivision. *J. ACM*, 38(1):18–73, January 1991. doi:10.1145/102782.102784.
- 31 Joseph S. B. Mitchell and Micha Sharir. New results on shortest paths in three dimensions. In *Proceedings of the 20th ACM Symposium on Computational Geometry, Brooklyn, New York, USA, June 8-11, 2004*, pages 124–133. ACM, 2004. doi:10.1145/997817.997839.
- 32 Eunjin Oh, Sang Won Bae, and Hee-Kap Ahn. Computing a geodesic two-center of points in a simple polygon. *Computational Geometry*, 82:45–59, 2019. doi:10.1016/j.comgeo.2019.04.003.
- 33 Joseph O’Rourke. *Art Gallery Theorems and Algorithms*. Oxford University Press, Inc., New York, NY, USA, 1987.
- 34 R. Pollack, M. Sharir, and G. Rote. Computing the geodesic center of a simple polygon. *Discrete & Computational Geometry*, 4(6):611–626, 1989. doi:10.1007/BF02187751.
- 35 Christian Rieck and Christian Scheffer. The Dispersive Art Gallery Problem. In Sang Won Bae and Heejin Park, editors, *33rd International Symposium on Algorithms and Computation (ISAAC 2022)*, volume 248 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 67:1–67:18, Dagstuhl, Germany, 2022. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.ISAAC.2022.67.
- 36 Jacob T Schwartz and Micha Sharir. On the “piano movers” problem. II. general techniques for computing topological properties of real algebraic manifolds. *Advances in Applied Mathematics*, 4(3):298–351, 1983. doi:10.1016/0196-8858(83)90014-3.
- 37 Jacob T. Schwartz and Micha Sharir. On the piano movers’ problem: III. coordinating the motion of several independent bodies: The special case of circular bodies moving amidst polygonal barriers. *The International Journal of Robotics Research*, 2(3):46–75, 1983. doi:10.1177/027836498300200304.

- 38 Micha Sharir and Amir Schorr. On shortest paths in polyhedral spaces. In *Proceedings of the Sixteenth Annual ACM Symposium on Theory of Computing*, STOC '84, pages 144–153, New York, NY, USA, 1984. Association for Computing Machinery. doi:10.1145/800057.808676.
- 39 G. T. Toussaint. An optimal algorithm for computing the relative convex hull of a set of points in a polygon. *EURASIP, Signal Processing III: Theories Applications, Part 2*, pages 853–856, 1986.
- 40 Haitao Wang. Shortest paths among obstacles in the plane revisited. In *Proceedings of the Thirty-Second Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '21, pages 810–821, USA, 2021. Society for Industrial and Applied Mathematics. doi:10.1137/1.9781611976465.51.
- 41 Glenn E. Wangdahl, Stephen M. Pollock, and John B. Woodward. Minimum-trajectory pipe routing. *Journal of Ship Research*, 18(01):46–49, March 1974. doi:10.5957/jsr.1974.18.1.46.