

How to Walk Your Dog in the Mountains with No Magic Leash

Sariel Har-Peled · Amir Nayyeri · Mohammad Salavatipour ·
Anastasios Sidiropoulos

the date of receipt and acceptance should be inserted later

Abstract We describe a $O(\log n)$ -approximation algorithm for computing the homotopic Fréchet distance between two polygonal curves that lie on the boundary of a triangulated topological disk. Prior to this work, algorithms were known only for curves on the Euclidean plane with polygonal obstacles.

A key technical ingredient in our analysis is a $O(\log n)$ -approximation algorithm for computing the minimum height of a homotopy between two curves. No algorithms were previously known for approximating this parameter. Surprisingly, it is not even known if computing either the homotopic Fréchet distance, or the minimum height of a homotopy, is in NP.

Keywords Fréchet distance, approximation algorithms, homotopy height

Mathematics Subject Classification (2000) 68W25

1 Introduction

Comparing the shapes of curves – or sequenced data in general – is a challenging task that arises in many different contexts. The *Fréchet distance* and its variants (e.g. dynamic time-warping [28]) have been used as a similarity measure in various applications such as matching of time series in databases [29], comparing melodies in music information retrieval [36], matching coastlines over time [32], as well as in map-matching of vehicle tracking data [5, 39], and moving objects analysis [7, 8]. See [2, 3] for algorithms for computing the Fréchet distance.

Informally, for a pair of such curves $f, g : [0, 1] \rightarrow \mathcal{D}$, for some ambient metric space (\mathcal{D}, d) , their Fréchet distance is the minimum length of a leash needed to traverse both curves in sync. To this end, imagine a person traversing f starting from $f(0)$, and a dog traversing g starting from $g(0)$, both traveling continuously along these curves without ever moving backwards. Then the Fréchet distance is the infimum over all possible traversals, of the maximum distance between the person and the dog. This notion can be formalized via a reparameterization: a continuous bijection $\phi : [0, 1] \rightarrow [0, 1]$. The *width* of ϕ , i.e., the longest leash needed by ϕ , is $\text{width}(\phi) = \sup_{t \in [0, 1]} d(f(t), g(\phi(t)))$, where $d(x, y)$ is the distance between x and y on \mathcal{D} . Consequently, the Fréchet distance between f and g is defined to be

$$d_{\mathcal{F}}(f, g) = \inf_{\phi: [0, 1] \rightarrow [0, 1]} \text{width}(\phi),$$

where ϕ ranges over all orientation-preserving homeomorphisms.

While this measure captures similarities between two curves when the underlying space is Euclidean, it is not as informative for more complicated underlying spaces such as a surface. For example, imagine walking a dog in the

A preliminary version of this paper appeared in SoCG 2012 [25].

S. Har-Peled

Department of Computer Science, University of Illinois, Urbana-Champaign; sariel@illinois.edu. Work on this paper was partially supported by NSF AF awards CCF-0915984, CCF-1421231, and CCF-1217462.

A. Nayyeri

School of Electrical Engineering & Computer Science, Oregon State University; nayyeria@eecs.oregonstate.edu.

M. Salavatipour

Department of Computing Science, University of Alberta; Supported by NSERC and Alberta Innovates. Part of this work was done while visiting Toyota Technological Institute at Chicago; mreza@cs.ualberta.ca.

A. Sidiropoulos

Department of Computer Science & Engineering, and Department of Mathematics, The Ohio State University; sidiropoulos.1@osu.edu. Research supported in part by the NSF grants CCF 1423230 and CAREER 1453472.




Fig. 1.1: (i) Two curves f and g , and (ii) the parameterization of their homotopic Fréchet distance.

woods. The leash might get tangled as the dog and the person walk on two different sides of a tree. Since the Fréchet distance cares only about the distance between the two moving points, the leash would “magically” jump over the tree. In reality, when there is no “magic” leash that jumps over a tree, one has to take into account the extra length needed (for the leash) to pass over such obstacles.

Homotopic Fréchet distance. To address this shortcoming, **homotopic Fréchet distance**, a natural extension of the above notion was introduced by Chambers *et al.* [10]. Informally, revisiting the above person-dog analogy, we consider the infimum over all possible traversals of the curves, but this time, we require that the person is connected to the dog via a leash, i.e., a curve that moves continuously over time. Furthermore, one keeps track of the leash during the motion, where the purpose is to minimize the maximum leash length needed.

To this end, consider a homotopy $h : [0, 1]^2 \rightarrow \mathcal{D}$, which can also be viewed as a homeomorphism between the unit square and \mathcal{D} . For parameters $\sigma, \tau \in [0, 1]$ consider the one dimensional functions $\ell(\tau) = h(\tau, \cdot) : [0, 1] \rightarrow \mathcal{D}$ and $\mu(\sigma) = h(\cdot, \sigma) : [0, 1] \rightarrow \mathcal{D}$. These are parameterized curves that are the natural restrictions of h into one dimension. We require that $\mu(0) = f$ and $\mu(1) = g$. The **homotopy width** of h is $\text{width}(h) = \sup_{\tau \in [0, 1]} \|\ell(\tau)\|$, and the **homotopic Fréchet distance** between f and g is

$$d_{\mathcal{H}}(f, g) = \inf_{h: [0, 1]^2 \rightarrow \mathcal{D}} \text{width}(h),$$

where the infimum is over all homeomorphisms h between $[0, 1]^2$ and \mathcal{D} , and $\|\cdot\|$ denotes the length of a curve. Note that $\ell(\cdot)$, in particular, specifies a reparametrization between the curves f and g .

Clearly, $d_{\mathcal{H}}(f, g) \geq d_{\mathcal{F}}(f, g)$ and, furthermore, $d_{\mathcal{H}}(f, g)$ can be arbitrarily larger than $d_{\mathcal{F}}(f, g)$. We remark that $d_{\mathcal{H}}(f, g) = d_{\mathcal{F}}(f, g)$ for any pair of curves in the Euclidean plane, as we can always pick the leash to be a straight line segment between the person and the dog. In other words, the map h in the definition of $d_{\mathcal{H}}$ can be obtained from the map ϕ in the definition of $d_{\mathcal{F}}$ via an appropriate affine extension. However, this is not true for general ambient spaces, where the leash might have to pass over obstacles, hills, or more generally regions of positive or negative curvature, etc. In particular, in the general settings, usually, the leash would not be a geodesic (i.e., a shortest path) during the motion.

The homotopic Fréchet distance is referred to as the **morph width** of f and g , and it bounds how far a point on f has to travel to its corresponding point in g under the morph of h [20]. The length of $\mu(\sigma)$ is the **height of the morph at time σ** , and the **height** of such a morph is $\text{height}(h) = \sup_{\sigma \in [0, 1]} \|\mu(\sigma)\|$. The **homotopy height** between f and g , bounded by given starting and ending leashes $\ell(0)$ and $\ell(1)$, is

$$h(f, g, \ell(0), \ell(1)) = \inf_h \text{height}(h),$$

where h varies over all possible morphs between f and g , such that each curve $\mu(\sigma)$ has one end on $\ell(0)$ and one end on $\ell(1)$. See Figure 1.1 for an example. Note that if we do not constrain the endpoints of the curves during the morph to stay on $\ell(0)$ and $\ell(1)$, the problem of computing the minimum height homotopy is trivial. One can contract f to a point, send it to a point in g , and then expand it to g . To keep the notation simple, we use $h(f, g)$ when f and g have common endpoints.

Intuitively, the homotopy height measures how long the curve has to become as it deforms from f to g , and it was introduced by Chambers and Letscher [11, 12] and Brightwell and Winkler [6]. Observe that if we are given the starting and ending leashes $\ell(0)$ and $\ell(1)$ then the homotopy height of f and g when restricted on homotopies that agree with $\ell(0)$ and $\ell(1)$ is the homotopic Fréchet distance between $\ell(0)$ and $\ell(1)$.

Here, we are interested in the problems of computing the homotopic Fréchet distance and the homotopy height between two simple polygonal curves that lie on the boundary of an arbitrary triangulated topological disk.

Why are these measures interesting? For the sake of the discussion here, assume that we know the starting and ending leash of the homotopy between f and g . The region bounded by the two curves and these leashes, forms a topological disk, and the mapping realizing the homotopic Fréchet distance specifies how to sweep over \mathcal{D} in a geometrically “efficient” way (especially if the leash does not sweep over the same point more than once), so that the leash (i.e., the sweeping curve) is never too long [20]. As a concrete example, consider the two curves as enclosing several mountains between them on the surface – computing the homotopic Fréchet distance corresponds to deciding which mountains to sweep first and in which order.

Furthermore, this mapping can be interpreted as surface parameterization [22,37] and can thus be used in applications such as texture mapping [4,35]. In the texture mapping problem, we wish to find a continuous and invertible mapping from the texture, usually a two-dimensional rectangular image, to the surface.

Another interesting interpretation is when f is a closed curve, and g is a point. Interpreting f as a rubber band in a 3d model, the homotopy height between f and g here is the minimum length the rubber band has to have so that it can be collapsed to a point through a continuous motion within the surface. In particular, a short closed curve with large homotopy height to any point in the surface is a “neck” in the 3d model.

To summarize, these measures seem to provide us with a fundamental understanding of the structure of the given surface/model.

Continuous vs. discrete. Here we are interested in two possible models. In the *continuous* settings, as described above, the leash moves continuously in the interior of the domain. In the *discrete* settings, the leash is restricted to the triangulation edges. As such, a transition of the leash corresponds to the leash “jumping” over a single face at each step. The two versions are similar in nature, but technically require somewhat different tools and insights. This issue is discussed more formally in Section 2.

Previous work. The problem of computing the (standard) Fréchet distance between two polygonal curves in the plane has been considered by Alt and Godau [3], who gave a polynomial time algorithm. Eiter and Mannila [21] studied the easier discrete version of this problem. Computing the Fréchet distance between surfaces [23], appears to be a much more difficult task, and its complexity is poorly understood. The problem has been shown to be NP-Hard by Godau [24], while the best algorithmic result is due to Alt and Buchin [2], who showed that it is upper semi-computable.

Efrat *et al.* [20] considered the Fréchet distance inside a simple polygon as a way to facilitate sweeping it efficiently. They also used the Fréchet distance with the underlying geodesic metric as a way to obtain a morph between two curves. For recent work on the Fréchet distance, see [17,13,26,16,19,18] and references therein.

Chambers *et al.* [10] gave a polynomial time algorithm to compute the homotopic Fréchet distance between two polygonal curves on the Euclidean plane with polygonal obstacles. Chambers and Letscher [11,12] and Brightwell and Winkler [6] considered the notion of minimum homotopy height, and proved structural properties for the case of a pair of paths on the boundary of a topological disk. We remark that in general, it is not known whether the optimum homotopy has polynomially long description. In particular, it is not known whether the problem is in NP.

Variants of the Fréchet distance for curves that are known to be computationally hard, include (i) the problem of finding the most similar *simple* (i.e., no self crossings) curve to a given curve on a surface [38], and (ii) computing the optimal Fréchet distance when allowing shortcuts anywhere on one of the curves [9]. Chambers and Wang [14] study a measure of similarity between curves that involves the minimum area spanned by a homotopy.

For a Riemannian 2-disk with boundary length L , diameter d and area $A \ll d$, Papasoglu [34] showed that there is a homotopy that contracts the disk to a point, such that the maximum length of the homotopy curve is at most $L + 2d + O(\sqrt{A})$. Chambers and Rotman [15] showed that given such a homotopy with maximum length L (i.e., any contraction of a disk to a point), one can modify it into a homotopy using only the loops of a base point p , contracting the disk into p , with maximum length $L + 2d + \varepsilon$, where $\varepsilon > 0$ is arbitrarily small.

Our results. In this paper, we consider the problems of computing the homotopic Fréchet distance and the homotopy height between two simple polygonal curves that lie on the boundary of a triangulated topological disk \mathcal{D} that is composed of n triangles. We give a polynomial time $O(\log n)$ -approximation algorithm for computing the homotopy height between f and g . Our approach is based on a simple, yet delicate divide and conquer approach.

We use the homotopy height algorithm as an ingredient for a $O(\log n)$ -approximation algorithm for the homotopic Fréchet distance problem. Here is a high-level description of our algorithm for approximating the homotopic Fréchet distance: We first guess (i.e., search over) the optimum (i.e., $d_{\mathcal{H}}(f, g)$). Using this guess, we classify parts of \mathcal{D} as “obstacles”, i.e., regions over which a short leash cannot pass. Consider the punctured disk obtained from \mathcal{D} after removing these obstacles. Intuitively, two leashes are homotopic if one can be continuously deformed to the other within the punctured disk, while its endpoints remain on the boundary during the deformation. Observe that the leashes of the optimum solution are homotopic. We describe a greedy algorithm to compute a “small” number of homotopy classes out of infinitely many choices. The homotopic Fréchet distance constrained to paths inside one of these classes is a polynomial approximation to the homotopic Fréchet distance in \mathcal{D} . We can then perform a binary search over this interval to acquire a better approximation. An extended version of the homotopy height algorithm is used in this algorithm in several places.

The $O(\log n)$ factor shows up in the homotopic Fréchet distance algorithm only because it uses the homotopy height as a subroutine. Thus, any constant factor approximation algorithm for the homotopy height problem implies a constant factor approximation algorithm for the homotopic Fréchet distance.

We also shortly sketch, in [Appendix A](#), an algorithm for sweeping the boundary of a convex polytope in three dimensions, by a guard that is connected by a continuously moving leash to a base point on the boundary of the polytope. This algorithm is a warm-up exercise for the more involved problem studied in this paper, and it might be of independent interest.

Organization. We provide basic definitions in [Section 2](#). Then we consider the discrete version of the homotopy height problem in [Section 3](#). Later, in [Section 4](#), we describe an algorithm to approximately find the shortest homotopy in the continuous settings. In [Section 5](#), we address the homotopic Fréchet distance, for both the discrete and the continuous cases. We conclude in [Section 6](#).

2 Background

2.1 Planar graphs

Let $G = (V, E)$ be a simple undirected graph with edge weights $w : E \rightarrow \mathbb{R}^+$. For any $u, v \in V$ we denote by $d_G(x, y)$ the shortest-path distance between u and v in G , where every edge e has length $w(e)$. An *embedding* of G in the plane maps the vertices of G to distinct points in the plane and its edges to disjoint paths except for the endpoints. The faces of an embedding are maximal connected subsets of the plane that are disjoint from the union of the (images of the) edges of the graph. The notation ∂f refers to the boundary of a single face f . The term *plane graph* refers to a graph together with its embedding in the plane.

The *dual* graph G^* of a plane graph G is the (multi-)graph whose vertices correspond to the faces of G , where two faces are joined by a (dual) edge if and only if their corresponding faces are separated by an edge of G . Thus, any edge e in G corresponds to a dual edge e^* in G^* , any vertex v in G corresponds to a face v^* in G^* and any face f in G^* corresponds to a vertex f^* in G .

A *walk* W in G is a sequence of vertices (v_1, v_2, \dots, v_k) such that each adjacent pair $e_i = (v_i, v_{i+1})$ is an edge in G . The length of W is $\|W\| = \sum_{i=1}^{k-1} w(e_i)$.

Let v_i and v_j be two vertices that appear on W . Here, $W[v_i, v_j]$ denotes the sub-walk of W that starts at the first appearance of v_i and ends at the first appearance of v_j after v_i on W . For two walks, $W_1 = (v_1, v_2, \dots, v_i)$ and $W_2 = (v_i, v_{i+1}, \dots, v_j)$, their *concatenation* is $W_1 \cdot W_2 = (v_1, v_2, \dots, v_i, v_{i+1}, \dots, v_j)$.

A walk with all the vertices being distinct is a *path*. The term (u, v) -walk refers to a walk that starts at u and ends in v , and (u, v) -path is defined similarly. A walk is closed if its first and last vertices are identical. A closed path is a *cycle*. Two walks *cross* if and only if their images cross in the plane. That is, no infinitesimal perturbation makes them disjoint.

2.2 Piecewise linear surfaces and geodesics

A *piecewise linear* surface is a 2-dimensional manifold composed of a finite number of Euclidean triangles by identifying pairs of equal length edges. In this paper, we work with piecewise linear surfaces that can be embedded in three dimensional space such that all triangles are flat and the surface does not cross itself. This assumption lets us exploit existing shortest paths algorithms for polyhedral surfaces^①.

A triangulated surface is *non-degenerate* if no interior vertex has curvature 0, i.e., when for every non-boundary vertex x , the sum of the angles of the triangles incident to x is not equal to 2π .

Assumption 1 We assume that the input surface is always non-degenerate. One can turn any triangulated surface into being non-degenerate by perturbing all edge lengths by a factor of at most $1 + \varepsilon$, for some $\varepsilon = O(1/n^2)$ (alternatively, one can perturb the vertices and edges). This changes the metric of the surface by at most a factor of $1 + 1/n$, and thus the minimum height of a homotopy. Such a factor will be negligible for our approximation guarantee.

A *path* γ on the surface \mathcal{D} is a function $\gamma : [0, 1] \rightarrow \mathcal{D}$; $\gamma(0)$ and $\gamma(1)$ are the endpoints of the path, and $\|\gamma\|$ denotes the length of γ . The path γ is *simple* if and only if it is bijective. A path is a *geodesic* if and only if it is locally a shortest path; i.e., it cannot be shortened by an infinitesimal perturbation. In particular, global shortest paths are geodesics. The terms path or *curve* are used interchangeably, and mean the same thing. A path or a curve is polygonal if it is composed of a finite number of line segments.

^①However, all of these algorithms should work verbatim even if the surface is not embedded in 3d, assuming it is an oriented and has the topology of a disk. Nevertheless, we keep this assumption to make the discussion more concrete and hopefully more intuitive.




Fig. 2.1: From left to right: face-flip, spike/reverse spike, person-move and dog-move.

2.2.1 Computing shortest paths on a polyhedral surface

Mitchell *et al.* [33] describe an algorithm to compute the shortest path distance from a single source to the whole surface in $O(n^2 \log n)$ time. Underlying this algorithm are the following two observations:

- (A) Shortest paths from the source s can not intersect in their interior.
- (B) Consider two points p and p' on an edge e of \mathcal{D} , and their two shortest paths ζ and ζ' , respectively, to the source s . Furthermore, assume that these shortest paths approach e from the same side. Then, all the shortest paths from s to the points on the edge e between p and p' (coming from the same side of e as ζ and ζ'), must lie inside the disk on \mathcal{D} having the boundary $s \cdot \zeta \cdot e[p, p'] \cdot \zeta'$. This property requires that the topology of the input surface to be either a disk or, more generally, a punctured disk.

These two observations still hold when the source is an edge instead of a point. As such, the algorithm of Mitchell *et al.* [33] can be adapted to compute the shortest path distance from an edge to the whole surface (with the same running time). This requires modifying the wavefront maintenance and propagation to be the distance from an edge instead from a point.

As such, by running this modified algorithm $O(n)$ times, one can compute, in $O(n^3 \log n)$ time, the shortest path from a set of $O(n)$ edges to the whole surface.

Signature & medial points. Any shortest path in \mathcal{D} is a polygonal line that intersects every edge at most once and passes through a face along a segment. Moreover, the shortest path crossing an edge looks locally like a straight line segment, if one rotates the adjacent faces so that they are coplanar. See [33] for more details.

Let S be a set of edges of \mathcal{D} , and let ζ be a shortest path from S to a point $p \in \mathcal{D}$. The *signature* of ζ is defined to be the ordered set of edges and vertices (crossed or used) by ζ . Since ζ is locally a straight line segment, we can rotate all faces that intersect ζ one by one so that ζ becomes a straight line. It follows that two geodesics with the same signature from p are identical. A point p on the surface is a *medial* point with respect to S if there are more than one shortest paths (with different signatures) from p to S .

Note that a shortest path has a vertex of \mathcal{D} in its interior, if and only if, the vertex is a boundary vertex, or the vertex has negative curvature (i.e., the total sum of the angles of the triangles adjacent to this vertex is larger than 2π). In particular, a vertex with positive curvature (i.e., total sum of angles $< 2\pi$), which is not on the boundary of \mathcal{D} , can not be in a signature of a shortest path, see also [Assumption 1](#).

2.3 Homotopy and leash function

Let L and R be two paths with the same endpoints s and t on a surface \mathcal{D} . A homotopy $h : [0, 1] \times [0, 1] \rightarrow \mathcal{D}$ is a continuous function, such that $h(\cdot, 0) = L$, $h(\cdot, 1) = R$, $h(0, \cdot) = s$ and $h(1, \cdot) = t$. So, for each $\tau \in [0, 1]$, $h(\cdot, \tau)$ is an (s, t) -path. The *height* of such a homotopy (as defined previously) is defined to be $\sup_{\tau \in [0, 1]} \|h(\cdot, \tau)\|$.

Let B and C be two disjoint curves. A curve connecting a point in B to a point in C is a (B, C) -*leash*. A (B, C) -*leash function* is a function f that maps every $\tau \in [0, 1]$ to a leash with endpoints $b(\tau) \in B$ and $c(\tau) \in C$ such that $b : [0, 1] \rightarrow B$ and $c : [0, 1] \rightarrow C$ are reparametrizations of B and C , respectively. A (B, C) -leash function f is *continuous* if the leash $f(\tau)$ varies continuously with τ . The *height* of a leash function f is $\sup_{\tau \in [0, 1]} \|f(\tau)\|$. Recall that the *Fréchet distance* between B and C is the height of the minimum height (B, C) -leash function. The *homotopic Fréchet distance* between B and C is the height of the minimum height continuous (B, C) -leash function.

2.3.1 The discrete version

Let W_1 be an (s, t) -walk and f be a face in an embedded planar graph G . Assume that α_1 is a subwalk of W_1 and $\partial f = \alpha_1 \cdot \alpha_2$, where α_1 and α_2 are walks that share endpoints u and v , such that u is closer to s on W_1 . The *face*

flip operation is the following: The walk $W_2 = W_1[s, u] \cdot \alpha_2 \cdot W_1[v, t]$ is the result of flipping W_1 over f . In this case, W_1 and W_2 are one face flip operation apart. See [Figure 2.1](#).

Now, let W_1 be an (s, t) -walk, and suppose $W_1 = W_1' \cdot W_1''$. Also, let u be the common endpoint of W_1' and W_1'' , and let $e = (u, v)$ be any edge in G . By applying a **spike** to W_1 we obtain $W_2 = W_1' \cdot (u, v) \cdot (v, u) \cdot W_1''$. Equivalently, we can obtain W_1 from W_2 via a **reverse spike**. In this case, W_1 and W_2 are one spike operation apart.

In general, W_1 and W_2 are one operation apart if one can transform one to the other using a single face flip, spike, or reverse spike. Chambers and Letscher [11, 12] introduce the same set of operations with the names: face lengthening, face shortening, spike and reverse spike.

Let L and R be two (s, t) -walks on the outer face of G . The sequence of walks $(L = W_0, W_1, \dots, W_m = R)$ is a (L, R) -**discrete homotopy** if, for $i = 1, \dots, m$, W_i and W_{i-1} are one operation apart. We may use the word homotopy as a short form of discrete homotopy when it is clear from the context. The height of the homotopy is defined to be the length of the longest walk in its sequence. The homotopy height between L and R , is the height of the shortest possible (L, R) -homotopy.

Definition 1 Let $B = (b_0, b_1, \dots, b_k)$ and $C = (c_0, c_1, \dots, c_{k'})$ be walks of G . The walk $W_1 = (b_i = w_1, w_2, \dots, w_k = c_j)$ changes to the walk $W_2 = (b_{i+1}, b_i = w_1, w_2, \dots, w_k)$ after a **person move**. Similarly, the walk $W_1 = (b_i = w_1, w_2, \dots, w_k = c_j)$ changes to the walk $W_2 = (w_1, w_2, \dots, w_k = c_j, c_{j+1})$ after a **dog move**. A **leash operation** is a dog move, a person move, a face flip, a spike or a reverse spike.

Definition 2 An (B, C) -walk is a walk that has one endpoint on B and one endpoint on C . A sequence of (B, C) -walks, (W_1, W_2, \dots, W_q) is called an (B, C) -**leash sequence** if

- (i) W_1 is a (b_0, c_0) -walk,
- (ii) W_q is a $(b_k, c_{k'})$ -walk, and
- (iii) we have that, for $i = 1, \dots, q-1$, W_i changes to W_{i+1} by performing a sequence of leash operations containing either at most one dog move and no person moves, or at most one person move and no dog moves.

The height of a leash sequence is the length of its longest walk.

Definition 3 The **discrete Fréchet distance** of B and C is the height of the minimum height (B, C) -leash sequence^②. The **homotopic discrete Fréchet distance** of B and C is the height of the minimum height (B, C) -leash sequence, where two consecutive walks differ by a single leash operation (and this is not required in the *discrete Fréchet distance*, where two consecutive walks might “jump”).

3 Approximating the height – the discrete case

In this section, we give an approximation algorithm for finding a discrete homotopy of minimum height in a topological disk \mathcal{D} , whose boundary is defined by two walks L and R that share their endpoints s and t . The disk \mathcal{D} is a triangulated edge-weighted planar graph. The ideas developed here are used later for the continuous case, see [Section 4](#).

3.1 Preliminaries

We are given an edge-weighted plane graph G all of whose faces (except possibly the outer face) are triangles. Let $s, t \in \partial G$ and L and R be two non-crossing (s, t) -walks on ∂G in counter-clockwise and clockwise order, respectively. We use \mathcal{D} to denote the topological disk enclosed by $L \cdot R$. The vertices of G (inside or on the boundary of \mathcal{D}) are also the vertices of \mathcal{D} . Our goal is to find a minimum height homotopy from L to R of non-crossing walks. Recall that a homotopy is a sequence of walks, where every two consecutive walks differ by either a triangle, or an edge (being traversed twice).

Lemma 1 *Let x and y be vertices of G that are at graph distance ρ . Then any discrete homotopy between L and R has height at least ρ .*

Proof Fix a homotopy of height δ . This homotopy contains an (s, t) -walk ω that passes through x , and an (s, t) -walk χ that passes through y . We have, by the triangle inequality, that $\rho = d_G(x, y) \leq \|\omega[s, x]\| + \|\chi[s, y]\|$, and, similarly, $\rho \leq \|\omega[x, t]\| + \|\chi[y, t]\|$. Therefore, $\rho \leq (\|\omega\| + \|\chi\|)/2 \leq \max(\|\omega\|, \|\chi\|) \leq \delta$, as required. \square

Lemma 2 *Suppose d_1 is the maximum distance of a vertex of G from L , d_2 is the largest edge weight, and let $d_L = \max\{d_1, d_2\}$. Furthermore, let \mathcal{D} , L , and R be defined as above. Then any discrete homotopy between L and R has height at least d_L .*

Proof The height is at least d_1 by [Lemma 1](#). On the other hand, for every homotopy between L and R , and for every edge e , there exists a walk in the homotopy that passes through e . Therefore, the height of the homotopy is at least d_2 . \square

^②The discrete Fréchet distance defined here is different than the more standard definition, which is usually defined over sequences of points.




Fig. 3.1: (I) Corridors edges. (II) Case (A). (III) Case (B).

3.2 The algorithm

Theorem 1 *Let \mathcal{D} be an edge-weighted triangulated topological disk with n faces such that its boundary is formed by two walks L and R that share endpoints s and t . Then one can compute, in $O(n \log n)$ time, a homotopy from L to R of height at most $\|L\| + \|R\| + O(d_L \log n)$, where d_L is the largest among (i) the maximum distance of a vertex of \mathcal{D} from L , and (ii) the maximum edge weight.*

In particular, the generated homotopy has height $O(h_{\text{opt}} \log n)$, where h_{opt} is the minimum homotopy height between L and R .

Proof We present a recursive algorithm that reduces the problem to sub-problems with a smaller number of triangles. The recursion might create instances where the boundary walks L and R are not interior disjoint. For such instances, it is immediate that one can obtain a solution by computing a homotopy independently for each maximal disk bounded by $L \cdot R$, and composing them to obtain the desired homotopy between L and R . We may therefore focus on the case where L and R are interior disjoint.

Let $b(\delta, d_L, n)$ be the maximum possible height of a homotopy obtained by our algorithm for any disk of perimeter δ that is composed of n faces and has d_L as defined in the statement of the theorem. We prove by induction that $b(\delta, d_L, n) \leq \delta + c_0 d_L \log n$ for some constant c_0 , implying the theorem statement. Note that the inductive argument implies that b is linear in δ .

The base case $n = 0$ is easy. Indeed, if we have two edges (u, v) and (v, u) consecutive in R (or in L) we can retract these two edges. By repeating this we arrive at both L and R being identical, and we are done. The case $n = 1$ is handled in a similar fashion. After one face flip, the problem reduces to the case $n = 0$. Hence, $b(\|L\| + \|R\|, d_L, 1) \leq \|L\| + \|R\| + d_L$.

For $n > 1$, compute for each vertex of G its shortest path to L , and consider the set of edges \mathcal{E} used by all these shortest paths. Clearly, these shortest paths can be chosen so that $L \cup \mathcal{E}$ form a tree. We consider each edge of R to be “thick” and have two sides (i.e., we think about these edges as being corridors – this is done to guarantee that in the recursive scheme, done below, there are exactly two subproblems to each instance). If \mathcal{E} uses an edge of R then it uses the inner copy of this edge, while R uses the outer side. Similarly, we consider each original vertex of R to be two vertices (one inside and the other one on the boundary R). The set \mathcal{E} would use only the inner vertices of R , while R would use only the outer vertices. To keep the graph triangulated we also arbitrarily triangulate inside each thick edge of R by adding corridor edges. Each corridor edge either connects two copies of a single vertex (thus has weight zero) or copies of two neighbors on R (and so has the same weight as the original edge).

Clearly, if we cut \mathcal{D} along the edges of \mathcal{E} , what remains is a simple triangulated polygon (it might have “thin” corridors along the edges of R). One can find a diagonal uv such that each side of the diagonal contains at least $\lceil n/3 \rceil$ triangles and at most $\lfloor 2n/3 \rfloor$ triangles of the *original* G . We emphasize that we count only the “real” triangles of G . This can be achieved as follows: We first assign weight zero to the faces within corridors and unit weight to all other faces. Then we find a diagonal uv such that each side contains faces with total weight at least $\lceil n/3 \rceil$. Furthermore, because the faces inside corridors have weight zero, we can ensure that if the separating edge uv is a corridor edge (i.e., corresponding to an edge e of R) then u and v are copies of the same vertex. Indeed, if not, we can change

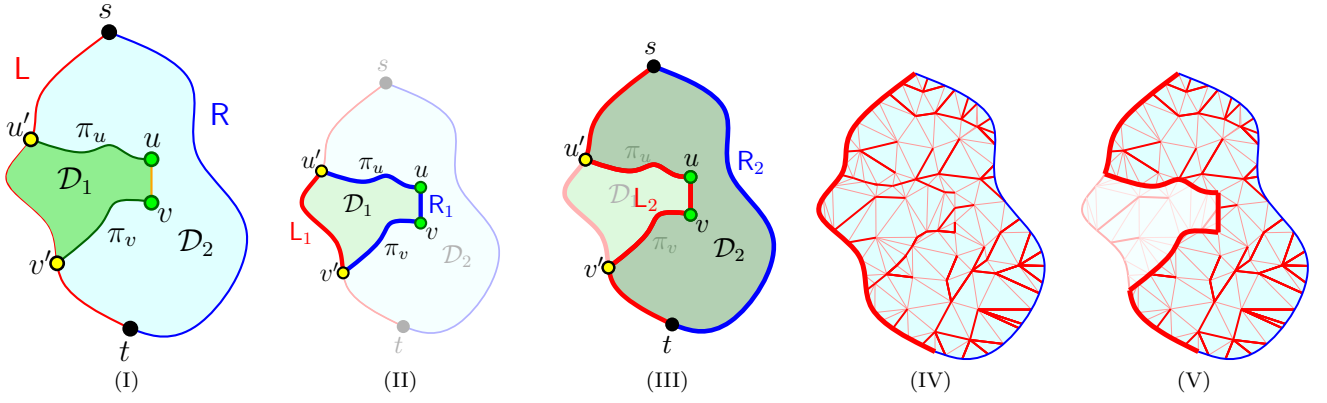


Fig. 3.2: (I–III) Illustration of case (C) in the proof [Theorem 1](#). (IV–V) How the shortest path tree get sent to the recursive subproblem.

the separating edge so this property holds, and the new separating edge separates regions with the same weight, see [Figure 3.1](#) (I). We use this property in the following case analysis³.

(A) Consider the case⁴ that u and v are both vertices of R . In this case, let $R[u, v]$ be the portion of R in between u and v , and let \mathcal{D}_2 be the disk having $R[u, v] \cdot uv$ as its outer boundary. Let \mathcal{D}_1 be the disk $\mathcal{D} \setminus \mathcal{D}_2$. Let $M = R[s, u] \cdot uv \cdot R[v, t]$, see [Figure 3.1](#) (II).

Clearly, the distance of any vertex of \mathcal{D}_1 from L is at most d_L . By induction, there is a homotopy of height $b(\|L\| + \|M\|, d_L, \lfloor 2n/3 \rfloor)$ from L to M . Similarly, the distance of any vertex of \mathcal{D}_2 from uv is at most its distance to L . Therefore, by induction, there is a homotopy between uv and $R[u, v]$ of height at most $b(\|R[u, v]\| + d_L, d_L, \lfloor 2n/3 \rfloor)$. Clearly, we can extend this to a homotopy of M to R of height $\|R[s, u]\| + b(\|R[u, v]\| + d_L, d_L, \lfloor 2n/3 \rfloor) + \|R[v, t]\|$ which using induction hypothesis is at most $\|R\| + d_L + c_0 d_L \log \lfloor 2n/3 \rfloor \leq \|R\| + c_0 d_L \log n$, for sufficiently large c_0 .

Putting these two homotopies together results in the desired homotopy from L to R .

(B) Consider the case that v is a vertex of \mathcal{E} and u is a vertex of R . So, v is an inner vertex of R (that belongs to \mathcal{E}) and u is an outer vertex of R . Recall that we can assume that v and u are inner and outer copies of the same vertex of R . Let π_v be the shortest path in \mathcal{D} from v to L , and let v' be its endpoint on L .

Consider the disk \mathcal{D}_1 having the “left” boundary $L_1 = L[s, v'] \cdot \pi_v \cdot vu$ and $R_1 = R[s, u]$ as its “right” boundary, see [Figure 3.1](#) (III). This disk contains at most $\lfloor 2n/3 \rfloor$ triangles, and by induction, it has a homotopy of height $b(\|L_1\| + \|R_1\|, d_L, \lfloor 2n/3 \rfloor)$. To see why we can apply the recursion, observe that u and v are copies of the same vertex of R . That is, all shortest paths of vertices inside \mathcal{D}_1 to L are completely inside \mathcal{D}_1 . As such, the distance of all vertices in \mathcal{D}_1 to L_1 are at most d_L .

Similarly, the topological disk \mathcal{D}_2 with the left boundary $L_2 = uv \cdot \pi_v \cdot L[v', t]$ and the right boundary $R_2 = R[u, t]$ has a homotopy of height $b(\|L_2\| + \|R_2\|, d_L, \lfloor 2n/3 \rfloor)$.

We combine these two homotopies as follows. Let L' be the walk obtained by concatenating L_1 and L_2 . Note that L' consists of a copy of L and two copies of π_v . Clearly, there exists a homotopy between L and L' of height at most $\|L\| + 2d_L$, which is obtained by a sequence of spike moves along π_v . We compose the resulting homotopy with the homotopy of \mathcal{D}_1 that moves L_1 to R_1 , followed by the homotopy of \mathcal{D}_2 that moves L_2 to R_2 . The result is a homotopy between L and R of height at most

$$\max \left(\begin{array}{c} \|L\| + 2d_L, \\ b(\|L_1\| + \|R_1\|, d_L, \lfloor 2n/3 \rfloor) + \|L_2\|, \\ \|R_1\| + b(\|L_2\| + \|R_2\|, d_L, \lfloor 2n/3 \rfloor) \end{array} \right).$$

If the first number is the maximum, we are done. Otherwise, using the induction hypothesis, the above value is at most $\|L\| + \|R\| + 2d_L + c_0 d_L \log \lfloor 2n/3 \rfloor$ which is at most $\|L\| + \|R\| + c_0 d_L \log n$ for sufficiently large c_0 .

(C) Here we handle the case that u and v are both vertices of $L \cup \mathcal{E}$. Then as before, let u' and v' be the closest points on L to u and v , respectively. Now, let π_u (resp. π_v) be the shortest path from u (resp. v) to u' (resp. v'). Note that we might have $u' = v'$.

Consider the disk \mathcal{D}_1 having $L_1 = L[u', v']$ as left boundary, and $R_1 = \pi_u \cdot uv \cdot \pi_v$ as right boundary, see [Figure 3.2](#). This disk contains between $n/3$ and $2n/3$ triangles of the original surface. The distance of any vertex of \mathcal{D}_1 to L_1

³Note, that the corridors were used only in generating this partition, and are an artifact that is not necessarily sent to the recursive subproblems. In particular, one can describe this partition scheme without using the corridors, but it seems somewhat messier and less intuitive.

⁴Strictly speaking this case is not possible because of the corridor diagonals. Nevertheless, it provides a good warm-up exercise for the followup cases which are more involved.

(when restricted to \mathcal{D}_1) is at most d_L , and therefore by induction, there is a homotopy from L_1 to R_1 of height at most $\alpha = b(\|L_1\| + \|R_1\|, d_L, \lfloor 2n/3 \rfloor) \leq \|L[u', v']\| + 3d_L + c_0 d_L \log \lfloor 2n/3 \rfloor$. This yields a homotopy of height $\alpha_1 = \|L[s, u']\| + \alpha + \|L[v', t]\|$, from L to $L_2 = L[s, u'] \cdot \pi_u \cdot uv \cdot \pi_v \cdot L[v', t]$. It is straightforward to check that $\alpha_1 \leq \|L\| + 3d_L + c_0 d_L \log \lfloor 2n/3 \rfloor$.

Next, let \mathcal{D}_2 be the disk with its left boundary being L_2 and its right boundary being $R_2 = R$. Observe, that as before, the maximum distance of any vertex of \mathcal{D}_2 to L_2 is at most d_L . As before, by induction, there is a homotopy from L_2 to R_2 of height $\alpha_2 = b(\|L_2\| + \|R_2\|, d_L, \lfloor 2n/3 \rfloor)$. Since $\|L_2\| \leq \|L\| + 3d$, we have $\alpha_2 \leq b(\|L\| + \|R\| + 3d_L, d_L, \lfloor 2n/3 \rfloor)$.

In all cases the length of the homotopy is at most

$$\|L\| + \|R\| + 3d_L + c_0 d_L \log \lfloor 2n/3 \rfloor \leq \|L\| + \|R\| + c_0 d_L \log n,$$

if we choose c_0 sufficiently large. The final guarantee of approximation follows as $d_L \leq h_{\text{opt}}$, by [Lemma 2](#).

We can compute the shortest path tree in linear time using the algorithm of Henzinger *et al.* [27]. The separating edge can also be found in linear time using DFS. So, the running time for a graph with n faces is $T(n) = T(n_1) + T(n_2) + O(n)$, where $n_1 + n_2 = n$ and $n_1, n_2 \leq (2/3)n$. It follows that $T(n) = O(n \log n)$. \square

Remark 1 (A) In the algorithm of [Theorem 1](#), it is not necessary that we have the shortest paths from L to all the vertices of \mathcal{D} . Instead, it is sufficient if we have a tree structure that provides paths from any vertex of \mathcal{D} to L of distance at most d_L in this tree, and we send the relevant portions of the tree into the recursive subproblems. We will use this property in the continuous case, where recomputing the shortest path tree is relatively expensive. This is demonstrated in [Figure 3.2](#) (IV–V).

(B) A more careful analysis shows that the height of the homotopy generated by [Theorem 1](#) is at most $\max(\|L\|, \|R\|) + O(d_L \log n)$.

(C) Note, that if $d_L = O(\max(\|L\|, \|R\|) / \log n)$ then [Theorem 1](#) provides a constant factor approximation. This is the situation when L and R are close to each other compared to their relative length.

(D) Note, that the $O(n \log n)$ running time algorithm cannot explicitly output the list of paths in the homotopy. Indeed, that list requires $O(n^2)$ space to be stored and so $O(n^2)$ time to output. The output of the algorithm of the above lemma is a shortest path tree T together with an ordered list of edges. Each edge $e = (u, v)$ in the list represents an (s, t) -walk $T[s, u] \cdot (u, v) \cdot T[v, t]$, where $T[s, u]$ and $T[v, t]$ are the unique (s, u) -path and (v, t) -path in T , respectively.

4 Approximating the height – the continuous case

In this section we extend the algorithm from [Section 3](#) to the continuous case. The continuous case is somewhat similar to the solution to the problem of sweeping over the boundary of a convex polytope in three dimensions from a base point. Since this is tangential for our main trust, we delegated describing this algorithm to [Appendix A](#), but the reader might still benefit from reading it first.

4.1 Preliminaries

We are given a piecewise linear triangulated topological disk, \mathcal{D} , with n triangles, and we consider the underlying metric to be the geodesic distance on this surface⁵. The boundary of \mathcal{D} is composed of two paths L and R with shared endpoints s and t , and the task at hand is to compute a morphing from L to R that minimizes the distance traversed by each point of L during this motion. See [Section 2.3](#) for the formal definition.

Here, we build a homotopy of height at most $\|L\| + \|R\| + O(d \log n)$, where d is the maximum distance of any point in \mathcal{D} from either L or R . We use the following observations (see [Section 2.2.1](#) for details):

- (A) The shortest path from a vertex to the whole surface can be computed in $O(n^2 \log n)$ time.
- (B) The shortest path from a set of $O(n)$ edges to the whole surface can be computed in $O(n^3 \log n)$ time.
- (C) A shortest path (i.e., a geodesic) intersects a face along a segment and it locally looks like a segment if the adjacent faces are rotated to be coplanar.

See [Section 2.2.1](#) for more details.

4.2 Homotopy height if edges are short

Similar to the discrete case, d_1 is the maximum distance for any point of \mathcal{D} from L , d_2 is the maximum length of any edge, and $d_L = \max(d_1, d_2)$. Here, we assume $d_2 \leq 2d_1$. In this case we can obtain the desired approximation algorithm via an argument that is similar to the one used in the discrete case.

As in the discrete case, let \mathcal{E} be the union of all the shortest paths from the vertices of \mathcal{D} to L (as before, we treat the edges and vertices of R as having infinitesimal thickness). For a vertex v of \mathcal{D} , its shortest path π_v is a

⁵Formally, for two points $p, q \in \mathcal{D}$, their *geodesic distance* is the length of the shortest path inside \mathcal{D} connecting p with q .




Fig. 4.1

polygonal path that crosses between faces (usually) in the middle of edges (it might also go to a vertex, merge with some other shortest paths and then follow a common shortest path back to L). In particular, each such shortest path might intersect a face of \mathcal{D} along a single segment. Thus, the polygon resulting from cutting \mathcal{D} along \mathcal{E} , call it P , is a polygon that has complexity $O(n^2)$. A face of P is a hexagon, a pentagon, a quadrilateral, or a triangle. However, each such face has at most three edges that are portions of the edges of \mathcal{D} . The degree of a face is i if it has i edges that are portions of the edges of \mathcal{D} . Observe that, each triangle of \mathcal{D} is now decomposed into a set of faces. Obviously, each triangle of \mathcal{D} contains at most one face of degree 3 in P . Overall, there are $O(n)$ faces of degree 3 in P .

Now consider C^* , the dual of the graph that is inside the polygon (ignore the edges on the boundary). More precisely, C^* has a vertex corresponding to each face inside the polygon P , let n_p be number of vertices of C^* . Two vertices of C^* are adjacent if and only if their corresponding faces share a portion of an edge of \mathcal{D} (this shared edge is a diagonal of P). Note that because P is simply connected C^* is a tree. Since the maximum degree of the tree C^* is 3, there is an edge that is a good separator (i.e., a separator that has at most $2/3$ of the faces on one side)[®]. Since the length of the edge is at most $2d_1$ it can be used in a similar fashion as the proof of [Theorem 1](#). However, in the recursion of the continuous case we avoid recomputing the shortest paths (i.e., we use the old shortest paths and distances computed in the original disk), see [Remark 1](#). So, we compute the shortest paths once in the beginning in $O(n^3 \log n)$ time. Then in each step we can find the separator in $O(n^2)$ time. Namely, the total time spent on computing the separators is $T(n_p) = T(n_1) + T(n_2) + O(n^2)$, where $n_1 + n_2 = n_p$ and $n_1, n_2 \leq (2/3)(n_1 + n_2)$; since $n_p = O(n^2)$, $T(n) = O(n^2 \log n)$. As such, the total running time is dominated by the computation of the shortest paths. The output is a list of $O(n^2)$ paths each of complexity $O(n)$, and so it can be explicitly presented in $O(n^3)$ time and space. Note that, we need a continuous deformation between any two consecutive paths in the list, which can be implicitly presented by a collection of functions in linear time and space (this is similar to what we describe below in the beginning of [Section 4.3](#)).

The proof of [Theorem 1](#) then goes through literally in this case. Since all the edges have length at most $2d_1$, by assumption, we obtain the following.

Lemma 3 *Let \mathcal{D} be a topological disk with n faces where every face is a triangle (here, the distance between any two points on the triangle is their Euclidean distance). Furthermore, the boundary of \mathcal{D} is formed by two walks L and R (that share two endpoints s, t). Let d_1 be the maximum distance of any point of \mathcal{D} from L. Finally, assume that all edges of \mathcal{D} have length at most $2d_1$. Then one can compute, in $O(n^3 \log n)$ time, a continuous homotopy from L to R of height at most $\|L\| + \|R\| + O(d_1 \log n)$.*

4.3 Homotopy height if there are long edges

4.3.1 Breaking the disk into strips, pockets and chunks

For any two points in \mathcal{D} consider a shortest path π connecting them. The signature of π is the ordered sequence of edges (crossed or used) and vertices used by π , see [Section 2](#). For a point $p \in R$, let $\text{sg}(p)$ denote the signature of the shortest path from p to L. The signature $\text{sg}(p)$ is well defined in R except for a finite set of *medial* points, where there are two (or more) distinct shortest paths from L to p . In particular, let Π_R be the set of all shortest paths from any medial point on R to L. Observe that, the medial points are the only points (on R) where the signature of the shortest path from R to L changes in any non-degenerate triangulation.

Cutting \mathcal{D} along the paths of Π_R breaks \mathcal{D} into corridors. If the intersection of a corridor with R is a point (resp. segment) then it is a *delta* (resp. *strip*), see [Figure 4.1](#) (I). In a strip C , all the shortest paths to L from the points in the interior of the segment $C \cap R$ have the same signature. Intuitively, strips have a natural way to morph from one side to the other. We further break each delta into chunks and pockets, as follows.

[®]The existence of such a tree edge separator is folklore – its proof is provided by Lewis *et al.* [30].




Fig. 4.2

Consider a delta C with an apex c (i.e., the point of R on the boundary of C). For a point $x \in L \cap C$, its signature (in relation to C), is the signature of the shortest path from x to c (restricted to lie inside C). Again, we partition $L \cap C$ into maximum intervals that have the same signature, and let P be the set of endpoints of these intervals. For each point $p \in P$, consider all the different shortest paths from c to p inside the delta C , and cut C along these paths. This breaks C into regions. If a newly created region has a single intersection point with both L and R , then it is a **pocket**, otherwise, it is a **chunk**. Clearly, this process decomposes C into pockets and chunks. See Figure 4.1 (II).

Applying the above partition scheme to all the deltas results in a decomposition of \mathcal{D} into strips, chunks and pockets.

Analysis. Recall, that d_1 is the maximum distance of any point of \mathcal{D} to L , and let d_3 be the maximum distance of any point of \mathcal{D} to R . Also, let

$$d = \max(d_1, d_3). \quad (4.1)$$

Now, consider a chunk C' . Its intersection with L is a segment, and its intersection with R is a point (i.e., the apex c of the delta).

Lemma 4 *A strip C cannot have any vertex of \mathcal{D} in its interior.*

Proof Let $e_L = L \cap C$ and $e_R = R \cap C$ be two edges bounding a strip. For two points p, p' in the interior of e_R , consider their corresponding shortest paths ζ and ζ' to L . By definition, these two paths have the same signature $\text{sg}(\zeta) = \text{sg}(\zeta')$. If not, then by a limit argument, there must be a point $p'' \in e_R$ in between these two points, which has two different shortest paths with different signature arriving to it; that is, p'' is a medial point, implying that e_R is broken into (at least) two edges, and it can not be the right side of a strip.

Now, for $i = 1, \dots, m$, let e_i be the i th edge of \mathcal{D} that intersect ζ , as we move from L to R along ζ . Observe that for any i , the edges e_i, e_{i+1} belong to some triangle Δ_i of \mathcal{D} , which ζ and ζ' goes through. In particular, being shortest paths, ζ and ζ' each intersect Δ_i along a segment. In particular, let B_i be the region of Δ_i bounded by ζ and ζ' . The region B_i does not contain any vertex of \mathcal{D} in its interior, and it thus follows that the region C enclosed between ζ and ζ' (i.e., $\bigcup_i B_i$) does not contain any vertex of \mathcal{D} . Now, applying this argument to a sequence of points (p, p') that converge to the endpoint of e_R , implies the claim. \square

Remark Lemma 4 testifies that no vertex of \mathcal{D} can be interior to a strip. However, strangely enough, a strip might be pinched together by some middle vertices. To see that, visualize a terrain with saddle points (i.e., passes high in the mountains), and the strip is made out of two triangle like shapes (with e_L and e_R as their respective bases), connected by the unique path between the two extreme saddle points⁷.

The somewhat more challenging case to handle is that of pockets. A pocket is a topological disk such that its intersections with L and R are both single points, and the two boundary paths between these intersections are of length at most $2d$. The overall perimeter of a pocket is of length at most $4d$, see Figure 4.2 (I). Pockets are handled by using the recursive scheme developed for the discrete case.

4.3.2 The algorithm in detail

We use the algorithm of Section 4.3.1 to break the given disk \mathcal{D} into strips, chunks and pockets (notice, that we assume nothing on the length of the edges). Next, order the resulting regions according to their order along L , and transform each one of them at time, such that starting with L we end up with R . In each such chunk or strip, the homotopy has height (roughly) proportional to its perimeter, while for a pocket the situation is more involved.

⁷Thus, a strip might look like a dissected butterfly. Sad indeed.

(A) **Morphing a chunk/strip S :** Let $\sigma_L = L \cap S$ and $\sigma_R = R \cap S$, and let π_t and π_b be the top and bottom paths forming the two sides of S . There is a natural homotopy from $\pi_t \cdot \sigma_L$ to $\sigma_R \cdot \pi_b$.

The strip/chunk S has no vertex of \mathcal{D} in its interior, and therefore it is formed by taking planar quadrilaterals and gluing them together along common edges. Observe that by the triangle inequality, all such edges of any of these quadrilaterals are of length at most $\max(\|\sigma_L\|, \|\sigma_R\|) + 4d$. It is now easy to check that we can collapse each such quadrilateral in turn to obtain the required homotopy. Since each of π_t and π_b is composed of two shortest paths (Figure 4.2 demonstrates why such a path can potentially be made of two shortest paths), there is a linear number of such quadrilaterals. See Figure 4.2 (II) for an example.

(B) **Morphing a pocket:** We apply the algorithm of Lemma 3 recursively to a pocket.

Specifically, the above decomposes \mathcal{D} into m chunk/strips/pockets: $\mathcal{D}_1, \dots, \mathcal{D}_m$ ordered by their intersection with L . Each such disk \mathcal{D}_i has a left (resp. right) subcurve $L_i = L \cap \mathcal{D}_i$, (resp. $R_i = L \cap \mathcal{D}_i$), and similarly, it has a top curve $T_i = \mathcal{D}_{i-1} \cap \mathcal{D}_i$ and a bottom curve $B_i = \mathcal{D}_i \cap \mathcal{D}_{i+1}$, for $i = 1, \dots, m$. In the end of the i th iteration, of this morphing process, the current curve is going to be

$$M_i = R_1 \cdot \dots \cdot R_i \cdot B_i \cdot L_{i+1} \cdot \dots \cdot L_m.$$

Specifically, at the i th iteration, the algorithm morph M_{i-1} to M_i , as described above (depending on what kind of region it is). In particular, initially $M_0 = L$ and in the end $M_m = R$. As such, this results in the desired homotopy.

4.3.3 Analysis

Why can we apply Lemma 3 to a pocket. A pocket has perimeter at most $4d$, and there is a point on its boundary, such that the distance of any point in it to this base point is at most $2d$. Indeed, the boundary of d in the worst case, is made out of four shortest paths in the original disk, and as such, its total length is at most $4d$, see Eq. (4.1). Furthermore, the distance of any point in the pocket to the apex on R , is at most $2d$.

Now, by the triangle inequality, we have that if in a topological disk \mathcal{D} all the points of \mathcal{D} are in distance at most $2d$ from some point c , then the longest edge in \mathcal{D} has length at most $4d$. Therefore, all the edges inside a pocket cannot be longer than $4d$.

Running time. The shortest paths from R to L can be computed in $O(n^3 \log n)$ time. The shortest paths inside a delta to its apex can be computed in $O(n^2 \log n)$ time. Since there is a linear number of deltas, the total running time for building the strips is $O(n^3 \log n)$.

Lemma 5 *The number of paths in Π_R is $O(|V(\mathcal{D})|)$, where $V(\mathcal{D})$ is the set of vertices of \mathcal{D} .*

In particular, the total number of parts (i.e., strips, chunks, and pockets) generated by the above decomposition is $O(|V(\mathcal{D})|)$.

Proof Let $\{\sigma_1, \sigma_2, \dots, \sigma_k\}$ be the paths in Π_R sorted by the order of their endpoints along R . Observe that these paths are geodesics and so one can assume that they are interior disjoint, or share a suffix (or a prefix). Now, if $l_i \in L$ and $r_i \in R$ are the endpoints of σ_i , for $i = 1, \dots, k$, then these endpoints are sorted along their respective curves. In particular, let \mathcal{D}_i be the disk with boundary $L[s, l_i] \cdot \sigma_{i+1} \cdot R[s, r_i]$. We have that $\mathcal{D}_1 \subseteq \mathcal{D}_2 \subseteq \dots \subseteq \mathcal{D}_k$. The signatures of σ_i and σ_{i+2} must be different as otherwise they would be consecutive. Furthermore, because of the inclusion property, if an edge or a vertex of \mathcal{D} intersects σ_i but does not intersect σ_{i+1} then it cannot intersect any later path. Therefore, every other path in Π_R can be charged to vertices or edges that are added or removed from the signature of the respective path. Since an edge or a vertex can be added at most once, and deleted at most once, this implies the desired bound on the number of paths.

The second claim follows readily by the above. □

The following bounds the quality of the morphing for a pocket or a chunk.

Lemma 6 *Consider a strip or a chunk S generated by the above partition of \mathcal{D} . Let $\sigma_L = L \cap S$ and $\sigma_R = R \cap S$. Let π_t and π_b be the top and bottom paths forming the two sides of S that do not lie on R or L .*

(A) *We have $\|\pi_b\| \leq 2d$ and $\|\pi_t\| \leq 2d$.*

(B) *If $\|\sigma_L\| > 0$ or $\|\sigma_R\| > 0$ then there is no vertex of \mathcal{D} in the interior of S .*

(C) *If $\|\sigma_L\| > 0$ or $\|\sigma_R\| > 0$ then there is a homotopy from $\pi_t \cdot \sigma_L$ to $\sigma_R \cdot \pi_b$ of height $\max(\|\sigma_L\|, \|\sigma_R\|) + 4d$. We can compute such a homotopy in linear time.*

Proof (A) If the strip was generated by the first stage of partitioning then the claim is immediate.

Otherwise, consider a delta C with an apex c . For any point $x \in L \cap C$ we claim that there is a path of length at most $2d$ to c . Indeed, consider the shortest path π_x from x to R in \mathcal{D} . If this path goes to c the claim holds immediately. Otherwise, the shortest path (that has length at most d) must cross either the top or bottom shortest path forming the boundary of C that are emanating from c . We can now modify π_x , so that after its intersection point with this




Fig. 5.1

shortest path, it follows it back to c . Clearly, the resulting path has length at most $2d$ and lies inside the resulting chunk.

(B) Follows readily from the argument of [Lemma 4](#).

(C) Immediate from the algorithm description. □

4.4 The result

Theorem 2 *Suppose that we are given a triangulated piecewise linear surface with the topology of a disk, such that its boundary is formed by two walks L and R . Then there is a continuous homotopy from L to R of height at most $\|L\| + \|R\| + O(d \log n)$, where d is the maximum geodesic distance of any point of \mathcal{D} from either L or R . This homotopy can be computed in $O(n^3 \log n)$ time, and it is a $O(\log n)$ approximation to the optimal minimum height homotopy.*

Proof The algorithm is described above. The quality of approximation (i.e., $O(\log n)$) follows from plugging in the above into the analysis of [Theorem 1](#). Indeed, the intermediate curves M_0, \dots, M_m have length at most $\|L\| + \|R\| + 2d$. The intermediate morphing of a strip or a chunk might result in a curve of length $\|L\| + \|R\| + 4d$, as can be verified easily. As such, any further expansion in the length needed is a result of the recursive morphing of a pocket, thus accounting for the additional $O(d \log n)$ term.

Note, that $\max(d/2, L, R)$ is a lower bound on the height of the optimal homotopy. □

5 Approximating the homotopic Fréchet distance

In this section, fix \mathcal{D} to be a triangulated topological disk with n faces. Let the boundary of \mathcal{D} be composed of T , R , B , and L , four internally disjoint walks appearing in clockwise order along the boundary. Also, let $t_l = L \cap T$, $b_l = L \cap B$, $t_r = R \cap T$, and $b_r = R \cap B$ [®]. See [Figure 5.1](#).

5.1 Approximating the regular Fréchet distance

5.1.1 The continuous case

Let $d_{\mathcal{F}}(T, B)$ (resp. $d_{\mathcal{H}}(T, B)$) be the regular (resp. homotopic) Fréchet distance between T and B . Clearly, $d_{\mathcal{F}}(T, B) \leq d_{\mathcal{H}}(T, B)$. The following lemma implies that the Fréchet distance can be approximated within a constant factor.

Lemma 7 *Let \mathcal{D} , n , L , T , R , and B be as above. Then, for the continuous case, one can compute, in $O(n^3 \log n)$ time, reparametrizations of T and B of width at most 2δ , where $\delta = d_{\mathcal{F}}(T, B)$.*

Proof In the following, consider \mathcal{D} to be the region bounded by these four curves L , T , R , and B . We decompose \mathcal{D} into strips, chunks and pockets using the algorithm of [Section 4.3.1](#). Let Π be the set of shortest paths from all points of T to the curve B . As in the algorithm of [Section 4.3.1](#), let Π_T be the set of all shortest paths from medial points on T to B . Arguing as in [Lemma 5](#), we have that the set Π_T is composed of a linear number of paths. The paths in Π_T do not cross and so partition \mathcal{D} into a set of regions. Each region is bounded by a portion of T , a portion of B and two paths in Π_T . A region is a *delta* if the two paths of Π_T in its boundary share a single endpoint (on T), it is a *pocket* if they share two endpoints (one on T and one on B), and it is *strip* if they share no endpoints.

Obviously, the (endpoints of the) paths in Π cover all of the vertices of T . The paths in Π also cover all of B except for the bases of deltas. Now, for each delta we compute the set of all shortest paths from the vertices of its base to its apex inside the delta. Let Π_B be the set of all such paths in all deltas. Clearly, the union of Π_B and Π_T is a set of non-crossing paths whose endpoints cover all the vertices of T and B .

The shortest path from any point of T to B is at most δ . So, all paths in Π have length at most δ . Similarly, the shortest path from a point of B to T is at most δ . Now, consider a delta C with apex c . Let b be a point on the base

[®]We use the same notation to argue about the discrete and continuous problems.

of C (and so on B). The shortest path π_b from b to T has length at most δ . Let x be the first point that π_b intersects a boundary path of C , π_C . Now, $\pi_b[b, x] \cdot \pi_C[x, c]$ has length at most 2δ and it is inside C . We conclude that all paths in Π_B have length at most 2δ .

The paths in $\Pi_B \cup \Pi_T$ decompose \mathcal{D} into strips and corridors. The left and right portions of a strip is of length at most 2δ , and its top and bottom sides have as such Fréchet distance at most 2δ from each other. Similarly, the leash can jump over a pocket from the left leash to the right leash. Doing this to all corridors and pockets, results in reparametrizations of L and R such that their maximum length of a leash for these reparametrizations are at most 2δ . This implies that the Fréchet distance is at most 2δ , and we have an explicit reparametrization that realizes this distance.

As for the running time, in $O(n^3 \log n)$ time, one can compute all shortest paths from T to the whole surface. Then one can, in $O(n^2 \log n)$ time, compute the shortest paths inside each of the linear number of deltas. It follows that the total running time is $O(n^3 \log n)$. \square

5.1.2 The discrete case

We can use a similar idea to the decomposition into deltas, pockets and strips as done in the proof of [Lemma 7](#).

Lemma 8 *Let \mathcal{D} be a triangulated topological disk with n faces, and T and B be two internally disjoint walks on the boundary of \mathcal{D} . Then, for the discrete case, one can compute, in $O(n)$ time, reparametrizations of T and B that approximate the discrete Fréchet distance between T and B . The computed reparametrizations have width at most 3δ , where δ is the maximum of the Fréchet distance between T and B , and the maximum length of an edge on*

Proof First, compute the set of shortest paths, $\Pi_T = \{\zeta_1, \zeta_2, \dots, \zeta_k\}$, from vertices of T to the path B . To this end, we (conceptually) collapse all the vertices of B into a single vertex, and compute the shortest path from this meta vertex to all the vertices in T . Let \mathcal{T} be the resulting shortest path tree.

Next, for $i = 1, \dots, k - 1$, let $e_i = t_i t_{i+1}$ be the i th edge of T , and let ζ_i be the shortest path from t_i to B , with b_i being its endpoint on B , and consider the region \mathcal{D}_i bounded by the curve $\zeta_i \cdot e_i \cdot \zeta_{i+1} \cdot B[b_{i+1}, b_i]$. Now, compute the shortest path tree \mathcal{T}_i inside \mathcal{D}_i , from the two vertices of e_i to all the other vertices of \mathcal{D}_i . For each internal vertex v of $B[b_{i+1}, b_i]$, the shortest path to either t_i or t_{i+1} inside \mathcal{D}_i can be retrieved from \mathcal{T}_i . Let Π_i be the set of all such shortest paths for internal vertices of $B[b_{i+1}, b_i]$, and let $\Pi = \Pi_T \cup \bigcup_i \Pi_i$.

As for the length of the paths in Π , observe that the shortest path ψ , in \mathcal{D} , from such a vertex v to T has length at most δ . If ψ wanders outside \mathcal{D}_i then one can modify it to lie in \mathcal{D}_i . Specifically, if this path intersect, say, ζ_i then we can modify it into a path from v to t_i , and the modified path has length $\leq \|\psi\| + \|\zeta_i\| \leq 2\delta$.

Now, every edge of T or B must be used by a valid leash sequence, see [Definition 2](#). As such, the height of any leash sequence is at least the length of the longest such edge. Note, that two consecutive paths in Π might be either share an endpoint or adjacent, in either the top or bottom curve. As such, the set Π can be turned into a valid leash sequence by adding at most two moves, in the worst case both a person and a dog move, between two such consecutive paths. Let Π' denote the resulting leash sequence. Now, every path in Π' has length at most $2\delta + \delta$, as the modified added paths are longer by at most the length of a single edge of T or B . Thus, the leash sequence Π' has height at most 3δ .

Using the algorithm of Henzinger *et al.* [27] to compute the shortest paths from B takes linear time. Since all the regions are disjoint, and every edge appears on the boundary of at most two regions, we can compute all the shortest paths inside all these regions to T in $O(n)$ time overall (this step requires careful implementation to achieve this running time). \square

Remark 2 (A) The paths realizing the Fréchet distance computed by [Lemma 8](#) are stored using an implicit data-structure (essentially shortest path trees that are intertwined). This is why the space used is linear and why it can be constructed in linear time. Of course, an explicit representation of the sequence of walks realizing the Fréchet distance might require quadratic space in the worst case.

(B) We emphasize that two consecutive paths of Π' , from the proof of [Lemma 8](#), might enclose a region that have (potentially) many interior vertices. Thus, the leash might “jump” over obstacles – the remainder of this section deals with removing this drawback.

5.2 Minimum reparameterization width if there are no mountains

The following lemma implies a $O(\log n)$ -approximation algorithm for the case that all vertices in \mathcal{D} are sufficiently close to both of the two curves.

Lemma 9 *Let \mathcal{D} be a triangulated topological disk with n faces, and T and B be two internally disjoint walks on the boundary of \mathcal{D} . Further, assume for all $p \in \mathcal{D}$, the distance between p and T is at most x , and the distance between p*




Fig. 5.2

and B is at most x . Then one can compute reparametrization of B of width $O(x \log n)$. The running time is $O(n^4 \log n)$ (resp. $O(n^2)$) in the continuous (resp. discrete) case.

In particular, if $x = O(d_{\mathcal{H}}(\mathbb{T}, \mathbb{B}))$ then this is an $O(\log n)$ -approximation to the optimal homotopic Fréchet distance.

Proof Consider the continuous case. Using the algorithm of Lemma 7 we compute a reparametrization of B of width δ , realizing approximately the regular Fréchet distance, where $\delta = O(x)$. Let $\ell(t)$ denote the leash at time t that we obtain from the reparametrization mentioned above. Note that the leash $\ell(\cdot)$ is not required to deform continuously in t . In particular, for a given time $t \in [0, 1]$, let $\ell^-(t) = \lim_{t' \rightarrow t^-} \ell(t')$ and $\ell^+(t) = \lim_{t' \rightarrow t^+} \ell(t')$, where $\lim_{t' \rightarrow t^-}$ and $\lim_{t' \rightarrow t^+}$ are the left-sided and right-sided limits, respectively. By definition, the leash is discontinuous at t if and only if $\ell^-(t) \neq \ell^+(t)$.

Naturally, the above reparameterization can be used as long as it is continuous. Whenever the leash jumps over a gap (i.e., the leash is discontinuous at this point in time), say at time t , we are going to replace this jump by a $(\ell^-(t), \ell^+(t))$ -homotopy between the two leashes. Clearly, this would result in the desired continuous homotopy.

To this end, observe that all the vertices inside the disk with boundary $\ell^-(t) \cdot \ell^+(t)$ have distance $O(x)$ to \mathbb{T} and B , and thus also to $\ell^-(t)$ and $\ell^+(t)$. Hence, using the algorithm of Theorem 2, compute an $(\ell^-(t), \ell^+(t))$ -homotopy with height $O(x \log n)$. Since a gap must contain a vertex there are $O(n)$ gaps, so this filling in is done at most $O(n)$ times. Computing the initial reparameterization takes $O(n^3 \log n)$ time. Each gap can be filled in $O(n^3 \log n)$ time.

The discrete case is similar. The Fréchet distance here can be computed in linear time using the algorithm of Lemma 8 (see also Remark 2). However, we can only obtain the value of the Fréchet distance as well as an implicit representation of the actual deformation in linear time. Indeed we can compute an explicit listing of the paths in $O(n^2)$ time. Each path in the list can be charged to a single face or edge of \mathcal{D} . It immediately follows that the number of paths is linear. For any two consecutive paths, π_i and π_{i+1} in the list, we can fill in the possible gap and compute the explicit solution in $O(n_i^2)$ time, where n_i is the number of faces between π_i and π_{i+1} , see Theorem 1 and Remark 1 (D). Since $\sum n_i = O(n)$ the total running time of the algorithm is $O(n^2)$. \square

The above lemma demonstrates that if the starting and ending leashes are known (i.e., the region of the disk \mathcal{D} swept over by the morph) then an approximation algorithm can be obtained. The challenge is that a priori, we do not know these two leashes, as the input is a topological disk \mathcal{D} with the two curves \mathbb{T} and \mathbb{B} on its boundary, and the start/end leashes might be curves that lie somewhere in the interior of \mathcal{D} .

5.3 A Decision Procedure for the Homotopic Fréchet distance in the presence of mountains

Here, we are handling both the discrete and continuous cases together.

For a parameter $\tau \geq 0$, a vertex $v \in V(\mathcal{D})$ is τ -tall if its distance to \mathbb{T} or \mathbb{B} is larger than τ (intuitively τ is a guess for the value of $d_{\mathcal{H}}(\mathbb{T}, \mathbb{B})$). Here, we consider the case where there are τ -tall vertices. Intuitively, one can think about tall vertices as insurmountable mountains. Thus, to find a good homotopy between \mathbb{T} and \mathbb{B} , we have to choose which “valleys” to use (i.e., what homotopy class the solution we compute belongs to if we think about tall vertices as punctures in the disk). As a concrete example, consider Figure 5.2 (I), where there are three tall vertices, and two possible solutions are being shown.

In the discrete case, we subdivide each edge in the beginning so that if an edge has length $> 2\tau$, then the vertex inserted in the middle of it is τ -tall. Observe that, if $\tau > d_{\mathcal{H}}(\mathbb{T}, \mathbb{B})$ then no leash of the optimum homotopic motion can afford to contain a τ -tall vertex. We use M^τ to denote the set of all τ -tall vertices in $V(\mathcal{D})$.

Now, let ω and ω' be two walks connecting points on \mathbb{T} and \mathbb{B} . The walks ω and ω' are *homotopic* in $\mathcal{D} \setminus M^\tau$ if and only if they are homotopic in $\mathcal{D} \setminus M^\tau$ after contracting \mathbb{T} and \mathbb{B} (each to a single point). Two *non-crossing* walks ω and ω' are homotopic if and only if $\mathbb{T} \cdot \mathbb{B} \cdot \omega \cdot \omega'$ contains no tall vertices. It is straightforward to check that homotopy is an equivalence relation. So it partitions (\mathbb{T}, \mathbb{B}) -paths into *homotopy classes*; we call each class a τ -homotopy class or simply a homotopy class (given that τ is fixed).




Fig. 5.3

For a homotopy class h , let $\pi_{L,h}$ (resp. $\pi_{R,h}$) be the *left geodesic* (resp. *right geodesic*); that is, $\pi_{L,h}$ denotes the shortest path in h from t_l to b_l (resp. from t_r to b_r).

Let ω be any walk in h from $b \in B$ to $t \in T$. The *left tall set* of h , denoted by $M_l(h) = M_l(\omega)$, is the set of all τ -tall vertices to the left of ω . Namely, $M_l(h)$ is the set of tall vertices inside the disk with boundary $L \cdot T[t_l, t_r] \cdot \omega \cdot B[b_l, b_r]$, where L is the “left” portion of the boundary of \mathcal{D} , having endpoints t_l and b_l . We similarly define the *right tall set* of h , $M_r(h) = M_r(\omega)$, to be the set of all τ -tall vertices to the right of ω . See Figure 5.2 (II).

Note that the sets $M_l(h)$ and $M_r(h)$ do not depend on the particular choice of ω , since all paths in h are homotopic and so have the same set of τ -tall vertices to their left and right side. However, we emphasize that the left and right tall sets do not identify homotopy classes. Figure 5.2 (III) demonstrates two non-homotopic paths with identical left and right tall sets.

The set h is *τ -extendable* from the left if and only if $\|\pi_{L,h}\| \leq \tau$ and there is a homotopy class h' , such that $\|\pi_{L,h'}\| \leq \tau$ and $M_l(h) \subset M_l(h')$. In particular, h is *τ -saturated* if it is not τ -extendable and $\|\pi_{L,h}\| \leq \tau$.

5.3.1 On the left and right geodesics

Lemma 10 *Let h be a τ -saturated homotopy class, where $\tau \geq d_{\mathcal{H}}(T, B)$. Then $\|\pi_{R,h}\| \leq 4\tau$.*

Proof Let h_{opt} be the homotopy class of the leashes in the optimum solution. Of course, no leash in the optimum solution contains a τ -tall vertex. Further, all leashes in the optimal solution are homotopic because there is a homotopy that contains all of them by definition.

Since h is saturated the set $M_l(h)$ is not a proper subset of $M_l(h_{\text{opt}})$. It follows that either $M_l(h) = M_l(h_{\text{opt}})$ or $M_l(h)$ intersects $M^{\tau} \setminus M_l(h_{\text{opt}}) = M_r(h_{\text{opt}})$.

If $M_l(h) = M_l(h_{\text{opt}})$ then either $h = h_{\text{opt}}$, and in particular $\|\pi_{R,h}\| = \|\pi_{R,h_{\text{opt}}}\| \leq \tau$, or $\pi_{L,h}$ crosses $\pi_{R,h_{\text{opt}}}$.

Otherwise, the set $M_l(h) \cap M_r(h_{\text{opt}})$ is not empty. Again, it follows that $\pi_{L,h}$ crosses $\pi_{R,h_{\text{opt}}}$.

Therefore, we only need to address the case that $\pi_{L,h}$ crosses $\pi_{R,h_{\text{opt}}}$.

Let x be the first intersection point between $\pi_{L,h}$ and $\pi_{R,h_{\text{opt}}}$, as one traverses $\pi_{L,h}$ from t_l to b_l . Let x' be the last intersection point of $\pi_{L,h}[t_l, x]$ with $\pi_{L,h_{\text{opt}}}$. Similarly, y is the last intersection point between $\pi_{L,h}$ and $\pi_{R,h_{\text{opt}}}$, and y' is the first intersection of $\pi_{L,h}[y, b_l]$ and $\pi_{L,h_{\text{opt}}}$. Observe that the interiors of $\pi_{L,h}[x', x]$ and $\pi_{L,h}[y, y']$ do not intersect the curves $\pi_{L,h_{\text{opt}}}$ and $\pi_{R,h_{\text{opt}}}$. See Figure 5.3 (I).

As the curves $\pi_{L,h}$ and $\pi_{R,h}$ are homotopic (by definition), the disk with the boundary $T \cdot \pi_{L,h} \cdot B \cdot \pi_{R,h}$ does not contain any tall vertex, and $T \cdot \pi_{L,h} \cdot B$ is homotopic to $\pi_{R,h}$.

Consider the walk $T' = \pi_{R,h_{\text{opt}}}[t_r, x] \cdot \pi_{L,h}[x, x'] \cdot \pi_{L,h_{\text{opt}}}[x', t_l]$, see Figure 5.3 (II). The walk T' is homotopic to T . Similarly, $B' = \pi_{L,h_{\text{opt}}}[b_l, y'] \cdot \pi_{L,h}[y', y] \cdot \pi_{R,h_{\text{opt}}}[y, b_r]$ is homotopic to B . It follows that $\pi_{R,h}$ is homotopic to $T' \cdot \pi_{L,h} \cdot B'$. As $\pi_{R,h}$ is the shortest path in its homotopy class with these endpoints, it follows that

$$\|\pi_{R,h}\| \leq \|T' \cdot \pi_{L,h} \cdot B'\| \leq \|\pi_{L,h}\| + (\|\pi_{L,h_{\text{opt}}}\| + \|\pi_{L,h}\| + \|\pi_{R,h_{\text{opt}}}\|) \leq 4\tau,$$

as T' and B' are disjoint, and $T' \cup B' \subseteq \pi_{R,h_{\text{opt}}} \cup \pi_{L,h_{\text{opt}}} \cup \pi_{L,h}$. \square

A region that contains no τ -tall vertices can still, potentially, contain τ -tall points (that are not vertices) on its edges or faces. We next prove that this does not happen in our setting.

Lemma 11 *For any $\tau \geq 0$, let h be a τ -homotopy class, such that $\max(\|\pi_{L,h}\|, \|\pi_{R,h}\|) \leq x$, where $x \geq \tau \geq d_{\mathcal{H}}(T, B)$. Let \mathcal{D}' be the disk with boundary $T \cdot \pi_{R,h} \cdot B \cdot \pi_{L,h}$. Then all the points inside \mathcal{D}' are within distance $O(x)$ to both T and B in \mathcal{D}' .*

Proof We first consider the continuous case. By the definition of τ -homotopy, the disk \mathcal{D}' has no τ -tall vertices. Furthermore, by the definition of x , we have that the distance of any point on T to B , restricted to paths in \mathcal{D}' is at most δ_1 , where $\delta_1 = x + d_{\mathcal{F}}(T, B) \leq 2x$. Indeed, the shortest path from any point on T to B in \mathcal{D}' , either stays inside \mathcal{D}' , or alternatively intersects either $\pi_{L,h}$ or $\pi_{R,h}$.

We can now deploy the decomposition of \mathcal{D}' into strips, pockets and chunks as done in [Section 4.3.1](#). Every strip (or a chunk) is being swept by a leash of length at most $\delta_2 = 2\delta_1 \leq 4x$ (the factor two is because a strip might rise out of a delta), and therefore the claim trivially holds for points inside such regions.

Every pocket P has perimeter of length at most $\|\partial P\| \leq \delta_3 = 2\delta_2 = 8x$ (the perimeter also contains two points of T and B and they are in distance at most δ_2 from each other in either direction along the perimeter). So, consider such a pocket P . Since \mathcal{D}' contains no τ -tall vertices, P does not contain any tall vertex. Let e be an edge in P (or a subedge if it intersects the boundary of P). The two endpoints of e are in P , and such an endpoint is either a (not tall) vertex or it is contained in ∂P . In either case, these endpoints are in distance at most x from ∂P , and so they are in distance at most $\delta_4 = 2x + \|\partial P\|/2 = 2x + \delta_2 \leq 6x$ from each other even if the geodesic distance is restricted to P . We conclude that $\|e\| \leq \delta_4$, and consequently, any point in e is in distance at most $\delta_5 = \|e\|/2 + x + \delta_2 \leq 3x + x + 8x \leq 12x$ from T and B .

Now, consider any point p in P , and consider the face F that contains it. Since the surface is triangulated, F is a triangle. Clipping F to P results in a planar region F' that has perimeter at most $\delta_6 = 3\delta_4 + \|\partial P\| \leq 3 \cdot 6x + \delta_3 \leq (18+8)x \leq 26x$ (note, that an edge might be fragmented into several subedges, but the distance between the furthest two points along a single edge is at most δ_4 using the same argument as above). Thus, the furthest a point of P can be from an edge of P is at most $\delta_7 = \delta_6/2\pi \leq 5x$. Hence, the maximum distance of a point of P from either T or B (inside \mathcal{D}') is at most $\delta_5 + \delta_7 \leq 12x + 5x = 17x$.

The discrete case is easy. Any edge of length $\geq 2\tau$ was split, by introducing a middle vertex, which must be τ -tall. So the claim immediately holds. \square

5.3.2 The decision algorithm

Lemma 12 *Let \mathcal{D} , n , T , L , B , R , t_l , b_l as in the first paragraph of [Section 5](#) and τ as in the previous subsection, and let $X \subseteq V(\mathcal{D})$ be a set of τ -tall vertices. Consider the shortest path σ_l (between t_l and b_l) that belongs to any homotopy class h such that $X \subseteq M_l(h)$. Then the path σ_l can be computed in $O(n^4 \log n)$ (resp. $O(n \log n)$) time in the continuous (resp. discrete) case.*

Proof For each vertex of $v \in X$, compute its shortest path ψ_v to L in \mathcal{D} . Cut the disk \mathcal{D} along these paths. The result is a topological disk \mathcal{D}' . Compute the shortest path ζ in \mathcal{D}' between t_l and b_l .

We claim that $\zeta = \sigma_l$. To this end, consider σ_l and any path ψ_v computed by the algorithm. We claim that σ_l and ψ_v do not cross in their interior. Indeed, if σ_l cross ψ_v an odd number of times, then v is inside the disk $\sigma_l \cdot T \cdot R \cdot B$, which contradicts the condition that $v \in X \subseteq M_l(h)$. Clearly, σ_l and ψ_v cannot cross in their interiors more than once, because otherwise, one can shorten one of them, which is a contradiction as they are both shortest paths. Thus, σ_l is a path in \mathcal{D}' connecting t_l to b_l , thus implying that ζ is σ_l .

As for the running time, each shortest path computation takes time $O(n^2 \log n)$, in the continuous (resp. discrete) case. The resulting disk has complexity $O(n^2)$, and computing a shortest path in it takes $O(n^4 \log n)$ time in the continuous case. In the discrete case, computing the paths can be done by collapsing L to a vertex, forbid the shortest path tree edges, and run a shortest path algorithm in the remaining graph. Clearly, this takes $O(n \log n)$ time. \square

Lemma 13 *Let \mathcal{D} be a triangulated topological disk with n faces, and T and B be two internally disjoint walks on \mathcal{D} 's boundary. Given $\tau > 0$, one can compute a τ -saturated homotopy class, in $O(n^5 \log n)$ (resp. $O(n^2 \log n)$) time, in the continuous (resp. discrete) case.*

Proof Start with an empty initial set $X = \emptyset$. At each iteration, try adding one of the τ -tall vertices $v \in M^\tau$ of \mathcal{D} to X , by using [Lemma 12](#). The algorithm of [Lemma 12](#) outputs a path σ between t_l and b_l and a set $X' \supset X \cup \{v\}$.

If σ is of length at most τ update X to be the new set X' , otherwise reject v . If v is rejected then the left geodesic of any superset of $X \cup \{v\}$ has length larger than τ . It follows that v cannot be accepted in any later iteration, so we do not need to reinspect it. Clearly, after trying all the vertices of M^τ , the set X defines the desired saturated class, which can be computed by using the algorithm of [Lemma 12](#). \square

Lemma 14 *Let \mathcal{D} be a triangulated topological disk with n faces, and T and B be two internally disjoint walks on the boundary of \mathcal{D} . Given a real number $x > 0$, one can either:*

- (A) *Compute a homotopy from T to B of width $O(x \log n)$, or*
- (B) *Return that $x < d_{\mathcal{H}}(T, B)$.*

The running time of this procedure is $O(n^5 \log n)$ (resp. $O(n^2 \log n)$) in the continuous (resp. discrete) case.

Proof Assume $x \geq \delta_H = d_{\mathcal{H}}(T, B)$, and we use x as a guess for this value δ_H . Using [Lemma 13](#), one can compute a x -saturated homotopy class, h . [Lemma 10](#) implies that both $\pi_{L,h}$ and $\pi_{R,h}$ are at most $4x$. Let $\mathcal{D}' \subseteq \mathcal{D}$ be the disk with boundary $T \cdot \pi_{L,h} \cdot B \cdot \pi_{R,h}$. By [Lemma 11](#), all the vertices in \mathcal{D}' are in distance $O(x)$ from T and B (this holds for all points in \mathcal{D}' in the continuous case). That is, there are no $O(x)$ -tall vertices in \mathcal{D}' . Finally, [Lemma 9](#) implies that a continuous leash sequence of height $\leq Z = O(x \log n)$ between T and B , inside \mathcal{D}' , can be computed.




Fig. 5.4

Thus, if x is larger than $d_{\mathcal{H}}(\mathsf{T}, \mathsf{B})$ then this algorithm returns the desired approximation; that is, a homotopy of width $\leq Z$. If the width of the generated homotopy is however larger than Z (a value that can be computed directly from x), then the value of x was too small. That is, the algorithm fails in this case only if $x < d_{\mathcal{H}}(\mathsf{T}, \mathsf{B})$. In the case of such failure, return that x is too small. \square

5.4 A strongly polynomial approximation algorithm

For a vertex $v \in \mathsf{V}(\mathcal{D})$, define $\text{cost}(v)$ to be the length of the shortest path between t_l and b_l that has v on its left side. Similarly, for a set of vertices $X \subseteq \mathsf{V}(\mathcal{D})$, let $\text{Cost}(X)$ be the length of the shortest path between t_l and b_l that has X on its left side. For a specific v or X , one can compute $\text{cost}(v)$ and $\text{Cost}(X)$ by invoking the algorithm of [Lemma 12](#) once.

5.4.1 The algorithm

- (I) **Identifying the tall vertices.** Observe that using the algorithm of [Lemma 14](#), we can decide given a candidate value δ_H for $d_{\mathcal{H}}(\mathsf{T}, \mathsf{B})$ if it is too large, too small, or leads to the desired approximation. Indeed, if the algorithm returns an approximation of values $O(\delta_H \log n)$ but fails for $\delta_H/2$, we know it is the desired approximation. For each vertex $v \in \mathsf{V}(\mathcal{D})$ let α_v be the maximum distance of v to either T or B . Note that v cannot be a -tall for any $a \geq \alpha_v$. Sort these values, and using binary search, compute the vertex w , with the minimum value α_w , such that [Lemma 14](#) returns a parameterization with homotopic Fréchet distance $O(\alpha_w \log n)$. If the algorithm of [Lemma 14](#) returns that α_w/n is too small of a guess, then $[\alpha_w/n, \alpha_w \log n]$ contains δ_H . In this case, we can use binary search to find an interval $[\gamma/2, \gamma]$ that contains δ_H and use [Lemma 14](#) to obtain the desired approximation. Similarly, if v is the tallest vertex shorter than w , then we can assume that $\alpha_v n$ is too small of a guess, otherwise we are again done as $[\alpha_v, \alpha_v n]$ contains δ_H . Therefore, in the following, we know that the desired distance δ_H lies in the interval $[x, y]$ where $x = \alpha_w/n$ and $y = \alpha_w \log n$, and for every vertex u of \mathcal{D} it holds that (i) $\alpha_u \leq x/n$, or (ii) $\alpha_u \geq yn$. Naturally, we consider all the vertices that satisfy (ii) as tall vertices, by setting $\tau = 2x/n$. In the following, let M denote the set of these τ -tall vertices.
- (II) **Computing candidate homotopy classes.** Start with $X_0 = \emptyset$. In the i th iteration, the algorithm computes the vertex $v_i \in \mathsf{M} \setminus X_{i-1}$, such that $\text{Cost}(X_{i-1} \cup \{v_i\})$ is minimized, and set $X_i = X_{i-1} \cup \{v_i\}$. Let h_i be the homotopy class having X_i on its left side, and $\mathsf{M} \setminus X_i$ on its right side.
- (III) **Binary search over candidates.** We approximate the homotopic Fréchet width of each one of the classes h_1, \dots, h_n . Let x be the minimum homotopic Fréchet width computed among these n candidates. Next, do a binary search in the interval $[x/n^2, x]$ for the homotopic Fréchet distance. We return the smallest width reparametrization computed as the desired approximation.

5.4.2 Analysis

Lemma 15 (i) For any $X' \subseteq X \subseteq \mathsf{V}(\mathcal{D})$, we have $\text{Cost}(X') \leq \text{Cost}(X)$.

(ii) For any $x \in X \subseteq \mathsf{V}(\mathcal{D})$, we have $\text{cost}(x) \leq \text{Cost}(X)$.

(iii) For $X, Y \subseteq \mathsf{V}(\mathcal{D})$, we have that $\text{Cost}(X \cup Y) \leq \text{Cost}(X) + \text{Cost}(Y)$.

Proof (i) Observe that the path realizing $\text{Cost}(X')$ is less constrained than the path realizing $\text{Cost}(X)$, therefore it might only be shorter.

(ii) Follows immediately from (i).

(iii) Consider the disk \mathcal{D} and the two paths σ_X and σ_Y realizing $\text{Cost}(X)$ and $\text{Cost}(Y)$, respectively. The close curves $\sigma_X \cdot \mathsf{L}$ and $\sigma_Y \cdot \mathsf{L}$ encloses two topological disks. Consider the union of these two disks, and its connected outer

boundary $\sigma_{XUY} \cup L$. Clearly, σ_{XUY} connects \mathbf{t}_i and \mathbf{b}_i , and it has all the points of X and Y on one side of it, and finally $\|\sigma_{XUY}\| \leq \|\sigma_X\| + \|\sigma_Y\|$ as $\sigma_{XUY} \subseteq \sigma_X \cup \sigma_Y$. See [Figure 5.4](#).

Lemma 16 *The cheapest homotopic Fréchet parameterization computed among $\mathbf{h}_1, \dots, \mathbf{h}_n$ has width $O(d_{\mathcal{H}}(\mathbf{T}, \mathbf{B})n \log n)$.*

Proof Consider the set Y that is the subset of tall vertices on the left side of the optimal solution. Let i be the first index such that $Y \subseteq X_i$ and $Y \not\subseteq X_{i-1}$. Let v be any vertex in $Y \setminus X_{i-1}$. By construction, we have that $\text{Cost}(X_i) \leq \text{Cost}(X_{i-1} \cup \{v\})$, and furthermore, for all $j \leq i$, we have that $\text{Cost}(X_j) \leq \text{Cost}(X_{j-1} \cup \{v\})$, by the greediness in the construction of X_1, \dots, X_i . Now, we have

$$\begin{aligned}
\text{Cost}(X_i) &\leq \text{Cost}(X_{i-1} \cup \{v\}) && \text{(by construction of } X_i) \\
&\leq \text{Cost}(X_{i-1}) + \text{cost}(v) && \text{(by Lemma 15 (iii))} \\
&\leq \text{Cost}(X_{i-1}) + \text{Cost}(Y) && \text{(by Lemma 15 (ii))} \\
&\leq (\text{Cost}(X_{i-2}) + \text{Cost}(Y)) + \text{Cost}(Y) && \text{(applying same argument to } X_{i-1}) \\
&= \text{Cost}(X_{i-2}) + 2\text{Cost}(Y) \\
&\leq \dots \leq i\text{Cost}(Y) \leq n\text{Cost}(Y).
\end{aligned}$$

Now, setting $\tau = \text{Cost}(X_i)$, it follows that X_i is τ -saturated. Applying [Lemma 10](#), implies that $\|\pi_{\mathbf{R}, \mathbf{h}_i}\| \leq 4\tau$. Observe, that the disk defined by $\mathbf{T}, \pi_{\mathbf{L}, \mathbf{h}_i}, \mathbf{B}, \pi_{\mathbf{R}, \mathbf{h}_i}$ cannot contain any tall vertex (by construction).

Now, plugging this into [Lemma 9](#) implies the homotopic Fréchet width of \mathbf{h}_i (starting with $\pi_{\mathbf{L}, \mathbf{h}_i}$ and ending up with $\pi_{\mathbf{R}, \mathbf{h}_i}$, so \mathcal{D} in [Lemma 9](#) is bounded by $\mathbf{T}, \mathbf{B}, \pi_{\mathbf{L}, \mathbf{h}_i}, \pi_{\mathbf{R}, \mathbf{h}_i}$) is $O(\tau \log n)$, which implies the claim since $\text{Cost}(X_i) \leq n\text{Cost}(Y) \leq nd_{\mathcal{H}}(\mathbf{T}, \mathbf{B})$. \square

5.4.3 The algorithm

Theorem 3 *Let \mathcal{D} be a triangulated topological disk with n faces, and \mathbf{T} and \mathbf{B} be two internally disjoint walks on the boundary of \mathcal{D} . One can compute a homotopic Fréchet parameterization of \mathbf{T} and \mathbf{B} of width $O(d_{\mathcal{H}}(\mathbf{T}, \mathbf{B}) \log n)$, where $d_{\mathcal{H}}(\mathbf{T}, \mathbf{B})$ is the homotopic Fréchet distance between \mathbf{T} and \mathbf{B} in \mathcal{D} .*

The running time of this procedure is $O(n^6 \log n)$ (resp. $O(n^3 \log n)$) in the continuous (resp. discrete) case.

Proof Consider the algorithm described in the previous subsection. For correctness, observe that the algorithm either found the desired value, or identified correctly the tall vertices. Next, by [Lemma 16](#), the range the algorithm searches over contains the desired value.

The algorithm requires $O(n^2)$ calls to [Lemma 12](#), which takes $O(n^6 \log n)$ (resp. $O(n^3 \log n)$) time in the continuous (resp. discrete) case. Then the algorithm requires the method of [Lemma 9](#) to compute the homotopic Fréchet distance of the classes $\mathbf{h}_1, \dots, \mathbf{h}_n$. The algorithm also performs $O(\log n)$ calls to the algorithm of [Lemma 14](#). \square

6 Conclusions

We presented a $O(\log n)$ approximation algorithm for approximating the homotopy height and the homotopic Fréchet distance between curves on piecewise linear surfaces. It seems quite believable that the approximation quality can be further improved, and we leave this as the main open problem. Since our algorithm works both for the continuous and discrete cases, it seems natural to conjecture that this algorithm should also work for more general surfaces and metrics.

Another problem for further research is to solve our main problem without the restriction that the two curves lie on the boundary of the disk.

Connection to planar separator. Our basic algorithm ([Theorem 1](#)) is inspired to some extent by the proof of the planar separator theorem [\[31\]](#). In particular, our result implies sufficient conditions to having a separator that can continuously deform from enclosing nothing in a planar graph, till it encloses the whole graph, without being too long at any point in time. As a result, our work can be viewed as extending the planar separator theorem. A natural open problem is to extend our work to graphs with higher genus.

Acknowledgments The authors thank Jeff Erickson and Gary Miller for their comments and suggestions. The authors also thank the anonymous referees for their detailed and insightful reviews.

References

1. P. K. Agarwal, B. Aronov, J. O'Rourke, and C. A. Schevon. Star unfolding of a polytope with applications. *SIAM J. Comput.*, 26:1679–1713, 1997.
2. H. Alt and M. Buchin. Semi-computability of the Fréchet distance between surfaces. In *Proc. 21st Euro. Workshop on Comput. Geom.*, pages 45–48, 2005.
3. H. Alt and M. Godau. Computing the Fréchet distance between two polygonal curves. *Internat. J. Comput. Geom. Appl.*, 5:75–91, 1995.
4. C. Bennis, J.-M. Vézien, G. Iglésias, and A. Gagalowicz. Piecewise surface flattening for non-distorted texture mapping. In Thomas W. Sederberg, editor, *Proc. SIGGRAPH '91*, volume 25, pages 237–246, 1991.
5. S. Brakatsoulas, D. Pfoser, R. Salas, and C. Wenk. On map-matching vehicle tracking data. In *Proc. 31st VLDB Conference*, pages 853–864. VLDB Endowment, 2005.
6. G. A. Brightwell and P. Winkler. Submodular percolation. *SIAM J. Discret. Math.*, 23(3):1149–1178, 2009.
7. K. Buchin, M. Buchin, and J. Gudmundsson. Detecting single file movement. In *Proc. 16th ACM SIGSPATIAL Int. Conf. Adv. GIS*, pages 288–297, 2008.
8. K. Buchin, M. Buchin, J. Gudmundsson, Maarten L., and J. Luo. Detecting commuting patterns by clustering subtrajectories. In *Proc. 19th Annu. Internat. Sympos. Algorithms Comput. (ISAAC)*, pages 644–655, 2008.
9. M. Buchin, A. Driemel, and B. Speckmann. Computing the Fréchet distance with shortcuts is np-hard. *CoRR*, abs/1307.2097, 2013.
10. E. W. Chambers, E. Colin de Verdière, J. Erickson, S. Lazard, F. Lazarus, and S. Thite. Homotopic fréchet distance between curves or, walking your dog in the woods in polynomial time. *Comput. Geom. Theory Appl.*, 43(3):295–311, 2010.
11. E. W. Chambers and D. Letscher. On the height of a homotopy. In *Proc. 21st Canad. Conf. Comput. Geom. (CCCG)*, 2009.
12. E. W. Chambers and D. Letscher. Erratum for on the height of a homotopy. <http://mathcs.slu.edu/~chambers/papers/herratum.pdf>, 2010.
13. E. W. Chambers, D. Letscher, T. Ju, and L. Liu. Isotopic Fréchet distance. In *Proc. 23rd Canad. Conf. Comput. Geom. (CCCG)*, 2011.
14. E. W. Chambers and Y. Wang. Measuring similarity between curves on 2-manifolds via homotopy area. In *Proc. 29th Annu. Sympos. Comput. Geom. (SoCG)*, pages 425–434, 2013.
15. G. R. Chambers and R. Rotman. Contracting loops on a Riemannian 2-surface. *ArXiv e-prints*, November 2013.
16. A. F. Cook, A. Driemel, S. Har-Peled, J. Sherette, and C. Wenk. Computing the Fréchet distance between folded polygons. In *Proc. 12th Workshop Algorithms Data Struct. (WADS)*, pages 267–278, 2011.
17. A. F. Cook and C. Wenk. Geodesic Fréchet distance inside a simple polygon. *ACM Trans. Algo.*, 7:9:1–9:19, 2010.
18. A. F. Cook and C. Wenk. Shortest path problems on a polyhedral surface. *Algorithmica*, 2012. to appear.
19. A. Driemel, S. Har-Peled, and C. Wenk. Approximating the Fréchet distance for realistic curves in near linear time. *Discrete Comput. Geom.*, 48:94–127, 2012.
20. A. Efrat, L. J. Guibas, S. Har-Peled, J. S.B. Mitchell, and T.M. Murali. New similarity measures between polylines with applications to morphing and polygon sweeping. *Discrete Comput. Geom.*, 28:535–569, 2002.
21. T. Eiter and H. Mannila. Computing discrete Fréchet distance. Tech. Report CD-TR 94/64, Christian Doppler Lab. Expert Sys., TU Vienna, Austria, 1994.
22. M. S. Floater. Parameterization and smooth approximation of surface triangulations. *Comput. Aided Geom. Design*, 14(4):231–250, 1997.
23. M. Fréchet. Sur la distance de deux surfaces. *Ann. Soc. Polonaise Math.*, 3:4–19, 1924.
24. M. Godau. *On the complexity of measuring the similarity between geometric objects in higher dimensions*. PhD thesis, Free University of Berlin, 1999.
25. S. Har-Peled, A. Nayyeri, M. Salavatipour, and A. Sidiropoulos. How to walk your dog in the mountains with no magic leash. In *Proc. 28th Annu. Sympos. Comput. Geom. (SoCG)*, pages 121–130, 2012.
26. S. Har-Peled and B. Raichel. The Fréchet distance revisited and extended. In *Proc. 27th Annu. Sympos. Comput. Geom. (SoCG)*, pages 448–457, New York, NY, USA, 2011. ACM. <http://sarielhp.org/papers/10/frechet3d/>.
27. M. R. Henzinger, P. Klein, S. Rao, and S. Subramanian. Faster shortest-path algorithms for planar graphs. *J. Comput. Sys. Sci.*, 55:3–23, August 1997.
28. E. J. Keogh and M. J. Pazzani. Scaling up dynamic time warping to massive dataset. In *Proc. of the Third Euro. Conf. Princip. Data Mining and Know. Disc.*, pages 1–11, 1999.
29. M.S. Kim, S.W. Kim, and M. Shin. Optimization of subsequence matching under time warping in time-series databases. In *Proc. ACM symp. Applied comput.*, pages 581–586, 2005.
30. P. M. Lewis II, R. E. Stearns, and J. Hartmanis. Memory bounds for recognition of context-free and context-sensitive languages. In *Proc. 6th Annu. IEEE Sympos. Found. Comput. Sci. (FOCS)*, pages 191–202, 1965.
31. R. J. Lipton and R. E. Tarjan. A separator theorem for planar graphs. *SIAM J. Appl. Math.*, 36:177–189, 1979.
32. A. Mascaret, T. Devogele, I. Le Berre, and A. Hénaff. Coastline matching process based on the discrete Fréchet distance. In *Proc. 12th Int. Sym. Spatial Data Handling*, pages 383–400, 2006.
33. J. S.B. Mitchell, D. M. Mount, and C. H. Papadimitriou. The discrete geodesic problem. *SIAM J. Comput.*, 16:647–668, 1987.
34. P. Papasoglu. Contracting thin disks. *ArXiv e-prints*, September 2013.
35. D. Piponi and G. Borshukov. Seamless texture mapping of subdivision surfaces by model pelting and texture blending. In *Proc. SIGGRAPH 2000*, pages 471–478, August 2000.
36. J. Serrà, E. Gómez, P. Herrera, and X. Serra. Chroma binary similarity and local alignment applied to cover song identification. *IEEE Transactions on Audio, Speech & Language Processing*, 16(6):1138–1151, 2008.
37. A. Sheffer and E. de Sturler. Surface parameterization for meshing by triangulation flattening. In *Proc. 9th International Meshing Roundtable*, pages 161–172, 2000.
38. J. Sherette and C. Wenk. Simple curve embedding. *CoRR*, abs/1303.0821, 2013.
39. C. Wenk, R. Salas, and D. Pfoser. Addressing the need for map-matching speed: Localizing global curve-matching algorithms. In *Proc. 18th Int. Conf. Sci. Statis. Database Manag.*, pages 879–888, 2006.

A Sweeping a convex polytope, star unfolding, and banana peels

Consider a convex polytope \mathcal{P} in three dimensions, a base point b on its boundary, and the problem of finding the minimum length leash needed for a guard that walks on the polytope such that the leash sweeps over all the points on the surface of the polytope. Specifically,

at any point in time, the guard maintains a connection to the base point b via a path (i.e., the leash) connecting it to the base point, and the leash has to move continuously as the guard moves around.

For a point p on the boundary of \mathcal{P} , let $d_{\mathcal{P}}(p)$ be the geodesic distance from b to p (i.e., the length shortest path ζ that lies on the boundary of \mathcal{P} connecting b to p). Let M be the *medial axis* of this distance – formally, a point p is on the medial axis if there are two distinct shortest paths ζ and ψ from b to p , such that $\|\zeta\| = \|\psi\| = d_{\mathcal{P}}(p)$. It is known that M is a tree in this case [1].

Now, let Π be the union of all the shortest paths from b to the vertices of \mathcal{P} (we assume that no vertex is on the medial axis, which holds under general position assumption). The set Π is also a tree. Surprisingly, if you cut $\partial\mathcal{P}$ along Π , then the resulting polygon can be flattened on the plane. Maybe even more surprisingly, this even holds if one cuts $\partial\mathcal{P}$ along M . This is known as *star unfolding* of a polytope, see Agarwal *et al.* [1] for details.

Consider cutting $\partial\mathcal{P}$ along both M and Π . This breaks \mathcal{P} into a collection of polygons \mathcal{Q} , where each polygon $Q \in \mathcal{Q}$, has no vertices of \mathcal{P} in its interior, and has b as a vertex. As such, one can unfold this Q into the plane. Here, the two paths of Π adjacent to b that belong to the boundary of Q maps in this unfolding to two straight edges. The rest of the boundary Q is a closed connected portion of M . One can think about Q as being a “leaf” in a decomposition of $\partial\mathcal{P}$ (i.e., think about the sides of a banana peel). Here, shortest paths from b to any point on $p \in \partial\mathcal{P}$ that belongs to Q results in a straight segment in (the planar embedded version of) the polygon Q . As such, the polygons of \mathcal{Q} completely capture the structure of all the shortest paths on $\partial\mathcal{P}$ to b .

Back to the problem of sweeping $\partial\mathcal{P}$. For the points of $M_Q = M \cap \partial Q$, we can sweep the region of ∂Q that corresponds to Q , by walking along the curve M_Q (say counterclockwise), and the leash being the shortest path in $\partial\mathcal{P}$ (that lies inside Q). This completely sweeps over the region of Q . We then continue this sweeping in the next polygon of \mathcal{Q} adjacent to Q around b . We continue in this fashion till all the boundary of the polytope is swept over. Note, that the leash is moving continuously, and during this motion, the maximum length of the leash is the distance to the furthest point on $\partial\mathcal{P}$ from b . We conclude that this is an optimal solution and using the known algorithms for computing shortest paths [1].

Let us recap the algorithm: We compute the medial axis M of b on $\partial\mathcal{P}$, under the shortest path distance on the boundary of the polytope. Next, the parameterize a point $p(t)$ to move continuously around the tree M (i.e., traversing along each edge twice, in both direction). At each point in time, the leash is connected via the shortest path to the base point b .

Lemma 17 *Given a convex polytope \mathcal{P} in three dimensions, and a base point $p \in \partial\mathcal{P}$, one can compute in polynomial time, a continuous motion of a point $p(t)$, $t \in [0, 1]$, and an associated leash $\ell(t)$ connecting $p(t)$ with b , such that (i) the leash sweeps over all the points of $\partial\mathcal{P}$, (ii) the leash moves continuously, (iii) a point of $\partial\mathcal{P}$ get swept over only once during this motion, and (iv) the maximum length of the leash is $\max_{p \in \partial\mathcal{P}} d_{\mathcal{P}}(p)$, which is optimal.*

The above is in sharp contrast to our original problem of computing the homotopy height, as the two ends of the leash must move along two prespecified curves L and R . Furthermore, because of that we no longer have the property that the leashes do not jump, as the leash head no longer moves along the medial axis, as the resulting paths might be too long (compared to the optimal morph). Nevertheless, the above captures our basic strategy of breaking the input disk into smaller slivers, induced by shortest paths, and solving the problem in each sliver separately, and gluing the solutions together.