

Can We Compute the Similarity between Surfaces?

Helmut Alt · Maïke Buchin

Received: 2 January 2008 / Revised: 23 February 2009 / Accepted: 26 February 2009 /
Published online: 17 March 2009
© Springer Science+Business Media, LLC 2009

Abstract A suitable measure for the similarity of shapes represented by parameterized curves or surfaces is the Fréchet distance. Whereas efficient algorithms are known for computing the Fréchet distance of polygonal curves, the same problem for triangulated surfaces is NP-hard. Furthermore, it remained open whether it is computable at all.

Using a discrete approximation, we show that it is *upper semi-computable*, i.e., there is a non-halting Turing machine which produces a decreasing sequence of rationals converging to the Fréchet distance. It follows that the decision problem, whether the Fréchet distance of two given surfaces lies below a specified value, is recursively enumerable.

Furthermore, we show that a relaxed version of the Fréchet distance, the *weak Fréchet distance* can be computed in polynomial time. For this, we give a computable characterization of the weak Fréchet distance in a geometric data structure called the *Free Space Diagram*.

Keywords Shape matching · Fréchet distance · Surfaces

This research was partially supported by the Deutsche Forschungsgemeinschaft (DFG) within the framework of the graduate program (Graduiertenkolleg) “Combinatorics, Geometry, and Computation,” No. GRK 588/3.


H. Alt (✉)

Institut für Informatik, Freie Universität Berlin, Takustr. 9, 14195 Berlin, Germany
e-mail: alt@inf.fu-berlin.de

M. Buchin

Department of Information and Computing Sciences, University of Utrecht, PO Box 80.089,
3508TB Utrecht, The Netherlands
e-mail: maïke@cs.uu.nl

Fig. 1 Nonsimilar polygonal chains with a small Hausdorff distance δ_H



1 Introduction

Suitable distance measures for comparing the similarity of shapes are an important issue in application areas such as computer vision and pattern recognition. The distance measure that comes to mind first is the Hausdorff distance which, between two compact sets $A, B \subset \mathbb{R}^d$, is defined as

$$\delta_H(A, B) = \max\left(\max_{a \in A} \min_{b \in B} \|a - b\|, \max_{b \in B} \min_{a \in A} \|a - b\|\right),$$

where $\|\cdot\|$ denotes the underlying norm on \mathbb{R}^d , for example, the Euclidean norm. So, the Hausdorff distance considers the maximum distance of a point in one set to the other set, and it gives reasonable results in many cases. Moreover, it can be computed efficiently in two dimensions if A and B are sets of line segments [1] or, more generally, in arbitrary but fixed dimension d where A and B are sets of k -dimensional simplices [2].

However, if shapes are modeled by curves or surfaces, there are examples of objects having little resemblance but a small Hausdorff distance, see Fig. 1.

In these cases, the *Fréchet distance* is more appropriate (see [3, 4]), which is a metric for parameterized geometric objects and was first introduced by Fréchet for curves [8] and later for surfaces [9]. The idea of the Fréchet distance is to take into account the “flow” of the curve or surface given by its parameterization. A popular illustration of the Fréchet distance of two curves is the following. Suppose a man is walking his dog on a leash. The man is walking on one curve, and the dog on the other. Both may stop but not walk backwards. Then the Fréchet distance is the shortest length of leash allowing them to walk on the two curves from start to end.

Formally the Fréchet distance is defined as follows.

Definition 1 Let f, g be parameterizations of curves or surfaces, i.e., continuous functions

$$f, g : [0, 1]^k \rightarrow \mathbb{R}^d, \quad k \in \{1, 2\}, d \geq k.$$

Then their *Fréchet distance* is

$$\delta_F(f, g) = \inf_{\sigma : [0, 1]^k \rightarrow [0, 1]^k} \max_{t \in [0, 1]^k} \|f(t) - g(\sigma(t))\|,$$

where the *reparameterization* σ ranges over all *orientation preserving* homeomorphisms.

Note that in the case $k = 1$, orientation preserving just means that σ is continuous and strictly increasing (rather than decreasing). In the case $k = 2$, it means that if any Jordan curve C (e.g., the boundary of $[0, 1]^2$) is traversed clockwise, the induced traversal of $\sigma(C)$ (in the example again the boundary) is clockwise as well.

In this paper, we assume that the norm $\|\cdot\|$ underlying the definition is the Euclidean norm, but the results hold for the L_1 - and L_∞ -norms and, with some modifications, can be generalized to other norms such as any L_p -norm, $p \in \mathbb{N}$.

For dimension $k = 1$ of the parameter space, in particular for polygonal curves, the Fréchet distance is known to be computable in polynomial time [4]. For two-dimensional surfaces, however, the computation of the Fréchet distance surprisingly seems to be much harder. In fact, Godau showed in his Ph.D. thesis [10] that computing the Fréchet distance between triangulated surfaces, even in two-dimensional space, is NP-hard. It remained open, however, how hard the problem really is, not even its computability could be shown.

For the special case that both surfaces are simple plane polygons (in a space of arbitrary dimension d), Buchin et al. [5] showed that the Fréchet distance can be computed in polynomial time.

The first part of this paper contains a partial result concerning the computability of the Fréchet distance. More specifically, we show that the Fréchet distance between triangulated surfaces is *upper semi-computable*, i.e., there is a non-halting Turing machine which produces a decreasing sequence of rationals converging to the result. It follows that the decision problem whether the Fréchet distance of two given surfaces lies below some specified value is *recursively enumerable*.

The computationally hard part of computing the Fréchet distance for dimension $k = 2$ seems to be that, according to the definition, the infimum over all homeomorphisms of the parameter space has to be taken. For dimension one, the orientation preserving homeomorphisms on $[0, 1]$ are “only” the continuous, onto, increasing functions on $[0, 1]$. For dimension two, the homeomorphisms can be much “wilder.”

We tackle this problem in our algorithm by approximating the homeomorphisms by discrete maps which are easier to handle. We do this by first approximating arbitrary homeomorphisms by piecewise linear homeomorphisms, which is a known result from topology. These piecewise linear homeomorphisms are then approximated by *mesh homeomorphisms*, i.e., ones that are compatible with certain subdivisions of the original triangulations of the parameter spaces. Finally, the distance obtained by considering only mesh homeomorphisms can be approximated for fine subdivisions by considering the distances of a finite number of points.

It remains open whether the Fréchet distance between triangulated surfaces is a computable function in the strong sense.

The second part of this paper considers a relaxed version of the Fréchet distance, which we call the *weak Fréchet distance* denoted by δ_{wF} . It is defined by modifying Definition 1 so that the reparameterization σ of g which is an orientation preserving homeomorphism is replaced by reparameterizations σ, τ of both surfaces that are *surjective* continuous maps. In the man–dog illustration for curves, this means that both can choose arbitrary starting and ending points, can move forward *and backward*, and each one has to traverse his curve completely.

We will show that the weak Fréchet distance between two given triangulated surfaces can be computed in polynomial time. This is done by first considering the corresponding decision problem, i.e., the question whether, for given f, g and $\varepsilon > 0$, the distance $\delta_{\text{wF}}(f, g) \leq \varepsilon$. The decision problem can be characterized by geometric properties of the so-called *free space diagram*, a data structure originally introduced

to solve the decision problem for polygonal curves [4]. This characterization leads to a polynomial-time algorithm for the decision problem. Since it can be shown that the exact distance must be one of a finite set of critical values, we can use the decision routine in a search strategy to actually compute the weak Fréchet distance.

2 Semi-Computability of the Fréchet Distance

We assume that the input to our algorithm are two triangulated surfaces in space \mathbb{R}^d , $d \geq 2$, which are represented by *piecewise linear* parameterizations $f, g : [0, 1]^2 \rightarrow \mathbb{R}^d$. For simplicity, we will denote the surfaces themselves by f and g as well.

Piecewise linear means that the parameter spaces of f and g are triangulated and on each triangle $\Delta = \langle u, v, w \rangle$, f and g , respectively, are linear maps in the sense that $f(\lambda_1 u + \lambda_2 v + \lambda_3 w) = \lambda_1 f(u) + \lambda_2 f(v) + \lambda_3 f(w)$ for all $\lambda_1, \lambda_2, \lambda_3 \geq 0$ with $\lambda_1 + \lambda_2 + \lambda_3 = 1$ and g has the same property.

We denote the triangulated parameter spaces of f and g by K and L , respectively. The vertices of the individual triangles have rational coordinates, and the coefficients describing the linear maps are rational as well. Thus, a problem instance has a canonical finite representation which can be given as input to a Turing machine (we will describe the algorithm in some high-level language, however.)

We will show that the Fréchet distance between triangulated surfaces is computable in a weak sense according to the following definition which has been considered in the research community concerned with the computability and complexity of real functions (see, e.g., [16]).

Definition 2 A function $\varphi : \mathbb{N} \rightarrow \mathbb{R}$ is called *upper (lower) semi-computable* iff there is a Turing machine which on input x outputs an infinite decreasing (increasing) sequence of rational numbers converging to $\varphi(x)$.

Now we can formulate the first *main result* of this article:

Theorem 1 *The Fréchet distance between two triangulated surfaces in space \mathbb{R}^d , $d \geq 2$, is upper semi-computable.*

Theorem 1 immediately implies the following corollary, where $\langle f, g, a \rangle$ denotes some standard encoding of a triple consisting of two triangulated surfaces f and g and some rational $a > 0$.

Corollary 1 *The set $\{\langle f, g, a \rangle \mid \delta_F(f, g) < a\}$, i.e., the decision problem for the Fréchet distance is recursively enumerable.*

Proof Consider the Turing machine producing a decreasing sequence converging to $\delta_F(f, g)$, which exists by Theorem 1. Stop this Turing machine as soon as it produces a value less than a . Thus, the algorithm will eventually halt for all triples $\langle f, g, a \rangle$ in the language, and it will run forever for the ones not in the language. \square

The computability of δ_F in the strong sense of computability theory of real functions (see, e.g., [15]) remains open, since the sequence produced by the algorithm proving Theorem 1 is not shown to *effectively* converge to $\delta_F(f, g)$. That is, we cannot compute an upper bound on the distance of the value produced by the algorithm after k steps to the real value $\delta_F(f, g)$.

Also observe that Corollary 1 cannot be deduced from Theorem 1 any more if we replace the $<$ -sign in the definition of the set by the \leq -sign.

Our proof can be modified to show a weaker form of Theorem 1 for more *general surfaces*. More precisely, if we just assume that the parameterizations f and g are *computable real functions*, it is still correct that there is an algorithm producing on input f, g (represented, say, by the Turing machines computing f and g) an infinite sequence of rational numbers converging to $\delta_F(f, g)$. However, this sequence is not necessarily decreasing, i.e., the corollary cannot be deduced any more.

2.1 Approximating the Homeomorphisms

We will prove the semi-computability of the Fréchet distance by showing that homeomorphisms can be approximated arbitrarily closely by so-called mesh homeomorphisms, which are compatible with finite subdivisions of the parameter space.

First, let us recall some standard definitions and notation from topology. For a triangulation T , let T^m denote its m th *barycentric subdivision*, where in one subdivision step, each triangle is subdivided into six triangles by the bisectors of its sides. By $|T|$ we denote the underlying space of T , i.e., the set of all points lying on triangles of T . In our case, $|K| = |L| = [0, 1]^2$. By $\text{mesh}(T)$ we denote the maximal diameter of triangles in T which gives a measure for the “fineness” of T .

Now we define mesh homeomorphisms.

Definition 3 Given two triangulations K and L , a piecewise linear homeomorphism $h : |K^m| \rightarrow |L^n|$ is called a *mesh homeomorphism* if it maps the edges of K^m to *edge chains* of L^n , i.e., polygonal chains consisting of edges of L^n .

We will show that any homeomorphism can be approximated arbitrarily closely by a mesh homeomorphism. In fact, we need only a weak form of closeness which is defined as follows.

Definition 4 Let K be a triangulation of the unit square, and let $h, h' : [0, 1]^2 \rightarrow [0, 1]^2$ be homeomorphisms on the unit square. Then we define

$$d_K(h, h') = \max_{\Delta \in K} \delta_H(h(\Delta), h'(\Delta)),$$

where $\Delta \in K$ ranges over all triangles in K , and δ_H denotes the Hausdorff distance.

Lemma 1 *Let K and L be triangulations, $\sigma : |K| \rightarrow |L|$ a homeomorphism, $m \in \mathbb{N}$, and $\varepsilon > 0$. Then there exist $n \in \mathbb{N}$ and a mesh homeomorphism $h : |K^m| \rightarrow |L^n|$ such that $d_K(\sigma, h) < \varepsilon$.*




Fig. 2 Approximating a homeomorphism by a mesh homeomorphism

Proof By a theorem from topology (see, e.g., Chap. 6 in the book by Moise [12]), σ can be approximated arbitrarily closely by a piecewise linear homeomorphism, i.e., for all $\varepsilon_1 > 0$, there exists a piecewise linear homeomorphism $h' : |K| \rightarrow |L|$ with $d_{K^m}(\sigma, h') < \varepsilon_1$. We use this fact as a first step, because piecewise linear homeomorphisms are much easier to handle than arbitrary homeomorphisms.

We will show that any piecewise linear homeomorphism h' can be approximated to any $\varepsilon_2 > 0$ (in the sense of Definition 4) by a mesh homeomorphism h , i.e., $d_{K^m}(h, h') < \varepsilon_2$. Choosing ε_1 and ε_2 such that $\varepsilon_1 + \varepsilon_2 < \varepsilon$ then proves the lemma, see Fig. 2.

In order to approximate h' , we first show how to find edge chains in L^n , for some $n \in \mathbb{N}$ large enough, that are close to the polygonal chains which are the images of edges of K^m under h' . Then we explain how this can be extended to a piecewise linear homeomorphism on the whole parameter space $|K^m|$.


In fact, the piecewise linear images of the edges of K^m under h' form a graph G' isomorphic to K^m embedded in $|L|$ where the edges of G' are polygonal chains in $|L|$. We want to modify G' to obtain an isomorphic graph G embedded in $|L^n|$ with edge chains of L^n as edges that have a distance smaller than ε_2 to the corresponding edges of G' .

We do this in two steps:

Step 1. We map the nodes of the graph G' and short initial segments of their incident edges to nearby vertices of $L^{n'}$ and (some of) their incident edges for some suitable $n' \in \mathbb{N}$.

More precisely, for mapping the nodes, we put a small circle of radius r around each node $h'(v)$ where $r < \varepsilon_2/2$ and less than half the smallest distance between two nodes $h'(v_1)$ and $h'(v_2)$. $r < \varepsilon_2/2$ yields disks of diameter less than ε_2 in which we may move nodes and edges freely. r less than half the smallest distance between two nodes ensures that disks around different nodes do not touch.

We choose $n' \in \mathbb{N}$ such that $\text{mesh}(L^{n'}) < r/2$ and, for all nodes $h'(v)$, there is a vertex $w \in L^{n'}$ of distance less than $r/2$ whose degree (in $L^{n'}$) is greater than the degree of $h'(v)$ (in G'). This is possible because $\text{mesh}(L^{n'})$ tends to zero for large n' and in the barycentric subdivision the degrees of vertices double in each subdivision step after their introduction. For each node $v \in K^m$, we choose $h(v) = w$ for such a w , i.e., satisfying $\|h'(v) - h(v)\| < r/2 < \varepsilon_2$ and $\text{deg}(h(v)) \geq \text{deg}(v)$. By this construction each $h(v)$ and all its incident edges of $L^{n'}$ lie in the disk of radius r around $h'(v)$, see Fig. 3(a).



(a) Construction in the neighborhood of a vertex

(b) Edge chain a_π approximating π

Fig. 3 Approximating a piecewise linear homeomorphism by a mesh homeomorphism

We start mapping the edges of G' by first choosing edges incident to $h(v)$ in $L^{n'}$ as initial segments. We do so maintaining the order given by G' , i.e., if the edges e_1, \dots, e_k leave the vertex $h'(v)$ of G' in clockwise order, we choose corresponding edges f_1, \dots, f_k of $L^{n'}$ leaving $h(v)$ in the same order (see Fig. 3(a)).

Next, we cut the edges e_1, \dots, e_k at the points p_1, \dots, p_k where they first leave the disk. Within the disk we connect the free endpoint of each f_i with the point p_i , $i = 1, \dots, k$, by nonintersecting polygonal chains, which replace the original polygonal chains from $h'(v)$ to p_i . This is possible because we chose the edges f_1, \dots, f_k in the same order as the cutting points and within a disk we are allowed to move freely, without violating the distance bound ε_2 (see Fig. 3(a)). Thus, we replaced all vertices $h'(v)$ by closely vertices $h(v)$ lying on the mesh $L^{n'}$.

Step 2. Next, we delete the nodes $h(v)$ together with the incident edges f_1, \dots, f_k from the scenery, leaving a finite set of pairwise disjoint polygonal chains which start and end in mesh points of $L^{n'}$. We show that they can be ε_2 -approximated by edge chains of L^n for suitable $n \geq n'$.

To achieve this, let η be the minimum distance between a vertex of a curve and a nonincident edge of a curve (possibly the same). For edges sharing an incident vertex, let θ be the minimum distance between the intersection points of the edges with a circle of radius $\varepsilon_2/2$ around their common vertex.

Then we subdivide $L^{n'}$ to L^n for a suitable $n \geq n'$ such that $\text{mesh}(L^n) < \min(\eta/2, \theta/2)$. Thus, L^n is so fine that each polygonal chain π traverses a sequence of triangles T_π so that T_π and even its neighboring triangles are not intersected by other polygonal chains (see Fig. 3(b)).

Now, we can approximate each polygonal chain π by any connected, not self-intersecting edge chain a_π of L^n that lies within T_π and connects the endpoints of π so that the Hausdorff distance between π and a_π is less than ε_2 .

Thus, we showed the existence of a polygonal chain $h(e)$ in L^n for each edge e of K^m which is arbitrarily close to $\sigma(e)$. The chains $h(e)$ and the vertices $h(v)$ form an embedded graph G isomorphic to G' and, therefore, to K^m . h can be extended to the interior points of each edge e to form a homeomorphism on e in a straightforward manner. Furthermore, the faces of G induce a partition of the set of triangles of L^n which is isomorphic to the triangulation K^m . To extend h to a piecewise linear homeomorphism on $|K^m|$, we subdivide each triangle Δ of K^m according to the triangulation of the associated set of triangles in the partition of L^n and extend h to the interior of Δ correspondingly. \square

The fact that homeomorphisms can be approximated arbitrarily closely by mesh homeomorphisms implies that the Fréchet distance can be approximated arbitrarily closely by considering only mesh homeomorphisms and distances between the vertices of the corresponding triangulations. More precisely, we obtain:

Lemma 2

$$\delta_F(f, g) = \inf_{m, n \in \mathbb{N}} \inf_{h: |K^m| \rightarrow |L^n|} \max_{\Delta \in K_T^m} \max_{\substack{v \in V_\Delta \\ w \in M_{h(\Delta)}^n}} \|f(v) - g(w)\|,$$

where h ranges over all orientation preserving mesh homeomorphisms, K_T^m is the set of triangles in K^m , V_Δ are the vertices of Δ , and $M_{h(\Delta)}^n$ is the set of vertices of L^n that lie in $h(\Delta)$.

Proof In order to prove Lemma 2, we define the right-hand side of the equation in the lemma as the *discrete Fréchet distance*¹ $\delta_{dF}(f, g)$ of surfaces and show that it is equal in value to the Fréchet distance.

We first show that the Fréchet distance is not larger than the discrete Fréchet distance.

Claim 1 $\delta_F \leq \delta_{dF}$.

Since any mesh homeomorphism is, in particular, a homeomorphism, it suffices to show that for a mesh homeomorphism $h : |K^m| \rightarrow |L^n|$, we can bound the pointwise maximum by the maximum taken at vertices, i.e., show that

$$\max_{t \in [0, 1]^2} \|f(t) - g(h(t))\| \leq \max_{\Delta \in K_T^m} \max_{\substack{v \in V_\Delta \\ w \in M_{h(\Delta)}^n}} \|f(v) - g(w)\|. \tag{1}$$

To see this, let $t \in [0, 1]^2$ be arbitrary. t lies in a triangle $\Delta = \langle v_1, v_2, v_3 \rangle$ of K^m , and $h(t)$ lies in a triangle $\Delta' = \langle w_1, w_2, w_3 \rangle$ of $h(\Delta) \subset L^n$. Since f and g are piecewise linear and K^m and L^n are refinements of the underlying triangulations of the parameter spaces, $f(\Delta)$ and $g(\Delta')$ are triangles as well, namely $\langle f(v_1), f(v_2), f(v_3) \rangle$

¹Observe that we use the term *discrete Fréchet distance* in a meaning different from the one occasionally used in the literature for curves.

and $\langle g(w_1), g(w_2), g(w_3) \rangle$, respectively. Consequently, since the maximum distance between points of two triangles in 3-space is attained between two corners, we have that $\|f(t) - g(h(t))\| \leq \|f(v_i) - g(w_j)\|$ for some i, j with $1 \leq i, j \leq 3$. Taking the maximum on both sides yields (1).

Now we show that also the discrete Fréchet distance is not larger than the Fréchet distance.

Claim 2 For every $\varepsilon > 0$, $\delta_{dF} \leq \delta_F + \varepsilon$.

The idea is that for any homeomorphism, there is an arbitrarily close mesh homeomorphism, and for the mesh homeomorphism, the distance computation at vertices comes arbitrarily close to the distance computation on all parameter values by sufficient subdivision of the domain complex.

Let σ be a homeomorphism close to realizing δ_F , i.e., $\max_t \|f(t) - g(\sigma(t))\| \leq \delta_F + \varepsilon_1$ for some small $\varepsilon_1 > 0$. By Lemma 1, for any $\varepsilon_2 > 0$ and any $m \in \mathbb{N}$, there is $n \in \mathbb{N}$ and a mesh homeomorphism $h : |K^m| \rightarrow |L^n|$ such that $d_{K^m}(\sigma, h) \leq \varepsilon_2$.

Next, we observe that f and g are Lipschitz continuous since they are continuous and piecewise linear. Let c_f and c_g denote the respective Lipschitz constants.

Let Δ be some triangle in $|K^m|$ and v one of its vertices. Since $d_{K^m}(\sigma, h) \leq \varepsilon_2$, for any $w \in h(\Delta) \subset L^n$, there is $x \in \sigma(\Delta)$ with $\|w - x\| < \varepsilon_2$. Using $t = \sigma^{-1}(x)$ and the Lipschitz-continuity of g , we get $\|g(w) - g(\sigma(t))\| < c_g \cdot \varepsilon_2$ for some $t \in \Delta$.

t and v lie in the same triangle $\Delta \in K^m$, so $\|v - t\| \leq \text{mesh}(K^m)$, and, thus, $\|f(v) - f(t)\| \leq c_f \cdot \text{mesh}(K^m)$.

Putting everything together and applying the triangle inequality repeatedly, we obtain

$$\begin{aligned} \delta_{dF} &\leq \max_{\Delta \in K_T^m} \max_{\substack{v \in V_\Delta \\ w \in M_{h(\Delta)}^n}} \|f(v) - g(w)\| \\ &\leq \max_{\Delta \in K_T^m} \max_{\substack{v \in V_\Delta \\ x \in \sigma(\Delta)}} \|f(v) - g(x)\| + c_g \cdot \varepsilon_2 \\ &\leq \max_{\Delta \in K_T^m} \max_{t \in \Delta} \|f(t) - g(\sigma(t))\| + c_g \cdot \varepsilon_2 + c_f \cdot \text{mesh}(K^m) \\ &\leq \delta_F + \varepsilon_1 + c_g \cdot \varepsilon_2 + c_f \cdot \text{mesh}(K^m). \end{aligned}$$

Since $\varepsilon_1, \varepsilon_2$, and, by choosing m large enough, $\text{mesh}(K^m)$ can be made arbitrarily small, Claim 2 is shown.

Claims 1 and 2 yield $\delta_F = \delta_{dF}$, which proves Lemma 2. □

2.2 Semi-Computing the Fréchet Distance

Using Lemma 2, we can now give an algorithm showing the upper semi-computability of the Fréchet distance between surfaces as claimed in Theorem 1. This algorithm will, on input f, g , run forever and produce a decreasing sequence of rational numbers converging to $\delta_F(f, g)$.

Algorithm 1: SemiComputeFréchet(f, g)

Input: Triangulated surfaces f, g , including triangulations K, L of the parameter spaces, in a finite description

Output: A monotone decreasing sequence of rational numbers converging to $\delta_F(f, g)$

```

1 set  $D = \infty$ 
2 forall  $(n, m) \in \mathbb{N} \times \mathbb{N}$  do
3   generate the barycentric subdivisions  $K^m$  of  $K$  and  $L^n$  of  $L$ 
4   let  $E = \{e_1, \dots, e_k\}$  be the set of edges in  $K^m$ 
5   forall  $k$ -tuples  $(\pi_1, \dots, \pi_k)$  of simple polygonal chains in  $L^n$  do
6     assign to the edge  $e_i$  the polygonal chain  $\pi_i$  for  $i = 1, \dots, k$ 
7     if this assignment results in an orientation preserving homeomorphic
      image of  $K^m$  then
8       set  $M = 0$ 
9       forall triangles  $\Delta$  of  $K^m$  do
10        let  $H_\Delta \subset |L^n|$  be the region in  $L^n$  assigned to  $\Delta$ 
11        forall vertices  $v$  of  $\Delta$  and vertices  $w$  of  $H_\Delta$  do
12          set  $M = \max(M, \|f(v) - g(w)\|)$ 
13        end
14      end
15      set  $D = \min(D, M)$ 
16      output  $D$ 
17    end
18  end
19 end

```

Line 2 can be realized by some standard enumeration method for pairs of integers. The number of k -tuples of polygonal chains of L^n checked in line 5 is finite. In fact, it is bounded by $(l!)^k$ where l is the number of edges in L^n , which itself is exponential in n , whereas k is exponential in m . However, efficiency is not the issue here.

In line 12 we assume that the norm $\|\cdot\|$ underlying the Fréchet distance can be evaluated by rational operations. This is correct for, e.g., the L_1 - or L_∞ -metric but not directly for the Euclidean metric L_2 . In that case, one should rather operate with the square of the distance in line 12 and output some suitable rational approximation of \sqrt{D} (which is possible) in line 16.

In line 7 we check whether an assignment of edges in K^m to polygonal chains in L^n results in an orientation preserving homeomorphic image of K^m by checking the following three conditions:

- The edges on the boundary of $|K^m|$ are mapped onto the boundary of $|L^n|$ preserving the orientation.
- If a set of edges in K^m share an endpoint, the corresponding chains in L^n do as well.
- Other than that, there are no intersection points between two chains.

Note that checking that the boundary of $|K^m|$ is mapped orientation preservingly onto the boundary of $|L^n|$ entails that the mesh homeomorphism is orientation preserving also on the interior.

For each pair $(m, n) \in \mathbb{N} \times \mathbb{N}$ all mesh homeomorphisms $h : K^m \rightarrow L^n$ are evaluated by Algorithm 1, i.e.,

$$\delta_{h,m,n} = \max_{\Delta \in K_T^m} \max_{\substack{v \in V_\Delta \\ w \in M_{h(\Delta)}^n}} \|f(v) - g(w)\|$$

(see Lemma 2) is computed.²

To see that Algorithm 1 produces values arbitrarily close to $\delta_F(f, g)$, observe that any neighborhood of $\delta_F(f, g)$ must, by Lemma 2, contain some value of the form $\delta_{h,m,n}$. The algorithm will eventually encounter that pair (m, n) and the subdivision corresponding to h and output $\delta_{h,m,n}$. By line 15 the output sequence is decreasing. Since for all triples (h, m, n) , by Lemma 2, $\delta_{h,m,n} \geq \delta_F(f, g)$, line 15 is justified. Therefore, Algorithm 1 arbitrarily closely approximates $\delta_F(f, g)$, which proves Theorem 1.

3 Computability of the Weak Fréchet Distance

In this section we give a polynomial-time algorithm for computing the weak Fréchet distance between triangulated surfaces. The weak Fréchet distance is a relaxed version of the Fréchet distance which does not require the reparameterizations of the curves or surfaces to be injective. Instead, it uses surjective continuous maps as reparameterizations.

Definition 5 The weak Fréchet distance between curves or surfaces given by continuous parameterizations $f, g : [0, 1]^k \rightarrow \mathbb{R}^d$ with $k \in \{1, 2\}, k \leq d$, is

$$\delta_{wF}(f, g) = \inf_{\alpha, \beta : [0, 1]^k \rightarrow [0, 1]^k} \max_{\mathbf{x} \in [0, 1]^k} \|f(\alpha(\mathbf{x})) - g(\beta(\mathbf{x}))\|,$$

where α and β range over all surjective continuous maps.

As was mentioned before, in the man–dog illustration for curves this means that both are allowed to choose their start and endpoints, can move forward and backward, and eventually must have traversed both curves completely. The weak Fréchet distance is always less than or equal to the Fréchet distance and greater than or equal to the Hausdorff distance:

$$\delta_H(f, g) \leq \delta_{wF}(f, g) \leq \delta_F(f, g).$$

²A more detailed analysis shows that, in fact, it suffices to consider only the pairs $(m, 2m), m \in \mathbb{N}$.

Fig. 4 Curves with small weak Fréchet distance and their free space diagram for parameter ε




Figure 4, e.g., shows two curves with a small weak Fréchet distance and a large Fréchet distance. For curves, this relation between the distance measures is reflected in the *free space diagram*, which is a geometric data structure for computing the Fréchet distance of polygonal curves [4].

For the sake of illustration, we will shortly repeat these ideas and concepts for curves before we extend them to surfaces. For two parameterized curves f, g and parameter ε , the *free space* is defined as

$$F_\varepsilon(f, g) = \{(x, y) \mid \|f(x) - g(y)\| \leq \varepsilon\}.$$

The free space diagram illustrates the free space, see Fig. 4, where the direction from left to right corresponds to the parameterization of f and the one from bottom to top to the parameterization of g . The white area inside the diagram corresponds to the free space for the indicated value of ε . For polygonal curves f, g of lengths n and m , respectively, the free space diagram consists of nm cells, each one corresponding to the free space between two segments of the curves.

It can easily be seen that $\delta_F(f, g) \leq \varepsilon$ if and only if there is a monotone path in the free space from the lower left corner $(0, 0)$ to the upper right corner $(1, 1)$ (corresponding to the reparameterization σ in Definition 1). This fact is used in [4] to compute the Fréchet distance.

However, also the fact that the Hausdorff distance or the weak Fréchet distance are less than or equal ε can be characterized by the free space diagram. Let us call the free space or any subset of it *extensive* if its projections onto both parameter spaces are surjective. Then:

Hausdorff distance less than or equal ε is equivalent to the free space being extensive.

Weak Fréchet distance less than or equal ε means that there is a curve inside the free space which is extensive. This fact is used in [4] to compute a restricted version of the weak Fréchet distance called the *nonmonotone Fréchet distance* of polygonal curves. An extensive curve exists exactly if there is an extensive connected component of the free space.

Essentially, this concept and this characterization of the weak Fréchet distance can be extended to surfaces (yielding a four-dimensional free space diagram). We will use this idea to find an efficient algorithm for computing the weak Fréchet distance of triangulated surfaces.

In fact, the second *main result* of this article is the following theorem.

Theorem 2 *The weak Fréchet distance between two triangulated surfaces with m, n triangles, respectively, can be computed in $O((m^2n + mn^2) \log^2(mn))$ time.*

The algorithms we give will need to compute the intersection of ellipses, circles, and line segments and compare such intersection points. These intersection points can all be described as roots of polynomials of degree up to four, which can be compared and computed exactly using constantly many arithmetic operations on rationals, see [7, 13].

3.1 Free Space Diagram of Triangulated Surfaces

As in Sect. 2, we assume that we are given piecewise linear parameterizations $f, g : [0, 1]^2 \rightarrow \mathbb{R}^d, d \geq 2$, of triangulated surfaces with underlying triangulations K, L , respectively. All coefficients of the maps and coordinates of the triangulations are assumed to be rational. In the following, we use \mathbf{x}, \mathbf{y} to denote points in parameter space, i.e., $\mathbf{x}, \mathbf{y} \in [0, 1]^2$. Let $F_\varepsilon(f, g) = \{(\mathbf{x}, \mathbf{y}) \mid \|f(\mathbf{x}) - g(\mathbf{y})\| \leq \varepsilon\}$ be the free space diagram of the surfaces. The free space diagram lies in the product of the parameter spaces of the surfaces, i.e., in the four-dimensional cube. As in the case of curves, we can partition the free space into cells.

Definition 6 Let $f, g : [0, 1]^2 \rightarrow \mathbb{R}^d, d \geq 2$, be piecewise linear surface parameterizations with underlying triangulations K and L , respectively. A *cell* of the free space is $C_\varepsilon(\Delta_K, \Delta_L) := F_\varepsilon(f, g) \cap (\Delta_K \times \Delta_L)$ for triangles Δ_K of K and Δ_L of L . A *boundary cell* is the cell of a triangle and an edge, defined analogously. Two cells are called *neighboring* if they share a nonempty boundary cell.

We will consider the projection of the free space onto the two parameter spaces, i.e., the image of $F_\varepsilon(f, g)$ under $\text{Proj}_K : [0, 1]^4 \rightarrow [0, 1]^2, (\mathbf{x}, \mathbf{y}) \mapsto \mathbf{x}$, and $\text{Proj}_L : [0, 1]^4 \rightarrow [0, 1]^2, (\mathbf{x}, \mathbf{y}) \mapsto \mathbf{y}$. We will compute the projections of the free space by computing it for its cells, i.e., we will compute for all triangles Δ_K of K and Δ_L of L

$$\text{Proj}_K(C_\varepsilon(\Delta_K, \Delta_L)) = \{\mathbf{x} \in \Delta_K \mid \exists \mathbf{y} \in \Delta_L : \|f(\mathbf{x}) - g(\mathbf{y})\| \leq \varepsilon\},$$

and analogously for the projection onto the parameter space L . Although the free space lies in the product of the parameter spaces, it is defined based on the distances in image space.

The combinatorial structure of the free space is captured by the graph whose vertices are the cells $C_\varepsilon(\Delta_K, \Delta_L)$ of the free space and edges exist between neighboring cells which share a nonempty boundary cell. For a free space diagram $F_\varepsilon(f, g)$, we denote the corresponding graph by G_ε . The graph G_ε has at most mn vertices and $O(mn)$ edges. For $\varepsilon_1 \leq \varepsilon_2$, the graph G_{ε_1} is a subgraph of G_{ε_2} . For varying ε , G_ε changes only at finitely many values, namely when a boundary cell becomes non-empty. This happens when ε equals the minimum distance between points on some edge e_K of K and on some triangle Δ_L of L or vice versa:

$$C_\varepsilon(e_K, \Delta_L) \neq \emptyset \iff \min\{\|x - y\| \mid x \in f(e_K), y \in g(\Delta_L)\} \leq \varepsilon. \tag{2}$$

The main idea of our algorithm is contained in the following lemma.

Lemma 3 *The weak Fréchet distance between two triangulated surfaces is less than or equal to ε if and only if there is an extensive connected component A in the free space for the parameter ε , i.e., $\text{Proj}_K(A) = \text{Proj}_L(A) = [0, 1]^2$.*

Proof Let α, β be two reparameterizations of f and g , respectively. We define the set $M_{\alpha,\beta} \subset [0, 1]^4$ by

$$M_{\alpha,\beta} = \{(\alpha(\mathbf{x}), \beta(\mathbf{x})) \mid \mathbf{x} \in [0, 1]^2\}.$$

The set $M_{\alpha,\beta}$ is connected because α and β are continuous. $M_{\alpha,\beta}$ is extensive because α and β are surjective. By definition of $M_{\alpha,\beta}$ the following equivalence holds:

$$M_{\alpha,\beta} \subseteq F_\varepsilon(f, g) \Leftrightarrow \max_{\mathbf{x} \in [0,1]^2} \|f(\alpha(\mathbf{x})) - g(\beta(\mathbf{x}))\| \leq \varepsilon. \tag{3}$$

Now we show the two directions of the lemma.

First, let us prove the only-if-direction.

In case there are two reparameterizations α and β for f and g , respectively, realizing that the weak Fréchet distance $\leq \varepsilon$, i.e., $\max_{\mathbf{x} \in [0,1]^2} \|f(\alpha(\mathbf{x})) - g(\beta(\mathbf{x}))\| \leq \varepsilon$, then $M_{\alpha,\beta}$ is an extensive, connected subset of $F_\varepsilon(f, g)$, so the connected component containing $M_{\alpha,\beta}$ is an extensive connected component in $F_\varepsilon(f, g)$.

A weak Fréchet distance of at most ε , however, does not imply that two realizing reparameterizations α, β exist, as it is defined as $\inf_{\alpha,\beta} \max_{\mathbf{x}} \|f(\alpha(\mathbf{x})) - g(\beta(\mathbf{x}))\| \leq \varepsilon$. This can be reformulated as

$$\forall \vartheta > \varepsilon \exists \alpha, \beta : \max_{\mathbf{x} \in [0,1]^2} \|f(\alpha(\mathbf{x})) - g(\beta(\mathbf{x}))\| \leq \vartheta.$$

By (3) this implies that $F_\vartheta(f, g)$ contains, for all $\vartheta > \varepsilon$, an extensive connected component. We will show that this property holds also for $F_\varepsilon(f, g)$.

As discussed previously, the combinatorial structure of the connected components of the free space changes only at finitely many values for ε , i.e., when a boundary cell becomes nonempty. Let

$$\eta = \arg \min \{ \vartheta \mid \vartheta > \varepsilon, \text{ and a boundary cell becomes nonempty in } F_\vartheta(f, g) \}.$$

Then the combinatorial structure of F_ϑ , i.e., the graph $G = G_\vartheta$, is the same for all $\vartheta \in [\varepsilon, \eta)$.

Let $(\varepsilon_n)_{n \in \mathbb{N}}$ be a sequence in $[\varepsilon, \eta)$ converging to ε . F_{ε_n} contains, for all n , an extensive connected component. Since G has only finitely many connected components, there must be at least one of them, say A , which corresponds to an extensive connected component A_{ε_n} in infinitely many of the F_{ε_n} . Since for any ε' , the projection of $A_{\varepsilon'}$, the connected component of $F_{\varepsilon'}$ corresponding to A , is a closed subset of $[0, 1]^2$ and changes continuously with the parameter ε' , the connected component A_ε in F_ε is also extensive.

To prove the if-direction of the lemma, suppose that there exists an extensive connected component of the free space $A \subset F_\varepsilon(f, g)$. We will construct reparameterizations α and β of f and g , respectively, such that $M_{\alpha,\beta} \subset A$ and $\text{Proj}_K(M_{\alpha,\beta}) = \text{Proj}_K(A) = [0, 1]^2$ and $\text{Proj}_L(M_{\alpha,\beta}) = \text{Proj}_L(A) = [0, 1]^2$, i.e., $M_{\alpha,\beta}$ is extensive. By (3) this will imply that α and β realize that the weak Fréchet distance is less than ε .

We proceed in two steps: first, we define for each cell $C_\varepsilon(\Delta_K, \Delta_L)$ contained in the connected component A , two parameterized surface patches $M_K(\Delta_K, \Delta_L)$ and $M_L(\Delta_K, \Delta_L)$ which fulfill $\text{Proj}_K(C_\varepsilon(\Delta_K, \Delta_L)) = \text{Proj}_K(M_K(\Delta_K, \Delta_L))$ and $\text{Proj}_L(C_\varepsilon(\Delta_K, \Delta_L)) = \text{Proj}_L(M_L(\Delta_K, \Delta_L))$. By a *surface patch* we mean a two-dimensional manifold with boundary in \mathbb{R}^4 . Then we show how to “glue” two such parameterized surfaces patches together. Gluing together the surface patches of all cells contained in the connected component will give the desired reparameterizations α and β .

Defining Parameterized Surface Patches in Each Cell. For each cell $C_\varepsilon(\Delta_K, \Delta_L)$ of the free space, we define two surface patches:

$$M_K(\Delta_K, \Delta_L) = \{(\mathbf{x}, g^{-1}(n(f(\mathbf{x}), g(\Delta_L)))) \mid \mathbf{x} \in \text{Proj}_K(C_\varepsilon(\Delta_K, \Delta_L))\},$$

$$M_L(\Delta_K, \Delta_L) = \{(f^{-1}(n(g(\mathbf{y}), f(\Delta_K)), \mathbf{y})) \mid \mathbf{y} \in \text{Proj}_L(C_\varepsilon(\Delta_K, \Delta_L))\},$$

where

$$n(x, A) = \arg \min_{y \in A} \|x - y\|$$

is the nearest neighbor of a point x in the set A . Observe that, since $g(\Delta_L)$ is a triangle, it contains a unique nearest neighbor to $f(\mathbf{x})$. By g^{-1} we mean the unique preimage of this point in Δ_L . Therefore, $M_K(\Delta_K, \Delta_L)$ and, analogously, $M_L(\Delta_K, \Delta_L)$ are well defined.

These surface patches have the following properties:

1. $M_K(\Delta_K, \Delta_L) \cup M_L(\Delta_K, \Delta_L) \subset C_\varepsilon(\Delta_K, \Delta_L)$.
2. $\text{Proj}_K(C_\varepsilon(\Delta_K, \Delta_L)) = \text{Proj}_K(M_K(\Delta_K, \Delta_L))$ and $\text{Proj}_L(C_\varepsilon(\Delta_K, \Delta_L)) = \text{Proj}_L(M_L(\Delta_K, \Delta_L))$.
3. $M_K(\Delta_K, \Delta_L) \cap M_L(\Delta_K, \Delta_L) \neq \emptyset$.
4. If the shared boundary cell of the cells $C_\varepsilon(\Delta_K, \Delta_L)$ and $C_\varepsilon(\Delta_K, \Delta'_L)$ is non-empty, then $M_L(\Delta_K, \Delta_L) \cap M_L(\Delta_K, \Delta'_L) \neq \emptyset$.
If the shared boundary cell of the cells $C_\varepsilon(\Delta_K, \Delta_L)$ and $C_\varepsilon(\Delta'_K, \Delta_L)$ is non-empty, then $M_K(\Delta_K, \Delta_L) \cap M_K(\Delta'_K, \Delta_L) \neq \emptyset$.
5. There exist parameterizations $\alpha_K, \alpha_L : [0, 1]^2 \rightarrow \Delta_K$ and $\beta_K, \beta_L : [0, 1]^2 \rightarrow \Delta_L$ such that $M_K(\Delta_K, \Delta_L) = M_{\alpha_K, \beta_K}$ and $M_L(\Delta_K, \Delta_L) = M_{\alpha_L, \beta_L}$.

For the first property, consider a tuple $(\mathbf{x}, \mathbf{y}) \in M_K(\Delta_K, \Delta_L)$. By definition the point \mathbf{x} is in $\text{Proj}_K(C_\varepsilon(\Delta_K, \Delta_L))$, i.e., there is a point \mathbf{y}' such that $(\mathbf{x}, \mathbf{y}') \in C_\varepsilon(\Delta_K, \Delta_L)$. Then also $(\mathbf{x}, \mathbf{y}) \in C_\varepsilon(\Delta_K, \Delta_L)$ because \mathbf{y} was chosen based on the nearest neighbor in image space.

The second property holds by definition, that is, $M_K(\Delta_K, \Delta_L)$ is defined by choosing a tuple for each point $\mathbf{x} \in \text{Proj}_K(C_\varepsilon(\Delta_K, \Delta_L))$.

The third property is equivalent to the existence of two points in Δ_K and Δ_L whose images under f and g , respectively, are each others nearest neighbor, i.e., a tuple $(\mathbf{x}, \mathbf{y}) \in \Delta_K \times \Delta_L$ such that $g(\mathbf{y}) = n(f(\mathbf{x}), g(\Delta_L))$ and $f(\mathbf{x}) = n(g(\mathbf{y}), f(\Delta_K))$. Points fulfilling this condition are points whose images have minimal distance in the free space cell, i.e., a tuple (\mathbf{x}, \mathbf{y}) such that $\|f(\mathbf{x}) - g(\mathbf{y})\| \leq \|f(\mathbf{x}') - g(\mathbf{y}')\|$ for all tuples $(\mathbf{x}', \mathbf{y}') \in \Delta_K \times \Delta_L$. Such a tuple exists because $\Delta_K \times \Delta_L$ is compact.

The fourth property holds because $M_L(\Delta_K, \Delta_L)$ and $M_L(\Delta_K, \Delta'_L)$ coincide on the intersection of $C_\varepsilon(\Delta_K, \Delta_L)$ and $C_\varepsilon(\Delta_K, \Delta'_L)$, and analogously for $M_K(\Delta_K, \Delta_L)$ and $M_K(\Delta'_K, \Delta_L)$.

The fifth property is fulfilled by the following natural parameterizations. Let σ_K and σ_L be homeomorphisms from $[0, 1]^2$ to Δ_K and Δ_L , respectively. Then $\alpha_K = \sigma_K$ and $\beta_K = g^{-1} \circ n(\cdot, g(\Delta_L)) \circ f \circ \sigma_K$ and $\alpha_L = f^{-1} \circ n(\cdot, f(\Delta_K)) \circ g \circ \sigma_L$ and $\beta_L = \sigma_L$ are parameterizations of $M_K(\Delta_K, \Delta_L)$ and $M_L(\Delta_K, \Delta_L)$, respectively. These reparameterizations are continuous because the identity map, f, g , and the nearest neighbor function on convex sets in \mathbb{R}^d are continuous.

Now we define M_A to be the union of all surface patches of cells in the connected component A , i.e.,

$$M_A = \bigcup_{C_\varepsilon(\Delta_K, \Delta_L) \subset A} (M_K(\Delta_K, \Delta_L) \cup M_L(\Delta_K, \Delta_L)).$$

The first two properties of the surface patches $M_K(\Delta_K, \Delta_L)$ and $M_L(\Delta_K, \Delta_L)$ imply the following properties of M_A :

1. $M_A \subset A$.
2. $\text{Proj}_K(A) = \text{Proj}_K(M_A)$ and $\text{Proj}_L(A) = \text{Proj}_L(M_A)$.

Gluing Together the Parameterizations of the Surface Patches. We first show that two reparameterizations can be joined to one and then obtain parameterizations α_A, β_A for M_A by successively gluing together parameterizations of the surface patches in M_A .


Claim 3 Given parameterizations $\alpha, \beta, \alpha', \beta'$ with $M_{\alpha, \beta} \cap M_{\alpha', \beta'} \neq \emptyset$, there exist parameterizations α'', β'' such that $M_{\alpha, \beta} \cup M_{\alpha', \beta'} = M_{\alpha'', \beta''}$.

Proof of Claim 3 Let $\mathbf{p} = (\mathbf{x}_p, \mathbf{y}_p)$ be an intersection point of $M_{\alpha, \beta}$ and $M_{\alpha', \beta'}$. Let $\mathbf{x}_0, \mathbf{x}_1$ be such that $\mathbf{p} = (\alpha(\mathbf{x}_0), \beta(\mathbf{x}_0)) = (\alpha'(\mathbf{x}_1), \beta'(\mathbf{x}_1))$.

For gluing together the parameterizations $\alpha, \beta, \alpha', \beta'$, we partition the parameter space as illustrated in Fig. 5(a). Choose a point \mathbf{c} in the parameter space and radii r, r' such that $r < r'$ and the disk of radius r' around the point \mathbf{c} is contained in $[0, 1]^2$. Let C, C' be the circles of radius r, r' around \mathbf{c} and D, D' the (closed) disks of radius r, r' around \mathbf{c} . Let R be the half-open ring with radii r, r' around \mathbf{c} , and let $S = [0, 1]^2 \setminus D'$.

We will glue the parameterizations α, β and α', β' along the circle C' . On the circle, α'', β'' will be constant, outside they will be defined based on α', β' , and inside based on α, β .

Fig. 5 Gluing together two parameterizations



We define homeomorphisms on the parts of the partition as illustrated in Fig. 5(b). Let σ_1 be a homeomorphism from D to the unit cube. Let σ_2 be a homeomorphism from R to the unit cube without the point x_0 . Choose σ_1 and σ_2 such that they coincide on the circle C . Let σ_3 be a homeomorphism from S to $[0, 1]^2 \setminus \{x_1\}$. Now we define $\alpha'', \beta'' : [0, 1]^2 \rightarrow [0, 1]^2$ as follows:

$$\alpha''(\mathbf{x}) = \begin{cases} \alpha(\sigma_1(\mathbf{x})), & \mathbf{x} \in D \\ \alpha(\sigma_2(\mathbf{x})), & \mathbf{x} \in R \\ \mathbf{x}_p, & \mathbf{x} \in C' \\ \alpha'(\sigma_3(\mathbf{x})), & \mathbf{x} \in S \end{cases}, \quad \beta''(\mathbf{x}) = \begin{cases} \beta(\sigma_1(\mathbf{x})), & \mathbf{x} \in D \\ \beta(\sigma_2(\mathbf{x})), & \mathbf{x} \in R \\ \mathbf{y}_p, & \mathbf{x} \in C' \\ \beta'(\sigma_3(\mathbf{x})), & \mathbf{x} \in S \end{cases}.$$

α'', β'' are continuous and fulfill $M_{\alpha, \beta} \cup M_{\alpha', \beta'} = M_{\alpha'', \beta''}$ by construction. They are continuous inside D, R , and S as concatenations of continuous functions. On the circle C they are continuous because σ_1 and σ_2 coincide on C . They are continuous on the circle C' because $\sigma_2(\mathbf{x})$ tends to x_0 for \mathbf{x} tending to $\mathbf{x}' \in C'$, and therefore $\alpha(\sigma_2(\mathbf{x}))$ and $\beta(\sigma_2(\mathbf{x}))$ tend to \mathbf{x}_p and \mathbf{y}_p , respectively, for \mathbf{x} tending to $\mathbf{x}' \in C'$. Similarly, $\sigma_3(\mathbf{x})$ tends to x_1 for \mathbf{x} tending to $\mathbf{x}' \in C'$, and therefore $\alpha'(\sigma_3(\mathbf{x}))$ and $\beta'(\sigma_3(\mathbf{x}))$ tend to \mathbf{x}_p and \mathbf{y}_p , respectively, for \mathbf{x} tending to $\mathbf{x}' \in C'$. Thus, Claim 3 is proven.

We can glue together all surface patches in M_A because they are connected by intersection—properties 3 and 4 of the surface patches. By gluing together all surface patches in M_A we get reparameterizations α_A, β_A for f, g which realize that the weak Fréchet distance between f and g is less than ε : α_A and β_A are continuous because the parameterizations of single surface patches are continuous and the gluing of parameterizations preserves continuity. α_A and β_A are surjective because the projection of $M_{\alpha_A, \beta_A} = M_A$ equals the projection of A —property 2 of M_A —and this was assumed to be $[0, 1]^2$. Finally, $\|f(\alpha_A(\mathbf{x})) - g(\beta_A(\mathbf{x}))\| \leq \varepsilon$ holds for all $\mathbf{x} \in [0, 1]^2$ because M_{α_A, β_A} lies completely in $F_\varepsilon(f, g)$. \square

3.2 Decision Problem for the Weak Fréchet Distance

By Lemma 3, deciding if the weak Fréchet distance is less than some given value ε is equivalent to determining if there is an extensive connected component in the

free space for the parameter ε . We will do this algorithmically in two steps: First, we compute all connected components of the free space; secondly, we test whether there is an extensive one.

For the first step, we compute the combinatorial graph of the free space as described in the previous section in $O(mn)$ time.

In the second step we need to determine if the projection of a connected component A completely covers both parameter spaces. This property for, e.g., the parameter space K is equivalent to each triangle Δ_K in K being covered by the projection of A . Equivalently, each triangle $f(\Delta_K)$ in image space must be completely contained in the union of the ε -neighborhoods $g(\Delta_L) \oplus B_\varepsilon$ of all triangles Δ_L such that the cell $C_\varepsilon(\Delta_K, \Delta_L)$ lies in A .

We can reduce this to a two-dimensional problem by intersecting the ε -neighborhoods $g(\Delta_L) \oplus B_\varepsilon$ with the plane E_{Δ_K} which contains the triangle $f(\Delta_K)$. That is, we decide for each triangle Δ_K in K whether

$$f(\Delta_K) \subset \bigcup_{\substack{\Delta_L \in L \\ C_\varepsilon(\Delta_K, \Delta_L) \in A}} ((g(\Delta_L) \oplus B_\varepsilon) \cap E_{\Delta_K}). \tag{4}$$

This property can be decided by computing for each triangle Δ_K the arrangement of the triangle $f(\Delta_K)$ and the set of boundary curves $\partial((g(\Delta_L) \oplus B_\varepsilon) \cap E_{\Delta_K})$ for all triangles $\Delta_L \in L$ with $C_\varepsilon(\Delta_K, \Delta_L) \in A$, where ∂ denotes the boundary operator. This computation can be done by a line sweep which also checks whether each cell of the arrangement covered by $f(\Delta_K)$ is also covered by the interior of at least one of the other curves. Property (4) holds exactly if this is the case.

For the metrics L_1 and L_∞ , the ε -neighborhoods can be described by linear equations. For the Euclidean metric, the ε -neighborhoods are described by quadratic equations: the ε -neighborhood of a triangle in 3-space is the union of ε -balls around the vertices, cylinders of radius ε around the edges, and a triangular prism of height 2ε . The boundaries of the intersections of these ε -neighborhoods with a plane are the union of a constant number of elliptic, circular, and straight segments. The event points of the sweep over the arrangement are all intersection points of a boundary curve with the triangle $f(\Delta_K)$ and intersection points between two boundary curves that lie inside $f(\Delta_K)$.

The overall time complexity of the second step of the algorithm adds up as follows: For each triangle Δ_K , we need to decide by which connected components of the free space it is completely covered. For one connected component, the arrangement for Δ_K has size $O(l^2)$ and can be swept in time $O(l^2 \log l)$, where l is the number of cells projecting onto Δ_K . Let s be the number of connected components and l_i , $1 \leq i \leq s$, the number of cells projecting onto Δ_K in the i th connected component. Because the connected components of the free space are disjoint, the sum of the l_i , $1 \leq i \leq s$, is at most m . Thus, for each triangle $\Delta_K \in K$, we can determine, in $O(\sum_{i=1}^s (l_i^2 \log l_i)) \in O(m^2 \log m)$ time, by which of the connected components of the free space it is covered. We have to apply the same procedure with K and L exchanged. This yields a total runtime of $O(nm^2 \log m + mn^2 \log n)$ for all triangles in both parameter spaces.

Summarizing, we obtain the following algorithm for the decision problem of the weak Fréchet distance.

Algorithm 2: DecideWeakFréchet(f, g, ε)**Input:** parametrized triangulated surfaces $f, g, \varepsilon > 0$ **Output:** Is $\delta_{\text{wF}}(f, g) \leq \varepsilon$?

-
- 1 compute the graph G_ε of the free space diagram $F_\varepsilon(f, g)$
 - 2 **forall** connected components A of G_ε **do**
 - 3 **forall** triangles Δ in either parameter space **do**
 - 4 using line sweep decide whether Δ is completely covered by the component A
 - 5 **end**
 - 6 **end**
 - 7 output *true* if a connected component covering all triangles has been found, else output *false*
-

Lemma 3 and the analysis above yield the following theorem:

Theorem 3 *Algorithm 2 decides whether the weak Fréchet distance between two triangulated surfaces with n and m triangles, respectively, is less than a given parameter ε in $O(nm^2 \log m + mn^2 \log n)$ time.*

3.3 Computing the Weak Fréchet Distance

The decision algorithm can be extended to a computation algorithm by searching a set of critical values for ε as is done for the Fréchet distance between polygonal curves [4]. We first characterize these critical values and then show how to compute the weak Fréchet distance by searching over the set of critical values.

By Lemma 3 the weak Fréchet distance corresponds to an extensive connected component of the free space. Thus, the weak Fréchet distance equals a value ε if and only if the free space F_ε contains such a connected component, and it does not for any smaller value of ε . We can characterize the critical values as follows.

1. The combinatorial structure of the connected components changes.
2. The projection of a connected component covers a parameter space completely, whereas for smaller values of ε , this is not the case although there is no combinatorial change.

We now describe this subdivision of the critical values more precisely:

Type 1 The combinatorial structure of the free space changes when a boundary cell, i.e., the cell of an edge and a triangle, becomes nonempty. This happens for ε equal to the distance of the edge and the triangle defining the boundary cell, as stated in (2). There are $O(mn)$ such values, each of which can be computed in constant time.

Type 2 Apart from the combinatorial changes, the projection of a connected component grows monotonously with increasing ε . We distinguish the critical values of this type by the last points in a certain neighborhood covered by the projection, i.e.,

points $\mathbf{x} \in [0, 1]^2$ which are in the projection of the free space $F_\varepsilon(f, g)$ and for any $\varepsilon' < \varepsilon$, there is a neighborhood of \mathbf{x} which is disjoint to $F_{\varepsilon'}(f, g)$.

A last point covered may be either a vertex (Type 2a), an interior point of an edge (Type 2b), or an interior point of a triangle (Type 2c) of the triangulation of the parameter space, or—if the triangles are not in general position—the last points covered may be a segment (Type 2d).

Type 2a If the last point covered is a vertex, then its image (under f or g) lies on the boundary of the ε -neighborhood of a triangle. Thus, its image has distance ε to the triangle. We can compute these critical values by computing for all triangles of the one surface their distance in image space to the vertices of the other surface. Each such value can be computed in constant time and there are $O(mn)$ such values.

Type 2b If the last point covered lies in the interior of an edge of one surface, then the boundaries of ε -neighborhoods of two triangles of the other surface intersect in the image of that point. That is, the point has equal distance to two triangles. There are $O(mn^2 + m^2n)$ such critical values for both parameter spaces, each of which can be computed in constant time.

Type 2c If the last point covered lies in the interior of a triangle, then the boundaries of ε -neighborhoods of three triangles intersect in the point. That is, the point has equal distance to three triangles. There are $O(mn^3 + m^3n)$ such critical values for both parameter spaces, each of which can be computed in constant time.

Type 2d If the triangles are not in general position, it may also happen that the last points covered form a line segment. A detailed analysis shows that this can occur only if two edges, or two triangles, or an edge and a triangle in image space are parallel and ε is the distance between such a parallel pair. There are $O(m^2 + n^2)$ such values, each of which can be computed in constant time.

Summarizing, we have $O(m^2n + n^2m)$ critical values of Types 1, 2a, 2b, and 2d and $O(m^3n + mn^3)$ critical values of Type 2c, each of which can be computed in constant time.

We could now compute the weak Fréchet distance by computing and sorting the critical values and then doing a binary search over the set of critical values, solving Algorithm 2 in each step. The runtime of this algorithm would be dominated by the sorting of the critical values, which takes $O((m^3n + mn^3) \log(mn))$ time.

We can improve the runtime applying parametric search [11] to the $O(m^3n + mn^3)$ critical values of type 2c where the boundaries of three projected cells intersect. For the $O(m^2n + mn^2)$ other critical values, we first run a binary search involving the decision Algorithm 2 in each step. This part takes $O((m^2n + mn^2) \log(mn))$ time.

More precisely, we have the following algorithm.

Algorithm 3: ComputeWeakFréchetDist(f, g)**Input:** Parametrized triangulated surfaces f, g **Output:** $\delta_{\text{wF}}(f, g)$

- 1 Compute the set C_1 of critical values of Types 1, 2a, 2b, and 2d
- 2 Sort the set C_1
- 3 Do a binary search over the set C_1 applying Algorithm 2 in each step and branching to larger values if the answer is false and to smaller values if the answer is true
- 4 Let c_i be the smallest value in C_1 for which Algorithm 2 outputs true
- 5 Do a parametric search over the set C_2 of critical values of Type 2c in the interval $(c_{i-1}, c_i]$, again applying Algorithm 2 in each step
- 6 Output the smallest value for which Algorithm 2 outputs true

For the parametric search, instead of giving a parallel algorithm for the decision problem, it suffices to use any parallel algorithm whose critical values include the critical values of the decision problem. In our case, we can use a parallel comparison-based sorting algorithm for sorting lexicographically the intersection points of the boundary curves of ε -neighborhoods of triangles of the one surface intersected with a plane containing a triangle of the other surface. Formally, we sort all points of the form

$$\partial((\Delta'_1 \oplus B_\varepsilon) \cap E_\Delta) \cap \partial((\Delta'_2 \oplus B_\varepsilon) \cap E_\Delta),$$

where Δ is the image of some triangle in one-parameter space and Δ'_1, Δ'_2 images of triangles of the other parameter space. These are $O(m^2n + n^2m)$ many values, and their critical values include the critical values of the decision algorithm.

More specifically, we use Cole's variant of parametric search based on sorting [6]. His technique yields a runtime of $O((k + T_{\text{dec}}) \log k)$, where T_{dec} is the runtime of the decision algorithm, and k is the number of values to be sorted. In our case, $k \in O(m^2n + n^2m)$ and $T_{\text{dec}} \in O(m^2n \log n + n^2m \log m)$. Thus we get a total runtime of $O((m^2n + mn^2) \log^2(mn))$, which proves Theorem 2.

In practice, it would make sense to check whether the simplified randomized variant of parametric search by van Oostrum and Veltkamp [14], which has already been applied successfully for computing the Fréchet distance of curves, leads to a more efficient implementation.

References

1. Alt, H., Behrends, B., Blömer, J.: Approximate matching of polygonal shapes. *Ann. Math. Artif. Intell.* **13**, 251–266 (1995)
2. Alt, H., Braß, P., Godau, M., Knauer, C., Wenk, C.: Computing the Hausdorff distance of geometric patterns and shapes. In: Aronov, B., Basu, S., Pach, J., Sharir, M. (eds.) *Discrete and Computational Geometry. The Goodman–Pollack Festschrift. Algorithms and Combinatorics*, vol. 25, pp. 65–76. Springer, Berlin (2003)
3. Alt, H., Efrat, A., Rote, G., Wenk, C.: Matching planar maps. *J. Algorithms* 262–283 (2003)
4. Alt, H., Godau, M.: Computing the Fréchet distance between two polygonal curves. *Int. J. Comput. Geom. Appl.* **5**, 75–91 (1995)

5. Buchin, K., Buchin, M., Wenk, C.: Computing the Fréchet distance between simple polygons in polynomial time. In: Proc. 22nd Annu. ACM Symp. Comput. Geom., pp. 80–87. ACM Press, New York (2006)
6. Cole, R.: Slowing down sorting networks to obtain faster sorting algorithms. *J. ACM* **34**(1), 200–208 (1987)
7. Emiris, I.Z., Tsigaridas, E.P.: Comparing real algebraic numbers of small degree. In: Proc. 12th Annu. European Symp. Algorithms. Lecture Notes Comput. Sci., vol. 3221, pp. 652–663. Springer, Berlin (2004)
8. Fréchet, M.: Sur quelques points du calcul fonctionnel. *Rend. Circ. Mat. Palermo* **22**, 1–74 (1906)
9. Fréchet, M.: Sur la distance de deux surfaces. *Ann. Soc. Pol. Math.* **3**, 4–19 (1924)
10. Godau, M.: On the complexity of measuring the similarity between geometric objects in higher dimensions. PhD thesis, Freie Universität Berlin, Germany (1998)
11. Megiddo, N.: Applying parallel computation algorithms in the design of serial algorithms. *J. ACM* **30**(4), 852–865 (1983)
12. Moise, E.E.: *Geometric Topology in Dimensions 2 and 3*. Graduate Texts in Mathematics, vol. 47. Springer, Berlin (1977)
13. Mourrain, B., Pion, S., Schmitt, S., Técourt, J.-P., Tsigaridas, E., Wolpert, N.: Algebraic issues in computational geometry. In: Boissonnat, J.-D., Teillaud, M. (eds.) *Effective Computational Geometry for Curves and Surfaces*. Mathematics and Visualization. Springer, Berlin (2006)
14. van Oostrum, R., Veltkamp, R.C.: Parametric search made practical. *Comput. Geom. Theory Appl.* **28**(2–3), 75–88 (2004)
15. Weihrauch, K.: *Computable Analysis*. Texts in Theoretical Computer Science. An EATCS Series. Springer, Berlin (2000)
16. Weihrauch, K., Zheng, X.: Computability on continuous, lower semi-continuous, and upper semi-continuous real functions. *Theor. Comput. Sci.* **234**, 109–133 (2000)