Chapter 3

Semi-Computability of the Fréchet Distance between Triangulated Surfaces

3.1 Introduction

In this chapter 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 monotone decreasing sequence of rationals converging to the result. It follows that the decision problem, whether the Fréchet distance between two given surfaces is smaller than a specified value ε, is *recursively enumerable*. That is, there is a Turing machine that for two given triangulated surfaces f, g and a value ε will halt if and only if the Fréchet distance between f and g is smaller than ε. Furthermore, we show that the Fréchet distance between computable surfaces is *computably approximable*, i.e., there is a non-halting Turing machine which produces a sequence of rationals converging to the result.

As discussed in Section 1.2, the computationally hard part of computing the Fréchet distance between surfaces is that – according to the definition – the infimum over all homeomorphisms on the parameter spaces has to be taken. For curves, the orientation-preserving homeomorphisms on the unit interval can be characterized as the continuous, onto, monotone increasing functions. For homeomorphisms on the unit k-cube, $k > 1$, no such characterization exists.

We tackle this problem in this chapter 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 subdivisions of the given triangulations of the parameter spaces. Furthermore, we approximate the distances over all points by distances only at vertices of the subdivisions, for arbitrary fine subdivisions.

As discussed in Section 2.1, we assume that the input to our algorithm are two triangulated surfaces in $\mathbb{R}^d, d \geq 2$, which are given by simplicial parameterizations f, g. We will use K, L to denote the underlying triangulations of the parameter spaces of f, g, respectively. As discussed in Section 2.4, we assume that the vertices of the triangulated surfaces and the coefficients of the simplicial parameterizations are rational. Thus, a problem instance for the Fréchet distance between triangulated surfaces has a canonical finite representation which can be given as input to a Turing machine (we will describe our algorithm in some high level language, however). For the case of computable, but not triangulated surfaces we assume to be given two computable parameterizations $f, g\colon [0,1]^2 \to \mathbb{R}^d$.

3.2 Computability of Real-valued Functions

The computability of real-valued functions is studied in *computable analysis* [45, 37, 53]. Definitions for the computability of a real-valued function are based on the computability of a real number. A sequence of rational numbers $(r_k)_{k\in\mathbb{N}}$ is *computable* if it can be computed by a Turing machine, i.e., there exist three recursive functions $a, b, s : \mathbb{N} \to \mathbb{N}$ such that $b(k) > 0$ for all k and $r_k = (-1)^{s_k} a(k)/b(k)$ for all k. A real number x is *computable* if there is a computable sequence of rational numbers *effectively converging* to x, i.e., there is a computable sequence $(r_k)_{k\in\mathbb{N}}$ such that $\|x - r_k\| \le 2^{-k}$ holds for all k. Furthermore, weaker notions of computability are considered [10]:

Definition 3.1.

1. *A real number is* computable *if there is a computable sequence of rational numbers effectively converging to it.*

2. *A real number is* upper (lower) semi-computable *if there is a computable sequence of rational numbers converging to it from below (above).*
   It is semi-computable *if it is upper or lower semi-computable.*

3. *A real number is* weakly computable *if it is the difference of two lower or upper semi-computable numbers.*

4. *A real number is* computably approximable *if there is a computable sequence of rational numbers converging to it.*

Note that for semi-computability the sequence does not need to *effectively* converge. Instead of *computably approximable* also the term *recursively approximable* is used. These definitions carry over to real-valued functions on the natural numbers: A function $\phi : \mathbb{N} \to \mathbb{R}$ is called computably approximable | lower semi-computable | computable if there is a Turing machine that on input x outputs a sequence of rational numbers converging | converging from below | converging effectively to $\phi(x)$.

For the computability of functions $\phi : \mathbb{R} \to \mathbb{R}$ several definitions exists [53] which use different *representations* of real numbers. For example, an effectively converging sequence of rational numbers may be used as a representation of a real number, and a function $\phi : \mathbb{R} \to \mathbb{R}$ may be called computable if there is a Turing machine mapping a representation of $x \in \mathbb{R}$ to a representation of $\phi(x) \in \mathbb{R}$ [37]. The notion of semi-computability may also be defined for functions $\phi : \mathbb{R} \to \mathbb{R}$ [54].

In this chapter, we will show that the Fréchet distance between triangulated surfaces is upper semi-computable and the Fréchet distance between computable surfaces is computably approximable. Since we assume that the triangulated surfaces are given by finite rational input, the Fréchet distance between triangulated surfaces can be interpreted (by appropriate coding of the input) as a real-valued on the natural numbers. For the Fréchet distance between computable surfaces, we will assume to be given two Turing machines T_f, T_g for the parameterizations f, g which we can use as black boxes in the following way: On input $q \in \mathbb{Q}$ and $k \in \mathbb{N}$, T_f outputs a value r_k s.t. $\|f(r) - r_k\| \le 2^{-k}$ and analogously for T_g. We will only evaluate the surface parameterizations for rational coordinates and therefore do not need to work with representations of the input values.

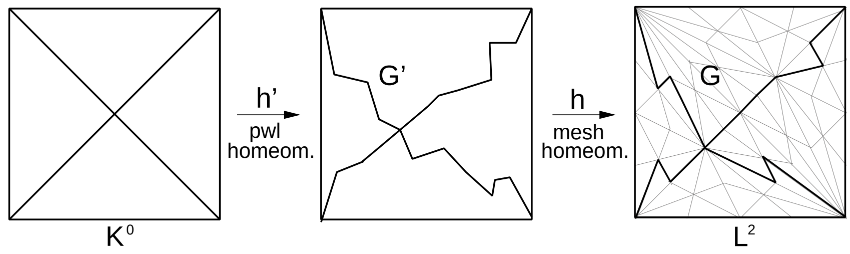

Figure 3.1: Approximating a homeomorphism by a mesh homeomorphism.

3.3 Approximating the Homeomorphisms

In this section, we show that homeomorphisms can be approximated arbitrarily closely by mesh homeomorphisms.

Let us recall some standard definitions and notations from topology. For a simplicial complex K – a triangulation in our case – its *mth barycentric subdivision* K^m is defined as follows. Let $\Delta_1, \ldots, \Delta_m$ be the simplices of K. The barycenter of a simplex $\Delta \in K$ with vertices $v_1, \ldots, v_k$ is $b(\Delta) = 1/k \sum_{i=1}^{k} v_i$. The first barycentric subdivision K' of K is the complex having vertices $b(\Delta_1), \ldots, b(\Delta_m)$ and simplices $(b(\Delta_{i_1}), \ldots, b(\Delta_{i_k}))$ for all sequences of simplices $\Delta_{i_1} \subset \ldots \subset \Delta_{i_n}$ in K. The mth barycentric subdivision is defined inductively by $K^m = (K^{m-1})'$. For triangulations, barycentric subdivision can be described geometrically by subdividing each triangle into six triangles by its angular bisectors. With $mesh(K)$ we denote the maximal diameter of simplices in K, again triangles in our case. Thus, $mesh(K)$ gives a measure for the fineness of K. For $m \to \infty$ the diameters of simplices of the mth barycentric subdivision tend to zero. The *underlying space of* K, denoted by $|K|$, is the set of all points lying in simplices of K. In our case the underlying space $|K|$ is always the unit square $[0,1]^2$.

Let us now define mesh homeomorphisms.

Definition 3.2. *Given two triangulations K and L, a piecewise linear homeomorphism $h : |K^m| \to |L^n|$ is called a* mesh homeomorphism *if it maps the edges of K^m to* edge chains *of L^n, i.e., polygonal chains made up of edges of L^n.*

Next 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 3.3. *Given two homeomorphisms $h, h' : |K| \to |L|$ on the underlying spaces of the triangulations K and L, let*

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

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

Lemma 3.1. *Let K and L be triangulations, $\sigma\colon |K| \to |L|$ a homeomorphism, $m \in \mathbb{N}$, and $\varepsilon > 0$. Then there exist $n \in \mathbb{N}$ and a mesh homeomorphism $h : |K^m| \to |L^n|$ such that $d_{K^m}(\sigma, h) < \varepsilon$.*

Proof. By a theorem from topology (see, e.g., Theorem 4 in Chapter 6 of the book by Moise [40]), the homeomorphism σ can be approximated arbitrarily closely by a piecewise linear homeomorphism, i.e., for all $\varepsilon_1 > 0$ there exists a piecewise linear homeomorphism $h' : |K| \to |L|$ with $d_K(\sigma, h') < \varepsilon_1$. We use this fact as a first step, because piecewise linear homeomorphisms are easier to handle than arbitrary homeomorphisms. In the second step, a piecewise linear homeomorphisms is approximated by a mesh homeomorphism, see Figure 3.1.

We will show that any piecewise linear homeomorphism h' can be approximated to any $\varepsilon_2 > 0$ in the sense of Definition 3.3 by a mesh homeomorphism h, i.e., $d_{K^m}(h, h') < \varepsilon_2$. Choosing ε_1 and ε_2 so that $\varepsilon_1 + \varepsilon_2 < \varepsilon$ then proves the lemma.

In order to approximate h' we first show how to find edge chains in L^n, for some large enough $n \in \mathbb{N}$, 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 will 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 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 the smallest distance between two nodes ensures that disks around different nodes do not touch.

We choose $n' \in \mathbb{N}$ such that $mesh(L^{n'}) < r/2$ 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 $mesh(L^{n'})$ 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 $\deg(h(v)) \geq \deg(v)$. 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 Figure 3.2(a). 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

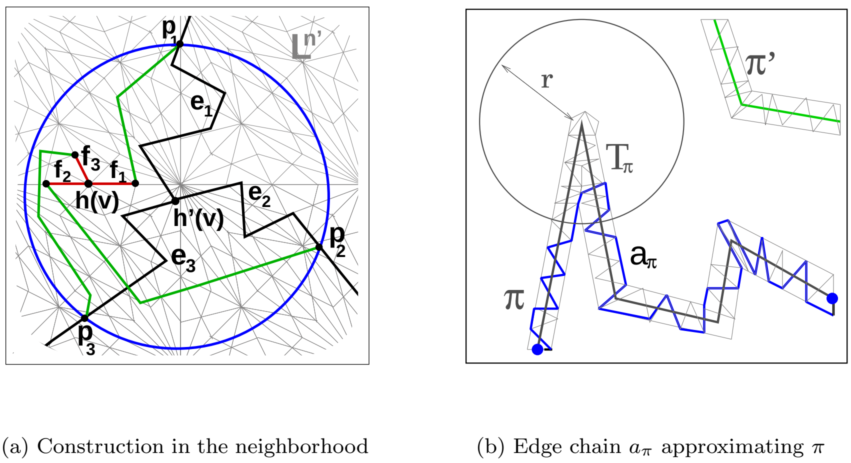

(a) Construction in the neighborhood of a vertex

(b) Edge chain a_π approximating π

Figure 3.2: Approximating a piecewise linear homeomorphism by a mesh homeomorphism.

given by G', i.e., if the edges $e_1, ..., e_l$ leave the vertex $h'(v)$ of G' in clockwise order, we choose corresponding edges $f_1, ..., f_l$ of $L^{n'}$ leaving $h(v)$ in the same order, see Figure 3.2(a).

Next we cut the edges $e_1, ..., e_l$ at the points $p_1, ..., p_l$ where they first leave the disk. Within the disk we connect the free endpoint of each f_i with the point p_i, $i = 1, ..., l$ by non-intersecting polygonal chains, which replace the original polygonal chains from $h'(v)$ to p_i. This is possible because we chose the edges $f_1, ..., f_l$ 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 Figure 3.2 (b). Thus, we replaced all vertices $h'(v)$ by close by vertices $h(v)$ lying on the mesh $L^{n'}$.

Step 2. Next we delete the nodes $h(v)$ together with the incident edges $f_1, ..., f_l$ 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 non-incident 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 $mesh(L^n) < \min(\eta/2, \theta/2)$. 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 Figure 3.2(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. □

Note that the mesh homeomorphism constructed in the proof of Lemma 3.1 is orientation-preserving if the original homeomorphism was. Thus Lemma 3.1 can be strengthened to state that an orientation preserving homeomorphism can be approximated arbitrarily closely by an orientation preserving mesh homeomorphism.

3.4 Discrete Fréchet Distance

In this section we define a discrete Fréchet distance for surfaces and show that it is equal in value to the Fréchet distance.

We define the discrete Fréchet distance between two surfaces by taking the infimum over all mesh homeomorphisms and for each mesh homeomorphism the maximum over distances at vertices. More formally, we define

Definition 3.4. *Let f, g be parameterized triangulated surfaces in $\mathbb{R}^d$, $d \geq 3$, with underlying triangulations K, L respectively, of the parameter space, I.e.,*

$$f\colon |K| \to \mathbb{R}^d, \quad g\colon |L| \to \mathbb{R}^d$$

are piecewise linear maps. Then their discrete Fréchet distance *is defined as*

$$\delta_{dF}(f,g) := \inf_{\substack{m,n \\ h:|K^m|\to|L^n|}} \max_{\Delta\in K_T^m} \max_{\substack{v\in V_\Delta \\ w\in M^n_{h(\Delta)}}} \|f(v)-g(w)\|$$

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 $M^n_{h(\Delta)}$ is the set of vertices of L^n that lie in $h(\Delta)$.

We first show that this definition yields a discrete Fréchet distance not smaller than the Fréchet distance.

Lemma 3.2. *The Fréchet distance between triangulated surfaces f, g is at most as large as their discrete Fréchet distance, i.e., $\delta_F(f,g) \le \delta_{dF}(f,g)$.*

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

$$\max_{t\in[0,1]^2} \|f(t)-g(h(t))\| \quad \le \quad \max_{\Delta\in K_T^m} \max_{\substack{v\in V_\Delta \\ w\in M^n_{h(\Delta)}}} \|f(v)-g(w)\|.$$

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 $\langle f(v_1), f(v_2), f(v_3)\rangle$ 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 $\|f(t) - g(h(t))\| \le \|f(v_i) - g(w_j)\|$ for some i, j with $1 \le i, j \le 3$. Taking the maximum on both sides yields the above equation. ☐

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

Lemma 3.3. *The discrete Fréchet distance between triangulated surfaces f, g is not larger than their Fréchet distance, i.e., $\delta_{dF}(f,g) \le \delta_F(f,g)$.*

Proof. We show that for all $\varepsilon > 0$, $\delta_{dF}(f,g) \le \delta_F(f,g) + \varepsilon$. The idea is that for any homeomorphism there is a mesh homeomorphism arbitrarily close 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. For this, we show that for any homeomorphism σ and point $t \in |K|$, there exists a mesh homeomorphism h and vertices $v \in \Delta \in K^m, w \in h(\Delta)$, s.t. $\|f(v) - g(w)\|$ is not much larger than $\|f(t) - g(\sigma(t))\|$. This is illustrated in Figure 3.3.

Let σ be a homeomorphism very close to realizing $\delta_F(f,g)$, i.e., $\max_t \|f(t) - g(\sigma(t))\| \le \delta_F(f,g) + \varepsilon_1$ for some small $\varepsilon_1 > 0$. By Lemma 3.1, for any $\varepsilon_2 > 0$ and any $m \in \mathbb{N}$ there is a mesh homeomorphism $h : |K^m| \to |L^n|$ such that $d_{K^m}(\sigma, h) \le \varepsilon_2$.

Let Δ be some triangle in $|K^m|$ and v one of its vertices. Since $d_{K^m}(\sigma, h) \le \varepsilon_2$, for any $w \in h(\Delta) \subset L^n$ there is an $x \in \sigma(\Delta)$ 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 $t \in \Delta$ where c_g denotes the Lipschitz constant of g.

t and v lie in the same triangle $\Delta \in K^m$, so $\|v - t\| \le \text{mesh}(K^m)$ and, therefore, $\|f(v) - f(t)\| \le c_f \cdot \text{mesh}(K^m)$ where c_f is the Lipschitz constant for f.

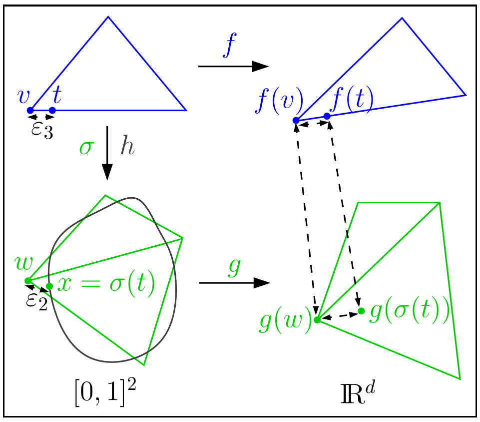

Figure 3.3: The discrete Fréchet distance is not larger than the Fréchet distance.

Putting everything together and applying the triangle inequality repeatedly we get

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

For any $\varepsilon > 0$ we can now choose $\varepsilon_1, \varepsilon_2$, and $\text{mesh}(K^m)$ such that $\varepsilon_1 + c_g \cdot \varepsilon_2 + c_f \cdot \text{mesh}(K^m) \leq \varepsilon$ holds. For instance, we can choose $\varepsilon_1 = \frac{\varepsilon}{3}$, $\varepsilon_2 = \frac{\varepsilon}{3c_g}$, and m large enough s.t. $\text{mesh}(K^m) \leq \frac{\varepsilon}{3c_f}$. Note that although for ε tending to zero, m, n will tend to infinity, for any fixed $\varepsilon > 0$ this yields finite values for m, n. □

Lemmas 3.2 and 3.3 yield the following corollary.

Corollary 3.1. *The Fréchet distance and discrete Fréchet distance between triangulated surfaces f, g are equal, i.e., $\delta_F(f,g) = \delta_{dF}(f,g)$.*

3.5 Semi-Computing the Fréchet Distance

Using the results of the previous sections we can now give an algorithm showing the upper semi-computability of the Fréchet distance between triangulated surfaces. This algorithm will, on input f, g run forever and produce a monotone decreasing sequence of rational numbers converging to $\delta_F(f,g)$.

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

We claim that this algorithm approximates the discrete Fréchet distance which is, by Corollary 3.1, the same as the Fréchet distance.

Theorem 3.1. *The Fréchet distance between two triangulated surfaces in space* $\mathbb{R}^d, d \geq 2$, *is upper semi-computable.*

Proof. We need to show that each step of the algorithm is finitely computable (except, of course, the loop over all pairs of natural numbers) and that the algorithm indeed computes the discrete Fréchet distance.

Line 2 can be realized by some standard enumeration method for pairs of integers. In fact, in the following observation we will see that it would also suffice to consider only the pairs $(m, 2m), m \in \mathbb{N}$. 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. But 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 d_1- or d_∞-metric but not directly for the Euclidean metric d_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 15.

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 preserving 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 $h : K^m \to L^n$ are evaluated by Algorithm 2, i.e.,

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

(see Definition 3.4) is computed.

To see that Algorithm 2 produces values arbitrarily close to $\delta_{dF}(f, g)$, observe that any neighborhood of $\delta_{dF}(f, g)$ must, by Definition 3.4, 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 14 the output sequence is monotone decreasing. Since for all triples (h, m, n), by Definition 3.4, $\delta_{h,m,n} \geq \delta_{dF}(f, g)$, line 14 is justified. Since by Corollary 3.1, $\delta_F = \delta_{dF}$ Algorithm 2 arbitrarily closely approximates $\delta_F(f, g)$ which proves Theorem 3.1. □

Observation 3.1. *In Algorithm 2 it would suffice in line* 2 *to loop over all tuples* $(m, 2m), m \in \mathbb{N}$.

Proof. Let $(m, n) \in \mathbb{N} \times \mathbb{N}$ be arbitrary. If $n \leq 2m$ then L^{2m} is a subdivision of L^n and any mesh homeomorphism $h : |K^m| \to |L^n|$ is also a mesh homeomorphism $h : |K^m| \to |L^{2m}|$.

If $n > 2m$ then we claim that any mesh homeomorphism $h : |K^m| \to |L^n|$ can be extended to a mesh homeomorphism $h' : |K^{n-m}| \to |L^{2(n-m)}|$. For this, observe that $n - m = m + (n - 2m)$ and $2(n - m) = n + (n - 2m)$, i.e., K^{n-m} and $L^{2(n-m)}$ can be obtained from K^m and L^n, respectively, by subdividing $(n - 2m)$ times. Let Δ be any triangle in K^m which is mapped to $h(\Delta) \subset L^n$. Let h' map Δ^{n-2m} mesh homeomorphic to $(h(\Delta))^{n-2m}$. This does not increase the distance achieved by h' compared to the distance achieved by h, i.e. $\delta_{h',n-m,2(n-m)} \leq \delta_{h,m,n}$.

To see this let v be any vertex of Δ^{n-2m} and w any vertex of a triangle $\Delta' \in (h(\Delta))^{n-2m}$. Let p_1, p_2, p_3 be the three vertices defining Δ and q_1, q_2, q_3 the three vertices defining Δ'. Thus, $v = \sum_{i=1}^3 \mu_i p_i$ and $w = \sum_{i=1}^3 \nu_i q_i$. Let $\lambda_i, i = 1, \ldots, l$, be a common subdivision of the μ_i and ν_i, i.e., we can write $v = \sum_{i=1}^l \lambda_i p_i$ and $w = \sum_{i=1}^l \lambda_i q_i$. Then

$$\|f(v) - g(w)\| \; = \; \|\sum_{i=1}^{l} \lambda_i (f(p_i) - g(q_i))\| \; \leq \; \sum \lambda_i \|f(p_i) - g(q_i)\|$$

where the first equality holds due to the linearity of f and g on Δ and Δ', respectively and the second inequality holds by the properties of a norm. □

Theorem 3.1 implies the following corollary, which states that the decision problem for the Fréchet distance between triangulated surfaces with a $<$-sign, i.e., the question "*Is* $\delta_F(f, g) < q$?", is recursively enumerable. Note that Corollary 3.2 cannot be deduced from Theorem 3.1 any more, if we replace the $<$-sign a $\leq$-sign. Let $<f, g, q>$ denote some standard encoding of a triple consisting of two triangulated surfaces f and g, and some rational $q > 0$.

Corollary 3.2. *The set* $\{<f, g, q> \mid \delta_F(f, g) < q\}$ *is recursively enumerable.*

Consider the Turing machine producing a monotone decreasing sequence converging to $\delta_F(f,g)$ which exists by Theorem 3.1. Stop this Turing machine as soon as it produces a value less than q. Thus, the algorithm will eventually halt for all triples $< f, g, q >$ in the language and it will run forever for the ones not in the language.

If we assume more *general surfaces* as input, that is we assume the parameterizations f and g to be *computable real functions*, then we can modify Algorithm 2 to show a weaker form of Theorem 3.1.

Theorem 3.2. *The Fréchet distance between two computable surfaces in space* $\mathbb{R}^d,\ d \geq 2$, *is computably approximable.*

Proof. We assume we are given two Turing machines T_f and T_g for the parameterized surfaces f and g, respectively. On input $(r \in \mathbb{Q}, k \in \mathbb{N})$, T_f, T_g output values $r_k, s_k \in \mathbb{Q}$, respectively, s.t. $\|f(r) - r_k\| \leq 2^{-k}$ and $\|g(r) - s_k\| \leq 2^{-k}$.

We will modify Algorithm 2 such that it can be applied also in this case. To run Algorithm 2 we can choose a triangulation of the vertices $(0,0), (0,1), (1,0), (1,1)$ as initial triangulations K and L. In each step, Algorithm 2 will refine these triangulations by barycentric subdivision. This implies that in each step the sampling density of the point set for which f and g, respectively, are evaluated is increased.

The more important change is that in each step we increase the precision and update previously computed values. We do this as follows: In the kth call of the loop in lines 2–18, Algorithm 2 computes the values for mesh homeomorphisms of the current refinement K^m, L^n. We change this to computing all values for previous refinements $K^{m'}, L^{n'},\ m' \leq m, n' \leq n$ at precision k. Using Observation 3.1 we can change the loop in lines 2–18 to a double loop over all natural numbers k and natural numbers $m \leq k$. In the inner loop over m we compute the mesh homeomorphisms for the subdivision $(m, 2m)$ and evaluate f, g at precision k, i.e., as value for $f(r)$ we use the output of T_f on input (r, k), and analogously for $g(r)$.

The sequence of rational numbers computed in this way will converge to the Fréchet distance by a similar argument as in Theorem 3.1: with m, n tending to infinity, we come infinitely close to all surface points and the discrete Fréchet distance equals the Fréchet distance.

The sequence of rational numbers computed in this way, will, however, not necessarily be monotone decreasing. This is because in any step of the algorithm, we do not have full information on the surfaces parameterized by f, g, but only on a finite sample of points in $[0,1]^2$. Thus in any step of the algorithm, we can still be missing some parts of the surfaces which may be far apart. ☐

Note that because of the missing monotonicity of the computed sequence, a similar corollary as Corollary 3.2 cannot be deduced from Theorem 3.2 for the Fréchet distance between computable surfaces.

3.6 Discussion

In this chapter we have shown the semi-computability of the Fréchet distance between triangulated surfaces. The computability of the Fréchet distance between surfaces in the strong sense of computability theory of real functions remains open, since the sequence produced by Algorithm 2 is not shown to converge effectively. That is, we cannot give an upper bound on the distance of the value produced by the algorithm after k steps to the real value. This is because in any step of the algorithm we consider mesh homeomorphisms $h : |K^m| \to |L^n|$ as approximation of homeomorphisms $\sigma\colon |K| \to |L|$. Although as shown in Lemma 3.1, mesh homeomorphism approximate homeomorphisms arbitrarily closely in the limit, for fixed m, n there is no approximation guarantee.

For a similar problem in topology, the *Simplicial Approximation Theorem* [14], also no approximation guarantee is known. Since in both approaches all homeomorphisms are approximated, it seems difficult to give such an approximation guarantee and thus to strengthen the result of this chapter to show computability in the strong sense. An approximation guarantee on the Fréchet distance might be achieved by a stronger restriction on the set of feasible homeomorphisms. In Chapter 5, we will show such a restriction for the special case of simple polygons. In fact, it seems likely that whether the Fréchet distance is computable is closely linked to the question whether such a restriction is possible.