Semi-Computability of the Fréchet Distance Between Surfaces

Helmut Alt　　Maike Buchin *

Abstract

The Fréchet distance is a distance measure for parameterized curves or surfaces. Using a discrete approximation, we show that for triangulated surfaces it 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 of two given surfaces lies below some specified value, is recursively enumerable.

1 Introduction

The *Fréchet distance* was first introduced by Fréchet for curves [Fré06] and later for surfaces [Fré24]. The idea of the Fréchet distance is to take into account the "flow" of the curve or surface given by its parameterization. In some cases, the Fréchet distance is a more suitable distance measure than the commonly used Hausdorff distance (see [AG95]).

Formally the Fréchet distance is defined as follows.

Definition 1 *Let f, g be parameterizations of k–dimensional surfaces, i.e., continuous functions*

$$f, g : [0,1]^k \to \mathbb{R}^d, \quad k \leq d.$$

Then their Fréchet distance is

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

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

The norm $||.||$ underlying the definition in this paper can be the L_1-, L_2-, or L_∞-norm as long as it can be computed or approximated by rational arithmetic.

For dimension $k = 1$ of the parameter space, in particular for polygonal curves, δ_F is known to be computable in polynomial time [AG95]. For two–dimensional surfaces, however, the computation of the Fréchet distance surprisingly seems to be much harder. In fact, Godau showed [God98] that computing the Fréchet distance between triangulated surfaces even in two–dimensional space is NP-hard. It remained open, how hard the problem really is, not even its computability could be shown.

In this paper, we present a partial result concerning the computability. More specifically, we will 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 of two given surfaces lies below some specified value is *recursively enumerable*.

The computationally hard part of computing the Fréchet distance for dimensions $k > 1$ seems to be, that according to the definition, the infimum over all homeomorphisms of the parameter space has to be taken. For dimension $k = 1$ the orientation-preserving homeomorphisms on $[0,1]$ are the continuous, onto, monotone increasing functions on $[0,1]$. For higher dimensions the homeomorphisms can be much "wilder".

We tackle this problem 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. The piecewise linear homeomorphisms are then approximated by *mesh homeomorphisms*, i.e., homeomorphisms that are compatible with certain subdivisions of the original triangulations of the parameter spaces. Finally, for mesh homeomorphisms on fine subdivisions the distance between the surfaces can be approximated by the distances at only a finite number of points.

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

2 Model of computation, main results

We assume that the input to our algorithm are two triangulated surfaces in space $\mathbb{R}^d, d \geq 2$, which are represented as *piecewise linear* parameterizations $f, g : [0,1]^2 \to \mathbb{R}^d$. 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 f and g are linear maps in the sense that for a triangle $\Delta = \langle u, v, w \rangle$ we have $f(\lambda_1 u + \lambda_2 v + \lambda_3 w) = \lambda_1 f(u) + \lambda_2 f(v) + \lambda_3 f(w)$ for all $\lambda_1, \lambda_2, \lambda_3$ with $\lambda_1 + \lambda_2 + \lambda_3 = 1$ and g has an analogous property.

*Freie Universität Berlin, Institut für Informatik, {alt,mbuchin}@inf.fu-berlin.de. This research was partially supported by the Deutsche Forschungsgemeinschaft within the European graduate program 'Combinatorics, Geometry, and Computation' (No. GRK 588/2).

We denote the triangulated parameter spaces of f and g by K and L. 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 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 complexity-of-real-functions community (see, e.g., [WZ00]).

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

Now we can formulate our main result:

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

In fact, consider the Turing machine producing a monotone 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. This algorithm will eventually halt for all triples $\langle f, g, a\rangle$ in the language and else will run forever.

The computability of δ_F in the strong sense of computability theory of real functions (see, e.g., [Wei00]) remains open, since the sequence produced by the algorithm in the proof of Theorem 1 is not shown to *effectively* converge to $\delta_F(f, g)$, i.e., we cannot give any estimate on the rate of convergence.

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 monotone decreasing, and the corollary cannot be deduced anymore.

3 Approximating the homeomorphisms

In this section, we approximate homeomorphisms arbitrarily closely by mesh homeomorphisms.

Let us first recall some standard definitions and notations from topology. For a simplicial complex K, a triangulation in our case, let K^m denote its m^{th} *barycentric subdivision*, where in one subdivision step the barycenters of the previous simplices are taken as vertices. *Mesh(K)* denotes the maximal diameter of simplices in K, again triangles in our case. The *underlying space of* K, denoted by $|K|$, is the set of all points lying in simplices of K. In our case $|K|$ is always the unit square $[0, 1]^2$.

We now define mesh homeomorphisms.

Definition 3 *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.*

For approximating homeomorphisms arbitrarily closely by mesh homeomorphisms, we need only a weak form of closeness which is defined as follows.

Definition 4 *Given two homeomorphisms $h, h' : |K| \to |L|$ on 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.

Now we can approximate homeomorphisms by mesh homeomorphisms.

Lemma 3 *Let K and L be triangulations, $\sigma : |K| \to |L|$ 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. We omit the details of this proof in this extended abstract but sketch the main idea.

By a theorem from topology (see, e.g., chapter 6 in [Moi77]), a homeomorphism can be approximated arbitrarily closely by a piecewise linear homeomorphism. We use this as a first step, because piecewise linear homeomorphisms are easier to handle than arbitrary homeomorphisms. For a piecewise linear homeomorphism, we see that it can be approximated arbitrarily closely (in the sense of Definition 4) by a mesh homeomorphism. Together this proves the lemma.

The idea of approximating piecewise linear homeomorphisms by mesh homeomorphisms, is to subdivide sufficiently using, e.g., barycentric subdivision. Because of growing degrees of vertices and growing fineness of the triangulations, we can find mesh homeomorphisms arbitrarily close to a piecewise linear homeomorphism. A simple example is shown in Figure 1.

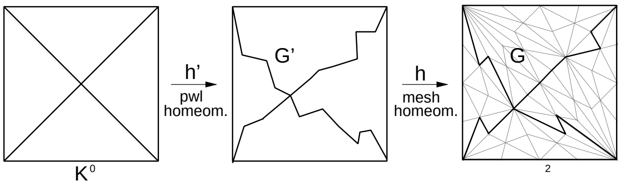

Figure 1: Approximating a piecewise linear homeomorphism by a mesh homeomorphism.

□

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 of two surfaces by taking the infimum over all mesh homeomorphisms and for each mesh homeomorphism taking the maximum over distances at vertices.

More formally, we define

Definition 5 *Let f, g be parametrized, triangulated two–dimensional surfaces in $\mathbb{R}^d$, $d \geq 2$, with underlying triangulations K, L respectively, of the parameter space, i.e.,*

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

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 $h(\Delta)$.

First we show that this definition yields a discrete Fréchet distance not smaller than the Fréchet distance.

Lemma 4 $\delta_F \leq \delta_{dF}$

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.,

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

To see this, let $t \in [0,1]^2$ be arbitrary. Then t lies in a triangle Δ of K^m and $h(t)$ lies in a triangle Δ' 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. Since the maximum distance between points of two triangles is attained between two corners, we have that $||f(t) - g(h(t))|| \leq ||f(v) - g(w)||$ for some $v \in \Delta, w \in \Delta'$. Taking the maximum on both sides yields equation (1). □

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

Lemma 5 *For all $\varepsilon > 0$, $\quad \delta_{dF} \leq \delta_F + \varepsilon$.*

Proof. 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.

Let σ be a homeomorphism close to realizing δ_F, i.e., $\max\limits_{t\in[0,1]^2} ||f(t) - g(\sigma(t))|| \leq \delta_F + \varepsilon_1$ for a small $\varepsilon_1 > 0$.

By Lemma 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 $d_{K^m}(\sigma, h) \leq \varepsilon_2$.

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 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.

Since t and v lie in the same triangle $\Delta \in K^m$, we have $||v - t|| \leq \text{mesh}(K^m)$ and $||f(v) - f(t)|| \leq c_f \cdot \text{mesh}(K^m)$ with c_f the Lipschitz constant of f.

Putting everything together and using the triangle inequality repeatedly we get

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

Since $\varepsilon_1, \varepsilon_2$, and mesh$(K^m)$ can be made arbitrarily small, this concludes the proof. □

Lemmas 2 and 3 yield the following corollary.

Corollary 6 $\delta_F = \delta_{dF}$

5 Semi-Computability of the Fréchet distance

We can now give an algorithm showing the upper semi-computability of the Fréchet distance between triangulated surfaces as claimed in Theorem 1. 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 CompFrec(f,g)

Input: Triangulated surfaces f, g, including triangulations K, L of the parameter spaces, in a finite description as explained in Section 2.

1 $D := \infty$;

2 for all $(m, n) \in \mathbb{N} \times \mathbb{N}$ do

2.1 generate the barycentric subdivisions K^m of K and L^n of L, let $E = \{e_1, ..., e_k\}$ be the set of edges in K^m;

2.2 for all k-tuples $(\pi_1, ..., \pi_k)$ of simple polygonal chains in L^n do

2.2.1 assign the polygonal chain π_i to the edge e_i for $i = 1, ..., k$ and check whether this assignment results in an orientation preserving homeomorphic image of K^m, i.e., whether

2.2.1.1 the edges on the boundary of $|K|$ are mapped onto the boundary of $|L|$ preserving the orientation; and

2.2.1.2 if a set of edges in K^m share an endpoint, the corresponding chains do, as well; and

2.2.1.3 other than that, there are no intersection points between two chains;

2.3 If the test in step 2.2.1 is passed, the chains form a subdivision of $|L|$ such that each triangle Δ of K^m has a corresponding area $H_\Delta \subset |L|$.

2.3.1 for each triangle Δ of K^m do

2.3.1.1 for all vertices v of Δ and all vertices w of L^n lying in H_Δ do compute $||f(v) - g(w)||$;

2.3.2 M := the maximum of all the values found in step 2.3.1.1;

2.3.3 $D := \min(D, M)$; output D;

In essence, algorithm *CompFrec* approximates the discrete Fréchet distance which is, by Section 4 the same as the Fréchet distance. Line 2 can be realized by some standard enumeration method for pairs of integers.

Observe, that the number of k-tuples of polygonal chains of L^n checked in step 2.2 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 step 2.3.1.1 we assume that the norm $||.||$ underlying the Fréchet distance can be evaluated by rational operations. This is correct for, e.g., the L_1- or L_∞-norm but not directly for L_2. In that case, one should rather operate with the square of the distance in line 2.3.1.1 and output some suitable rational approximation of $\sqrt{D}$ (which is possible) in line 2.3.3.

Note that checking that the boundary of $|K|$ is mapped orientation preserving onto the boundary of $|L|$ in step 2.2.1.1, 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, 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 5) is computed[1].

To see that the algorithm produces values arbitrarily close to $\delta_{dF}(f, g)$, observe that any neighborhood of $\delta_{dF}(f, g)$ must, by Definition 5, 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 2.3.3 the output sequence is monotone decreasing. Since for all triples (h, m, n) by Definition 5, $\delta_{h,m,n} \geq \delta_{dF}(f, g)$, line 2.3.3 is justified.

Since by Corollary 6 $\delta_F = \delta_{dF}$, algorithm *CompFrec* arbitrarily closely approximates $\delta_F(f, g)$ which proves Theorem 1.

References

[AG95] H. Alt and M. Godau. Computing the Fréchet distance between two polygonal curves. *Internat. J. Comput. Geom. Appl.*, 5:75–91, 1995.

[Fré06] M. Fréchet. Sur quelques points du calcul fonctionnel. *Rendic. Circ. Mat. Palermo*, 22:1–74, 1906.

[Fré24] M. Fréchet. Sur la distance de deux surfaces. *Ann. Soc. Polonaise Math.*, 3:4–19, 1924.

[God98] M. Godau. *On the complexity of measuring the similarity between geometric objects in higher dimensions.* PhD thesis, Freie Universität Berlin, 1998.

[Moi77] E. E. Moise. *Geometric Topology in Dimensions 2 and 3*, volume 47 of *Graduate Texts in Mathematics.* Springer-Verlag, New York, Heidelberg, Berlin, 1977.

[Wei00] K. Weihrauch. *Computable Analysis.* Texts in Theoretical Computer Science. An EATCS Series. Springer-Verlag, Berlin, Heidelberg, 2000.

[WZ00] K. Weihrauch and X. Zheng. Computability on continuous, lower semi-continuous, and upper semi-continuous real functions. *Theoretical Computer Science*, 234:109–133, 2000.

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