Chapter 4

Computability of the Weak Fréchet Distance between Triangulated Surfaces

4.1 Introduction

In this chapter 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 using surjective continuous maps as reparameterizations of the surfaces. Its definition is given in Section 4.2.

In Section 4.3 we analyze the free space diagram in higher dimension. In Section 4.4 we show that deciding the weak Fréchet distance is equivalent to determining whether the free space contains a connected component whose projection onto both parameter spaces covers the whole parameter space. We give a polynomial time algorithm for determining if such a connected component exists in Section 4.5.

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 in constant time [23, 42].

4.2 Weak Fréchet Distance

The weak Fréchet distance does not require the reparameterizations of the curves or surfaces to be injective. Instead it uses surjective continuous maps as reparameterizations. Furthermore, one can distinguish between the weak Fréchet distance with and without boundary condition. The weak Fréchet distance with boundary condition requires that the boundary of one parameter space is mapped onto the boundary of the other parameter space.

Definition 4.1. *The weak Fréchet distance between two k-dimensional surfaces given by continuous parameterizations $f, g\colon [0,1]^k \to \mathbb{R}^d$ with $k \le d$ is*

$$\delta_{wF}(f,g) := \inf_{\substack{\alpha,\beta\colon\ [0,1]^k \to [0,1]^k \\ \text{surj. cont.}}} \ \max_{\mathbf{x}\in[0,1]^k} \ \|f(\alpha(\mathbf{x})) - g(\beta(\mathbf{x}))\| \ ,$$

where α and β range over all surjective continuous maps on the unit k-cube.

The weak Fréchet distance with boundary condition is

$$\delta_{wFb}(f,g) := \inf_{\substack{\alpha,\beta:\ [0,1]^k\to[0,1]^k \\ \textit{surj. cont.} \\ \textit{boundary cond.}}} \max_{\mathbf{x}\in[0,1]^k} \|f(\alpha(\mathbf{x})) - g(\beta(\mathbf{x}))\| ,$$

where α and β range over all surjective continuous maps on the unit k-cube that map boundary to boundary.

For the weak Fréchet distance in contrast to the Fréchet distance it is necessary to reparameterize both curves or surfaces. For the Fréchet distance it suffices to reparameterize one curve or surface because the inverse of an orientation-preserving homeomorphism is again an orientation-preserving homeomorphism. Thus, instead of reparameterizing both curves or surfaces by α, β we can reparameterize the second curve or surface by $\beta^{-1} \circ \alpha$.

In the man-dog illustration of the weak Fréchet distance the man and dog are allowed to walk also backwards. For the weak Fréchet distance without boundary condition the man and dog may also choose their starting and endpoints.

The weak Fréchet distance with and without boundary condition are (by definition) larger or equal to the Hausdorff distance and smaller or equal the Fréchet distance. That is, the following relation holds between the different distance measures:

$$\delta_H(f,g) \;\le\; \delta_{wF}(f,g) \;\le\; \delta_{wFb}(f,g) \;\le\; \delta_F(f,g).$$

For curves, this relation is also reflected in the free space diagram, i.e., in the geometric data structure for computing the Fréchet distance between polygonal curves [6] which we reviewed in Section 2.3.1. The Hausdorff and the weak Fréchet distance can be characterized by (subsets of) the free space, whose projections onto the parameter spaces completely cover both parameter spaces. We will call such subsets *extensive*, i.e., a subset $A \subset F_\varepsilon(f,g)$ is called extensive if the projection of A onto the parameter space of f equals the parameter space of f and analogously for the parameter space of g.

1. Hausdorff distance less than ε is equivalent to the free space being extensive.
2. Weak Fréchet distance less than ε is equivalent to one connected component of the free space being extensive.
3. Weak Fréchet distance with boundary condition less than ε is equivalent to one connected component of the free space containing both corner vertices $(0,0)$ and $(1,1)$.
4. Fréchet distance less than ε is equivalent to a connected component of the free space containing a monotone path from $(0,0)$ to $(1,1)$.

These characterizations are illustrated for curves in Figure 4.1. The characterization of the Hausdorff distance holds also for surfaces. In this chapter we will show that the characterization of the weak Fréchet distance holds for triangulated surfaces. This characterization will allow us to compute the weak Fréchet distance between triangulated surfaces in polynomial time.

Thus the value of the weak Fréchet distance lies between the value of the Hausdorff distance, which can be computed in polynomial time but is a coarser distance measure, and the value of the Fréchet distance, which is a finer distance measure but not known to be computable in higher dimensions.

A polynomial time algorithm for computing the weak Fréchet distance with boundary condition between polygonal curves has been developed by Alt and Godau [6] (there it is called the *non-monotone Fréchet distance*).

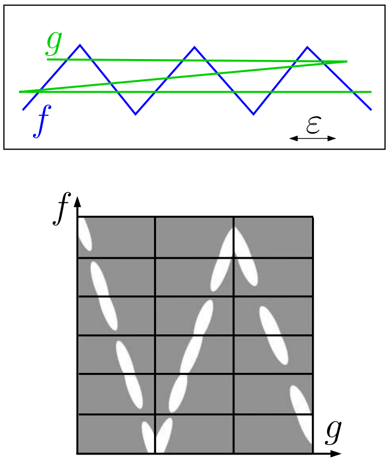

(a) small Hausdorff distance

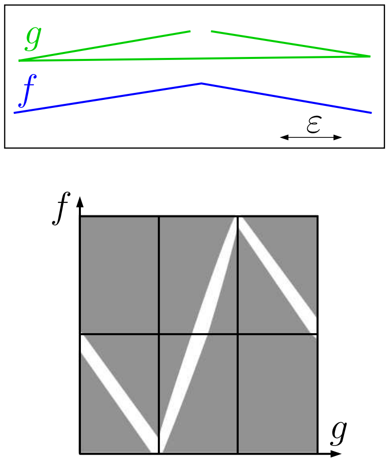

(b) small weak Fréchet distance

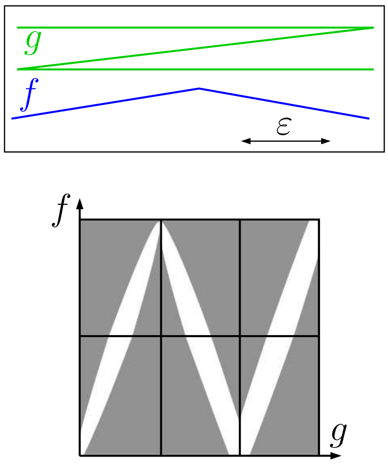

(c) small weak Fréchet distance with boundary condition

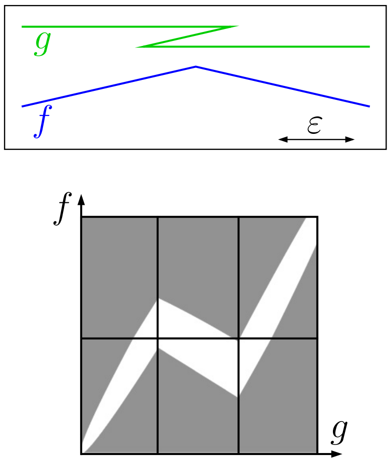

(d) small Fréchet distance

Figure 4.1: Examples of pairs of curves that have (a) a small Hausdorff distance but large weak Fréchet distance, (b) a small weak Fréchet distance but large weak Fréchet distance with boundary condition, (c) a small weak Fréchet distance with boundary condition but large Fréchet distance, and (d) a small Fréchet distance.

4.2.1 Is the Triangle Inequality fulfilled?

The weak Fréchet distance with boundary condition is a pseudo-metric for curves [28], i.e., it fulfills the triangle inequality but curves with distance zero may be distinct. For the weak Fréchet distance between surfaces it is not known whether the triangle inequality holds.

To see why the proof of the triangle inequality does not generalize to higher dimensions, consider three curves or surfaces given by parameterizations f, g, h. For the triangle inequality one needs to show

$$\delta_{wF}(f,h) \;\leq\; \delta_{wF}(f,g) \;+\; \delta_{wF}(g,h).$$

Let $\varepsilon > 0$, then there exist reparameterizations $\alpha, \beta, \gamma, \delta$ s.t. holds for all t

$$\begin{aligned} \|f \circ \alpha(t) - g \circ \beta(t)\| &\leq \delta_{wF}(f,g) + \varepsilon \quad \text{and} \\ \|g \circ \gamma(t) - h \circ \zeta(t)\| &\leq \delta_{wF}(g,h) + \varepsilon. \end{aligned}$$

For curves, the mountain climbing theorem [38] states that there are reparameterizations η, ϑ s.t. $\beta \circ \eta = \gamma \circ \vartheta$. Thus, we can reparameterize f by $\alpha \circ \eta$ and h by $\zeta \circ \vartheta$, and get

$$\begin{aligned} &\|f \circ \alpha \circ \eta(t) - h \circ \zeta \circ \vartheta(t)\| \\ \leq\;& \|f \circ \alpha \circ \eta(t) - g \circ \beta \circ \eta(t)\| \;+\; \|g \circ \gamma \circ \vartheta(t) - h \circ \zeta \circ \vartheta(t)\| \\ \leq\;& \delta_{wF}(f,g) \;+\; \varepsilon \;+\; \delta_{wF}(g,h) \;+\; \varepsilon. \end{aligned}$$

In the first step we use the triangle inequality of the underlying metric in image space and the equality $\beta \circ \eta = \gamma \circ \vartheta$.

For surfaces, however, the mountain climbing theorem cannot be applied because it is not known to hold in higher dimensions. We could try instead to reparameterize f by $\alpha \circ \gamma$ and h by $\beta \circ \zeta$, which yields

$$\begin{aligned} &\|f \circ \alpha \circ \gamma(t) - h \circ \beta(t) \circ \zeta(t)\| \\ \leq\;& \|f \circ \alpha \circ \gamma(t) - g \circ \beta \circ \gamma\| \;+\; \|g \circ \beta \circ \gamma(t) - h \circ \beta \circ \zeta(t)\| \\ \leq\;& \delta_{wF}(f,g) \;+\; \varepsilon \;+\; \|g \circ \beta \circ \gamma(t) - h \circ \beta \circ \zeta(t)\|. \end{aligned}$$

The term $\|g \circ \beta \circ \gamma(t) - h \circ \beta \circ \zeta(t)\|$ however can be much larger than $\delta_{wF}(f,g)$.

4.3 Free Space Diagram of Triangulated Surfaces

In this section we analyze the free space diagram of triangulated surfaces. We consider parameterized triangulated surfaces given by simplicial maps $f, g\colon [0,1]^2 \to \mathbb{R}^d,\ d \geq 2$. In the following, we will always use $\mathbf{x}, \mathbf{y}$ to denote points in parameter space, i.e., $\mathbf{x}, \mathbf{y} \in [0,1]^2$. Let $F_\varepsilon(f,g) := \big\{(\mathbf{x},\mathbf{y}) \;\big|\; \|f(\mathbf{x}) - g(\mathbf{y})\| \leq \varepsilon\big\}$ be the free space diagram, as defined by Alt and Godau [6] and reviewed in Section 2.3.1. The free space diagram lies in the product of the parameter spaces of the two surfaces, i.e., for two-dimensional surfaces it lies in the four-dimensional cube. As in the case of curves, we can partition the free space into cells and show that it is convex inside cells.

4.3.1 Cells of the Free Space Diagram

Definition 4.2. *Given simplicial maps $f, g\colon [0,1]^2 \to \mathbb{R}^d,\ d \geq 2$, with underlying triangulations K and L, respectively. I.e., $|K| = |L| = [0,1]^2$ and f and g map triangles of K and L, respectively, linearly into $\mathbb{R}^d$. A cell of the free space is the* *free space of f and g restricted to a simplex of K and L, respectively. Let Δ_K be a simplex of K and Δ_L a simplex of L. Then we define a cell of the free space as*

$$\begin{aligned} C_\varepsilon(\Delta_K,\Delta_L) &:= F_\varepsilon(f,g)\cap(\Delta_K\times\Delta_L) \\ &= \big\{(\mathbf{x},\mathbf{y})\in\Delta_K\times\Delta_L \;\big|\; \|f(\mathbf{x})-g(\mathbf{y})\| \le \varepsilon\big\}. \end{aligned}$$

The dimension of a cell is the sum of the dimensions of the simplices Δ_K and Δ_L. Two l-dimensional cells are called neighboring *if they share a $l-1$-dimensional cell.*

We will be interested mostly in cells where both simplices are triangles. These cells are four-dimensional and two such cells are neighboring if they share a three-dimensional cell, which is the cell of a triangle and an edge. More explicitly, two four-dimensional cells $C_\varepsilon(\Delta_K,\Delta_L)$ and $C_\varepsilon(\Delta'_K,\Delta'_L)$ are neighboring if and only if either $\Delta_K=\Delta'_K$ and $\Delta_L\cap\Delta'_L$ is an edge of the triangulation L or $\Delta_L=\Delta'_L$ and $\Delta_K\cap\Delta'_K$ is an edge of the triangulation K. The shared three-dimensional cell is the cell of the shared triangle and the shared edge. In the following, with a *cell* of the free space we will mean – unless stated otherwise – a four-dimensional cell, i.e., the cell of two triangles, and with a *boundary cell* a three-dimensional cell, i.e., the cell of a triangle and an edge.

As for curves the cells of the free space are convex. In fact, this holds for parameterized simplicial shapes of any dimension:

Lemma 4.1. *Cells of the free space diagram of k-dimensional parameterized simplicial shapes are convex.*

Proof. This follows from the simpliciality of the parameterizations and the triangle inequality of the underlying metric in $\mathbb{R}^d$. Let f and g be two simplicial parameterizations of two k-dimensional shapes. Let x,x' be two points in the same simplex of the underlying simplicial complex of f and y,y' two points in the same simplex of the underlying simplicial complex of g. Furthermore, let both (x,y) and (x',y') lie in free space, i.e., $\|f(x)-g(y)\|<\varepsilon$ and $\|f(x')-g(y')\|<\varepsilon$. Let $0\le\lambda\le1$. Then also the point $\lambda(x,y)+(1-\lambda)(x',y')$ lies in the free space cell:

$$\begin{aligned} &\|f\big(\lambda x+(1-\lambda)x'\big)-g\big(\lambda y+(1-\lambda)y'\big)\| \\ =\;& \|\lambda f(x)+(1-\lambda)f(x')-\lambda g(y)+(1-\lambda)g(y')\| \\ \le\;& \lambda\|f(x)-g(y)\|+(1-\lambda)\|f(x')-g(y')\| \\ \le\;& \lambda\varepsilon+(1-\lambda)\varepsilon \;=\; \varepsilon \end{aligned}$$

The first equality holds because x,x' and y,y' come from the same simplex on which f and g, respectively, are linear. The second inequality holds because of the triangle inequality of the underlying metric in $\mathbb{R}^d$. The third inequality holds by the assumption that (x,y) and (x',y') lie in free space. □

4.3.2 Projections of the Free Space Diagram

We will consider the projection of the free space onto the two parameter spaces, i.e., the projection of $F_\varepsilon(f,g)$ under

$$\begin{aligned} \mathrm{Proj}_K &: [0,1]^4\to[0,1]^2, \qquad (\mathbf{x},\mathbf{y})\mapsto\mathbf{x}, \\ \mathrm{Proj}_L &: [0,1]^4\to[0,1]^2, \qquad (\mathbf{x},\mathbf{y})\mapsto\mathbf{y}. \end{aligned}$$

We will compute the projections of the free space by computing it for its cells, i.e., we will compute

$$\begin{aligned} \mathrm{Proj}_K\big(C_\varepsilon(\Delta_K,\Delta_L)\big) &= \big\{\mathbf{x}\in\Delta_K \;\big|\; \exists\mathbf{y}\in\Delta_L : \|f(\mathbf{x})-g(\mathbf{y})\| \le \varepsilon\big\}, \\ \mathrm{Proj}_L\big(C_\varepsilon(\Delta_K,\Delta_L)\big) &= \big\{\mathbf{y}\in\Delta_L \;\big|\; \exists\mathbf{x}\in\Delta_K : \|f(\mathbf{x})-g(\mathbf{y})\| \le \varepsilon\big\}. \end{aligned}$$

Note that the free space lies in the product of the parameter spaces but is defined based on the distances in image space. That is, the projection of a cell of the free space contains all those points, which when mapped into image space are close to a point on the other surface in image space. More formally,

$$\begin{aligned} \mathrm{Proj}_K\big(C_\varepsilon(\Delta_K,\Delta_L)\big) &= f^{-1}\big(f(\Delta_K)\cap(g(\Delta_L)\oplus B_\varepsilon)\big) \\ \mathrm{Proj}_L\big(C_\varepsilon(\Delta_K,\Delta_L)\big) &= g^{-1}\big(g(\Delta_L)\cap(f(\Delta_K)\oplus B_\varepsilon)\big), \end{aligned} \tag{4.3.1}$$

where $\Delta\oplus B_\varepsilon$ denotes the Minkowski sum of the simplex Δ with a ball of radius ε.

4.3.3 Combinatorial Structure of the Free Space Diagram

The combinatorial structure of the free space is captured by the graph whose vertices are the cells of the free space and edges exist between neighboring cells which share a non-empty boundary face. For a free space diagram $F_\varepsilon(f,g)$ we define

$$\begin{aligned} G_\varepsilon &= (V,E_\varepsilon) \quad \text{where} \\ V &= \big\{(\Delta_K,\Delta_L)\;\big|\;\Delta_K \text{ triangle in } K,\ \Delta_L \text{ triangle in } L\big\} \quad \text{and} \\ E_\varepsilon &= \big\{\big((\Delta_K,\Delta_L),(\Delta_K,\Delta'_L)\big)\,\big|\,e_L=\Delta_L\cap\Delta'_L \text{ is an edge of } L\wedge C_\varepsilon(\Delta_K,e_L)\neq\emptyset\big\} \\ &\cup\ \big\{\big((\Delta_K,\Delta_L),(\Delta'_K,\Delta_L)\big)\,\big|\,e_K=\Delta_K\cap\Delta'_K \text{ is an edge of } K\wedge C_\varepsilon(e_K,\Delta_L)\neq\emptyset\big\}. \end{aligned}$$

The graph G_ε has mn vertices. The number of edges is at most the number of triangles in K times the number of edges in L plus the number of edges in K times the number of triangles in L. In a triangulation, the number of edges is at most $2n+1$ if n is the number of triangles. Thus, the number of edges is $O(mn)$.

Note that for $\varepsilon_1\le\varepsilon_2$ the graph G_{ε_1} is a subgraph of G_{ε_2}. Furthermore, G_ε changes for varying ε only at finitely many values, namely when a boundary face becomes non-empty. The number of boundary cells is finite, it is $m'n+mn'\in O(mn)$ where m,m' and n,n' are the number of triangles and edges in K and L, respectively.

A boundary cell $C_\varepsilon(e_K,\Delta_L)$ is non-empty if and only if the distance between the edge e_K and the triangle Δ_L is less than ε. In formulas,

$$C_\varepsilon(e_K,\Delta_L)\neq\emptyset \quad\Leftrightarrow\quad \mathrm{dist}\big(f(e_K),g(\Delta_L)\big)\le\varepsilon, \tag{4.3.2}$$

where the distance between an edge and a triangle is defined as the minimum distance between a point on the edge and a point on the triangle, i.e., $\mathrm{dist}(e,\Delta)=\min_{x\in e,y\in\Delta}\|x-y\|$.

We will search for extensive connected components in the free space. A necessary condition for a connected component to be extensive is that it contains a cell $C_\varepsilon(\Delta_K,\Delta_L)$ for all triangles Δ_K in K as well as for all triangles Δ_L in L. In particular, the connected component needs to have a size of at least $m+n-1$. We will call connected components of this size *large*. There are at most $\frac{mn}{m+n-1}\le\min(m,n)$ large connected components in a free space of size mn.

4.4 Characterizing the Weak Fréchet Distance

In this section we show that we can decide the weak Fréchet distance by searching for an extensive connected component in free space.

Lemma 4.2. *The weak Fréchet distance between two triangulated surfaces is less than ε if and only if there is an extensive connected component $A\subset F_\varepsilon(f,g)$ in the free space for the parameter ε, i.e., $Proj_K(A)=Proj_L(A)=[0,1]^2$.*

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

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

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:

$$M_{\alpha,\beta} \subseteq F_\varepsilon(f,g) \;\Leftrightarrow\; \max_{\mathbf{x}\in[0,1]^2} \|f(\alpha(\mathbf{x})) - g(\beta(\mathbf{x}))\| \le \varepsilon. \tag{4.4.1}$$

Now we show the two directions of the lemma.

$\Rightarrow$: If α, β are two reparameterizations for f and g, respectively, realizing the weak Fréchet distance, i.e., $\max_{\mathbf{x}\in[0,1]^2} \|f(\alpha(\mathbf{x})) - g(\beta(\mathbf{x}))\| = \varepsilon$ then $M_{\alpha,\beta}$ is an extensive connected component in $F_\varepsilon(f,g)$

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

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

By equation 4.4.1 this implies that $F_\vartheta(f,g)$ contains for all $\vartheta > \varepsilon$ an extensive connected component. We will show that this holds also for $F_\varepsilon(f,g)$.

As discussed in 4.3.3 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 non-empty. Let

$$\eta := \arg\min\{\vartheta \mid \vartheta > \varepsilon \wedge \text{a boundary cell becomes non-empty in } F_\vartheta(f,g)\}.$$

Then the combinatorial structure of F_ϑ 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. A free space diagram has only finitely many connected components, and therefore there must be one connected component that is extensive in infinitely many of the F_{ε_n}. Although the free space does not change combinatorially for any of the ε_n, its projection shrinks with decreasing ε. Therefore the same connected component that is extensive in F_{ε_m} need not be extensive in F_{ε_k} for $m < k$.

We claim however, that the connected component which is extensive in infinitely many of the F_{ε_n} is also extensive in F_ε. Assume this is not the case, i.e., there is a point in parameter space that does not lie in the projection of the connected component in F_ε. Then because the projection of the free space is closed this must have already been the case for a small neighborhood $[\varepsilon, \eta') \subset [\varepsilon, \eta)$, i.e., for all but finitely many of the $\varepsilon_n, n \in \mathbb{N}$. This is a contradiction to the assumption.

$\Leftarrow$: Let $A \subset F_\varepsilon(f,g)$ be an extensive connected component of the free space. We will construct reparameterizations α and β of f and g, respectively, such that $M_{\alpha,\beta} \subset A$ and $\mathrm{Proj}_K(M_{\alpha,\beta}) = \mathrm{Proj}_K(A)$ and $\mathrm{Proj}_L(M_{\alpha,\beta}) = \mathrm{Proj}_L(A)$. This will imply that α and β realize a weak Fréchet distance less than ε.

We do this in two steps: first we define for each cell contained in the connected component A two parameterized surface patches $M_K(\Delta_K, \Delta_L)$ and $M_L(\Delta_K, \Delta_L)$. With a *surface patch* we mean a two-dimensional manifold with boundary in $\mathbb{R}^4$. Next we show how to "glue" two parameterized surfaces patches together. Gluing together the surface patches of all cells contained in the connected component will give the desired reparameterizations.

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

$$\begin{aligned} M_K(\Delta_K, \Delta_L) &:= \big\{\big(\mathbf{x},\, g^{-1}\big(n\big(f(\mathbf{x}),\, \Delta_L\big)\big)\big) \,\big|\, \mathbf{x} \in \mathrm{Proj}_K(C_\varepsilon(\Delta_K, \Delta_L))\big\}, \\ M_L(\Delta_K, \Delta_L) &:= \big\{\big(f^{-1}\big(n\big(g(\mathbf{y}), \Delta_K\big), \mathbf{y}\big)\big) \,\big|\, \mathbf{y} \in \mathrm{Proj}_L\,(C_\varepsilon(\Delta_K, \Delta_L))\big\}, \end{aligned}$$

where

$$n(x, A) \quad := \quad \operatorname*{arg\,min}_{y \in A} \|x - y\|$$

is the nearest neighbor of a point x in a metric space A.

For the surface patches $M_K(\Delta_K, \Delta_L)$ and $M_L(\Delta_K, \Delta_L)$ natural parameterizations α_K, β_K and α_L, β_L exist which fulfill $M_{\alpha_K,\beta_K} = M_K(\Delta_K, \Delta_L)$ and $M_{\alpha_L,\beta_L} = M_L(\Delta_K, \Delta_L)$. 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 $\beta_K = \sigma_K \circ f \circ n(\cdot, \Delta_L) \circ g^{-1}$ and $\alpha_L = \sigma_L \circ g \circ n(\cdot, \Delta_K) \circ f^{-1}$ 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.,

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

M_A has the following properties:

1. $M_A \subset A$
2. $\mathrm{Proj}_K(A) = \mathrm{Proj}_K(M_A)$ and $\mathrm{Proj}_L(A) = \mathrm{Proj}_L(M_A)$.
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 nonempty, then $M_K(\Delta_K, \Delta_L) \cap M_K(\Delta_K, \Delta'_L) \neq \emptyset$.
   If the shared boundary cell of the cells $C_\varepsilon(\Delta_K, \Delta_L)$ and $C_\varepsilon(\Delta'_K, \Delta_L)$ is nonempty, then $M_L(\Delta_K, \Delta_L) \cap M_L(\Delta'_K, \Delta_L) \neq \emptyset$.

The first two properties hold by definition. 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 in the free space cell, i.e.,

$$\exists (\mathbf{x}, \mathbf{y}) \in \Delta_K \times \Delta_L : \; g(\mathbf{y}) = n\big(f(\mathbf{x}),\, \Delta_L\big) \;\wedge\; f(\mathbf{x}) = n\big(g(\mathbf{y}),\, \Delta_K\big).$$

Points fulfilling this condition are points whose images have minimal distance in the free space cell, i.e.,

$$(\mathbf{x}, \mathbf{y}) \;=\; \operatorname*{arg\,min}_{(\mathbf{x}', \mathbf{y}') \in \Delta_K \times \Delta_L} \|f(\mathbf{x}') - g(\mathbf{y}')\|.$$

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

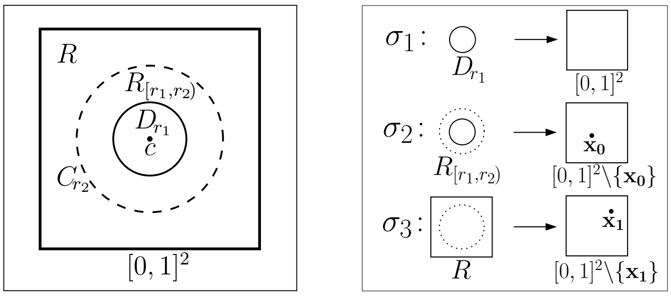

(a) Partitioning of the parameter space

(b) Homeomorphisms on the parts

Figure 4.2: Gluing together two parameterizations

Gluing together the Parameterizations of the Surface Patches We show that two parameterizations can be joined to one in the following sense:

$$\forall \alpha, \beta, \alpha', \beta' : \; M_{\alpha,\beta} \cap M_{\alpha',\beta'} \neq \emptyset \;\Rightarrow\; \exists \alpha'', \beta'' : \; M_{\alpha,\beta} \cup M_{\alpha',\beta'} = M_{\alpha'',\beta''} \quad (4.4.2)$$

where $\alpha, \beta, \alpha', \beta', \alpha'', \beta''$ are continuous functions on $[0,1]^2$. We will then obtain parameterizations α_A, β_A for M_A by successively gluing together parameterizations of the surface patches in M_A.

Proof of 4.4.2: 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 $\mathbf{p} = \big(\alpha(\mathbf{x_0}), \beta(\mathbf{x_0})\big)$ and $\mathbf{p} = \big(\alpha'(\mathbf{x_1}), \beta'(\mathbf{x_1})\big)$.

For gluing together the parameterizations $\alpha, \beta, \alpha', \beta'$, we partition the parameter space as illustrated in Figure 4.2 (a). Choose a point $\mathbf{c}$ in parameter space and radii r_1, r_2 such that $r_1 < r_2$ and the disc of radius r_2 around the point $\mathbf{c}$ is contained in $[0,1]^2$. Let D_r be the (closed) disc of radius r around $\mathbf{c}$ and C_r the circle of radius r around $\mathbf{c}$. Let $R_{[r_1,r_2)}$ be the half-open ring with radii r_1, r_2 around $\mathbf{c}$ and let $R = [0,1]^2 \backslash D_{r_2}$.

We will glue the parameterizations α, β and α', β' along the circle C_{r_2}. On the circle α'', β'' will be constant, outside they will be defined based on α', β' and inside based on α, β.

We define homeomorphisms on the parts of the partition as illustrated in Figure 4.2 (b). Let σ_1 be a homeomorphism from D_{r_1} to the unit cube. Let σ_2 be a homeomorphism from $R_{[r_1,r_2)}$ to the unit cube without the point $\mathbf{x_0}$. Choose σ_1 and σ_2 s.t. they coincide on the circle C_{r_1}. Let σ_3 be a homeomorphism from R to $[0,1]^2 \backslash \{\mathbf{x_1}\}$. Now we define $\alpha'', \beta'' \colon [0,1]^2 \to [0,1]^2$ as follows:

$$\alpha''(\mathbf{x}) := \begin{cases} \alpha(\sigma_1(\mathbf{x})), & \mathbf{x} \in D_{r_1} \\ \alpha(\sigma_2(\mathbf{x})), & \mathbf{x} \in R_{[r_1,r_2)} \\ \mathbf{x_p}, & \mathbf{x} \in C_{r_2} \\ \alpha'(\sigma_3(\mathbf{x})), & \mathbf{x} \in R \end{cases}, \quad \beta''(\mathbf{x}) := \begin{cases} \beta(\sigma_1(\mathbf{x})), & \mathbf{x} \in D_{r_1} \\ \beta(\sigma_2(\mathbf{x})), & \mathbf{x} \in R_{[r_1,r_2)} \\ \mathbf{y_p}, & \mathbf{x} \in C_{r_2} \\ \beta'(\sigma_3(\mathbf{x})), & \mathbf{x} \in R \end{cases}.$$

α'', β'' are continuous and fulfill $M_{\alpha,\beta} \cup M_{\alpha',\beta'} = M_{\alpha'',\beta''}$ by construction. They are continuous inside $D_{r_1}, R_{[r_1,r_2)}$, and R as concatenations of continuous functions. On the circle C_{r_1} they are continuous because σ_1 and σ_2 coincide on C_{r_1}. They are continuous on the circle C_{r_2} because $\sigma_2(\mathbf{x})$ tends to $\mathbf{x_0}$ for $\mathbf{x}$ tending to $\mathbf{x}' \in C_{r_2}$ 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 $\mathbf{x}' \in C_{r_2}$. Similarly, $\sigma_3(\mathbf{x})$ tends to $\mathbf{x_1}$ for $\mathbf{x}$ tending to $\mathbf{x}' \in C_{r_2}$ 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 $\mathbf{x}' \in C_{r_2}$.

We can glue together all surface patches in M_A because they are connected by intersection – properties 3. and 4. of M_A. By gluing together all surface patches in M_A we get reparameterizations α_A, β_A for f, g which realize a weak Fréchet distance between f and g which is less than ε: α_A and β_A are continuous because the parameterizations of single surface patches and the gluing of parameterizations are continuous. 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 $F_\varepsilon(f,g)$. □

4.5 Deciding the Weak Fréchet Distance

By the result of the previous section, Lemma 4.2, deciding the weak Fréchet distance is equivalent to determining if there is an extensive connected component in free space. We will do this algorithmically in two steps:

1. determine the connected components of the free space
2. for all large connected components, test if their projections completely cover both parameter spaces.

For the first step we compute the combinatorial graph of the free space as described in Section 4.3.3. For this, we need to decide which boundary cells are non-empty. By equation 4.3.2 we can do this for each boundary cell in constant time by computing the distance of the edge and the triangle defining the boundary cell. This yields a run time of $O(mn)$ for the first step.

In the second step we need to determine if the projection of a large connected component completely covers a parameter space. For this, we choose w.l.o.g. the parameter space K, i.e., we want to decide if $\mathrm{Proj}_K(A) = |K|$ holds for a connected component A of the free space.

The triangulation K is the union of its triangles and therefore K is completely contained in the projection of A if this holds for every triangle in K. Since each cell projects only onto the triangles defining it, we need to consider for each triangle Δ_K in K only the cells $C_\varepsilon(\Delta_K, \Delta_L)$ in A for triangles Δ_L in L. Using the characterization of a cell stated in equation 4.3.1 we get the following equivalence.

$$\begin{aligned} \Delta_K \subset \mathrm{Proj}_K(A) \quad &\Leftrightarrow \quad \Delta_K \subset \bigcup_{\substack{\Delta_L \in L \\ C_\varepsilon(\Delta_K,\Delta_L)\in A}} \mathrm{Proj}_K\big(C_\varepsilon(\Delta_K,\Delta_L)\big) \\ &\Leftrightarrow \quad f(\Delta_K) \subset \bigcup_{\substack{\Delta_L \in L \\ C_\varepsilon(\Delta_K,\Delta_L)\in A}} \big(g(\Delta_L) \oplus B_\varepsilon\big) \qquad (4.5.1) \end{aligned}$$

Note that we are now concerned only with the triangles in image space and not in parameter space. We can reformulate equivalence 4.5.1 as a decision problem for the directed Hausdorff distance, namely it is equivalent to the directed Hausdorff distance of the triangle $f(\Delta_K)$ to the set of triangles $g(\Delta_L)$, where $C_\varepsilon(\Delta_K, \Delta_L)$ is a cell in A, being less or equal than ε. In formulas,

$$\delta_{H,dir}\Big(f(\Delta_K), \bigcup_{\substack{\Delta_L \in L \\ C_\varepsilon(\Delta_K,\Delta_L)\in A}} g(\Delta_L)\Big) \leq \varepsilon. \qquad (4.5.2)$$

By a result of Alt et al. [4] the directed Hausdorff distance of n k-dimensional simplices to m k-dimensional simplices can be computed in $O(nm^{k+2})$ time. Thus,

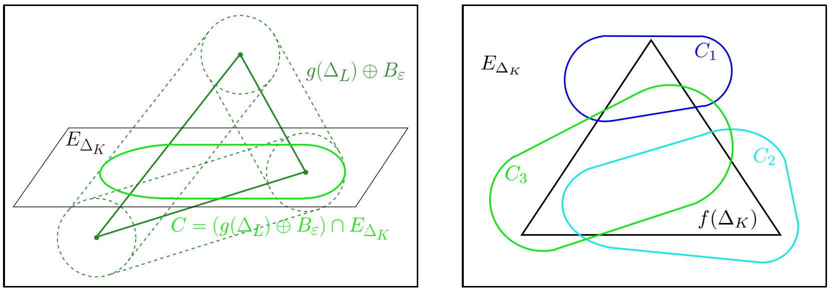

(a) ε-neighborhood of a triangle intersected by a plane

(b) arrangement of boundary curves covering a triangle

Figure 4.3: Computing the arrangement of the cells projected onto the plane.

equation 4.5.2 can be solved in $O(m_i^4)$ time for triangles, where m_i is the number of cells $C_\varepsilon(\Delta_K, \Delta_L) \in A$.

For triangles, we can improve this run time. If we define E_{Δ_K} to be the plane containing $f(\Delta_K)$ then we can modify the last expression of equivalence 4.5.1 to

$$f(\Delta_K) \subset \bigcup_{\substack{\Delta_L \in L \\ C_\varepsilon(\Delta_K, \Delta_L) \in A}} \Big(\big(g(\Delta_L) \oplus B_\varepsilon\big) \cap E_{\Delta_K} \Big). \tag{4.5.3}$$

Intersecting with the plane E_{Δ_K} makes the problem two-dimensional. We solve equation 4.5.3 by computing for each triangle Δ_K the arrangement of the triangle $f(\Delta_K)$ and the set of boundary curves $\partial\big(\big(g(\Delta_L) \oplus B_\varepsilon\big)\big) \cap E_{\Delta_K}\big)$ for all triangles $\Delta_L \in L$ fulfilling $C_\varepsilon(\Delta_K, \Delta_L) \in A$. With ∂ we denote the boundary operator. See Figure 4.3 (b) for an example. We sweep over the part of the arrangement contained in the triangle $f(\Delta_K)$. $f(\Delta_K)$ is contained in the union of the $\big(g(\Delta_L) \oplus B_\varepsilon\big)\big) \cap E_{\Delta_K}$ if and only if there is no empty cell in this part of the arrangement.

For the metrics d_1 and d_∞, 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 triangle prism of height ε. The boundaries of the intersections of these ε-neighborhoods with a plane are the union of a constant number of half-ellipses, circles or half-circles and straight line segments. See Figure 4.3 (a) for an illustration. 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 in which connected components it is contained. 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 ith connected component. Because the connected components 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 $O(\sum_{i=1}^{s}(m_i^2 \log m_i)) \in O(m^2 \log m)$ 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 run time of $O(nm^2 \log m + mn^2 \log n)$ for all triangles in both parameter spaces.

In the above analysis, we solve equation 4.5.3 in $O(m^2 \log m)$ time. The problem Triangles-cover-Triangle, which asks whether the union of a set of triangles in the plane contains another given triangle, is 3SUM-hard [27]. This implies that equation 4.5.3 is not likely to be solvable in subquadratic time.

Summarizing, we obtain the following algorithm for deciding the weak Fréchet distance.

Algorithm 3: DecideWeakFrechet(f, g, ε)

Input: parametrized triangulated surfaces f, g, $\varepsilon > 0$
Output: Is $\delta_{wF}(f, g) \leq \varepsilon$?

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

Lemma 4.2 and the analysis above yield the following theorem:

Theorem 4.1. *Algorithm 3 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.*

4.6 Computing the Weak Fréchet Distance

The decision algorithm can be extended to a computation algorithm by searching a set of critical values. We first give these critical values and show how to compute them. Then we show how to search the set of critical values in order to compute the weak Fréchet distance.

4.6.1 Critical Values

By Lemma 4.2 the weak Fréchet distance equals a value ε if and only if the free space $F_\varepsilon(f, g)$ contains an extensive connected component and it does not for any smaller values of ε. We can characterize the critical value where this may happen 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 the free space does not change combinatorially in ε.

We show that these are polynomially many critical values and each value can be computed in constant time.

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 non-empty. This happens for ε equal to the distance of the edge and the triangle defining the boundary cell, as stated in equation 4.3.2. These are $O(mn)$ values, each of which can be computed in constant time.

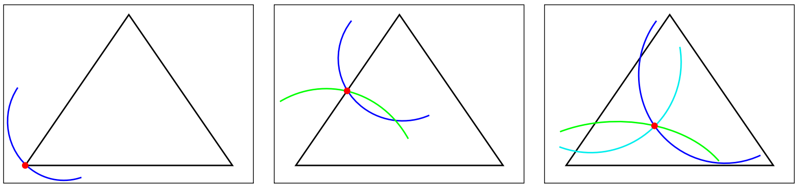

(a) vertex (b) edge (c) interior

Figure 4.4: Critical values of type 2 distinguished by the last point covered.

Type 2 Apart from the combinatorial changes, the projection of a connected component grows monotonously with increasing ε. We distinguish these critical values by the *last point* covered by the projection. A last point covered is a point x in parameter space, i.e., $x \in [0,1]^2$, fulfilling

$$x \in \text{Proj}(F_\varepsilon) \ \wedge \ (\forall \varepsilon' < \varepsilon)\ x \notin \text{Proj}(F_{\varepsilon'}). \tag{4.6.1}$$

This may be either (a) a vertex, (b) a point on an edge, or (c) an interior point of a triangle of the parameter space.

In the following we will assume that the points of the triangulated surfaces lie in general position, i.e., that there are no parallel edges or triangles. If this is not the case, i.e., if there are parallel edges or triangles, then the last points covered may form a *last segment*. Such a last segment covered can be the overlap of the boundaries of ε-neighborhoods of two triangles which lie in parallel planes or have two parallel edges. Or the ε-neighborhood of a triangle may overlap with the boundary of a triangle of the other surface in a last segment covered. We can compute a superset of these additional critical values for the degenerate case by computing all parallel edges and planes and adding the distance and half the distance of each pair of parallel edges or planes to the set of critical values. These are $O(m^2 + n^2)$ values, each of which can be computed in constant time.

Now we describe the critical values for the non-degenerate case. If the last point covered is a vertex, then equation 4.6.1 holds if this point lies on the boundary of $g(\Delta) \oplus B_\varepsilon$ for a triangle Δ in the other parameter space, as in Figure 4.4 (a). If the last point covered lies on the boundary of the parameter space, then two boundaries of ε-neighborhoods of $g(\Delta_L) \oplus B_\varepsilon$ intersect in this point, as in Figure 4.4 (b). If the last point covered lies in the interior of the parameter space, then three boundaries intersect in it, as in Figure 4.4 (c).

Type 2 a A point x lies on the boundary of the ε-neighborhood of a triangle if and only if it has distance ε to the triangle. In formulas,

$$x \in \partial\big(g(\Delta) \oplus B_\varepsilon\big) \ \Leftrightarrow\ \text{dist}(x, g(\Delta)) = \varepsilon \tag{4.6.2}$$

where ∂ denotes the boundary operator and $\text{dist}(x, g(\Delta))$, the distance of the point x to the triangle $g(\Delta)$ is defined as the minimal distance between x and a point on the triangle. That is, $\text{dist}(x, g(\Delta)) = \min_{y \in g(\Delta)} \|x - y\|$.

Thus we can compute these critical values by computing the distances between any vertex of the one surface in image space and a triangle of the other surface in image space. These are $O(mn)$ values, each of which can be computed in constant time.

Type 2 b By equation 4.6.2, a point lies on the boundary of two ε-neighborhoods of triangles if and only if it has equal distance to the two triangles. We can compute these critical values by solving for each triple consisting of an edge of the one triangulated surface in image space and two triangles of the other triangulated surface in image space the set of distance equations. This can be simplified by computing a superset of critical values, namely all equal distances between an edge of the one surface in image space and two "components" of triangles of the other surface, namely a vertex, a line defined by an edge or the plane defined by the triangle. These are one linear and two linear or quadratic equations, which we assume that we can solve in constant time (see Section 2.4). The number of critical values of this type is $O(mn^2 + m^2n)$.

Type 2 c These critical values are similar to those of type 2 *b*. Now we determine points inside one triangle with equal distance to three other triangles. For this we again solve the set of distance equations. These are $O(mn^3 + m^3n)$ critical values.

Summarizing, we get the following lemma.

Lemma 4.3. *For the weak Fréchet distance between triangulated surfaces there are $O(m^3n + mn^3)$ critical values, each of which can be computed in constant time.*

4.6.2 Searching the Set of Critical Values

We can compute the weak Fréchet distance by computing and sorting the critical values and then doing a binary or median search over the set of critical values, solving Algorithm 3 in each step. The run time of this algorithm is dominated by computing and sorting of the critical values which takes $O(m^3n + mn^3)$ and $O((m^3n + mn^3)\log(mn))$ time, respectively.

The run time can be improved using Meggido's parametric search [39], and Cole's trick for parametric search based on sorting [17]. We apply the parametric search to the critical values of type 2 *c*, of which there are $O(m^3n + mn^3)$ many. The other critical values, of which there are $O(m^2n + mn^2)$ many, we search with a binary search in $O((m^2n + mn^2)\log^2(mn))$ time or with a median search in $O((m^2n + mn^2)\log(mn))$ time. Thus, our final algorithm has two steps:

1. binary or median search on the $O(m^2n + mn^2)$ critical values
2. parametric search on the $O(m^3n + mn^3)$ critical values.

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 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. I.e., we want to sort all values

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

where Δ is the image of a triangle in one parameter space and Δ'_1, Δ'_2 images of triangles of the other parameter space. The critical values of type 2 *c* are the ε values where two such intersection points coincide, which share one triangle. Sorting all intersection points gives a superset of critical values, namely also those do not share a triangle.

We can simplify the computation of the intersection points by writing $\Delta' \oplus B_\varepsilon$ as the union of the balls of radius ε around its vertices, the cylinders of radius ε around its edges and a triangle pyramid of height ε. Then we sort the intersection points $\partial(P_{1\varepsilon} \cap E_\Delta) \;\cap\; \partial(P_{2\varepsilon} \cap E_\Delta)$, where $P_{1\varepsilon}, P_{2\varepsilon}$ are two such "components" of ε-neighborhoods of images of triangles Δ'_1, Δ'_2. These intersection points are a superset of the intersection points we need. There are $O(m^2n + n^2m)$ many of these, which is asymptotically the same as without the decomposition of the ε-neighborhoods into balls, cylinders, and pyramids.

For the parametric search, we can sort these (two-dimensional) points lexicographically or by one of the two coordinates. In both cases, the critical values for sorting will be a superset of the critical values for the Fréchet distance.

We use Cole's variant of parametric search based on sorting [17]. His technique yields a run time of $O((k + T_{\text{dec}}) \log k)$ where T_{dec} is the run time 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 run time of $O((m^2n + mn^2) \log^2(mn))$.

Theorem 4.2. *The weak Fréchet distance between two triangulated surfaces with m and n triangles, respectively, can be computed in $O((m^2n+mn^2)\log^2(mn))$ time.*

4.6.3 Non-parameterized Curves and Surfaces

All computations for computing the weak Fréchet distance are done on the surfaces in image space and not in parameter space. The parameter space is used only for the connectivity of the free space. Instead, we could also define the connectivity of the free space based on the connectivity in image space (assuming this to be given). Then we can run the same algorithms for non-parameterized triangulated surfaces.

Let us call the space in which the free space diagram lies the *configuration space*. For parameterized curves and surfaces, the configuration space is the product of the parameter spaces of the curves or surfaces, respectively. If we instead define the free space based on the curves and surfaces in image space, then the configuration space will be the product of the two shapes.

Let us consider the example of closed curves. For computing the Fréchet distance of closed curves Alt and Godau [6] use the double free space diagram, i.e., two concatenated copies of a single free space diagram. If we instead define the free space based on the connectivity in image space, then the configuration space of two closed curves will be homeomorphic to the product of two circles, which is a torus.

For curves and surfaces which are homeomorphic to their parameter spaces, i.e., the unit interval and unit square, respectively, the topology of the configuration space is the same whether we define it based on the connectivity in parameter or image space. If the curves and surfaces are not homeomorphic to the their parameter spaces, however, as in the case of closed curves parameterized over the unit interval, the topology of the configuration space will differ. An interesting open problem is to further investigate this definition of a "generalized" weak Fréchet distance for non-parameterized curves and surfaces.

4.7 Discussion

In this chapter we have shown that the weak Fréchet distance between triangulated surfaces is polynomial time computable. For this, we showed that the free space diagram in higher dimension is computable, and used that the weak Fréchet distance is a relaxed version of the Fréchet distance. Thus, we now have a gap between the NP-hardness of deciding the Fréchet distance and a polynomial time algorithm for deciding the weak Fréchet distance between triangulated surfaces. An immediate open question is to learn more about this gap. Another interesting open question is whether the weak Fréchet distance between triangulated surfaces fulfills the triangle inequality and therefore is a metric as in the case of curves.