

Computing the Fréchet Distance between Real-Valued Surfaces

Kevin Buchin *

Tim Ophelders *

Bettina Speckmann *

1 Introduction

The problem of measuring the similarity between shapes has recently gained much attention. While many measures have been defined, algorithms to compute such measures have been found for only some of them. We consider the problem of comparing real-valued functions $f: M \rightarrow \mathbb{R}$ on surfaces, focusing in particular on spheres and disks of constant boundary, i.e., $f(x) = f(x')$ for all $x, x' \in \partial M$. The kind of similarity we investigate is that under continuous deformations of surfaces, such as in Figure 1. Here shapes that can be deformed into each other have distance 0, otherwise, shapes have some meaningful positive distance. Two natural computational problems arise for each measure, namely deciding whether two images have distance 0, and the more general problem of computing the distance between two images.




Figure 1: Pictures that can deform into each other.

Major applications of computing such measures are in the comparison of medical imagery. For example, when comparing two MRI or CT scans of lungs, the images are often not aligned due to breathing and gravity. It is important to align the images through deformation to locate differences.

Definitions, background and results. Given two functions $f: M \rightarrow \mathbb{R}^k$ and $g: M \rightarrow \mathbb{R}^k$ with common parameter space M , their *Fréchet distance* is defined by Equation 1, where $\mu: M \rightarrow M$ ranges over orientation preserving homeomorphisms and $d(\cdot, \cdot)$ is the underlying norm of \mathbb{R}^k . Essentially, the Fréchet distance captures the similarity between two functions by realigning their parameter spaces to minimize the maximum difference in function value of aligned points. We assume that f and g are piecewise-linear.

$$d_F(f, g) = \inf_{\mu: M \rightarrow M} \sup_{x \in M} d(f(x), g \circ \mu(x)). \quad (1)$$

*Department of Mathematics and Computer Science, TU Eindhoven, the Netherlands,
[k.a.buchin|t.a.e.ophelders|b.speckmann]@tue.nl

Efficient algorithms for computing $d_F(f, g)$ exist for L^p norms if f and g are polylines [3], so if $M = [0, 1]$ or $M = S^1$ for closed polylines. The computational complexity in the case that f and $g: M \rightarrow \mathbb{R}^k$ are (triangulated) surfaces is much less understood. The problem is known to be NP-hard [7] also when $k = 2$ [4, 5]. But it is not known whether it is actually in NP, in fact it is only known to be upper semi-computable for surfaces in \mathbb{R}^k [2].

We show that the problem is in NP for $k = 1$ if M is a topological sphere or disk with constant boundary. Additionally, we show that even for $k = 1$, computing a factor $2 - \varepsilon$ approximation of the Fréchet distance is NP-hard. We achieve our results on surfaces (Section 3) by first defining a suitable similarity measure between contour trees, which we show to be NP-complete to approximate as well (Section 2).

In previous work, a few variants on the comparison of surfaces under the Fréchet distance have been investigated. For instance, there are efficient algorithms for computing the Fréchet distance with certain constraints on the homeomorphisms μ [5] and for computing the *weak Fréchet distance* [2] between triangulated surfaces homeomorphic to the disk.

2 Contour tree distance

The Reeb graph [8] of a function $f: M \rightarrow \mathbb{R}$ is the quotient space M/\sim_f (endowed with the quotient topology) where $a \sim_f b$ if and only if a and b are in the same connected component of the level set $f^{-1}(f(a))$. Denote by \mathcal{R}_f the corresponding quotient map. Because f associates a single real number to each equivalence class of \sim_f , the resulting Reeb graph has a natural \mathbb{R} -valued function associated with it, namely the (unique) function $f': M/\sim_f \rightarrow \mathbb{R}$ satisfying $f' \circ \mathcal{R}_f = f$. If M is the disk or the 2-sphere, the Reeb graph forms a tree called a contour tree.

For the sake of compact representation, in this paper we assume each surface to be triangulated to form a simplicial 2-complex. Furthermore, we assume function values along edges of Reeb graphs to be linearly interpolated between the values of the vertices at their endpoints. In this representation, the contour tree of a surface with n faces has complexity $O(n)$ and can be computed in $O(n \log n)$ time [11].

Based on the Fréchet distance between f and g , we derive a computationally simpler measure that abstracts from the realizability of the matching μ be-

tween spheres or disks. Throughout this paper, we use the notation $\mathbb{X} = M/\sim_f$ and $\mathbb{Y} = M/\sim_g$ for the contour trees of f and g , respectively. We shall denote the vertex set of \mathbb{X} by $V(\mathbb{X})$ (that is, the saddle points and minima and maxima of f) and its edge set by $E(\mathbb{X})$. With slight abuse of notation, we reuse function names f and g for the natural \mathbb{R} -valued functions associated with the contour trees \mathbb{X} and \mathbb{Y} . Our distance measure d_C compares the contour trees \mathbb{X} and \mathbb{Y} of f and g . We define the contour tree distance d_C between $f: \mathbb{X} \rightarrow \mathbb{R}$ and $g: \mathbb{Y} \rightarrow \mathbb{R}$ as

$$d_C(f, g) = \inf_{\tau \in \mathcal{M}(\mathbb{X}, \mathbb{Y})} \sup_{(x, y) \in \tau} |f(x) - g(y)|,$$

where $\tau \subseteq \mathbb{X} \times \mathbb{Y}$ is drawn from a specific class of matchings $\mathcal{M}(\mathbb{X}, \mathbb{Y})$, defined below. So τ defines a correspondence between contour trees, such that $(x, y) \in \tau$ if some points on contours x and y are matched by a corresponding matching μ on M . Denote $\tau(x) = \{y \mid (x, y) \in \tau\}$ and $\tau^{-1}(y) = \{x \mid (x, y) \in \tau\}$. The class $\mathcal{M}(\mathbb{X}, \mathbb{Y})$ captures the essential (but not all) properties of an orientation preserving matching μ . We define $\mathcal{M}(\mathbb{X}, \mathbb{Y})$ as the set of matchings τ for which the following properties hold:

1. τ is a connected subset of $\mathbb{X} \times \mathbb{Y}$;
2. $\tau(x)$ is a nonempty subtree of \mathbb{Y} for each $x \in \mathbb{X}$;
3. $\tau^{-1}(y)$ is a nonempty subtree of \mathbb{X} for each $y \in \mathbb{Y}$.




Figure 2: Two trees (left) and a matching.

Here, the term *subtree* is used to denote a connected subset of a tree, not necessarily containing leaves of that tree. By Conditions 2 and 3, each connected set matches with a connected set, and Condition 1 ensures continuity.

Figure 2 shows an example of a matching between trees. The two-dimensional patch of this matching represents

a many-to-many correspondence. For a matching $\mu: M \rightarrow M$ between surfaces f and g , define $T(\mu)$ to be the corresponding matching between the Reeb graphs of f and g :

$$T(\mu) = \{(\mathcal{R}_f(x), \mathcal{R}_g \circ \mu(x)) \mid x \in M\} \quad (2)$$

Lemma 1 *If $\mu: M \rightarrow M$ is orientation preserving and $\tau = T(\mu)$, then $\tau \in \mathcal{M}(\mathbb{X}, \mathbb{Y})$.*

Proof. Consider such a matching μ and the correspondence τ between \mathbb{X} and \mathbb{Y} . We show all three conditions on τ hold. The set $\{(x, y) \mid \mu(x) = y\}$ is a connected subset of $S^2 \times S^2$, and hence its image under the quotient map $(x, y) \mapsto (\mathcal{R}_f(x), \mathcal{R}_g(y))$

is connected, so Condition 1 holds. Because μ preserves orientation, $\tau(x)$ is connected, and by surjectivity nonempty. Hence, $\tau(x)$ and symmetrically $\tau^{-1}(y)$ is a nonempty subtree of \mathbb{Y} and \mathbb{X} , respectively. \square

Corollary 2 $d_C(f, g) \leq d_F(f, g)$.

Lemma 3 *Computing $d_C(f, g)$ is in NP.*

By Lemma 1 we have for each orientation preserving homeomorphism μ , that the matching $\tau = T(\mu)$ satisfies $\tau \in \mathcal{M}(\mathbb{X}, \mathbb{Y})$. Hence, $d_C(f, g) \leq d_F(f, g)$. On the other hand, a matching $\tau \in \mathcal{M}(\mathbb{X}, \mathbb{Y})$ does not need to correspond to an orientation preserving homeomorphism on M , as illustrated in Section 3.1.

To test whether the contour tree distance between two trees is zero, one needs to test only whether the trees are equal. We represent trees canonically by exhaustively removing degree 2 vertices that lie on the segment connecting the two adjacent vertices, and replacing them by a single edge between those vertices. This reduces the problem to labeled unordered unrooted tree isomorphism, solvable in linear time [1].

Computing the contour tree distance and Fréchet distance between trees are different problems. In fact, one major limitation of the Fréchet distance for trees is that non-homeomorphic trees have infinite Fréchet distance. Nonetheless, algorithms for computing the Fréchet distance between trees have been investigated before, yielding an $O(n^{5/2})$ time algorithm [4].

2.1 NP-hardness

We show that approximating the contour tree distance between \mathbb{R} -valued trees within factor 2 is NP-hard by a reduction from the NP-hard problem EXACT COVER BY 3-SETS [6].

Definition 1 EXACT COVER BY 3-SETS (X3C)

Input: A set \mathcal{S} of m subsets of $\{1, \dots, k\}$ of size 3.

Output: Is there a subset of \mathcal{S} consisting of $k/3$ disjoint triples whose union is $\{1, \dots, k\}$?

We introduce the gadgets used in our reduction in Figure 3. Gadget Y^* is a long segment from position 0 to position $6k + 6$. Gadget Y_l ($l \in \{1, \dots, k\}$) is a path from position 1 to $6k + 6$ with a single zig-zag of radius 2 around position $6l$. Similarly, gadget $X_{i,j}$ ($i \in \{1, \dots, m\}, j \in \{1, 2, 3\}$) has a single zig-zag around position $6 \cdot s(i, j)$, but with radius 1.




Figure 3: Gadgets.

The function s aligns the center of the zig-zag of $X_{i,j}$ with that of $Y_{s(i,j)}$, such that gadget $X_{i,j}$ has a contour tree

distance of 1 to Y^* and $Y_{s(i,j)}$, but a contour tree distance of 2 to any gadget Y_l with $l \neq s(i,j)$.




Figure 4: Connecting gadgets into trees f and g .

The function s can be configured such that each triple of gadgets $(X_{i,1}, X_{i,2}, X_{i,3})$ corresponds to one of the m subsets of \mathcal{S} . We connect the three elements of each triple at a common vertex at position 1, and finally connect all triples at a common vertex at position 2 (blue in Figure 4) to form tree $f: \mathbb{X} \rightarrow \mathbb{R}$.

Similarly, all gadgets Y_l correspond to an element of $\{1, \dots, k\}$, and all Y_l are connected to a common vertex at position 1. To obtain a low contour tree distance, $k/3$ triples of f must match the Y_l gadgets exactly; then what remains in f are $m - k/3$ triples that must be matched elsewhere. Each such unmatched triple of f is then forced to match with three copies of Y^* , connected at a vertex at position 0 to form a so called Y^* -triple. We use $m - k/3$ such Y^* -triples, each connected to the Y_l gadgets at position 1 to form tree $g: \mathbb{Y} \rightarrow \mathbb{R}$. We use a solution to X3C to derive a matching using only many-to-one correspondences between f and g , even though $\mathcal{M}(\mathbb{X}, \mathbb{Y})$ also permits many-to-many correspondences. In the full paper we also show that any many-to-one matching can be realized as an orientation preserving homeomorphism on the sphere, such that computing the Fréchet distance between \mathbb{R} -valued spheres is NP-hard, see Theorem 5.

Theorem 4 *Computing a $(2 - \varepsilon)$ -approximation of the contour tree distance is NP-complete.*

Theorem 5 *Computing a $(2 - \varepsilon)$ -approximation of the Fréchet distance of \mathbb{R} -valued spheres is NP-hard.*

3 Surfaces

We consider two surfaces: the disk $M = [0, 1]^2$ and the sphere S^2 . Not all matchings between the contour trees \mathbb{X} and \mathbb{Y} can be realized as orientation preserving homeomorphisms on the sphere as illustrated in Section 3.1. In the case of the disk, the boundaries must also be matched, which imposes additional constraints on the matching of the interiors. We prove that the Fréchet distance between \mathbb{R} -valued spheres or disks with constant boundary is in NP in Section 3.2. For this we use properties of Euler diagrams.

3.1 An unrepresentable matching

Consider the two rooted trees \mathbb{X} and \mathbb{Y} of Figure 5. The leaves of \mathbb{X} are labeled $x_{i,j}$ ($i \in \{1, \dots, 6\}$, $j \in \{1, 2, 3\}$) and the leaves of \mathbb{Y} are labeled $y_{k,l}$ ($k \in \{1, \dots, 9\}$, $l \in \{1, 2\}$). For both trees, leaves with the same i or k are grouped in subtrees. Based on the complete bipartite graph $K_{3,3}$ with vertices v_1, \dots, v_6 and edges e_1, \dots, e_9 , we construct a matching τ between those subtrees as follows. For an edge $e_k = (v_i, v_{i'})$ of $K_{3,3}$, match the path from $y_{k,1}$ to $y_{k,2}$ with the path between unused vertices $x_{i,j}$ and $x_{i',j'}$. Match the edge from the root to group i of \mathbb{X} with the edges of \mathbb{Y} from the root to the three groups that match with $x_{i,1}$, $x_{i,2}$ and $x_{i,3}$. Then $\tau \in \mathcal{M}(\mathbb{X}, \mathbb{Y})$ does not match any path from $y_{k,1}$ to $y_{k,2}$ (of edge $e_k = (v_i, v_{i'})$) with any group of \mathbb{X} not containing any $x_{i,j}$ or $x_{i',j'}$. However, because $K_{3,3}$ is not a planar graph, this matching cannot be realized on the sphere, as illustrated in Figure 5.

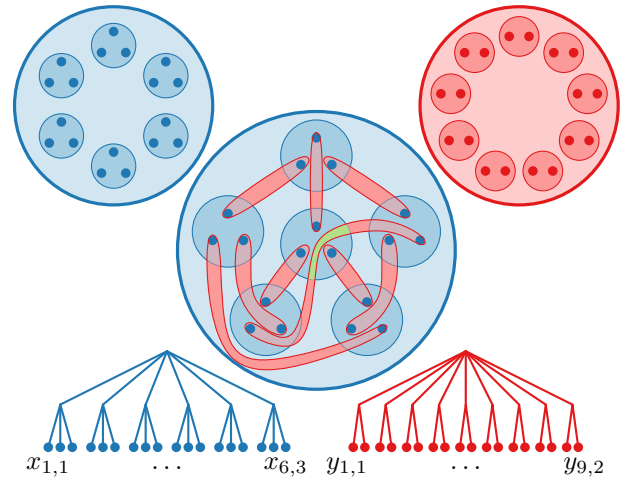


Figure 5: Top: surfaces. Bottom: trees \mathbb{X} and \mathbb{Y} . Middle: a matching in which a subtree of \mathbb{Y} must intersect an additional subtree of \mathbb{X} .

3.2 Fréchet distance in NP

An Euler diagram is a set of topological disks, drawn in the plane to capture relations such as overlap or containment between them. There are eight such relations: disjoint, equal, inside, contains, covered, cover, meet, overlap [10]. For a set D of n disks and a relation $\mathcal{P}(a, b)$ between each pair $a, b \in D$ of those disks, the tuple (D, \mathcal{P}) is called a *topological expression*. It was shown by Schaefer, Sedgwick and Štefankovič [9] that it is in NP to decide whether D can be drawn in the plane to satisfy all relations of \mathcal{P} .

We show that deciding whether the Fréchet distance between \mathbb{R} -valued surfaces f and $g: M \rightarrow \mathbb{R}$ is at most ε is in NP if M is a sphere. First consider the case where $M = [0, 1]^2$ is a topological disk

where f and g have constant boundary: $f(p) = f(q)$ and $g(p) = g(q)$ for all $p, q \in \partial M$.

We can represent the contour trees \mathbb{X} and \mathbb{Y} as rooted trees, and represent each subtree as a disk. A vertex is represented by a disk with punctures (one per edge to a subtree), and an edge e is represented by two nested disks, whose difference is an annulus $A(e)$. Let $W(\cdot)$ denote the disks of a tree.




Figure 6: \mathbb{X}' (green) refining \mathbb{X} (blue) for a single $d \in W(\mathbb{Y})$.

Define matching $\mu: M \rightarrow M$ to be an ε -matching if and only if $|f(x) - g(\mu(x))| \leq \varepsilon$ for all x . Consider any ε -matching μ between \mathbb{X} and \mathbb{Y} . A polynomial amount of information about μ is used to derive a topological expression that captures the relations between the disks (of \mathbb{X} and \mathbb{Y}) as drawn by μ . We construct trees \mathbb{X}' and \mathbb{Y}' that have additional vertices along the edges. We include for each disk d of \mathbb{X} the extremes p on \mathbb{Y} of the image of the boundary of d as vertices of \mathbb{Y}' . That is, $p \in \partial\mathcal{R}_g(\mu(\partial d)) \subseteq V(\mathbb{Y}')$, and we require the boundary of these disks to touch but not intersect ∂d . Symmetrically, refine \mathbb{X} into \mathbb{X}' and let D be the disks of $W(\mathbb{X}') \supseteq W(\mathbb{X})$, and $W(\mathbb{Y}') \supseteq W(\mathbb{Y})$, see Figure 6.

Consider any drawing of D that satisfies the constraints imposed by the topological expression derived from μ . We claim that if μ was an ε -matching, we can parameterize f and g such that they form an ε -matching μ' . Where boundaries of disks of \mathbb{X}' and \mathbb{Y}' intersect, both function values are fixed and differ by at most ε . Consider any face F in the drawing $\bigcup_{d \in D} \partial d$ of the boundaries of disks. It is bounded by (a subset of) boundaries of either a vertex or edge of \mathbb{X}' and those of a vertex or edge of \mathbb{Y}' . By construction, F is either entirely contained in the image of a vertex, or disk boundaries of different surfaces intersect on ∂F . In the former case, function values are fixed for ∂F and differ by at most ε . In the latter case, linearly interpolate the function values between those at intersections. The function value (of f and g) on each boundary of F is then either constant, or (only if F lies in the intersection of two annuli and has genus 0) two linear interpolations (back and forth) between the values at the boundary of an annulus.

If either function is constant along ∂F , assign the same constant to the interior, such that the other function can be interpolated arbitrarily. If neither function is constant, $F \subseteq A(e) \cap A(e')$ lies in the intersection of annuli ($e \in E(\mathbb{X}')$ and $e' \in E(\mathbb{Y}')$). If F is homeomorphic to an annulus, we interpolate f and g in its interior linearly between the boundaries. In the final case where both f and g are interpolated along the boundary of F , we separate the regions where the interpolation of f occurs from that of g , such that the

function values at any point differ by at most ε .

Deciding whether the topological expression (constructed using a polynomial amount of information about μ) has a realization in the plane is decidable in nondeterministic polynomial time. For any ε -matching μ , this topological expression has a solution, and each such solution encodes an ε -matching. Theorems 6 and 7 follow.

Theorem 6 *Deciding whether $d_F(f, g) \leq \varepsilon$ for disks f and g with constant boundary is in NP.*

Theorem 7 *Deciding whether $d_F(f, g) \leq \varepsilon$ for spheres f and g is in NP.*

Acknowledgements. K. Buchin, T. Ophelders, and B. Speckmann are supported by the Netherlands Organisation for Scientific Research (NWO) under project no. 612.001.207 (K. Buchin) and no. 639.023.208 (T. Ophelders & B. Speckmann).

References

- [1] A. V. Aho, J. E. Hopcroft, and J. D. Ullman. *The Design and Analysis of Computer Algorithms*. Addison-Wesley, 1974.
- [2] H. Alt and M. Buchin. Can we compute the similarity between surfaces? *DCG*, 43(1):78–99, 2010.
- [3] H. Alt and M. Godau. Computing the Fréchet distance between two polygonal curves. *IJCGA*, 5(01n02):75–91, 1995.
- [4] K. Buchin, M. Buchin, and A. Schulz. Fréchet distance of surfaces: Some simple hard cases. In *Proc. 18th ESA*, pages 63–74, 2010.
- [5] K. Buchin, T. Ophelders, and B. Speckmann. Computing the similarity between moving curves. In *Proc. 23rd ESA*, pages 928–940, 2015.
- [6] M. R. Garey and D. S. Johnson. *Computers and intractability: a guide to the theory of NP-completeness*. WH Freeman & Co., San Francisco, 1979.
- [7] M. Godau. *On the complexity of measuring the similarity between geometric objects in higher dimensions*. PhD thesis, Freie Universität Berlin, Germany, 1998.
- [8] G. Reeb. Sur les points singuliers d'une forme de Pfaff complètement intégrable ou d'une fonction numérique. *C. R. de L'Académie des Sciences*, 222:847–849, 1946.
- [9] M. Schaefer, E. Sedgwick, and D. Štefankovič. Recognizing string graphs in NP. *Journal of Computer and System Sciences*, 67(2):365–380, 2003.
- [10] M. Schaefer and D. Štefankovič. Decidability of string graphs. In *Proc. 33rd STOC*, pages 241–246, 2001.
- [11] M. Van Kreveld, R. van Oostrum, C. Bajaj, V. Pascucci, and D. Schikore. Contour trees and small seed sets for isosurface traversal. In *Proc. 13th SoCG*, pages 212–220, 1997.