Computing the similarity between moving curves

Computing the Similarity Between Moving Curves*

Kevin Buchin† Tim Ophelders† Bettina Speckmann†

1 Introduction

A significant amount of algorithmic research in recent years has focused on the analysis of trajectories: sequences of time-stamped points which represent the movement of objects over time. Not all moving objects, however, can be reasonably represented as points. Here we hence go beyond this basic setting, by studying moving complex, non-point objects. Specifically, we focus on similarity measures for moving curves which can, for example, model changing coastlines, retreating glacier termini, or slithering snakes.

We base the similarity measures between moving curves on the Fréchet distance. We model a moving curve as a sequence of $T+1$ polylines, each of $P+1$ vertices. Consecutive polylines are interpolated to form a quadrilateral mesh of $P \times T$ quadrilaterals with parameters $(p,t) \in [0,P] \times [0,T]$.

The Fréchet distance is commonly used to determine the similarity between curves A and $B : [0,1] \to \mathbb{R}^n$. A natural generalization to more complex shapes uses the definition of Eq. 1 where A and B have type $X \to \mathbb{R}^n$.

$$D_{\mathrm{fd}}(A,B) = \inf_{\mu : X \to X} \sup_{x \in X} \|A(x) - B(\mu(x))\| \qquad (1)$$

Here, $\|\cdot\| : \mathbb{R}^n \to \mathbb{R}$ is a norm such as the Euclidean norm (L^2) or the Manhattan norm (L^1). The *matching* μ ranges over orientation-preserving homeomorphisms (possibly with additional constraints) between the parameter spaces of the shapes compared; as such, it defines a correspondence between the points of the compared shapes. Given one such matching we obtain a distance between A and B by taking the largest distance between any two corresponding points of A and B. The Fréchet distance is then the infimum of these distances taken over all possible matchings. For moving points or static curves, we have as parameter space $X = [0,1]$ and for moving curves or static surfaces, we have $X = [0,1]^2$. We define various similarity measures by imposing further restrictions on μ.

Related work. The Fréchet distance between point trajectories or polygonal curves can be computed in near-quadratic time [2]. The natural generalization to moving (parameterized) curves is to interpret the curves as surfaces parameterized over time and over the curve parameter. The Fréchet distance between surfaces is NP-hard [5], even for terrains [4]. In terms of positive algorithmic results for general surfaces the Fréchet distance is only known to be semi-computable [1].

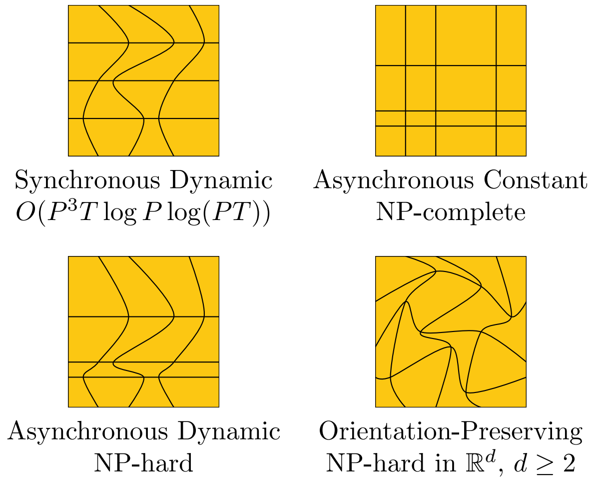

Figure 1: Time complexities; classes of matchings are illustrated as images of regular grids.

When interpreting moving curves as surfaces it is important to take the different roles of the two surface parameters into account: the first is inherently linked to time while the other is linked to space. Depending on the application we do not want to treat these parameters equally. This naturally leads to restricted versions of the Fréchet distance of surfaces. For curves, restricted versions of the Fréchet distance have been previously considered [3, 6]. For surfaces we are not aware of similar results.

Results. We refine the Fréchet distance between surfaces to meaningfully compare moving curves. To do so, we restrict matchings to be one of several suitable classes. Here we often separate the matching into positional and temporal matchings. Representative matchings and running times for the classes considered are illustrated in Fig. 1.

2 Synchronous Dynamic Matchings

Synchronous dynamic matchings align timestamps under the identity matching, but the matching of positions may change continuously over time. Specifi-

*K. Buchin, T. Ophelders, and B. Speckmann are supported by the Netherlands Organisation for Scientific Research (NWO) under project no. 612.001.106 (K. Buchin) and no. 639.023.208 (T. Ophelders & B. Speckmann)

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

cally, the matching is defined as $\mu(p, t) = (\pi_t(p), t)$. Here, $\mu(p, t) : [0, P] \times [0, T] \rightarrow [0, P] \times [0, T]$ is continuous, and for any t the matching $\pi_t : [0, P] \rightarrow [0, P]$ between the two curves is a nondecreasing surjection.

2.1 Freespace

Define the 3D freespace $\mathcal{F}_\varepsilon^{3D} \subseteq [0, P] \times [0, P] \times [0, T]$ by Eq. 2. Then the Fréchet distance is at most ε if and only if for some matching μ of the considered class, all points (x, y, t) with $\mu(x, t) = (y, t)$ lie in $\mathcal{F}_\varepsilon^{3D}$.

$$(x, y, t) \in \mathcal{F}_\varepsilon^{3D} \Leftrightarrow \|A(x, t) - B(y, t)\| \leq \varepsilon \quad (2)$$

Define cells $C_{x,y,t}$ of the freespace with $(x, y, t) \in \mathbb{N}^3$ by Eq. 3 as the freespace between two quadrilaterals.

$$C_{x,y,t} = [x, x + 1] \times [y, y + 1] \times [t, t + 1] \cap \mathcal{F}_\varepsilon^{3D} \quad (3)$$

To determine the conditions under which some matching lies in the freespace, we derive some properties of freespace cells in Lemma 1.

Lemma 1 *Any cell $C_{x,y,t}$ has a convex intersection with any line parallel to the xy-plane or the t-axis.*

2.2 Freespace Partitions in 2D

Whereas previous algorithms for the decision problem of the Fréchet distance between curves compute a path through the freespace, we use a dual problem that extends to moving curves. We illustrate this dual approach in the fictional 2D freespace of Fig. 2. Here, any matching—such as the red path—must be an x- and y-monotone path from the bottom left to the top right corner and this matching must avoid all obstacles. Therefore each such matching divides the obstacles in two sets: those above, and those below the matching.

Suppose we are allowed to draw a directed edge from an obstacle a to an obstacle b if and only if any matching that goes over a must necessarily go over b. The key observation is that no matching exists if and only if such edges can form a path from the lower-right boundary to the upper-left boundary of the freespace. A few of these edges are drawn in black and gray. In the example, observe that if all obstacles were slightly larger, an edge could connect a blue and green obstacle, connecting the boundaries by the black edges.

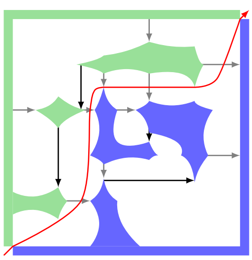

Figure 2: μ in 2D.

In contrast to the 2D freespace where the matching is a path, matchings form surfaces in the case of the 3D freespace. Such a surface again divides the obstacles in the freespace in two sets and can be punctured by a path connecting two boundaries. We shall formalize this approach for the 3D freespace ($\mathcal{F}_\varepsilon^{3D}$).

2.3 Freespace Partitions in 3D

Observe that any matching μ partitions obstacles into two sets, namely those above, and those below μ. Let O be the complete set of obstacles and $D \subseteq O$ be the obstacles below the matching. Then the upper boundary u of the freespace is never in D and the lower boundary d of the freespace is in D for any μ, see Fig. 3. The boundaries lie just outside the freespace. Let O' consist of all obstacles between the boundaries and let each obstacle be a connected subset of $\mathbb{R}^3$.

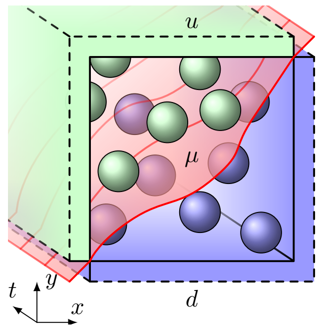

Figure 3: u, d and μ.

$$O = \{u, d\} \cup O' \text{ where } \bigcup O' = F \setminus \mathcal{F}_\varepsilon;$$

$$F = [0, P] \times [0, P] \times [0, T];$$

$$u = \{(x, y, t) \mid (x < 0 \wedge y > 0) \vee (x < P \wedge y > P)\};$$

$$d = \{(x, y, t) \mid (x > 0 \wedge y < 0) \vee (x > P \wedge y < P)\}.$$

Here, we use axes (x, y, t) and say that a point is below some other point if it has a smaller y-coordinate. Because each obstacle is a connected set and μ cannot intersect obstacles, a single obstacle cannot lie on both sides of the same matching. Because all matchings have $u \notin D$ and $d \in D$, a matching exists if and only if $\neg(d \in D \Rightarrow u \in D)$.

We compute a relation $\triangleright$ of elementary dependencies between obstacles, such that its transitive closure $\circledR$ has $d \circledR u$ if and only if $d \in D \Rightarrow u \in D$. Let $a \triangleright b$ if and only if $a \cup b$ is connected (a touches b) or there exists some point $(x_a, y_a, t_a) \in a$ and $(x_b, y_b, t_b) \in b$ with $x_a \leq x_b$, $y_a \geq y_b$ and $t_a = t_b$. By Lemmas 2 and 3, this choice of $\triangleright$ satisfies the required properties and by Theorem 4 we can use the transitive closure $\circledR$ of $\triangleright$ to solve the decision problem of the Fréchet distance.

Lemma 2 *If $a \circledR b$, then $a \in D \Rightarrow b \in D$.*

Lemma 3 *If $d \in D \Rightarrow u \in D$, then $d \circledR u$.*

Theorem 4 *The Fréchet distance is greater than ε if and only if $d \circledR u$ for ε.*

We choose the set of obstacles O' such that $\bigcup O' = F \setminus \mathcal{F}_\varepsilon$ and the relation $\triangleright$ is easily computable. Note that due to Lemma 1, each connected component contains a corner of a cell, therefore any cell in the freespace contains constantly many (up to eight) components of $F \setminus \mathcal{F}_\varepsilon$. As such, we can index the obstacles in O' by a grid point $(x, y, t) \in \mathbb{N}^3$ combined with one of the adjacent cells (with (x, y, t) as a corner).

$O' = \bigcup_{(x,y,t)\in\mathbb{N}^3\cap(F\setminus\mathcal{F}_\varepsilon)} O'_{x,y,t}$ where

$O'_{x,y,t} = \{o_{x,y,t,C} \mid \text{cell } C \text{ has } (x,y,t) \text{ as a corner}\}$,

$o_{x,y,t,C}$ is the maximal connected set with $(x,y,t) \in o_{x,y,t,C} \subseteq (F \setminus \mathcal{F}_\varepsilon) \cap C$.

Since obstacles in $O'_{x,y,t}$ touch at grid point (x,y,t), we treat them as a single obstacle $o_{x,y,t} = \bigcup O'_{x,y,t}$. Two obstacles $o_{x,y,t,C}$ and $o_{x',y',t',C}$ represent the same set of points if (x,y,t) is connected to (x',y',t') within C, but treat two such obstacles as distinct obstacles.

Each of the $O(P^2T)$ obstacles is now defined by a constant number of vertices. We therefore assume that for each pair of obstacles $(a,b) \in O^2$, we can decide in constant time whether $a \rhd b$. For each obstacle a in a cell $C_{x,y,t}$, there can only be $O(P^2)$ obstacles b for which $a \rhd b$; namely obstacles u, d, and those in cells $C_{x',y',t'}$ with $t' \in \{t-1, t, t+1\}$. Therefore we can compute the relation $\rhd$ in $O(P^4T)$ time.

Testing whether $d \circledast u$ is equivalent to testing whether there exists a path from d to u in the directed graph $(O, \rhd)$, which can be decided in $O(|\rhd|)$ time using a depth first search. Thus, the decision problem for the Fréchet distance is solved in $O(|\rhd|) = O(P^4T)$ time. There are many unnecessary edges in $\rhd$ which we do not have to compute (see Theorem 5). To compute the exact Fréchet distance, the parametric search of Section 2.4 is applied to the decision problem.

Theorem 5 *The decision problem for the synchronous dynamic Fréchet distance is solvable in $O(P^3T \log P)$ time.*

2.4 Parametric Search

To give an idea of what the 3D freespace looks like, we have drawn the obstacles of the eight cells of the freespace between two quadrilateral meshes of size $P \times T = 2 \times 2$ in Fig. 4. Cells of the 3D freespace lie within cubes and have six faces and twelve edges. According to the axis to which they are parallel, we denote such edges by x-, y- or t-edges.

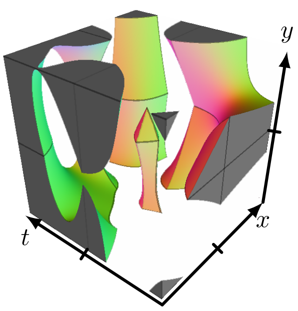

Figure 4: $[0,2]^3 \setminus \mathcal{F}^{3D}_\varepsilon$

We are looking for the minimum value of ε for which a matching exists. When increasing the value of ε, the relation $\rhd$ becomes smaller since obstacles shrink. Critical values of ε occur when $\rhd$ becomes smaller. Due to Lemma 1, all critical values involve an edge or an xt-face or yt-face of a cell, but never the internal volume, so the following critical values cover all cases.

a) The minimal ε such that $(0,0,t) \in \mathcal{F}^{3D}_\varepsilon$ and $(P,P,t) \in \mathcal{F}^{3D}_\varepsilon$ for all t.

b) An edge of $C_{x,y,t}$ becomes nonempty.

c) The endpoints of two y-edges (or two x-edges) of $C_{x,y,t}$ and $C_{x+i,y,t}$ (or $C_{x,y-j,t}$) align.

d) An endpoint of a t-edge of $C_{x,y,t}$ aligns with an endpoint of a t-edge of $C_{x+i,y-j,t}$.

e) An obstacle in $C_{x,y,t}$ stops overlapping with an obstacle in $C_{x+i,y,t}$ or $C_{x,y-j,t}$ when projected along the x- or y-axis.

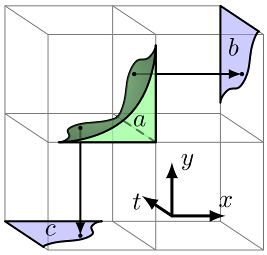

Figure 5: $a \rhd b$ and $a \rhd c$

We illustrate the need for critical values of type e) in Fig. 5. Here obstacle a overlaps with both obstacles b and c while the overlap in edges does not contribute to $\rhd$. The critical values of types a), b) and c) resemble those in the paper by Alt and Godau [2]. The endpoints involved in the critical values of type a), b), c) and d) can be captured in $O(P^2T)$ functions.

We apply a parametric search [7] on them to find the minimum critical value $\varepsilon_{\text{abcd}}$ of type a), b), c) or d) for which a matching exists. This parametric search takes $O((P^2T + \mathit{time}_{\text{dec}})\log(PT))$ time where $\mathit{time}_{\text{dec}} = O(P^3T\log P)$ is given by Theorem 5.

It is unclear how critical values of type e) can be incorporated in the parametric search directly. Instead, we enumerate and sort the $O(P^3T)$ critical values of type e) in $O(P^3T\log(PT))$ time. Using $O(\log(PT))$ calls to the decision algorithm, we apply a binary search to find the minimum critical value ε_{e} of type e) for which a matching exists. Finding ε_{e} then takes $O((P^3T + \mathit{time}_{\text{dec}})\log(PT))$ time.

The synchronous dynamic Fréchet distance is then the minimum of $\varepsilon_{\text{abcd}}$ and ε_{e}. Because the decision problem takes $\mathit{time}_{\text{dec}} = O(P^3T\log P)$ time, the running time of Theorem 6 is achieved for the exact Fréchet distance.

Theorem 6 *The synchronous dynamic Fréchet distance can be computed in $O(P^3T\log P\log(PT))$ time.*

3 Hardness

We extend the synchronous dynamic class of matchings to the asynchronous dynamic class by allowing realignments of timestamps. Matchings of this class have the form $\mu(p,t) = (\pi_t(p), \tau(t))$ where π and τ are realign positions and timestamps and the positional matching π_t changes over time. In the more restricted asynchronous constant class $\mu(p,t) = (\pi(p), \tau(t))$ the positional matching cannot change over time. The Fréchet distance is in NP for this class because piecewise linear π and τ exist whenever a matching exist.

Theorem 7 *Computing the Fréchet distance is in NP for the asynchronous constant class of matchings.*

Due to critical values of type e), it is unclear whether every asynchronous dynamic matching admits a piecewise-linear matching τ^* of polynomial size, which would mean that the asynchronous dynamic Fréchet distance is also in NP.

Computing the Fréchet distance is NP-hard for both classes by a reduction from 3-SAT. The idea behind the construction is illustrated in the two height maps of Fig. 6. The height maps represent quadrilateral meshes embedded in $\mathbb{R}^1$ and correspond to a single clause in a 3-CNF formula of four variables.

We distinguish valleys (dark), peaks (white on A, yellow on B) and ridges (denoted X_i, F_i and T_i). Observe that to obtain a low Fréchet distance of $\varepsilon < 3$, the n-th valley of A must be matched with the n-th valley of B. Moreover, each ridge X_i must be matched with F_i or T_i and each peak of A must be matched to a peak of B. Note that even for asynchronous dynamic matchings, if X_i is matched to F_i it cannot be matched to T_i and vice-versa because the (red) valley separating F_i and T_i has distance 3 from X_i.

Consider a 3-CNF formula with n variables and m clauses, then A and B consist of m clauses along the t-axis and n variables ($X_1 \dots X_n$ and $F_1, T_1 \dots F_n, T_n$) along the p-axis. The k-th clause of A is matched to the k-th clause of B due to the elevation pattern on the far left ($p = 0$). This means that the peaks of A are matched with peaks of the same clause on B and for these peaks have the same timestamp because $\tau(t)$ does not depend on p. For each clause, there are three rows (timestamps) of B with peaks on the ridges. On each such timestamp, exactly one ridge (depending on the disjuncts of the clause) does not have a peak. Specifically, if a clause has X_i or $\neg X_i$ as its k-th disjunct, then the k-th row of that clause has no peak on ridge F_i or T_i, respectively. By Theorem 10, it is then NP-hard to approximate the Fréchet distance within a factor 1.5.

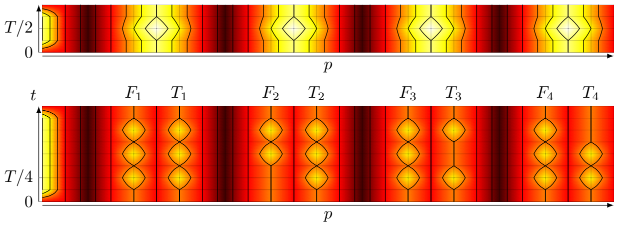

Figure 6: Meshes A (top) and B (bottom) in $\mathbb{R}^1$ (indicated by color). Their Fréchet distance is two isolines if $(X_2 \vee \neg X_3 \vee \neg X_4)$ is satisfiable and three otherwise.

Under the class of orientation-preserving homeomorphisms (restricted by aligning the four corners between of the meshes), we can embed the meshes in $\mathbb{R}^2$ and ensure that all points on A of the same timestamp are matched to similar timestamps of B and Theorem 11 follows.

Lemma 8 *The Fréchet distance between two such moving curves is at least* 3 *if the corresponding 3-CNF formula is unsatisfiable.*

Lemma 9 *The Fréchet distance between two such moving curves is at most* 2 *if the corresponding 3-CNF formula is satisfiable.*

Theorem 10 *No polynomial time algorithm can approximate the asynchronous constant or asynchronous dynamic Fréchet distance between two quadrilateral meshes in* $\mathbb{R}^1$ *within a factor* 1.5 *unless P=NP.*

Theorem 11 *No polynomial time algorithm can approximate the orientation-preserving Fréchet distance between quadrilateral meshes in* $\mathbb{R}^2$ *under the maximum norm within a factor* 1.5 *unless P=NP.*

References

[1] H. Alt and M. Buchin. Can we compute the similarity between surfaces? *Discrete Comput. Geom.*, 43(1):78–99, 2010.

[2] H. Alt and M. Godau. Computing the Fréchet distance between two polygonal curves. *Internat. J. Comput. Geom. Appl.*, 5(01n02):75–91, 1995.

[3] K. Buchin, M. Buchin, and J. Gudmundsson. Constrained free space diagrams: a tool for trajectory analysis. *International Journal of Geographical Information Science*, 24:1101–1125, 2010.

[4] K. Buchin, M. Buchin, and A. Schulz. Fréchet distance of surfaces: Some simple hard cases. In *Algorithms–ESA 2010*, pages 63–74. Springer, 2010.

[5] M. Godau. *On the complexity of measuring the similarity between geometric objects in higher dimensions.* PhD thesis, Berlin, Freie Universität Berlin, Diss., 1998, 1999.

[6] A. Maheshwari, J.-R. Sack, K. Shahbaz, and H. Zarrabi-Zadeh. Fréchet distance with speed limits. *Comput. Geom. Theory Appl.*, 44(2):110–120, 2011.

[7] N. Megiddo. Applying parallel computation algorithms in the design of serial algorithms. *J. ACM*, 30(4):852–865, 1983.