

Matching polygonal curves with respect to the Fréchet distance

Helmut Alt, Christian Knauer, and Carola Wenk*

Institut für Informatik, Freie Universität Berlin, Takustraße 9, D-14195 Berlin,
Germany.

{alt,knauer,wenk}@inf.fu-berlin.de

Abstract. We provide the first algorithm for matching two polygonal curves P and Q under translations with respect to the Fréchet distance. If P and Q consist of m and n segments, respectively, the algorithm has runtime $\mathcal{O}((mn)^3(m+n)^2 \log(m+n))$. We also present an algorithm giving an approximate solution as an alternative. To this end, we generalize the notion of a reference point and observe that all reference points for the Hausdorff distance are also reference points for the Fréchet distance. Furthermore we give a new reference point that is substantially better than all known reference points for the Hausdorff distance. These results yield a $(1 + \epsilon)$ -approximation algorithm for the matching problem that has runtime $\mathcal{O}(\epsilon^{-2}mn)$.

Keywords: Computational geometry, Shape matching, Fréchet distance, Parametric search, Approximation algorithm, Reference point, Steiner point.

1 Introduction

The task of comparing two two-dimensional shapes arises naturally in many applications, e.g., in computer graphics, computer vision and computer aided design. Often two-dimensional shapes are given by the planar curves forming their boundaries which directly leads to the problem of comparing two planar curves. There are several possible distance measures to assess the ‘resemblance’ of the shapes, and there are also different kinds of transformations that are allowed to match them, see [5] for a survey. We will focus here on the *Fréchet distance* δ_F for polygonal curves, and we will search for a *translation* which, when applied to the first curve, minimizes the Fréchet distance to the second one. In [4] it is shown how to compute the Fréchet distance for two polygonal curves.

The only algorithm we know of that decides whether there is a transformation that, when applied to the first curve, results in a Fréchet distance less or equal than some given parameter ϵ (this is called the *decision problem*, see Problem 2 below) is presented in [10], where the admissible transformations are translations

* This research was supported by the Deutsche Forschungsgemeinschaft under Grant No. AL 253/4-3.

in a fixed direction. But to our knowledge there is no algorithm which actually computes the Fréchet distance under a non-trivial class of transformations¹.

In the following we will adopt some basic definitions and results from [4] on which we will subsequently build up.

Definition 1 (Polygonal curve) A continuous mapping $f : [a, b] \rightarrow \mathbb{R}^2$ with $a, b \in \mathbb{R}$ and $a < b$ is called a curve. A polygonal curve is a curve $P : [0, n] \rightarrow \mathbb{R}^2$ with $n \in \mathbb{N}$, such that for all $i \in \{0, 1, \dots, n-1\}$ each $P_i := P|_{[i, i+1]}$ is affine, i.e., $P(i + \lambda) = (1 - \lambda)P(i) + \lambda P(i + 1)$ for all $\lambda \in [0, 1]$.

Definition 2 (Fréchet distance) Let $f : [a, a'] \rightarrow \mathbb{R}^2$ and $g : [b, b'] \rightarrow \mathbb{R}^2$ be curves. Then $\delta_F(f, g)$ denotes their Fréchet distance, defined as

$$\delta_F(f, g) := \inf_{\substack{\alpha : [0, 1] \rightarrow [a, a'] \\ \beta : [0, 1] \rightarrow [b, b']}} \max_{t \in [0, 1]} \|f(\alpha(t)) - g(\beta(t))\|.$$

where $\|\cdot\|$ denotes the L_2 norm, and α, β range over continuous and increasing functions with $\alpha(0) = a, \alpha(1) = a', \beta(0) = b$ and $\beta(1) = b'$ only.

As a popular illustration of the Fréchet-metric suppose a man is walking his dog, he is walking on the one curve the dog on the other. Both are allowed to control their speed but are not allowed to go backwards. Then the Fréchet distance of the curves is the minimal length of a leash that is necessary.

In the rest of the paper we will develop algorithms for the following two problems:

Problem 1 (δ_F – optimization problem)

Given two polygonal curves P, Q , and a class of transformations \mathcal{T} .

Find a $\tau \in \mathcal{T}$ such that $\delta_F(\tau(P), Q)$ is as small as possible.

Similar to [4] we will first consider the decision problem which we will afterwards optimize applying Megiddo’s parametric search technique, c.f. [8]. The decision problem in our setting is the following:

Problem 2 (δ_F – decision problem)

Given two polygonal curves P, Q , a class of transformations \mathcal{T} , and $\epsilon \geq 0$.

Decide, whether there exists a $\tau \in \mathcal{T}$ such that $\delta_F(\tau(P), Q) \leq \epsilon$.

We will show that in the case of translations we can solve the decision problem in $\mathcal{O}((mn)^3(m+n)^2)$ time. The parametric search adds only a logarithmic overhead, since we can apply Cole’s trick for parametric search based on sorting, so we can solve the optimization problem in $\mathcal{O}((mn)^3(m+n)^2 \log(m+n))$ time.

¹ We recently learned that Efrat et al. [7] have independently developed an algorithm for the decision problem under translations. However, the runtime they achieve is by a quadratic factor slower than ours, and their result is rather complicated and relies on complex data structures.

2 Computing the Fréchet distance

Throughout the rest of the paper let $P : [0, m] \rightarrow \mathbb{R}^2$ and $Q : [0, n] \rightarrow \mathbb{R}^2$ be polygonal curves. Unless stated otherwise $\epsilon \geq 0$ is a fixed real parameter. In the sequel we will use the notion of a *free space* which was introduced in [4]:

Definition 3 (Free space, [4]) *The set $F_\epsilon(P, Q) := \{(s, t) \in [0, m] \times [0, n] \mid \|P(s) - Q(t)\| \leq \epsilon\}$, or F_ϵ for short, denotes the free space of P and Q .*

Sometimes we refer to $[0, m] \times [0, n]$ as the *free space diagram*; the *feasible* points $p \in F_\epsilon$ will be called ‘white’ and the *infeasible* points $p \in [0, m] \times [0, n] - F_\epsilon$ will be called ‘black’ (for obvious reasons, c.f. Figure 1). Consider $[0, m] \times [0, n]$ as composed of the mn cells $C_{i,j} := [i-1, i] \times [j-1, j]$ $1 \leq i \leq n$, $1 \leq j \leq m$. Then $F_\epsilon(P, Q)$ is composed of the mn free spaces for each pair of edges $F_\epsilon(P_{i-1}, Q_{j-1}) = F_\epsilon(P, Q) \cap C_{i,j}$.

The following results from [4] describe the structure of the free space and link it to the problem of computing δ_F .

Lemma 4 (Alt/Godau, [4]) *The free space of two line segments is the intersection of the unit square with an affine image of the unit disk, i.e., with an ellipse, possibly degenerated to the space between two parallel lines.*

Lemma 5 (Alt/Godau, [4]) *For polygonal curves P and Q we have $\delta_F(P, Q) \leq \epsilon$, exactly if there exists a curve within $F_\epsilon(P, Q)$ from $(0, 0)$ to (m, n) which is monotone in both coordinates.*

For proofs of the above two Lemmas see [4]. Figure 1 shows polygonal curves P, Q , a distance ϵ , and the corresponding diagram of cells $C_{i,j}$ with the free space F_ϵ . Observe that the curve as a continuous mapping from $[0, 1]$ to $[0, m] \times [0, n]$ directly gives feasible reparametrizations, i.e., two reparametrizations α and β , such that $\max_{t \in [0, 1]} \|f(\alpha(t)) - g(\beta(t))\| \leq \epsilon$.




Fig. 1. Two polygonal curves P and Q and their free space diagram for a given ϵ . An example monotone curve in the free space (c.f. Lemma 5) is drawn bold.

For $(i, j) \in \{1, \dots, m\} \times \{1, \dots, n\}$ let $L_{i,j}^F := \{i-1\} \times [a_{i,j}, b_{i,j}]$ (or $B_{i,j}^F := [c_{i,j}, d_{i,j}] \times \{j-1\}$) be the left (or bottom) line segment bounding $C_{i,j} \cap F_\epsilon$ (see Figure 2).




Fig. 2. Intervals of the free space on the boundary of a cell.

By induction it can easily be seen that those parts of the segments $L_{i,j}^F$ and $B_{i,j}^F$ which are reachable from $(0,0)$ by a monotone path in F_ϵ are also line segments. Using a dynamic programming approach one can compute them, and thus decide if $\delta_F(P, Q) \leq \epsilon$. For details we refer the reader to the proof of the following theorem in [4]:

Theorem 6 (Alt/Godau, [4]) *For given polygonal curves P, Q and $\epsilon \geq 0$ one can decide in $\mathcal{O}(mn)$ time, whether $\delta_F(P, Q) \leq \epsilon$.*

Now let us observe a continuity property of F_ϵ : As we have already mentioned, each (possibly clipped ellipse) in F_ϵ is the affine image of a unit disk. Thus each ellipse in F_ϵ varies continuously in ϵ . This implies the following observation:

Observation 7 (see [4]) *If $\epsilon = \delta_F(P, Q)$, then F_ϵ contains at least one monotone path from $(0,0)$ to (m,n) and for each such path π one of the following cases occurs:*

- a) $L_{i,j}^F$ or $B_{i,j}^F$ is a single point on π for some pair (i, j) . (The path passes through a passage between two neighboring cells that consists of a single point.)
- b) $a_{i,j} = b_{k,j}$ (or $c_{i,j} = d_{i,k}$) for some i, j, k and π passes through $(i, a_{i,j})$ and $(k, b_{k,j})$ (or π passes through $(c_{i,j}, j)$ and $(d_{i,k}, k)$). (The path contains a ‘clamped’ horizontal or vertical passage, see Figure 3.)

Figure 4 shows the geometric situations that correspond to these two cases. In case a) the reparametrization maps the point $P(i-1)$ to the only point on the edge Q_j that has distance ϵ from $P(i-1)$. In case b) it maps the part of P between $P(i-1)$ and $P(k-1)$ to the only point on the edge Q_j that has distance ϵ from $P(i-1)$ and $P(k-1)$. This situation covers the case of horizontally clamped paths. The geometric situations that involve a vertically clamped passage are similar, with the roles of P and Q interchanged. Note that we can actually view case a) as a special case of case b) with $i = k$.




Fig. 3. The path contains a ‘clamped’ horizontal passage in the j -th row between the spikes A and B .




Fig. 4. The geometric situations corresponding to a horizontally clamped path.

3 Minimizing the Fréchet distance

First we give a rough sketch of the basic idea of our algorithm: Assume that there is at least one translation that moves P to a Fréchet distance at most ϵ to Q . Then we can move P to a position $\tau_{=}$ where the Fréchet distance to Q is exactly ϵ . According to Observation 7 the free space diagram $F_{\epsilon}(\tau_{=}(P), Q)$ then contains at least one clamped path. As a consequence, one of the geometric situations from Figure 4 must occur. Therefore the set of translations that attain a Fréchet distance of exactly ϵ is a subset of the set of translations that realize at least one of those geometric situations. The set of translations that create a geometric situation involving the two different vertices $P(i-1)$ and $P(k-1)$ from P and the edge Q_j from Q consist of two segments in transformation space, i.e., it can be described geometrically.

Now assume that the geometric situation from above is specified by the two vertices $P(i-1)$ and $P(k-1)$ and the edge Q_j . When we move P in such a way that $P(i-1)$ and $P(k-1)$ remain at distance ϵ from a common point on an edge of Q (i.e., we shift P ‘along’ Q), we will preserve one geometric situation (namely the one involving $P(i-1)$ and $P(k-1)$ and some edge of Q). At some point however, we will reach a placement where the Fréchet distance becomes larger than ϵ . This means that immediately before that point it was exactly ϵ , so there is a second placement $\tau'_{=}$ such that $F_{\epsilon}(\tau'_{=}(P), Q)$ contains at least one clamped path and consequently another geometric situation must occur. So the

set of translations that attain a Fréchet distance of exactly ϵ is a subset of the set of translations that realize at least *two* such geometric situations.

After this informal description of the basic ideas let us go into more detail now:

Convention: *In this section \mathcal{T}_2 denotes the group of planar translations, P and Q are polygonal curves with m and n vertices, respectively and $\epsilon \geq 0$ is a real parameter.*

A translation $\tau = \langle (x, y) \mapsto (x + \delta_x, y + \delta_y) \rangle \in \mathcal{T}_2$ can be specified by the pair $(\delta_x, \delta_y) \in \mathbb{R}^2$ of *parameters*. The set of parameters of all translations in \mathcal{T}_2 is called the parameter space of \mathcal{T}_2 , or *translation space* for short, and we identify \mathcal{T}_2 with its parameter space.

Let us now take a look at the free space $F_\epsilon(\tau(P), Q)$ when τ varies over \mathcal{T}_2 . Now we show that each of the $\mathcal{O}(mn)$ ellipses (and thus also each clipped ellipse) varies continuously in $\tau \in \mathcal{T}_2$. In fact we consider all mn ellipses, even those that have an empty intersection with their corresponding square in the diagram (let us call these 'invisible'). Note that an ellipse is generated by two linearly independent line segments; one from P and one from Q . Parallel line segments generate only a 'degenerate ellipse', namely the space between two parallel lines. So if we fix a translation $\tau \in \mathcal{T}_2$ it is easy to see that each (possibly invisible) ellipse in $F_\epsilon(\tau(P), Q)$ is a translation of the corresponding ellipse in $F_\epsilon(P, Q)$. In fact, the translation is $(-\lambda, \mu)$ where λ and μ are the coefficients which are obtained by representing τ as a linear combination of the direction unit vectors of the line segments. Thus each ellipse varies continuously in $\tau \in \mathcal{T}_2$. Note that this is still true if an ellipse is visible but its translate is invisible or vice versa. A similar argument holds for degenerate ellipses.

Definition 8 (Configuration) *A triple (p, p', s) that consists of two (not necessarily different) vertices p and p' of P and an edge s of Q is called an h-configuration. v-configurations are defined analogously with the roles of P and Q exchanged. A configuration is an h- or v-configuration.*

Definition 9 (Critical translations) *Let $c = (x, y, s)$ be an h-configuration and $c' = (x', y', s')$ be a v-configuration. The sets*

$$\begin{aligned} T_{crit}(c) &:= \{\tau \in \mathcal{T}_2 \mid \exists z \in s : \|\tau(x) - z\| = \|\tau(y) - z\| = \epsilon\} \quad \text{and} \\ T_{crit}(c') &:= \{\tau \in \mathcal{T}_2 \mid \exists z' \in s' : \|x' - \tau(z')\| = \|y' - \tau(z')\| = \epsilon\} \end{aligned}$$

are called the sets of critical translations for c and c' . A translation is called critical if it is critical for some configuration.

Lemma 10 *If $\delta_F(\tau(P), Q) = \epsilon$, then τ is critical.*

Proof. By Observation 7 there is a path π in $F_\epsilon(\tau(P), Q)$ for which case a) or b) occurs. If the corresponding geometric situation (c.f. Figure 4) involves the vertices $\tau(P(i-1))$ and $\tau(P(k-1))$ on $\tau(P)$ and a point on the edge Q_j then

the translation τ is critical for the h-configuration $(P(i-1), P(k-1), Q_j)$. If the geometric situation involves vertices from Q and a segment of $\tau(P)$, the same argument yields a v-configuration. \square

Note that the condition in Lemma 10 is only necessary but not sufficient, i.e., there are indeed critical translations τ with $\delta_F(\tau(P), Q) \neq \epsilon$. This is because a critical translation for a configuration (x, y, s) does not even have to map the part of the curve between x and y within distance ϵ to the corresponding point on s .

Let us now take a closer look at the critical translations in \mathcal{T}_2 : For a given configuration (x, y, s) with two different vertices (which corresponds to case (b) in Figure 4) the set of critical translations is described by two parallel line segments in translation space, where each line segment is a translate of s . If the two vertices in the configuration are the same (which is case (a) in Figure 4) the set of critical translations is described by a 'racetrack' in translation space, which is the locus of points having distance ϵ to a translate of s . Note that a 'racetrack' consists of line segments and circular arcs.

We call the arrangement in translation space consisting of the curves describing all critical translations of all configurations the *arrangement of critical translations*. There are $\mathcal{O}(mn(m+n))$ different configurations, so the combinatorial complexity of the arrangement of critical translations (i.e., the number of vertices line segments and circular arcs) is $\mathcal{O}((mn(m+n))^2)$.

Lemma 11 *If there is a translation $\tau_{\leq} \in \mathcal{T}_2$ such that $\delta_F(\tau_{\leq}(P), Q) \leq \epsilon$ then there is a translation $\tau_{=} \in \mathcal{T}_2$ that is critical such that $\delta_F(\tau_{=}(P), Q) = \epsilon$.*

Proof. Pick any translation $\tau_{>} \in \mathcal{T}_2$ such that $\delta_F(\tau_{>}(P), Q) > \epsilon$. By continuity, there exists a translation $\tau_{=}$ on any curve between τ_{\leq} and $\tau_{>}$ in translation space such that $\delta_F(\tau_{=}(P), Q) = \epsilon$. By Lemma 10 the translation $\tau_{=}$ is critical. \square

This result states that, whenever there is some translation τ_{\leq} that moves P into Fréchet distance at most ϵ to Q , there is also a 'canonical' translation $\tau_{=}$ that results in a Fréchet distance exactly ϵ and that lies on the arrangement of critical translations. So in order to check if there is a translation that moves P into Fréchet distance at most ϵ to Q , it is sufficient to check all translations on the arrangement of critical translations.

However, since the translation space has more than one degree of freedom, the arrangement of critical translations contains an infinite number of translations. So our observation does not help from an algorithmic point of view. Lemma 14 shows that we can restrict our attention to the zero-dimensional parts of the arrangement, i.e., intersection points and endpoints of the curves describing the critical translations. First we need the following two observations:

Observation 12 *Let $c = (P(i-1), P(k-1), Q_j)$, with $i \neq k$, be an h-configuration. Then $a_{i,j} = b_{k,j}$ in $F_{\epsilon}(\tau(P), Q)$ for all $\tau \in T_{crit}(c)$, i.e., the relative position of the two spikes stays the same for all $\tau \in T_{crit}(c)$ (c.f. Figures 3 and 4).*

Observation 13 Let $c = (P(i-1), P(k-1), Q_j)$, with $i \neq k$, be an h -configuration. Now we consider a feasible reparametrization for some $\tau_{=} \in T_{crit}(c)$, that maps the part of $\tau_{=}(P)$ between $\tau_{=}(P(i-1))$ and $\tau_{=}(P(k-1))$ to a point on Q_j . This corresponds to a path $\pi_{=}$ in $F_{\epsilon}(\tau_{=}(P), Q)$ that is clamped between the two corresponding vertical spikes in cell (i, j) and (k, j) of $F_{\epsilon}(\tau_{=}(P), Q)$. Now from Observation 12 it follows that for each $\tau \in T_{crit}(c)$ the relative position of the spikes does not change, i.e., we cannot ‘destroy’ $\pi_{=}$ locally by moving along $T_{crit}(c)$.

Of course both observations remain true if we consider h -configurations that correspond to case a) of Observation 7 (where $i = k$) or v -configurations (where the roles of P and Q are interchanged).

Lemma 14 If there is a translation $\tau_{\leq} \in \mathcal{T}_2$ such that $\delta_F(\tau_{\leq}(P), Q) \leq \epsilon$ then there is a translation $\tau_{=} \in \mathcal{T}_2$ that lies on a vertex of the arrangement of critical translations such that $\delta_F(\tau_{=}(P), Q) = \epsilon$.

Proof. Suppose all vertices of the arrangement of critical translations yield a Fréchet distance greater than ϵ . By Lemma 11 there is a critical translation $\tau_{=} \in \mathcal{T}_2$ with $\delta_F(\tau_{=}(P), Q) = \epsilon$, and by definitions 8 and 9 there is a configuration c such that $\tau_{=} \in T_{crit}(c)$. Now pick any translation $\tau_{>} \in \mathcal{T}_2$ such that $\delta_F(\tau_{>}(P), Q) > \epsilon$. We can assume without loss of generality that $\tau_{=}$ lies in an ‘extreme’ position on $T_{crit}(c)$, which means that $\delta_F(\tau(P), Q) > \epsilon$ for every $\tau \in T_{crit}(c)$ that lies ‘between’ $\tau_{=}$ and $\tau_{>}$. Considering the free space diagram this means that $F_{\epsilon}(\tau_{=}(P), Q)$ contains a monotone path, but $F_{\epsilon}(\tau(P), Q)$ does not contain this or any other monotone path anymore. By continuity this can only happen, if each monotone path in $F_{\epsilon}(\tau_{=}(P), Q)$ is ‘clamped’ between two ‘spikes’ which close the narrow passage in the free space when moving from $\tau_{=}$ to $\tau_{>}$ on $T_{crit}(c)$.

But according to Observation 12 this cannot be true for the spikes corresponding to the critical translations $T_{crit}(c)$. Thus there must be another configuration c' such that $\tau_{=} \in T_{crit}(c')$, and close to $\tau_{=}$ the curve describing $T_{crit}(c')$ differs from $T_{crit}(c)$. Since both curves are algebraic, $T_{crit}(c) \cap T_{crit}(c')$ is zero-dimensional, and thus a vertex of the arrangement. \square

So in order to solve the decision problem for a given ϵ it is sufficient to check for all translations τ that correspond to vertices of the arrangement of critical translations whether $\delta_F(\tau(P), Q) \leq \epsilon$. We thus have altogether $\mathcal{O}((mn)^2(m+n)^2)$ translations for each of which we check in $\mathcal{O}(mn)$ time if it brings P into distance at most ϵ to Q , which solves Problem 4 for the case of translations and yields the following theorem:

Theorem 15 For given polygonal curves P, Q and $\epsilon \geq 0$ one can decide in $\mathcal{O}((mn)^3(m+n)^2)$ time whether there is a translation $\tau \in \mathcal{T}_2$ such that $\delta_F(\tau(P), Q) \leq \epsilon$.

In order to find a translation that minimizes the Fréchet distance between the two polygonal curves we apply the parametric search paradigm. For this

we generalize the approach of [4]. Remember that for a given configuration $c = (x, y, s)$ the set of critical translations $T_{crit}(c)$ is described by two parallel line segments or by a 'racetrack' in translation space. Now when we let ϵ vary $T_{crit}(c)$ changes accordingly, namely the distance between the parallel line segments or the radius of the 'racetrack' varies depending on ϵ . Note that for small ϵ , $T_{crit}(c)$ might even be empty, which happens for example when $\|x - y\| < \epsilon$.

For a given ϵ let $S(\epsilon)$ be the set of $\mathcal{O}((mn(m+n))^2)$ vertices of the arrangement of critical translations. In fact, one can track each vertex in $S(\epsilon)$ for varying ϵ , i.e., one can interpret each vertex in $S(\epsilon)$ as a function of ϵ . Let S be the set of these vertex-functions. Note that the vertex-functions in S might not be defined for small ϵ . For each of the $\mathcal{O}((mn(m+n))^2)$ translation functions $\tau(\epsilon)$ in S we compute the free space $F_\epsilon(\tau(\epsilon)(P), Q)$ depending on ϵ . In fact we only compute all $a_{i,j}(\tau, \epsilon)$, $b_{i,j}(\tau, \epsilon)$, $c_{i,j}(\tau, \epsilon)$, and $d_{i,j}(\tau, \epsilon)$ which depend on ϵ and τ , and of which there are $\mathcal{O}((mn)^3(m+n)^2)$. For the parametric search an ϵ is critical if two of these functions have the same value (for the same translation function τ). A parametric search over all $\mathcal{O}((mn)^3(m+n)^2)$ values of $a_{i,j}(\tau, \epsilon)$, $b_{i,j}(\tau, \epsilon)$, $c_{i,j}(\tau, \epsilon)$, and $d_{i,j}(\tau, \epsilon)$ thus yields an optimum ϵ together with an optimum translation. As in [4] we apply a parallel sorting algorithm which generates a superset of the critical values of ϵ we need. By utilizing Cole's trick [6] for parametric search based on sorting, which in general yields a runtime of $\mathcal{O}((k + T_{seq}) \log k)$ where T_{seq} is the sequential runtime for the decision problem and k is the number of values to be sorted, we obtain a runtime of $\mathcal{O}((mn)^3(m+n)^2 \log(m+n))$. This solves Problem 3 for the case of translations and proves the following theorem:

Theorem 16 *For given polygonal curves P, Q one can compute a translation τ_{min} in $\mathcal{O}((mn)^3(m+n)^2 \log(m+n))$ time, such that $\delta_F(\tau_{min}(P), Q) = \min_{\tau \in \mathcal{T}_2} \delta_F(\tau(P), Q)$.*

3.1 Other transformation classes

We are currently investigating the application of the techniques from above to other classes of transformations, such as translations in a fixed direction, rotations around a fixed center, rigid motions, and arbitrary affine maps, for matching curves in two and higher dimensions.

In the parameter space of the transformation class under consideration the set of critical transformations for a configuration is a semi-algebraic set in general, which is defined by a constant number of polynomials of bounded degree. Therefore we can define the arrangement of critical transformations in the same way as before.

A suitable generalization of Lemma 14 should imply that only the zero-dimensional pieces of this arrangement have to be considered as candidates for a successful match. This immediately yields an algorithm with a runtime that depends on the complexity of the arrangement of critical transformations, which in turn depends on the dimension of the parameter space as well as on the dimension of the underlying Euclidean space.

4 Approximately minimizing the Fréchet distance

The algorithms we described so far cannot be considered to be efficient. To remedy this situation, we present approximation algorithms which do not necessarily compute the optimal transformation, but one that yields a Fréchet distance which differs from the optimum value by a constant factor only. To this end, we generalize the notion of a *reference point*, c.f. [2] and [1], to the Fréchet metric and observe that all reference points for the Hausdorff distance are also reference points for the Fréchet distance.

We first need the concept of a *reference point* that was introduced in [1]. A reference point of a figure is a characteristic point with the property that similar figures have reference points that are close to each other. Therefore we get a reasonable matching of two figures if we simply align their reference points.

Definition 17 (Reference point, [1]) *Let \mathcal{K} be a set of compact subsets of \mathbb{R}^2 and δ be a metric on \mathcal{K} . A mapping $\mathbf{r} : \mathcal{K} \rightarrow \mathbb{R}^2$ is called a δ -reference point for \mathcal{K} of quality $c > 0$ with respect to a set of transformations \mathcal{T} on \mathcal{K} , if the following holds for any two sets $P, Q \in \mathcal{K}$ and each transformation $\tau \in \mathcal{T}$:*

$$\text{(Equivariance)} \quad \mathbf{r}(\tau(P)) = \tau(\mathbf{r}(P)) \quad (1)$$

$$\text{(Lipschitz continuity)} \quad \|\mathbf{r}(P) - \mathbf{r}(Q)\| \leq c \cdot \delta(P, Q). \quad (2)$$

In other words a reference point is a Lipschitz-continuous mapping between the metric spaces (\mathcal{K}, δ) and $(\mathbb{R}^2, \|\cdot\|)$ with Lipschitz constant c , which is equivariant under \mathcal{T} . Various reference points are known for a variety of distance measures and classes of transformations, like, e.g., the centroid of a convex polygon which is a reference point of quality $11/3$ for translations, using the area of the symmetric difference as a distance measure, see [3]. However, most work on reference points has focused on the Hausdorff distance, see [1].

Definition 18 (Hausdorff distance) *Let P and Q be curves. Then $\delta_H(P, Q)$ denotes their Hausdorff distance, defined as*

$$\delta_H(P, Q) := \max(\tilde{\delta}_H(P, Q), \tilde{\delta}_H(Q, P)), \quad \text{with}$$

$$\tilde{\delta}_H(X, Y) := \sup_{x \in X} \inf_{y \in Y} \|x - y\|, \quad \text{the one-sided Hausdorff distance from } X \text{ to } Y.$$

We will only mention the following result that provides a δ_H -reference point for polygonal curves with respect to similarities, the so called *Steiner point*. The Steiner point of a polygonal curve is the weighted average of the vertices of the convex hull of the curve, where each vertex is weighted by its exterior angle divided by 2π .

Theorem 19 (Aichholzer et al., [1]) *The Steiner point is a δ_H -reference point with respect to similarities of quality $4/\pi$. It can be computed in linear time.*

Note that the Steiner point is an *optimal* δ_H -reference point with respect to similarities, i.e., the quality of any δ_H -reference point for that transformation class is at least $4/\pi$, see [1].

Two feasible reparametrizations α and β of P and Q demonstrate, that for each point $P(\alpha(t))$ there is a point $Q(\beta(t))$ with $\|P(\alpha(t)) - Q(\beta(t))\| \leq \epsilon$ (and vice versa), thus $\delta_H(P, Q) \leq \delta_F(P, Q)$. This shows the following observation:

Observation 20 *Let $c > 0$ be a constant and \mathcal{T} be a set of transformations on \mathcal{K} . Then each δ_H -reference point with respect to \mathcal{T} is also a δ_F -reference point with respect to \mathcal{T} of the same quality.*

This shows that we can use the known δ_H -reference points to obtain δ_F -reference points. However, since each reparametrization has to map $P(0)$ to $Q(0)$, the distance $\|P(0) - Q(0)\|$ is a lower bound for $\delta_F(P, Q)$. So we get a new reference point that is substantially better than all known reference points for the Hausdorff distance.

Observation 21 *Let \mathcal{C}_o be the set of all planar curves. The mapping*

$$\mathbf{r}_o : \begin{cases} \mathcal{C}_o \rightarrow \mathbb{R}^2 \\ P \mapsto P(0) \end{cases}$$

is a δ_F -reference point for curves of quality 1 with respect to translations.

The quality of this reference point, i.e., 1, is better than the quality of the Steiner point, which is $4/\pi$. Since the latter is an optimal reference point for the Hausdorff distance, this shows that for the Fréchet distance substantially better reference points exist. For closed curves however \mathbf{r}_o is not defined at all.

Based on the existence of a δ_F -reference point for \mathcal{T}_2 we obtain the following algorithm for approximate matchings with respect to the Fréchet distance under the group of translations, which is the same procedure as already used in [1] for the Hausdorff distance.

Algorithm T: Compute $\mathbf{r}(P)$ and $\mathbf{r}(Q)$, translate P by $\tau := \mathbf{r}(Q) - \mathbf{r}(P)$, and output this matching as the approximate solution, together with $\delta_F(\tau(P), Q)$.

Theorem 22 *Suppose that \mathbf{r} is a δ_F -reference point of quality c with respect to translations that can be computed in $\mathcal{O}(T_{\mathbf{r}}(n))$ time. Then algorithm T produces a $(c + 1)$ -approximation to Problem 1 in $\mathcal{O}(mn + T_{\mathbf{r}}(m) + T_{\mathbf{r}}(n))$ time.*

Proof. Let τ_{opt} be a translation, such that $\min_{\tau} \delta_F(\tau(P), Q) = \delta_F(\tau_{opt}(P), Q)$. Then

$$\|\mathbf{r}(\tau_{opt}(P)) - \mathbf{r}(Q)\| \leq c \cdot \delta_F(\tau_{opt}(P), Q).$$

Let $\tau_{diff} := \mathbf{r}(\tau_{opt}(P)) - \mathbf{r}(Q) \in \mathcal{T}_2$; then

$$\tau_{approx} := \tau_{diff} \circ \tau_{opt}$$

maps $\mathbf{r}(P)$ onto $\mathbf{r}(Q)$ and

$$\delta_F(\tau_{\text{approx}}(P), Q) \leq \delta_F(\tau_{\text{opt}}(P), Q) + \|\tau_{\text{diff}}\| \leq (c + 1) \cdot \delta_F(\tau_{\text{opt}}(P), Q).$$

The proof of the claimed time bound is obvious. \square

Note that with an idea from [9] it is possible to reduce the approximation constant for reference point based matching to $(1 + \epsilon)$ for any $\epsilon > 0$; the idea places a sufficiently small grid of size $\mathcal{O}(1/\epsilon^2)$ around the reference point of Q and checks each grid point as a potential image point for the reference point of P . The runtime increases by a factor proportional to the grid size.

Acknowledgements

We would like to thank Günter Rote and an anonymous referee for helpful comments on an earlier version of this paper.

References

1. O. Aichholzer, H. Alt, and G. Rote. Matching shapes with a reference point. *Internat. J. Comput. Geom. Appl.*, 7:349–363, 1997.
2. H. Alt, B. Behrends, and J. Blömer. Approximate matching of polygonal shapes. *Ann. Math. Artif. Intell.*, 13:251–266, 1995.
3. H. Alt, U. Fuchs, G. Rote, and G. Weber. Matching convex shapes with respect to the symmetric difference. *Algorithmica*, 21:89–103, 1998.
4. H. Alt and M. Godau. Computing the Fréchet distance between two polygonal curves. *Internat. J. Comput. Geom. Appl.*, 5:75–91, 1995.
5. H. Alt and L. J. Guibas. Discrete geometric shapes: Matching, interpolation, and approximation. In J.-R. Sack and J. Urrutia, editors, *Handbook of Computational Geometry*, pages 121–153. Elsevier Science Publishers B.V. North-Holland, Amsterdam, 2000.
6. R. Cole. Slowing down sorting networks to obtain faster sorting algorithms. *Journal of the ACM*, 34(1):200–208, 1987.
7. A. Efrat, P. Indyk, and S. Venkatasubramanian. Pattern matching for sets of segments, September 2000. Manuscript, accepted to the 12th Symposium on Discrete Algorithms, 2001.
8. N. Megiddo. Applying parallel computation algorithms in the design of serial algorithms. *J. ACM*, 30(4):852–865, 1983.
9. S. Schirra. Über die Bitkomplexität der ϵ -Kongruenz. Master's thesis, Fachbereich Informatik, Universität des Saarlandes, 1988.
10. S. Venkatasubramanian. *Geometric Shape Matching and Drug Design*. PhD thesis, Department of Computer Science, Stanford University, August 1999.