Approximately Matching Polygonal Curves under Translation, Rotation and Scaling with Respect to the Fréchet-Distance

Michael Clausen* and Axel Mosig
Institut für Informatik III, Universität Bonn

January 3, 2003

Abstract

Let $P : [0, m] \to \mathbb{R}^2$ and $Q : [0, n] \to \mathbb{R}^2$ be polygonal curves in the plane, G a subgroup of the affine group $\mathrm{AGL}(2, \mathbb{R})$, and $\varepsilon \geq 0$. By definition, a transformation $g \in G$ yields a (G, ε)-Fréchet-match of P and Q if the Fréchet-distance of P and the transformed version gQ of Q is at most ε. In this paper we design a c-approximation algorithm, $c > 2$, that constructs such $(G, c\varepsilon)$-Fréchet-matches for both the group G_r of rigid motions and the group G_s generated by translations and uniform scalings. We associate to P, Q and ε a certain acyclic digraph $\mathcal{M}_{m,n}$, see Fig. 1, whose edges are either weighted by closed intervals in $\mathbb{R}_{>}0$ ($G = G_s$) or by circular arcs ($G = G_r$). All maximal paths in $\mathcal{M}_{m,n}$ correspond to discrete reparametrization pairs; such a pair yields a c-approximate solution if the intervals assigned to the edges along the path have a non-empty intersection. To decide whether such a path exists, we use a dynamic programming approach, whose time complexity is $O(m^2n^2)$. There is related work dealing with smaller subgroups of $\mathrm{AGL}(2, \mathbb{R})$: Alt and Godau [1] investigated the case $G = \{1\}$, whereas both Alt, Knauer and Wenk [2] and Efrat, Indyk and Venkatasubramanian [4] studied the case $G = T_2$, which denotes the group of translations.

1 Fréchet-Matches

A *polygonal curve of length* $m \in \mathbb{N}$ in $\mathbb{R}^2$ is defined as a continuous mapping $P : [0, m] \to \mathbb{R}^2$ with the property that for all $i \in [0 : m - 1] := \{0, 1, \ldots, m - 1\}$ the curve $P|_{[i,i+1]}$ is affine, i.e., $P(i + \lambda) = (1 - \lambda)P(i) + \lambda P(i + 1)$ for $\lambda \in [0, 1]$. A polygonal curve P is completely described by the sequence of its vertices $\langle p_0, \ldots, p_m \rangle$, where $p_i := P(i)$. For real numbers $x < y$ and $x' < y'$, let $\mathrm{Mon}([x, y], [x', y'])$ denote the set of all continuous, weakly increasing and surjective functions $\varphi : [x, y] \to [x', y']$; note that the surjectivity implies $\varphi(x) = x'$ and $\varphi(y) = y'$ for all $\varphi \in \mathrm{Mon}([x, y], [x', y'])$. Let P and Q be polygonal curves of lengths m and n, respectively. The *Fréchet-distance* $d_F(P, Q)$ of P and Q is defined as $d_F(P, Q) := \inf_{\alpha,\beta} \max_{t \in [0,1]} d(v(\alpha(t)), w(\beta(t)))$, where the infimum is taken over all $\alpha \in \mathrm{Mon}([0, 1], [0, m])$ and $\beta \in \mathrm{Mon}([0, 1], [0, n])$. Any subgroup G of $\mathrm{AGL}(2, \mathbb{R})$ acts on $\mathbb{R}^2$ as well as on the set of all polygonal curves P by $(gP)(t) := gP(t)$ for $g \in G$. Moreover, if $P = \langle p_0, \ldots, p_m \rangle$, then $gP = \langle gp_0, \ldots, gp_m \rangle$.

We are now ready to describe a typical question in pattern matching: given a subgroup G of $\mathrm{AGL}(2, \mathbb{R})$, two polygonal curves P and Q, and $\varepsilon \geq 0$, is there a $g \in G$ such that $d_F(P, gQ) \leq \varepsilon$? Motivated by this question, we define the set of all (G, ε)*-Fréchet-matches of* P *and* Q as

$$\mathcal{F}_G^{\varepsilon}(P, Q) := \{g \in G \mid d_F(P, gQ) \leq \varepsilon\}. \tag{1}$$

Given two polygonal curves P and Q and $\varepsilon \geq 0$, a decision algorithm for this pattern matching task outputs 1 if $\mathcal{F}_G^{\varepsilon}(P, Q) \neq \emptyset$ and 0 otherwise. Letting $c > 1$, a c*-approximation algorithm* also outputs 1 if $\mathcal{F}_G^{\varepsilon}(P, Q) \neq \emptyset$. However, the output is guaranteed to be 0 only if $\mathcal{F}_G^{c\varepsilon}(P, Q) = \emptyset$. In case $\mathcal{F}_G^{c\varepsilon}(P, Q) \setminus \mathcal{F}_G^{\varepsilon}(P, Q) \neq \emptyset$, the algorithm may answer either 0 or 1. The algorithm proposed in this paper yields c-approximate solutions for arbitrary $c > 2$. The algorithm will also be able to compute specific elements $g \in \mathcal{F}_G^{c\varepsilon}(P, Q)$ in case of output 1.

*This work was supported in part by Deutsche Forschungsgemeinschaft under grant CL 64/3

2 Approximating Fréchet-Matches by Transporter Sets

In this section $P = \langle p_0, \dots, p_m \rangle$ and $Q = \langle q_0, \dots, q_n \rangle$ will always denote polygonal curves of lengths m and n, respectively. P is called *reducible* if and only if, for some i, the vertex p_i is contained in the line segment $[p_{i-1}, p_{i+1}]$. Eliminating p_i from the sequence yields another curve P' with $d_F(P, P') = 0$. This elimination process finally yields an irreducible curve. In general, two polygonal curves P and P' are called *equivalent* if and only if their Fréchet-distance is zero. The Fréchet-distance defines a metric on the equivalence classes of polygonal curves. Obviously, in each equivalence class there is a unique irreducible curve. All other members of this class can be viewed as *oversamplings* of this irreducible version. In what follows, oversampling will play a crucial role. Let $\delta > 0$. A polygonal curve P is said to be *δ-sampled* if and only if $d(p_{i-1}, p_i) \le 2\delta$, for all $i \in [1 : m]$. Given a polygonal curve P, an equivalent, δ-sampled curve P' can be constructed in an obvious way.

The notion of δ-sampled curves is a first step towards discretizing the reparametrizations α and β. We will replace $(\alpha, \beta) \in \mathrm{Mon}([0,1],[0,m]) \times \mathrm{Mon}([0,1],[0,n])$ by discrete reparametrizations $(\kappa, \lambda) \in \mathcal{I}_{m,n}$, where

$$\mathcal{I}_{m,n} := \{(\kappa,\lambda) \mid \begin{array}{l} \kappa : [0 : m+n] \to [0 : m] \text{ and } \lambda : [0 : m+n] \to [0 : m] \\ \text{are both weakly increasing and surjective}\}. \end{array} \tag{2}$$

From the facts that the index sequences κ and λ are surjective and weakly increasing, we may conclude that $\{\kappa_{s+1} - \kappa_s, \lambda_{s+1} - \lambda_s\} \subseteq \{0, 1\}$ for all $s \in [0 : n + m - 1]$. To approximate sets of Fréchet-matches we use certain transporter subsets of the group G. If d denotes the Euclidean distance in $\mathbb{R}^2$ and if P and Q have equal length, then $\tau^{G,\varepsilon}_{P,Q} := \{g \in G \mid \max_i d(p_i, gq_i) \le \varepsilon\}$ is called the *(G, ε)-transporter of Q to P*. Similarly, $\tau^{G,\varepsilon}_{p,q} := \{g \in G \mid d(p, gq) \le \varepsilon\}$ denotes the (G, ε)-transporter of $q \in \mathbb{R}^2$ to $p \in \mathbb{R}^2$. Obviously, $\tau^{G,\varepsilon}_{P,Q} = \cap_i \tau^{G,\varepsilon}_{p_i,q_i}$.

Theorem 2.1 *Let P and Q be δ-sampled polygonal curves of lengths m and n, respectively, with $\mathcal{F}^{\varepsilon}_{G}(P, Q) \neq \emptyset$. Then there exists a pair $(\kappa, \lambda) \in \mathcal{I}_{m,n}$ such that $\emptyset \neq \tau^{G,\varepsilon+\delta}_{P\circ\kappa,Q\circ\lambda} \subseteq \mathcal{F}^{\varepsilon+\delta}_{G}(P, Q)$.*

Unfortunately, our proof of the theorem is not completely constructive, since we require some $g \in \mathcal{F}^{\varepsilon}_{G}(P, Q)$ for computing (κ, λ). When matching two curves, however, such a g is what we are looking for. Thus, our algorithm for deciding whether $\mathcal{F}^{\varepsilon}_{G}(P, Q)$ is non-empty has to find suitable integer sequences κ and λ in a different way. A naive method that enumerates all surjective and weakly increasing candidate sequences and checks if $\tau^{G,\varepsilon+\delta}_{P\circ\kappa,Q\circ\lambda}$ is non-empty for each candidate sequence only yields an exponential-time algorithm.

3 Intersecting Projected Transporter Sets

Regarding the last theorem, both $P \circ \kappa$ and $Q \circ \lambda$ are in $(\mathbb{R}^2)^{m+n+1}$. Thus $\tau^{G,\varepsilon+\delta}_{P\circ\kappa,Q\circ\lambda}$ is the intersection of $m + n + 1$ individual transporters of the form $\tau^{G,\varepsilon+\delta}_{P(\kappa_s),Q(\lambda_s)}$. Unfortunately, these individual transporters have a rather complicated structure. In order to simplify the intersection problem we use the fact that our groups are semidirect products: $G = T_2 \rtimes H$ with $H = H_r := \mathrm{SO}(2)$ for $G = G_r$ and $H = H_s := \{\sigma E_2 \mid \sigma > 0\}$ for $G = G_s$, where E_2 denotes the 2×2 unit matrix. Thus the projection η of G onto H with kernel T_2, i.e., $\eta(th) := h$, for $t \in T_2$ and $h \in H$, is well-defined. Instead of $\tau^{G,\varepsilon}_{P,Q}$ we work with its η-image:

$$\eta^{G,\varepsilon}_{P,Q} := \eta[\tau^{G,\varepsilon}_{P,Q}] = \{h \in H \mid \exists t \in T_2 : th \in \tau^{G,\varepsilon}_{P,Q}\}.$$

Note that an analogous statement to $\tau^{G,\varepsilon}_{P,Q} = \cap_i \tau^{G,\varepsilon}_{p_i,q_i}$ does not hold for the η-images. Furthermore, for $p, q \in \mathbb{R}^2$ we always have $\eta^{G,\varepsilon}_{p,q} = H$, thus to obtain non-trivial transporters we use projected transporter sets of the form $\eta^{G,\varepsilon}_{\langle p_0,p_1\rangle,\langle q_0,q_1\rangle}$ as building blocks. To simplify notation, we let $k := m + n + 1$ and write P and Q instead of $P \circ \kappa$ and $Q \circ \lambda$.

Theorem 3.1 *Let $G \in \{G_r, G_s\}$, $k \in \mathbb{N}$, and $\varepsilon \geq 0$. For $P, Q \in (\mathbb{R}^2)^k$ and every $j \in [1:k-1]$ define*

$$H_j := \bigcap_{i \in [1:j]} \eta^{G,\varepsilon}_{\langle p_{i-1}, p_i \rangle, \langle q_{i-1}, q_i \rangle} \cap \eta^{G,\varepsilon}_{\langle p_0, p_i \rangle, \langle q_0, q_i \rangle}.$$

Then $\tau^{G,\varepsilon}_{P,Q} \neq \emptyset$ implies $H_{k-1} \neq \emptyset$, whereas $\tau^{G,2\varepsilon}_{P,Q} = \emptyset$ forces $H_{k-1} = \emptyset$.

Expressed in simple terms, the preceding result states that deciding whether the intersection H_{k-1} is non-empty yields an approximate solution to the decision problem asking if the intersection $\tau^{G,\varepsilon}_{P,Q} = \cap_i \tau^{G,\varepsilon}_{p_i,q_i}$ is non-empty. Next we take a closer look at the projected (G, ε)-transporters $\eta^{G,\varepsilon}_{\langle p_0, p_1 \rangle, \langle q_0, q_1 \rangle}$ for $G \in \{G_r, G_s\}$.

Theorem 3.2 *Each $\eta^{G_r,\varepsilon}_{\langle p_0, p_1 \rangle, \langle q_0, q_1 \rangle}$ can be viewed as a circular arc on the unit circle S^1, whereas each $\eta^{G_s,\varepsilon}_{\langle p_0, p_1 \rangle, \langle q_0, q_1 \rangle}$ can be regarded as a closed interval on $\mathbb{R}_{>0}$.*

Sketch of Proof. One shows that it suffices to compute $\eta^{G,\varepsilon}_{\langle p_0, p_1 \rangle, \langle q_0, q_1 \rangle}$ for line segments $\langle p_0, p_1 \rangle$ and $\langle q_0, q_1 \rangle$ centered at the origin. The construction of the circular arc and the interval is illustrated in the figure below. □

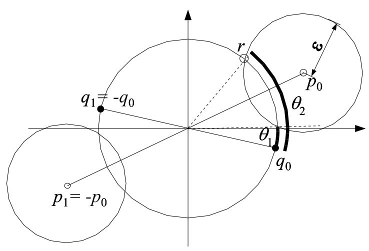

$G = G_r$. Obviously, $\eta^{G,\varepsilon}_{P,Q} = [\theta_1, \theta_2]$, where $\theta_2 - \theta_1 = 2\sphericalangle(r, p_0)$ and $\theta_2 = \sphericalangle(r, q_0)$. Hence, we get $\theta_1 = \theta_2 - (\theta_2 - \theta_1) = \sphericalangle(r, q_0) - 2\sphericalangle(r, p_0)$.

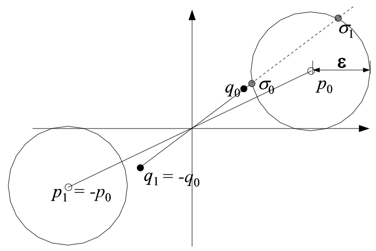

$G = G_s$. The intersection of the ray Hq_0 with the disc $U_\varepsilon(p_0)$ is a line segment, and thus $\eta^{G,\varepsilon}_{P,Q}$ can be identified with the closed interval $[\|\sigma_0\|_2/\|q_0\|_2, \|\sigma_1\|_2/\|q_0\|_2]$.

In case $G = G_s$, we can easily decide whether the intersection of finitely many projected transporters is non-empty as $\cap_i[x_i, y_i] \neq \emptyset$ iff $\max_i x_i \leq \min_j y_j$.
For $G = G_r$, the projected (G, ε)-transporters are circular arcs, hence intervals on the unit circle. Such intervals differ in some respects from real intervals. For example, the intersection of two intervals on S^1 may consist of up to two disjoint intervals. An easy way to avoid the difficulties in conjunction with circular-arc intersections is to unroll S^1 — and intervals on S^1 — to the interval $[0, 2\pi]$. Unrolling an interval that covers the angle 0 requires the interval to be split into two intervals on $[0, 2\pi]$. Thus, unrolling $\eta^{G,\varepsilon}_{\langle p_0, p_i \rangle, \langle q_0, q_i \rangle} \cap \eta^{G,\varepsilon}_{\langle p_{i-1}, p_i \rangle, \langle q_{i-1}, q_i \rangle}$ yields up to three (disjoint) intervals in $[0, 2\pi]$.

4 An Efficient Approximate Matching Algorithm

We are now prepared to design an efficient algorithm for approximately matching two polygonal curves with respect to the Fréchet-distance under a transformation group $G \in \{G_r, G_s\}$. To this end, we introduce for δ-sampled polygonal curves P and Q of lengths m and n, respectively, the acyclic digraph $\mathcal{M}_{m,n} := (V_{m,n}, E_{m,n})$ together with a function that assigns a real interval ($G = G_s$) or up to two circular arcs ($G = G_r$) to each edge of the graph. (Efrat et al. [4] also use a graph for finding paths in free-space. However, our graph differs substantially from their construction.) The digraph $\mathcal{M}_{m,n}$, defined by

$$V_{m,n} := [0:m] \times [0:n] \quad \text{and} \quad E_{m,n} := \{((a,b),(c,d)) \in V^2_{m,n} \mid \{1\} \subseteq \{c-a, d-b\} \subseteq \{0,1\}\},$$

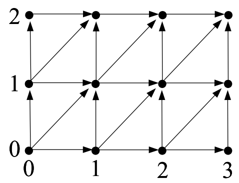

Figure 1: The digraph $\mathcal{M}_{3,2}$.

only depends on m and n, see Fig. 1, whereas the weight of the edge $e = ((a,b),(c,d)) \in E_{m,n}$ depends on $P, Q, \delta, \varepsilon$, and G:

$$I_{P,Q,\delta,\varepsilon,G}((a,b),(c,d)) := \eta^{G,\varepsilon+\delta}_{\langle p_a,p_c\rangle,\langle q_b,q_d\rangle} \cap \eta^{G,\varepsilon+\delta}_{\langle p_0,p_c\rangle,\langle q_0,q_d\rangle}.$$

Obviously, every pair $(\kappa,\lambda) \in \mathcal{I}_{m,n}$ defines (after eliminating loops) a path in $\mathcal{M}_{m,n}$ with source $(0,0)$ and sink (m,n). Conversely, every such path can be transformed into an element of $\mathcal{I}_{m,n}$ as follows: turn every vertex along the path into a pair (κ_s,λ_s) and repeat the target node of every diagonal edge, i.e., replace $(a,b) \to (a+1,b+1)$ by the subsequence $(a,b),(a+1,b+1),(a+1,b+1)$.

Now we take a closer look at the case $G = G_s$. (The case $G = G_r$ is similar, but a little bit more technical.) Here, each edge e is assigned the empty set or a closed real interval $[\ell_e, r_e]$, see Theorem 3.2. We have to find a path from $(0,0)$ to (m,n) in $\mathcal{M}_{m,n}$ such that the intersection of the involved intervals is non-empty. To decide whether such a path exists, we use the facts that $\cap_{i=1}^N [x_i, y_i] = [\max_i x_i, \min_i y_i]$ and that $\max_i x_i \in \{x_1, \ldots, x_N\}$. In particular, as $\mathcal{M}_{m,n}$ has $3mn + m + n$ edges, we have at most $3mn + m + n$ different left borders to consider. For each possible left border ℓ we define a new $0-1$ weight on the edges: edge e has weight 1 iff ℓ is contained in $[\ell_e, r_e]$. By dynamic programming one can test in time $O(mn)$ whether there is a path from $(0,0)$ to (m,n) involving only edges with weight 1.

Theorem 4.1 *For $G = G_s$, there is an algorithm that on input $P, Q, m, n, \delta, \varepsilon$ (with the above meaning) computes an element $g \in \mathcal{F}_G^{2(\varepsilon+\delta)}(P,Q)$ if $\mathcal{F}_G^{\varepsilon}(P,Q) \neq \emptyset$ and computes the output 0 if $\mathcal{F}_G^{2(\varepsilon+\delta)}(P,Q) = \emptyset$. Its running time is $O(m^2n^2)$.*

Thus there is a c-approximation algorithm for determining whether the set of (G_s,ε)-Fréchet-matches is non-empty, for $c = 2(1+\delta/\varepsilon)$. The same result holds for $G = G_r$.

5 Final Remarks and Future Work

The systematic use of group transporter sets is the basis of a new technique that generalizes the concept of inverted files from full-text retrieval. It has been successfully applied to content-based multimedia retrieval, see [3]. In the present work this concept has been extended to (G,ε)-transporters. We are currently investigating variants of the described algorithm, including matching curves partially as well as matching under other subgroups of $\mathrm{AGL}(2,\mathbb{R})$, in particular the group of similarity transformations generated by translations, rotations and uniform scalings.

References

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

[2] H. Alt, C. Knauer, and C. Wenk. Matching polygonal curves with respect to the Fréchet distance. In *Proc. 18th Int. Symp. on Theoretical Aspects of Computer Science*, pages 63–74, 2001.

[3] M. Clausen and F. Kurth. A Unified Approach to Content-Based and Fault Tolerant Music Recognition. *IEEE Transactions on Multimedia*, 2003. To appear.

[4] A. Efrat, P. Indyk, and S. Venkatasubramanian. Pattern matching with sets of segments. In *Proc. 12th ACM Symp. on Discrete Algorithms*, January 2001.