Approximately Matching Polygonal Curves with Respect to the Fréchet Distance

Axel Mosig

Lehrstuhl für Bioinformatik, Universität Leipzig, D-04103 Leipzig

Michael Clausen

Institut für Informatik III, Universität Bonn, D-53117 Bonn

Abstract

In this paper we present approximate algorithms for matching two polygonal curves with respect to the Fréchet distance. We define a discrete version of the Fréchet distance as a distance measure between polygonal curves and show that this discrete version is bounded by the continuous version of the Fréchet distance.

For the task of matching with respect to the discrete Fréchet distance, we develop an algorithm that is based on intersecting certain subsets of the transformation group under consideration. Our algorithm for matching two point sequences of lengths m and n under the group of rigid motions has a time complexity of $O(m^2n^2)$ for matching under the discrete Fréchet distance and can be modified for matching subcurves, closed curves and finding longest common subcurves. Group theoretical considerations allow us to eliminate translation components of affine transformations and to consider matching under arbitrary linear algebraic groups.

---

1 Introduction

A typical scenario in geometric pattern matching is as follows: we are given two geometric objects P and Q as well as a group G of admissible transformations and a distance measure d for computing the resemblance of P and Q. The matching task, stated as a decision problem, is to determine whether there exists a transformation $g \in G$ that brings an object Q close to another object P so that $d(P, gQ) \leq \varepsilon$; here, gQ denotes the object Q transformed by g. Sometimes, one is also interested in the *optimization problem* of finding a transformation g that minimizes $d(P, gQ)$. Typical applications range from computer vision and image retrieval to computer aided drug design. For a survey on geometric pattern matching, we refer to [1].

Preprint submitted to Elsevier Science *30 April 2004*

In our case, the objects under consideration are polygonal curves in some real vector space V; the transformation groups studied are affine transformations, in particular translations, rotations and scalings, while the distance measure considered is (a discrete version of) the Fréchet distance.

Many aspects of the Fréchet distance as a distance measure between polygonal curves have recently been examined in the field of computational geometry. Introduced in [2], algorithms for computing the Fréchet distance were developed. Several authors address the problem of *matching* curves with respect to the Fréchet distance: Efrat et al. [3] as well as Alt, Knauer and Wenk [4] designed polynomial time algorithms for matching under the group of arbitrary two-dimensional translations. In [5], the idea from [4] is generalized to larger transformation groups using techniques from real algebraic geometry.

We define a discrete version of the Fréchet distance and show that the continuous Fréchet distance is bounded by this discrete version. As a consequence, it suffices to design algorithms for matching with respect to the discrete version of the Fréchet distance in order to obtain algorithms for matching approximately with respect to the continuous Fréchet distance. Letting m and n denote the number of vertices of the two polygonal curves to be matched, our algorithm's running time for matching approximately under rigid motions is bounded by $O(m^2n^2)$ under the discrete version of the Fréchet distance. Improving the quality of approximation for matching under the continuous version, however, results in a running time that depends on the Euclidean length of the two polygonal curves to be matched. This compares to a running time of $O(n^{11})$ (where $m \leq n$) for the algorithm proposed in [5] for solving the matching problem under the continuous version exactly.

Our algorithms for matching under rigid motions rely on elementary algorithmic and geometric computations and can hence be implemented easily. Note that the algorithm from [5] for matching under rigid motions in the plane as well as the algorithms we propose for some larger transformation groups rely on techniques from real algebraic geometry. Generally, algorithms relying on such techniques can be considered as being difficult to implement.

All matching problems we consider are based on intersecting certain subsets of the underlying transformation groups, which is motivated by the technique described in [6,7]. This leads to some group theoretical considerations, which are the subject of Section 3.

2 The Discrete Fréchet Distance

We first introduce some notation. Let $[x, y]$ denote the compact real interval between x and y; moreover, for integers a and b, let $[a : b]$ denote the set $\{a, a+1, \ldots, b\}$ of all integers between a and b. Given two sets X and Y, Y^X denotes the set of all mappings from X to Y; for $f \in Y^X$ and $I \subseteq X$, we denote $f[I] := \{f(x) \mid x \in I\}$. Analogously, we denote $f^{-1}[J] := \{x \in X \mid f(x) \in J\}$ for $J \subseteq Y$. Since $x \in X^{[a:b]}$ is completely described by a sequence of $b - a + 1$ values in X, we also write $x = \langle x_a, \ldots, x_b \rangle \in X^{[a:b]}$. Let $V = \mathbb{R}^k$ denote a Euclidean vector space with the Euclidean norm $\|.\| := \|.\|_2$. A *curve* in V is a continuous mapping $f \in V^{[a,b]}$ with $a, b \in \mathbb{R}$; a *polygonal curve of length* $m \in \mathbb{N}$, is a mapping $P \in V^{[0,m]}$, such that for all $i \in [0 : m-1]$, $P|_{[i,i+1]}$ is *affine*, i.e., $P(i + \lambda) = (1 - \lambda)P(i) + \lambda P(i + 1)$ for all $\lambda \in [0, 1]$.

For $f \in V^I$, let $\|f\|_\infty := \sup_{t \in I} \|f(t)\|$. The *Fréchet distance* between $P \in V^{[0,m]}$ and $Q \in V^{[0,n]}$ (for some $m, n > 0$) is defined as $d_{\mathrm{F}}(P, Q) = \min_{(\alpha,\beta)} \|P \circ \alpha - Q \circ \beta\|_\infty$, where (α, β) ranges over all continuous, weakly increasing and surjective mappings $\alpha \in [0, m]^{[0,1]}$ and $\beta \in [0, n]^{[0,1]}$. In the sequel, we denote the set of all continuous, weakly increasing and surjective mappings from a set X to another set Y by $\mathrm{Mon}(X, Y)$ and write $\mathrm{Mon}_{m,n} := \mathrm{Mon}([0, 1], [0, m]) \times \mathrm{Mon}([0, 1], [0, n])$. In case $m \in \mathbb{N}$ and $P \in V^{[0,m]}$ denotes a polygonal curve, we can identify P with the mapping $[0 : m] \ni i \mapsto P(i) =: p_i$, and hence we also write $P \in V^{[0:m]}$.

2.1 Definition and Basic Properties

Given two polygonal curves $P \in V^{[0:m]}$ and $Q \in V^{[0:n]}$, we define the *discrete Fréchet distance* as $\mathbf{d}_{\mathrm{F}}(P, Q) = \min_{(\kappa,\lambda)} \|P \circ \kappa - Q \circ \lambda\|_\infty$, where the pairs (κ, λ) range over the set $\mathbf{Mon}_{m,n} := \mathrm{Mon}([1 : m+n], [0 : m]) \times \mathrm{Mon}([1 : m+n], [0 : n])$. Correspondingly, one can define $\mathbf{d}_{\mathrm{F}}$ for polygonal curves $P \in V^{[a:b]}$ and $Q \in V^{[c:d]}$ for integers a, b, c, d by adapting domain and range of the reparametrizations. Note that the integer interval $[1 : m+n]$ substitutes the real interval $[0, 1]$ as the common domain for the two reparametrizations. The discrete Fréchet distance is similar to the *dynamic time warping distance* that is defined as $\mathbf{d}_{\mathrm{W}}(P, Q) := \min_{\kappa,\lambda} \|P \circ \kappa - Q \circ \lambda\|_2$, where (κ, λ) range over $\cup_{K \in [\max(m,n), m+n]} \mathrm{Mon}([0 : K], [0 : m]) \times \mathrm{Mon}([0 : K], [0 : n])$, respectively, and $\|f\|_2 := (\| \sum_{i \in I} (f(t))^2 \|)^{1/2}$ for $f \in V^I$ with $|I| < \infty$.

Dynamic time warping has been considered in the context of speech signal processing and time series databases [8], in both cases for $V = \mathbb{R}$. More recently, dynamic time warping has been used for matching polygonal curves in the plane under the group of translations [9]. The results presented in the sequel can be seen as a bridge between these works and the results obtained in the area of computational geometry.

We can compute the discrete Fréchet distance between $P \in V^{[0:m]}$ and $Q \in V^{[0:n]}$ in a straightforward way. Defining $d_{i,j} := \mathbf{d}_{\mathrm{F}}(P|_{[0:i]}, Q|_{[0:j]})$, we compute $d_{m,n}$ in $O(mn)$ time using dynamic programming:

$d_{0,0} := \|p_0 - q_0\|$;
for $j := 1$ **to** n **do** $d_{0,j} := \max\{d_{0,j-1}, \|p_0 - q_j\|\}$;
for $i := 1$ **to** m
  $d_{i,0} := \max\{d_{i-1,0}, \|p_i - q_0\|\}$
  for $j := 1$ **to** n
    $d_{i,j} := \max(\min\{d_{i,j-1}, d_{i-1,j}, d_{i-1,j-1}\}, \|p_i - q_j\|)$;
  end
end
return $d_{m,n}$.

2.2 Bounding $\mathbf{d}_{\mathrm{F}}$ *by* d_{F}

A major property of the discrete Fréchet distance is that it is bounded by the continuous Fréchet distance. Before we can state this bounding property, we require the notion of *sampling* and *oversampling* polygonal curves. We say that $P \in V^{[0:m]}$ is *δ-sampled* if $\|p_i - p_{i-1}\| \leq \delta$ for all $i \in [1 : m]$.

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. We say that a polygonal curve P is *reducible* if and only if, for some i, the vertex p_i is contained in the line segment $\langle p_{i-1}, p_{i+1}\rangle$. Eliminating p_i from the sequence yields another curve P' with $d_{\mathrm{F}}(P, P') = 0$. This elimination process finally yields a curve that cannot be reduced any further. Obviously, in each equivalence class there is one unique such irreducible curve. All other members of this class can be viewed as *oversamplings* of this irreducible version. Also, we can produce *oversampled* versions of a polygonal curve by inserting additional vertices that leave d_{F} unchanged. As can be seen easily, oversampling decreases $\mathbf{d}_{\mathrm{F}}$, while d_{F} is left unchanged. This fact clearly suggests that $\mathbf{d}_{\mathrm{F}}$ can be bounded by d_{F}. However, there are no obvious *tight* bounds. In this section, we provide distance bounds between d_{F} and $\mathbf{d}_{\mathrm{F}}$ that are tight.

Theorem 1 *Let $P \in V^{[0:m]}$ and $Q \in V^{[0:n]}$ be δ-sampled polygonal curves. Then,*

$$d_{\mathrm{F}}(P, Q) \quad \leq \quad \mathbf{d}_{\mathrm{F}}(P, Q) \quad \leq \quad d_{\mathrm{F}}(P, Q) + \delta/2.$$

PROOF. We start with the proof of the first inequality. Let $(\kappa, \lambda) \in \mathbf{Mon}_{m,n}$ be optimal in the sense that $\mathbf{d}_{\mathrm{F}}(P, Q) = \|P \circ \kappa - Q \circ \lambda\|_\infty$. By affine interpolation, one obtains $(\alpha, \beta) \in \mathrm{Mon}_{m,n}$ with $\alpha(\frac{i}{m+n}) = \kappa_i$ and $\beta(\frac{i}{m+n}) = \lambda_i$. Then,

$$
\begin{aligned}
d_{\mathrm{F}}(P, Q) &= \min_{\alpha', \beta'} \max_{t \in [0,1]} \|P(\alpha'(t)) - Q(\beta'(t))\| \\
&\leq \max_{t \in [0,1]} \|P(\alpha(t)) - Q(\beta(t))\| \\
&= \max_{s \in [1:m+n]} \|P(\kappa(s)) - Q(\lambda(s))\| = \mathbf{d}_{\mathrm{F}}(P, Q),
\end{aligned}
$$

where the last but one equality follows from the fact that for line segments $L = \langle L_0, L_1 \rangle$ and $L' = \langle L_0', L_1' \rangle$, $d_{\mathrm{F}}(L, L') = \max\{\|L_0 - L_0'\|, \|L_1 - L_1'\|\}$. This proves the first inequality.

For the proof of the second inequality, let $(\alpha, \beta) \in \mathrm{Mon}_{m,n}$ be optimal, i.e., $d_{\mathrm{F}}(P, Q) = \|P \circ \alpha - Q \circ \beta\|_\infty$. Define

$$
\mu_j := \min \beta^{-1}[j] \quad \text{and} \quad \nu_i := \min \alpha^{-1}[i] \tag{1}
$$

for $i \in [0 : m - 1]$ and $j \in [0 : n - 1]$ as well as $\mu_n := 1$ and $\nu_m := 1$.

Since α and β are weakly increasing, we have $\mu_{j-1} \leq \mu_j$ and $\nu_{i-1} \leq \nu_i$ for all i and j.

Now, let $\theta_0 \leq \theta_1 \leq \cdots \leq \theta_{m+n+1}$ denote the ordered sequence of all $m+n+2$ values μ_i and ν_j, including multiplicities. This allows us to define

$$
\kappa_s := \lfloor \alpha(\theta_s) \rceil \quad \text{and} \quad \lambda_s := \lfloor \beta(\theta_s) \rceil, \tag{2}
$$

where $\lfloor x \rceil$ assigns the nearest integer to $x \in \mathbb{R}$, i.e., $\lfloor x \rceil := \lfloor x \rfloor$ if $x - \lfloor x \rfloor < \frac{1}{2}$ and $\lfloor x \rceil := \lceil x \rceil$ if $\lceil x \rceil - x \leq \frac{1}{2}$.

The sequences $\kappa := \langle \kappa_1, \ldots, \kappa_{m+n} \rangle \in [0 : m]^{[1:m+n]}$ and $\lambda := \langle \lambda_1, \ldots, \lambda_{m+n} \rangle \in [0 : n]^{[1:m+n]}$ are weakly increasing, since the sequence $\theta := \langle \theta_0, \ldots, \theta_{m+n+1} \rangle$ is weakly increasing and the mappings α, β as well as $\lfloor . \rceil$ are order preserving.

The surjectivity of λ follows from the fact that the sequence $\langle \mu_0, \ldots, \mu_n \rangle$ is contained in θ as a subsequence, and $\lfloor \beta(\mu_j) \rceil = \lfloor \beta(\min \beta^{-1}[j]) \rceil = \lfloor j \rceil = j$; note that since $\theta_0 = \theta_1 = 0$ and $\theta_{m+n} = \theta_{m+n+1} = 1$, we can omit λ_0 and λ_{m+n+1} without losing surjectivity. The surjectivity of κ follows analogously.

So far, we know that $(\kappa, \lambda) \in \mathbf{Mon}_{m,n}$. It remains to be shown that $\|p_{\kappa_s} - q_{\lambda_s}\| \leq d_{\mathrm{F}}(P, Q) + \delta/2$ for all $s \in [1 : m + n]$. To this end, observe that since P and Q are δ-sampled, we have $\|P(x) - P(\lfloor x \rceil)\| \leq \delta/2$. Additionally, we have $\|P(\alpha(\theta_s)) - Q(\beta(\theta_s))\| \leq d_{\mathrm{F}}(P,Q)$ for all $s \in [1 : m+n]$. For each $s \in [1 : m+n]$, either $\theta_s = \nu_i$ or $\theta_s = \mu_j$ for some $i \in [0 : m]$ or $j \in [0 : n]$. If $\theta_s = \mu_j$, we have $\lambda_s = \lfloor \beta(\mu_j) \rceil = \lfloor \beta(\min \beta^{-1}[j]) \rceil = \lfloor j \rceil = j = \beta(\mu_j)$, in other words, $\|Q(\beta(\mu_j)) - Q(\lfloor \beta(\mu_j) \rceil)\| = 0$. If $\theta_s = \nu_i$ for some i, we similarly obtain $\|P(\alpha(\nu_i) - P(\lfloor \alpha(\nu_i) \rceil)\| = 0$.

Now, using the triangle inequality, we get:

$$\begin{aligned}\|P(\kappa_s) - Q(\lambda_s)\| \leq & \|P(\kappa_s) - P(\alpha(\theta_s))\| + \|P(\alpha(\theta_s)) - Q(\beta(\theta_s))\| \\ & + \|Q(\beta(\theta_s)) - Q(\lambda_s)\| \\ \leq & \, \delta/2 + d_{\mathrm{F}}(P,Q).\end{aligned}$$

Altogether, we have constructed $(\kappa, \lambda) \in \mathbf{Mon}_{m,n}$ such that

$$\begin{aligned}d_{\mathrm{F}}(P,Q) + \delta/2 \geq & \max_{s \in [1:m+n]} \|P(\kappa_s) - Q(\lambda_s)\| \\ \geq & \min_{\kappa', \lambda'} \max_{s \in [1:m+n]} \|P(\kappa'_s) - Q(\lambda'_s)\| \\ = & \, \mathbf{d}_{\mathrm{F}}(P,Q).\end{aligned}$$

□

The bounds stated are tight, as the following examples for $V = \mathbb{R}^2$ show: For $P = Q$, we have $d_{\mathrm{F}}(P,Q) = \mathbf{d}_{\mathrm{F}}(P,Q) = 0$, so that the first bound ist tight. For the second bound, set $P := \langle (0,0), (0,1) \rangle$ and $Q := \langle (0,0), (0,1), (0,0), (0,1) \rangle$. Both P and Q are 1-sampled, and we have $d_{\mathrm{F}}(P,Q) = \frac{1}{2}$ as well as $\mathbf{d}_{\mathrm{F}}(P,Q) = 1$, so that $\mathbf{d}_{\mathrm{F}}(P,Q) = d_{\mathrm{F}}(P,Q) + \frac{1}{2}$.

3 Pattern Matching via Transporter Sets

In the last section, we have seen bounds between d_{F} and $\mathbf{d}_{\mathrm{F}}$. In this section, we let a group G act on the polygonal curves. The bounds from the last section carry into distance bounds between a polygonal curve P and the G-orbit of a second curve Q, so that algorithms for matching with respect to $\mathbf{d}_{\mathrm{F}}$ yield approximate algorithms for matching with respect to d_{F}.

Let G denote a subgroup of $\mathrm{AGL}(k)$, the group of all affine transformations in $V = \mathbb{R}^k$. Since G acts on V, G also acts on the set of all finite sequences of points in V. Furthermore, G acts on the set of all polygonal curves. This motivates us to write $gP := \langle gp_0, \ldots, gp_m \rangle$ for $g \in G$ and

$P \in V^{[0:m]}$. We now define the set of all $(G, \varepsilon, \mathbf{d}_{\mathrm{F}})$-*matches of* Q *with respect to* P as $G(P, Q, \varepsilon, \mathbf{d}_{\mathrm{F}}) := \{g \in G \mid \mathbf{d}_{\mathrm{F}}(P, gQ) \leq \varepsilon\}$. Analogously, $G(P, Q, \varepsilon, d_{\mathrm{F}}) := \{g \in G \mid d_{\mathrm{F}}(P, gQ) \leq \varepsilon\}$. The matching task we deal with in the sequel can now be stated as the following decision problem: Given $P \in V^{[0:m]}, Q \in V^{[0:n]}$ and $\varepsilon \geq 0$, determine whether $G(P, Q, \varepsilon, \mathbf{d}_{\mathrm{F}})$ is empty or not.

The bounding property of $\mathbf{d}_{\mathrm{F}}$ and d_{F} from Theorem 1 immediately yields a relation between matches with respect to $\mathbf{d}_{\mathrm{F}}$ and matches with respect to d_{F}:

Corollary 2 *Let* $P \in V^{[0:m]}$ *and* $Q \in V^{[0:n]}$ *be* δ*-sampled polygonal curves. Then, for* $G \leq \mathrm{AGL}(k)$*, we have* $G(P, Q, \varepsilon, d_{\mathrm{F}}) \subseteq G(P, Q, \varepsilon + \delta/2, \mathbf{d}_{\mathrm{F}}) \subseteq G(P, Q, \varepsilon + \delta/2, d_{\mathrm{F}})$.

3.1 A Basic Matching Algorithm

Our approach for solving the decision problem is based on considering (G, ε)-transporter sets for points $p, q \in V$ defined as

$$\tau_{p,q}^{G,\varepsilon} := \{g \in G \mid \|p - gq\| \leq \varepsilon\}. \tag{3}$$

The next remark shows the close relation of transporter sets to $(G, \varepsilon, \mathbf{d}_{\mathrm{F}})$-matches.

Remark 3 *Let* $P \in V^{[0:m]}$ *and* $Q \in V^{[0:n]}$*. Then,*

$$G(P, Q, \varepsilon, \mathbf{d}_{\mathrm{F}}) = \bigcup_{(\kappa,\lambda) \in \mathbf{Mon}_{m,n}} \bigcap_{s \in [1:m+n]} \tau_{p_{\kappa(s)}, q_{\lambda(s)}}^{G,\varepsilon}.$$

In particular, $G(P, Q, \varepsilon, \mathbf{d}_{\mathrm{F}}) \neq \emptyset$ *iff* $\bigcap_{s \in [1:m+n]} \tau_{p_{\kappa(s)}, q_{\lambda(s)}}^{G,\varepsilon} \neq \emptyset$ *for some* $(\kappa, \lambda) \in \mathbf{Mon}_{m,n}$.

For $P \in V^{[0:m]}$ and $Q \in V^{[0:n]}$, we consider the family $(\tau_{p_i,q_j}^{G,\varepsilon})_{i \in [0:m], j \in [0:n]}$ of $(m+1)(n+1)$ transporter sets. According to the above remark, we would like to decide whether at least one of the intersections $\bigcap_{s \in [1:m+n]} \tau_{p_{\kappa(s)}, q_{\lambda(s)}}^{G,\varepsilon}$ is non-empty. To this end, we define an equivalence relation on G such that *every* intersection of transporters is a union of equivalence classes. If we compute a subset C of G containing from each equivalence class at least one element, then $G(P, Q, \varepsilon, \mathbf{d}_{\mathrm{F}}) \neq \emptyset$ iff $C \cap G(P, Q, \varepsilon, \mathbf{d}_{\mathrm{F}}) \neq \emptyset$ (such a C will be called a (P, Q, G, ε)-*suptransversal*). Thus, for deciding $G(P, Q, \varepsilon, \mathbf{d}_{\mathrm{F}}) \neq \emptyset$, it suffices to test each $g \in C$ for membership in $G(P, Q, \varepsilon, \mathbf{d}_{\mathrm{F}})$.

The announced equivalence relation on G is defined by $g \sim_{P,Q,G,\varepsilon} g'$ if and only if for all i, j we have $g \in \tau_{p_i,q_j}^{G,\varepsilon} \Leftrightarrow g' \in \tau_{p_i,q_j}^{G,\varepsilon}$, i.e., g is contained in exactly the same transporter sets as g'. We call the $\sim_{P,Q,G,\varepsilon}$-equivalence class the (P, Q, G, ε)-*cell of* g.

Algorithm 4
INPUT: $P \in V^{[0:m]}, Q \in V^{[0:n]}, G \leq \mathrm{AGL}(k), \varepsilon \geq 0$
OUTPUT: $\begin{cases} g \in G(P, Q, \varepsilon, \mathbf{d}_\mathrm{F}) & \textit{if } G(P, Q, \varepsilon, \mathbf{d}_\mathrm{F}) \neq \emptyset \\ \textbf{false} & \textit{otherwise.} \end{cases}$

Match(P, Q, G, ε)
  compute a (P, Q, G, ε)*-suptransversal* $C \subseteq G$
  for each $g \in C$ **do**
    if $\mathbf{d}_\mathrm{F}(P, gQ) \leq \varepsilon$ **then return** g
  end
  return false
end.

The complexity of the above algorithm mainly depends on the size of the suptransversal C and the time it takes to compute C. We now study an example for the case $V = \mathbb{R}^2$ where a cell enumeration can be done with elementary geometric computations: let $\mathrm{SC}(2)$ denote the group of all uniform scalings (without reflections) in the plane, i.e., $\mathrm{SC}(2)$ is the matrix group $\{\lambda\, \mathrm{id}_2 \mid \lambda \in \mathbb{R}_{>0}\}$, where id_2 denotes the 2×2 unit matrix.

Figure 1 provides a geometric construction of an $(\mathrm{SC}(2), \varepsilon)$-transporter, showing that each such transporter can be characterized as a (possibly empty or unbounded) closed interval on the real line. In order to use Algorithm 4, we need to enumerate one representative from each cell defined by a set of closed intervals. Note that the cells are closed intervals also, and the border of each cell belongs to the border of at least one transporter $\tau_{p_i,q_j}^{\mathrm{SC}(2),\varepsilon}$. Thus, it suffices to compute all (upper and lower) borders of the $(m + 1)(n + 1)$ intervals in the parameter space. This can be done in $O(mn)$ time (since computing a transporter's interval borders as in Figure 1 takes $O(1)$ time), yielding a total running time of $O(m^2n^2)$ for matching with respect to $\mathbf{d}_\mathrm{F}$ under the group $\mathrm{SC}(2)$.

Matching with respect to the group $\mathrm{SO}(2)$, i.e., the group of rotations around the origin, works very similar. The group $\mathrm{SO}(2)$ can be parametrized by the unit circle. As shown in Figure 1, a single transporter can be characterized as a circular arc in this parameter space. Now, the cells defined by a set of circular arcs are also circular arcs. Each border of a single cell corresponds to the border of (at least) one transporter $\tau_{p_i,q_j}^{\mathrm{SC}(2),\varepsilon}$, and just as for the case $G = \mathrm{SC}(2)$, we can compute a suptransversal by enumerating all transporters'

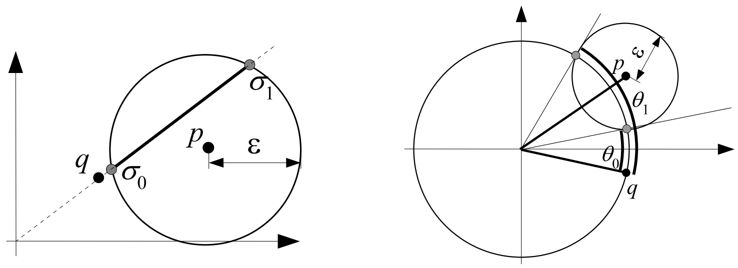

Fig. 1. Construction of $\tau_{p,q}^{G,\varepsilon} \equiv \{\lambda \mid \|\sigma_0\|/\|q\| \leq \lambda \leq \|\sigma_1\|/\|q\|\}$ for $G = \mathrm{SC}(2)$ (left) and $\tau_{p,q}^{G,\varepsilon} \equiv \sphericalangle(\theta_0, \theta_1)$ for $G = \mathrm{SO}(2)$ (right).

borders. Computing these takes $O(mn)$ time, and the total time complexity obtained for Algorithm 4 amounts to $O(m^2n^2)$ as well.

3.2 Projecting Transporter Sets

In this section, we present some group theoretical considerations in order to decrease the computational complexity of matching tasks. As demonstrated in [4] and [3], translating the starting point of Q (which is a reference points in the sense of [10]) onto the starting point of P can easily be shown to yield an approximate solution for matching under translations with respect to d_F. We generalize this result (for $\mathbf{d}_\mathrm{F}$) by showing that the starting points of P and Q can be used to *eliminate translation components* of the transformation group. This is related to a result from [10]. In this work, Alt, Aichholzer and Rothe demonstrate that reference points for the Hausdorff distance can be used to eliminate translation components of the group of similarity motions. Our group theorical point of view allows us to state results for $\mathbf{d}_\mathrm{F}$ that hold for arbitrary subgroups of affine motions in $\mathbb{R}^k$, for any $k > 0$.

A group G is called the *semidirect product* of its subgroup H and its normal subgroup N if $G = \{nh \mid n \in N, h \in H\}$ and $N \cap H = \{1\}$; in this case we sometimes write $G = N \rtimes H$. The most important example of a semidirect product used in the sequel is the *affine general linear group* $\mathrm{AGL}(k) = T(k) \rtimes \mathrm{GL}(k)$, where $T(k)$ denotes the group of all translations in $\mathbb{R}^k$.

The groups $\mathrm{SC}(2)$ and $\mathrm{SO}(2)$ that we have studied so far are both subgroups of $\mathrm{GL}(2)$. Transformation groups that are relevant in practical applications – rigid motions, homothetic motions or similarity motions – are usually affine linear groups, i.e., they additionally contain the subgroup of translations. In the sequel, we study the case that $G = T(k) \rtimes H$ for some $H \leq \mathrm{GL}(k)$ in more detail. Since the Euclidean distance is *translation invariant*, i.e., $\|x - y\| = \|tx - ty\|$ for any $t \in T(k)$ and $x, y \in \mathbb{R}^k$, the following Lemma will be of some use later on:

Lemma 5 *Let N be an abelian group acting transitively on V, i.e., for each pair $(v, w) \in V^2$ there exists an $n \in N$ with $nv = w$. Furthermore, let d be an N-invariant metric on V. Then $d(nx, n'x) = d(ny, n'y)$, for arbitrary $n, n' \in N$ and $x, y \in V$.*

PROOF. Since N acts transitively on V, we can write $x = ty$ for some $t \in N$. As N is abelian and d is N-invariant, we get $d(nx, n'x) = d(nty, n'ty) = d(tny, tn'y) = d(ny, n'y)$. □

According to Remark 3, Algorithm 4 can be seen as an algorithm that decides whether certain intersections of transporters are empty, presuming we can compute a (P, Q, G, ε)-suptransversal for the group under consideration. For the groups SO(2) and SC(2), computing a (P, Q, G, ε)-suptransversal could be done by elementary geometric computations. For most other groups, however, there is no obvious way to compute such suptransversal.

Given a group $G = T(k) \rtimes H$, $H \leq \mathrm{GL}(k)$, we show how to reduce the problem of matching with respect to G to the problem of matching with respect to H. As a result, we will obtain matching algorithms for matching with respect to rigid motions (in place of SO(2)) and homothetic motions (in place of SC(2)).

To this end, we apply the projection η of G onto H with kernel $T(k)$, i.e., $\eta(th) := h$, for $t \in T(k)$ and $h \in H$. This projection is well defined since for every $g \in G$, there is a unique $t \in T(k)$ and $h \in H$ so that $g = th$, which is due to the fact that G is the semidirect product of $T(k)$ and H. Instead of a set $A \subseteq G$, we work with its η-image:

$$\eta[A] := \{h \in H \mid \exists t \in T(k)\colon th \in A\}.$$

Theorem 6 *Let $V = \mathbb{R}^k$ and $G = T(k) \rtimes H$ for some $H \leq \mathrm{GL}(k)$. Given $P \in V^{[0:m]}$ and $Q \in V^{[0:n]}$ as well as $(\kappa, \lambda) \in \mathbf{Mon}_{m,n}$, define*

$$\begin{aligned} &\tilde{p}_i := \tfrac{1}{2}(p_i - p_0),\ \tilde{q}_j := \tfrac{1}{2}(q_j - q_0) \quad \textit{and} \\ &H_{s,\varepsilon} := \tau^{H,\varepsilon}_{\tilde{p}_{\kappa(s)}, \tilde{q}_{\lambda(s)}} \quad \textit{as well as} \\ &G_{s,\varepsilon} := \tau^{G,\varepsilon}_{p_{\kappa(s)}, q_{\lambda(s)}} \end{aligned} \tag{4}$$

Then, we have

(a) $\bigcap_{s\in[1:m+n]} H_{s,\varepsilon} = \emptyset \implies \bigcap_{s\in[1:m+n]} G_{s,\varepsilon} = \emptyset.$
(b) $\bigcap_{s\in[1:m+n]} H_{s,\varepsilon} \neq \emptyset \implies \bigcap_{s\in[1:m+n]} G_{s,2\varepsilon} \neq \emptyset.$

We prepare for the proof of the theorem.

Lemma 7 *Let $V=\mathbb{R}^k$ for some $k>0$ and $P,Q\in V^{[0:1]}$ be two line segements in V, and let $\tilde{P}=\langle -\tilde{p},\tilde{p}\rangle,\tilde{Q}=\langle -\tilde{q},\tilde{q}\rangle\in V^{[0,1]}$ denote the centered versions of P and Q, respectively, so that $\tilde{p}=\frac{1}{2}(p_1-p_0)$ and $\tilde{q}=\frac{1}{2}(q_1-q_0)$. Moreover, let $G=T(k)\rtimes H$ for some $H\leq \mathrm{GL}(k)$. Then, the following holds:*

(a) $\|\tilde{P}-\tilde{Q}\|_\infty\leq\|\tilde{P}-n\tilde{Q}\|_\infty$ *for any* $n\in T(k)$.
(b) $\eta[\tau^{G,\varepsilon}_{p_0,q_0}\cap\tau^{G,\varepsilon}_{p_1,q_1}]=\tau^{H,\varepsilon}_{-\tilde{p},-\tilde{q}}=\tau^{H,\varepsilon}_{\tilde{p},\tilde{q}}$.

PROOF. (a) We have $\|\tilde{P}-\tilde{Q}\|_\infty=\max\{\|-\tilde{p}-(-\tilde{q})\|,\|\tilde{p}-\tilde{q}\|\}=\|\tilde{p}-\tilde{q}\|$ and $\|\tilde{P}-(\tilde{Q}+\langle n,n\rangle)\|_\infty=\max\{\|\tilde{p}-\tilde{q}+n)\|,\|\tilde{p}-\tilde{q}-n\|\}$. As for any $a\in V$,

$$\|a\|=\|\frac{1}{2}(a+n)+\frac{1}{2}(a-n)\|\leq\frac{1}{2}(\|a+n\|+\|a-n\|)\leq\frac{1}{2}\cdot 2\max\{\|a+n\|,\|a-n\|\},$$

our claim follows with $a=\tilde{p}-\tilde{q}$.
(b) We start with the second equality. Since for any $h\in\mathrm{GL}(k)$, we have $-h\tilde{q}=h(-\tilde{q})$, the equality follows from $\|\tilde{p}-h\tilde{q}\|=\|(-\tilde{p})-h(-\tilde{q})\|$, for all $h\in H$.

We get to the proof of the first equality. Note that $\eta[\tau^{G,\varepsilon}_{-\tilde{p},-\tilde{q}}\cap\tau^{G,\varepsilon}_{\tilde{p},\tilde{q}}]=\eta[\tau^{G,\varepsilon}_{p_0,q_0}\cap\tau^{G,\varepsilon}_{p_1,q_1}]$, since $\tilde{P}$ and $\tilde{Q}$ are translated versions of P and Q, respectively. Now, it suffices to prove $\eta[\tau^{G,\varepsilon}_{-\tilde{p},-\tilde{q}}\cap\tau^{G,\varepsilon}_{\tilde{p},\tilde{q}}]=\tau^{H,\varepsilon}_{\tilde{p},\tilde{q}}$.

$\tau^{H,\varepsilon}_{\tilde{p},\tilde{q}}\subseteq\eta[\tau^{G,\varepsilon}_{-\tilde{p},-\tilde{q}}\cap\tau^{G,\varepsilon}_{\tilde{p},\tilde{q}}]$ follows immediately from $\|\tilde{P}-h\tilde{Q}\|_\infty\leq\varepsilon$ for any $h\in\tau^{H,\varepsilon}_{\tilde{p},\tilde{q}}$, and it remains to show the reverse inclusion $\tau^{H,\varepsilon}_{\tilde{p},\tilde{q}}\supseteq\eta[\tau^{G,\varepsilon}_{-\tilde{p},-\tilde{q}}\cap\tau^{G,\varepsilon}_{\tilde{p},\tilde{q}}]$. To this end, let $h\in\eta[\tau^{G,\varepsilon}_{-\tilde{p},-\tilde{q}}\cap\tau^{G,\varepsilon}_{\tilde{p},\tilde{q}}]$. By definition of η, there is a translation t such that $th\in\tau^{G,\varepsilon}_{-\tilde{p},-\tilde{q}}\cap\tau^{G,\varepsilon}_{\tilde{p},\tilde{q}}$, in other words, $\|\tilde{P}-th\tilde{Q}\|_\infty\leq\varepsilon$. From part (a), we get $\|\tilde{P}-h\tilde{Q}\|_\infty\leq\|\tilde{P}-th\tilde{Q}\|_\infty\leq\varepsilon$, so that $h\in\tau^{H,\varepsilon}_{-\tilde{p},-\tilde{q}}\cap\tau^{H,\varepsilon}_{\tilde{p},\tilde{q}}$, and in particular, $h\in\tau^{H,\varepsilon}_{\tilde{p},\tilde{q}}$. □

PROOF of Theorem 6. To prove (a), it suffices to show that

$$\bigcap_{s\in[1:m+n]}\tau^{G,\varepsilon}_{p_{\kappa(s)},q_{\lambda(s)}}\neq\emptyset\implies\bigcap_{s\in[1:m+n]}\tau^{H,\varepsilon}_{\tilde{p}_{\kappa(s)},\tilde{q}_{\lambda(s)}}\neq\emptyset.$$

To begin with, let $g\in\cap_s\tau^{G,\varepsilon}_{p_{\kappa(s)},q_{\lambda(s)}}$. Now, $(\kappa,\lambda)\in\mathbf{Mon}_{m,n}$ implies $g\in G(P,Q,\varepsilon,\mathbf{d}_\mathrm{F})$. Since $\kappa(1)=\lambda(1)=0$, we get $g\in\tau^{G,\varepsilon}_{p_0,q_0}$. This implies

$$g\in\cap_s(\tau^{G,\varepsilon}_{p_{\kappa(s)},q_{\lambda(s)}}\cap\tau^{G,\varepsilon}_{p_0,q_0}).$$

Since $G=T(k)\rtimes H$, we can write $g=th$ for uniquely defined $t\in T(k)$ and $h\in H$, yielding for all $s\in[1:m+n]$

$$h\in\eta[\tau^{G,\varepsilon}_{p_{\kappa(s)},q_{\lambda(s)}}\cap\tau^{G,\varepsilon}_{p_0,q_0}]=\tau^{H,\varepsilon}_{\tilde{p}_{\kappa(s)},\tilde{q}_{\lambda(s)}},$$

where the last equality follows from Lemma 7.(b). This proves implication (a).

For the proof of (b), let $h \in \cap_s \tau^{H,\varepsilon}_{\tilde{p}_{\kappa(s)},\tilde{q}_{\lambda(s)}}$. From Lemma 7.(b) and the definition of $\tilde{p}_i$ and $\tilde{q}_j$, we know that for all $s \in [1 : m+n]$

$$h \in \tau^{H,\varepsilon}_{\tilde{p}_{\kappa(s)},\tilde{q}_{\lambda(s)}} = \eta[\tau^{G,\varepsilon}_{p_0,q_0} \cap \tau^{G,\varepsilon}_{p_{\kappa(s)},q_{\lambda(s)}}]. \tag{5}$$

We claim that the group element $g := t_0 h$ with $t_0 := p_0 - hq_0$ is contained in $\cap_s \tau^{G,2\varepsilon}_{p_{\kappa(s)},q_{\lambda(s)}}$.

First, we observe that $gq_0 = p_0$. Furthermore, due to Eq. (5), we get for all $s \in [1 : m+n]$:

$$\exists n_s \in T(k)\colon\, g_s := n_s h \in \tau^{G,\varepsilon}_{p_0,q_0} \cap \tau^{G,\varepsilon}_{p_{\kappa(s)},q_{\lambda(s)}}. \tag{6}$$

Using Eq. (6), the triangle inequality and Lemma 5, we get

$$\begin{aligned}
\|p_{\kappa(s)} - gq_{\lambda(s)}\| &\le \|gq_{\lambda(s)} - g_s q_{\lambda(s)}\| + \|g_s q_{\lambda(s)} - p_{\kappa(s)}\| \\
&\le \|t_0 h q_{\lambda(s)} - n_s h q_{\lambda(s)}\| + \varepsilon \\
&\le \|t_0 h q_0 - n_s h q_0\| + \varepsilon \\
&= \|p_0 - g_s q_0\| + \varepsilon \ \le\ 2\varepsilon.
\end{aligned}$$

□

With Theorem 6, we get an approximate algorithm for matching polygonal curves with respect to $\mathbf{d}_\mathrm{F}$ by computing $\tilde{P} := \langle \tilde{p}_0, \ldots, \tilde{p}_m \rangle$ and $\tilde{Q} := \langle \tilde{q}_0, \ldots, \tilde{q}_n \rangle$ as defined in Eq. (4) and then match $\tilde{P}$ and $\tilde{Q}$ using Algorithm 4 with respect to SO(2) or SC(2). Since computing $\tilde{P}$ and $\tilde{Q}$ takes $O(m+n)$ time, the following algorithm runs in the same asymptotical time as Algorithm 4.

Algorithm 8
INPUT: $P \in V^{[0:m]}, Q \in V^{[0:n]}; p, q \in V; G = T(k) \rtimes H,\ H \le \mathrm{GL}(k); \varepsilon \ge 0$.
OUTPUT: **Projection-Match**$(P, Q, p_0, q_0, G, \varepsilon) =$

$$\begin{cases}
g \in G(P,Q,2\varepsilon,\mathbf{d}_\mathrm{F}) & \textit{if } G(P,Q,\varepsilon,\mathbf{d}_\mathrm{F}) \neq \emptyset \\
\textbf{false} & \textit{if } G(P,Q,2\varepsilon,\mathbf{d}_\mathrm{F}) = \emptyset \\
g \in G(P,Q,2\varepsilon,\mathbf{d}_\mathrm{F}) \textit{ or } \textbf{false} & \textit{otherwise.}
\end{cases}$$

Projection-Match$(P, Q, p, q, G, \varepsilon)$

$\tilde{P} := \frac{1}{2}(P - p)$ *and* $\tilde{Q} := \frac{1}{2}(Q - q)$;

$h := \textbf{Match}(\tilde{P}, \tilde{Q}, H, \varepsilon)$;

if $h \neq$ **false then** $t_0 := p_0 - hq_0$; **return** $t_0 h$ **else return false**;

end.

We now study the use of this algorithm for matching with respect to $\mathbf{d}_{\mathrm{F}}$ under two subgroups of AGL(2). Let $\mathrm{RM}(k) := T(k) \rtimes \mathrm{SO}(k)$ denote the group of rigid motions and $\mathrm{HM}(k) := T(k) \rtimes \mathrm{SC}(k)$ the group of homothetic motions in the plane. Then, we can use Algorithm 8 for matching with respect to RM(2) and HM(2); the time bounds we obtain are exactly the same as for matching under SO(2) or SC(2). The only price for matching under HM(2) instead of SC(2) is that Algorithm 8 has an indecision interval of size ε. The indecision interval for matching with respect to d_{F} stated in Corollary 2 increases by a factor of 2 correspondingly.

3.3 Matching Subcurves and Closed Curves

We now turn to the *partial Fréchet distance* $\vec{\mathbf{d}}_{\mathrm{F}}$ for measuring resemblance of $Q \in V^{[0:m]}$ as a subcurve of $P \in V^{[0:n]}$ and the discrete Fréchet distance for closed polygonal curves, $\mathbf{d}_{\mathrm{F}}^{\circ}$. As for the discrete Fréchet distance, we first show how to compute $\vec{\mathbf{d}}_{\mathrm{F}}$ as well as $\mathbf{d}_{\mathrm{F}}^{\circ}$, and then propose algorithms for matching with respect to these distance measures.

For measuring whether $Q \in V^{[0:n]}$ is a subcurve of $P \in V^{[0:m]}$, we define the *partial Fréchet distance* as

$$\vec{\mathbf{d}}_{\mathrm{F}}(P, Q) := \min_{[a:b] \subseteq [0:m]} \mathbf{d}_{\mathrm{F}}(P|_{[a:b]}, Q).$$

In order to adapt the discrete Fréchet distance to closed curves, we view P as cyclically continued, i.e., for $i > m$ we let $P(i) := P(i \mod m + 1)$. In analogy to the continuous Fréchet distance for closed curves in [2], we define

$$\mathbf{d}_{\mathrm{F}}^{\circ}(P, Q) := \min_{a \in [0:m]} \mathbf{d}_{\mathrm{F}}(P|_{[a:a+m]}, Q).$$

An important concept we use for deciding $\vec{\mathbf{d}}_{\mathrm{F}}(P, Q) \leq \varepsilon$ and $\mathbf{d}_{\mathrm{F}}^{\circ}(P, Q) \leq \varepsilon$ is the *discrete ε-free space* [2] of two polygonal curves, defined as $\mathbf{F}_{\varepsilon}(P, Q) = \{(i, j) \in \mathbb{Z} \times \mathbb{Z} \mid \|p_i - q_j\| \leq \varepsilon\}$. Defining a *monotonic path of length K* as a mapping $\pi \in (\mathbb{Z} \times \mathbb{Z})^{[0:K]}$ with the property that $\pi(i) - \pi(i-1) \in \{(1,0), (0,1), (1,1)\}$ for all $i \in [1:K]$, we can state a basic property of $\mathbf{d}_{\mathrm{F}}$:

Theorem 9 *Let $P \in V^{[0:m]}$ and $Q \in V^{[0:n]}$ be polygonal curves and let $\varepsilon \geq 0$. Then, we have $\mathbf{d}_{\mathrm{F}}(P,Q) \leq \varepsilon$ iff there is a monotonic path of length $K \leq m+n$ within $\mathbf{F}_\varepsilon(P,Q)$ that starts at $(0,0)$ and ends at (m,n).*

SKETCH OF PROOF. Let $\mathbf{d}_{\mathrm{F}}(P,Q) = \|P \circ \kappa - Q \circ \lambda\|_\infty \leq \varepsilon$, for suitable reparametrizations $(\kappa, \lambda) \in \mathbf{Mon}_{m,n}$. Then the monotonic curve is given by all pairs $(\kappa(i), \lambda(i))$, omitting pairs that yield loops (i.e. $(\kappa(i), \lambda(i)) = (\kappa(i-1), \lambda(i-1))$). Conversely, given a monotonic curve, we obtain suitable reparametrizations by introducing a loop for every diagonal step of the curve (i.e., $(\kappa(i), \lambda(i)) = (\kappa(i-1)+1, \lambda(i-1)+1)$). □

Carrying this result to algorithms for deciding $\vec{\mathbf{d}}_{\mathrm{F}}(P,Q) \leq \varepsilon$ and $\mathbf{d}^\circ_{\mathrm{F}}(P,Q) \leq \varepsilon$, we need to determine whether there is a monotonic path from $(a,0)$ to (b,n) for some $[a:b] \subseteq [0:m]$ (in case of $\vec{\mathbf{d}}_{\mathrm{F}}$) or whether there is a monotonic path from $(a,0)$ to $(a+m,n)$ for some a (in case of $\mathbf{d}^\circ_{\mathrm{F}}$). This idea is crucial for the decision algorithms we propose.

Lemma 10 *Let $P \in V^{[0:m]}$,$Q \in V^{[0:n]}$ and $\varepsilon \geq 0$. Then, $\vec{\mathbf{d}}_{\mathrm{F}}(P,Q) \leq \varepsilon$ as well as $\mathbf{d}^\circ_{\mathrm{F}}(P,Q) \leq \varepsilon$ can be decided in $O(mn)$ time.*

PROOF. We start with a decision algorithm for $\vec{\mathbf{d}}_{\mathrm{F}}(P,Q) \leq \varepsilon$. Consider the following algorithm: for $i \in [0:m]$ and $j \in [0:n]$, define $r_{i,j} := \max\{a \in [0:i] \mid \mathbf{d}_{\mathrm{F}}(P|_{[a:i]}, Q|_{[0:j]}) \leq \varepsilon\}$ (and $r_{i,j} := -\infty$ if no such a exists). Using dynamic programming as in the algorithm from Section 2 for computing $\mathbf{d}_{\mathrm{F}}$, we can compute in $O(mn)$ time the values $r_{i,n}$ for each $i \in [0:m]$. We have $\vec{\mathbf{d}}_{\mathrm{F}}(P,Q) \leq \varepsilon$ if and only if $r_{i,n} > -\infty$, for some $i \in [0:2m]$.

Deciding $\mathbf{d}^\circ_{\mathrm{F}}(P,Q) \leq \varepsilon$ works similar as deciding $\vec{\mathbf{d}}_{\mathrm{F}}(P,Q) \leq \varepsilon$. We compute $r_{i,j}$ as well as $R_{i,j} := \max\{b \in [i:2m] \mid \mathbf{d}_{\mathrm{F}}(P|_{[i:b]}, Q|_{[j:n]}) \leq \varepsilon\}$ for all $i \in [0:2m]$ and $j \in [0:n]$, which can also be done in $O(mn)$ time using dynamic programming as follows: We start with computing $R_{m,n}$; then, the algorithm continues as the algorithm for computing $\mathbf{d}_{\mathrm{F}}(P,Q)$, except for the **for**-loops: these run from m downto 0 rather than running from 0 to m and, analogously, from n downto 0 rather than from 0 to n.

We claim that $\mathbf{d}^\circ_{\mathrm{F}}(P,Q) \leq \varepsilon$ if and only if for some i, $r_{i+m,n} \geq i$ and $R_{i,0} \geq i+m$. The necessity of this condition follows immediately from the definitions of $\mathbf{d}^\circ_{\mathrm{F}}(P,Q)$, $r_{i,j}$ and $R_{i,j}$; for the proof that the condition is sufficient, we follow the construction shown in Figure 2:

By definition of $\mathbf{d}^\circ_{\mathrm{F}}$, it suffices to regard the free space restricted to $[0:2m] \times [0:n]$ instead of the complete free space in $\mathbb{Z} \times \mathbb{Z}$. We have two paths contained in $\mathbf{F}_\varepsilon(P,Q)$, one from $(i+c,0)$ to $(i+m,n)$ for some $c \geq 0$ (since $r_{i+m,n} \geq i$) and one from $(i,0)$ to $(i+m+e,n)$ for some $e \geq 0$ (since $R_{i,0} \geq i+m$). As demonstrated in Figure 2, these two paths intersect in some point (i',j'). Hence we can construct a path from $(i,0)$ via (i',j') to $(i+m,n)$ that is completely contained in $\mathbf{F}_\varepsilon(P,Q)$, which proves the claim. $\square$

The stated upper bounds of $O(mn)$ for deciding $\vec{\mathbf{d}}_{\mathrm{F}}(P,Q) \leq \varepsilon$ and $\mathbf{d}^{\circ}_{\mathrm{F}}(P,Q) \leq \varepsilon$ are slightly smaller than the upper bounds of $O(mn\log(mn))$ from [2] for deciding whether the continuous Fréchet distance for closed curves or partial correspondences is at most ε.

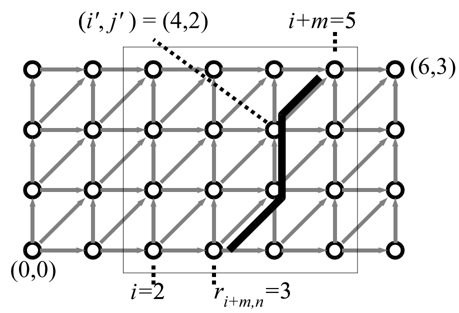

Fig. 2. Deciding $\mathbf{d}^{\circ}_{\mathrm{F}}(P,Q) \leq \varepsilon$ can be done using $r_{i+m,n}$ and $R_{i,0}$ for $i \in [0:m]$. In the example shown, we have $m = n = 3$ and $r_{i+m,n} \geq i$ (left) as well as $R_{i,0} \geq i+m$ (right) for $i = 2$.

Next, we consider matching with respect to $\vec{\mathbf{d}}_{\mathrm{F}}$ and $\mathbf{d}^{\circ}_{\mathrm{F}}$. For matching under the groups SO(2) and SC(2), we can use the idea of enumerating (P,Q,G,ε)-cells as in the first section. Since both $G(P,Q,\varepsilon,\vec{\mathbf{d}}_{\mathrm{F}})$ and $G(P,Q,\varepsilon,\mathbf{d}^{\circ}_{\mathrm{F}})$ are unions of (P,Q,G,ε)-cells, we can apply Algorithm 4; instead of testing $\mathbf{d}_{\mathrm{F}}(P,gQ) \leq \varepsilon$ for each cell, we test $\vec{\mathbf{d}}_{\mathrm{F}}(P,gQ) \leq \varepsilon$ or $\mathbf{d}^{\circ}_{\mathrm{F}}(P,gQ) \leq \varepsilon$, respectively. Hence, we obtain exactly the same time bounds as for matching with respect to $\mathbf{d}_{\mathrm{F}}$.

We now apply the technique of projecting transporters for matching approximately with respect to $\vec{\mathbf{d}}_{\mathrm{F}}$ and $\mathbf{d}^{\circ}_{\mathrm{F}}$. Let $G = T(k) \rtimes H$, $H \leq \mathrm{GL}(k)$, be a transformation group. Our goal is to decide approximately whether $G(P,Q,\varepsilon,\mathbf{d}) \neq \emptyset$, where $\mathbf{d} \in \{\vec{\mathbf{d}}_{\mathrm{F}}, \mathbf{d}^{\circ}_{\mathrm{F}}\}$. We need to modify the algorithms from Section 3.2 slightly, since the proof of Theorem 6 (and hence Algorithm 8) relies on the fact that $\kappa(1) = \lambda(1) = 0$, so that q_0 is always matched with p_0. This only holds for $\mathbf{d}_{\mathrm{F}}$, not for $\vec{\mathbf{d}}_{\mathrm{F}}$ or $\mathbf{d}^{\circ}_{\mathrm{F}}$. For the latter distance measures, we only know that q_0 is matched with some vertex p_a. Hence, we try each of the m vertices of P if it can be matched with q_0 by computing **Projection-Match**$(P,Q,p_a,q_0,G,\varepsilon)$ for each $a \in [0:m]$. This introduces an extra factor of m to the time complexity of the resulting matching algorithm. Thus, matching approximately with respect to $\vec{\mathbf{d}}_{\mathrm{F}}$ or $\mathbf{d}^{\circ}_{\mathrm{F}}$ under RM(2) or HM(2) can be done in $O(m^3n^2)$ time; the resulting approximation property reads analogous to Theorem 6.

Finally, we propose a method for finding common subcurves of P and Q. We restrict our considerations to curves that are not cyclically continued and define

$$\mathrm{LCSC}(P,Q,\varepsilon) := \max\{b \in [0:m] \mid \exists\, a,c,d \in \mathbb{N}\colon a+b \le m, c+d \le n, \\ \mathbf{d}_{\mathrm{F}}(P|_{[a:a+b]}, Q|_{[c:c+d]}) \le \varepsilon\}$$

as the *length of the longest common subcurve of P and Q*. Stated as a matching problem, we want to find $\mathrm{LCSC}(P,Q,G,\varepsilon) := \max_{g\in G} \mathrm{LCSC}(P,gQ,\varepsilon)$. $\mathrm{LCSC}(P,Q,\varepsilon)$ can be computed in $O(mn)$ time: we define $L_{i,j} := \max\{b \in [0:i] \mid \exists\, c \in \mathbb{N}\colon \mathbf{d}_{\mathrm{F}}(P|_{[i-b:i]}, Q|_{[c:j]}) \le \varepsilon\}$ if such b exists and $L_{i,j} := -\infty$ otherwise. Using dynamic programming, we can compute $L_{i,j}$ for all $(i,j) \in [0:m] \times [0:n]$ in $O(mn)$ time. Determining the maximum $L_{i,j}$ yields $\mathrm{LCSC}(P,Q,\varepsilon)$.

Observe that every (P,Q,G,ε)-cell is $\mathrm{LCSC}(P,Q,\varepsilon)$-invariant in the sense that for $g \sim_{P,Q,G,\varepsilon} g'$, we have $\mathrm{LCSC}(P,gQ,\varepsilon) = \mathrm{LCSC}(P,g'Q,\varepsilon)$; in fact, we could use any mapping $L\colon V^{[0:m]} \times V^{[0:n]} \to R$, where R is some totally ordered set and we have the property that $g \sim_{P,Q,G,\varepsilon} g'$ implies $L(P,gQ) = L(P,g'Q)$. Due to this invariance property of each cell, we can apply Algorithm 4 again by computing $\mathrm{LCSC}(P,gQ)$ (or, $L(P,gQ)$ in general) instead of deciding $\mathbf{d}_{\mathrm{F}}(P,gQ) \le \varepsilon$ for each cell representative g; the maximum $\mathrm{LCSC}(P,gQ)$ over all $g \in C$ yields the longest common subcurve's length. Since $\mathrm{LCSC}(P,Q,\varepsilon)$ can be computed in $O(mn)$ time using dynamic programming, we obtain a running time of $O(m^2n^2)$ for $G = \mathrm{SC}(2)$ and $G = \mathrm{SO}(2)$.

Applying transporter projection for computing $\mathrm{LCSC}(P,Q,G,\varepsilon)$ gets one order of magnitude more complex than matching with respect to $\mathbf{d}_{\mathrm{F}}^{\circ}$ or $\vec{\mathbf{d}}_{\mathrm{F}}$, since for the latter distance measures, we used the fact that q_0 is matched with some vertex p_a. However, q_0 is not necessarily part of the longest common subcurve. All we know is that some vertex q_c is matched with some vertex p_a. Hence, we try all $(m+1)(n+1)$ possible combinations of vertices p_a and q_c as a substitute for p_0 and q_0 in Algorithm 8. I.e., we compute **ProjectionMatch**$(P,Q,p_a,q_c,G,\varepsilon)$ for each $a \in [0:m]$ and $c \in [0:n]$. This results in a total time complexity of $O(m^3n^3)$ for finding longest common subcurves under the groups $\mathrm{RM}(2)$ or $\mathrm{HM}(2)$.

3.4 Other Transformation Groups and Distance Measures

The algorithms proposed so far rely on the fact that enumerating all (P,Q,G,ε)-cells can be done efficiently using only basic geometric calculations. This applies to the groups $\mathrm{SO}(2)$ and $\mathrm{SC}(2)$. For larger groups and transformations in higher dimensional spaces, cell enumeration can be done using methods from algebraic geometry.

We consider the case that $G \leq \mathrm{GL}(k)$ is a *linear algebraic group* [11]. In this situation, the group G also is an algebraic subset of $\mathbb{R}^K$, for some $K > 0$. In addition, the transporter sets are semialgebraic subsets of G: let $p, q \in \mathbb{R}^k$ and $g \in G$. As G acts rationally on $\mathbb{R}^k$, each coordinate of $p - gq$ is a rational function in the K coordinates of g; the coefficients of this rational function depend on the coordinates of p and q. Hence, the condition $\|p - gq\|^2 \leq \varepsilon^2$ can be described by one polynomial inequality $u_{p,q,G,\varepsilon} \leq 0$, with a suitable polynomial $u_{p,q,G,\varepsilon} \in \mathbb{R}[X_1, \ldots, X_K]$. Consequently, the family $(\tau^{G,\varepsilon}_{p_i,q_j})_{i\in[0:m],j\in[0:n]}$ of $(m+1)(n+1)$ transporters is described by the family $U_{P,Q,G,\varepsilon} := (u_{p_i,q_j,G,\varepsilon})_{i\in[0:m],j\in[0:n]}$ of $(m+1)(n+1)$ polynomials. For computing a (P, Q, G, ε)-suptransversal, we use the following result by Basu, Pollack and Roy [12,13]:

Theorem 11 *Let $d > 0$, $W \in \mathbb{R}[X_1, \ldots, X_d]$ and let d' denote the real dimension of the variety $\mathbf{V} = \{x \in \mathbb{R}^d \mid W(x) = 0\}$. Furthermore, let U denote a subset of $\mathbb{R}[X_1, \ldots, X_d]$ with cardinality $\ell < \infty$. Define an equivalence relation on $\mathbf{V}$ by $x \sim_{U,\mathbf{V}} y$ iff for all $u \in U$ $\mathrm{sign}(u(x)) = \mathrm{sign}(u(y))$. If all $u \in U$ have degree at most D, then a $(U, \mathbf{V})$-suptransversal C can be computed in $O(\ell^{d'+1} D^{O(d)})$ time. Furthermore, $|C|$ is bounded by $O(\ell^{d'} O(D)^d)$.*

We use this result as follows: We set $\mathbf{V} := G$ and $U := U_{P,Q,G,\varepsilon}$, so that $d = K$ and d' is the real dimension of the group variety G. Now, $g \sim_{U,\mathbf{V}} g'$ implies $g \sim_{P,Q,G,\varepsilon} g'$, so that a $(U, \mathbf{V})$-suptransversal also is a (P, Q, G, ε)-suptransversal. Hence, a (P, Q, G, ε)-suptransversal can be computed in the time bounds stated in Theorem 11.

For a fixed linear algebraic group $G \subseteq \mathbb{R}^K$, the time for computing a (P, Q, G, ε)-suptransversal is $O((mn)^{d'+1})$, since the degree of a polynomial $u_{p,q,G,\varepsilon}$ is bounded by some $D > 0$, independent of m and n. Since $d = K$ is a constant for a fixed group G, the factor $D^{O(d)}$ is constant as well. For the same reason, the cardinality of the suptransversal is bounded by $O((mn)^{d'})$. As a result, the running time of Algorithm 4 equipped with the above mentioned technique for computing a (P, Q, G, ε)-suptransversal is $O((mn)^{d'+1})$. Using the technique of transporter projection from Section 3.2, we obtain the same running times for groups $G = T(k) \rtimes H$ for matching under H with respect to $\mathbf{d}_{\mathrm{F}}$. For matching with respect to $\vec{\mathbf{d}}_{\mathrm{F}}$ and $\mathbf{d}^{\circ}_{\mathrm{F}}$, the results from Section 3.2 yield a running time of $O(m(mn)^{d'+1})$. Computing $\mathrm{LCSC}(P, Q)$ requires $O((mn)^{d'+2})$ time. The implications of these general running times for some common transformation groups are shown in Figure 3.

Finally, it should be mentioned that Algorithm 4 in combination with the technique of cell enumeration can be applied to other distance measures between point sets, such as the directed or the undirected Hausdorff distance as well as the bottleneck distance; the only requirement a distance measure $\mathbf{d}$ needs to satisfy for the correctness of Algorithm 4 is that every set

	SO(2), SC(2)	RM(2), HM(2)	RM(3)
$\mathbf{d}_{\mathrm{F}}$	$O(m^2n^2)$	$O(m^2n^2)$	$O(m^4n^4)$
$\vec{\mathbf{d}}_{\mathrm{F}}$,$\mathbf{d}_{\mathrm{F}}^{\circ}$	$O(m^2n^2)$	$O(m^3n^2)$	$O(m^5n^4)$
LCSC	$O(m^2n^2)$	$O(m^3n^3)$	$O(m^5n^5)$

Fig. 3. Running times obtained by our algorithms for the different distance measures proposed and some typical transformation groups. The running times stated for the groups RM(2), HM(2) and RM(3) refer to approximate matching algorithms using the technique of transporter projection.

of $(G, \varepsilon, \mathbf{d})$-matches is a union of intersections of (P, Q, G, ε)-cells. Furthermore, largest common subcurve computation can be generalized as follows. Given $f\colon V^{[0:m]} \times V^{[0:n]} \to \mathbb{R}$ with the property that $g \sim_{P,Q,G,\varepsilon} g'$ implies $f(P, gQ) = f(P, g'Q)$, for all $P \in V^{[0:m]}$ and $Q \in V^{[0:n]}$. Then, we can solve the maximization problem $(P, Q) \mapsto \max_{g\in G} f(P, gQ)$ by computing $f(P, gQ)$ for each g contained in a (P, Q, G, ε)-suptransversal. This way, we obtain approximate algorithms for finding largest common point sets with respect to the bottleneck distance, as studied in [14]. For details on this generalized scenario, including a generalized result on eliminating translation components based on reference points, we refer to [15].

References

[1] H. Alt, L. J. Guibas, Discrete geometric shapes: Matching, interpolation, and approximation, in: J.-R. Sack, J. Urrutia (Eds.), Handbook of Computational Geometry, Elsevier Science Publishers B.V. North-Holland, Amsterdam, 2000, pp. 121–153.

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

[3] A. Efrat, P. Indyk, S. Venkatasubramanian, Pattern matching with sets of segments, in: Proc. 12th ACM Symp. on Discrete Algorithms, 2001.

[4] H. Alt, C. Knauer, C. Wenk, Matching polygonal curves with respect to the Fréchet distance, in: Proc. 18th Int. Symp. on Theoretical Aspects of Computer Science, 2001, pp. 63–74.

[5] C. Wenk, Shape matching in higher dimensions, Ph.D. thesis, Freie Universität Berlin (2003).

[6] M. Clausen, F. Kurth, Content-based information retrieval by group theoretical methods, in: Proceedings of the NATO Advanced Study Institute on Computational Noncommutative Algebra and Applications, Kluwer Academic Publishers, Dordrecht, NL, 2003, to appear.

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

[8] L. Rabiner, B. Juang, Fundamentals of Speech Recognition, Prentic Hall, Englewood Cliffs, N. J., 1993.

[9] M. Munich, P. Perona, Continuous dynamic time warping for translation-invariant curve alignment with applications to signature verification, in: IEEE International Conference on Computer Vision, Vol. 1, Corfu, Greece, 1999, pp. 108–.

[10] H. Alt, O. Aichholzer, G. Rote, Matching shapes with a reference point, Internat. J. Comput. Geom. Appl. 7 (1997) 349–363.

[11] J. Humphreys, Linear Algebraic Groups, Springer, New York, Heidelberg, Berlin, 1975.

[12] S. Basu, R. Pollack, M.-F. Roy, On the number of cells defined by a family of polynomials on a variety, Mathematika 43 (1996) 120–126.

[13] S. Basu, R. Pollack, M.-F. Roy, On computing a set of points meeting every semi-algebraically connected component of a family of polynomials on a variety, Journal of Complexity 13 (1) (1997) 28–37.

[14] C. Ambühl, S. Chakraborty, B. Gärtner, Computing Largest Common Point Sets under Approximate Congruence, in: Proc. Europ. Symp. Alg., 2000.

[15] A. Mosig, Efficient algorithms for shape and pattern matching, Ph.D. thesis, Institut für Informatik III, Universität Bonn (2004).