Exact and approximate Geometric Pattern Matching for point sets in the plane under similarity transformations *

Dror Aiger† Klara Kedem‡

Abstract

We consider the following geometric pattern matching problem: Given two sets of points in the plane, P and Q, and some $\delta > 0$, find a similarity transformation (translation, rotation and scale) such that $h(T(P), Q) < \delta$, where $h(.,.)$ is the directional Hausdorff distance with L_∞ as the underlying metric. Similarity transformations have not been dealt with in the context of the directional Hausdorff distance and we fill the gap here. We present efficient, exact and approximate algorithms for this problem imposing a reasonable separation restriction on the set Q. For the exact case if the minimum L_∞ distance between every pair of points in Q is 8δ then the problem can be solved in $O(n^2 m \log n)$ time where m and n are the number of points in P and Q respectively. If the points in Q are just $c\delta$ apart from each other for any $0 < c < 1$ we get a randomized approximate solution with expected runtime $O(n^2 c^{-4} \varepsilon^{-8} \log^4 mn)$, where $\varepsilon > 0$ controls the approximation and the answer is correct with high probability.

1 Introduction

A central problem in pattern recognition, computer vision, and robotics, is the question of whether two point sets P and Q in the plane *resemble* each other. The Hausdorff distance (HD) between two point sets P and Q is defined as

$$H(P, Q) = \max(h(P, Q), h(Q, P))$$

where $h(P, Q)$ is the directional Hausdorff distance from P to Q:

$$h(P, Q) = \max_{p \in P} \min_{q \in Q} d(p, q).$$

Here, $d(\cdot, \cdot)$ represents a more familiar metric on points; for instance, the standard Euclidean metric (L_2) or the L_∞ metric.

Algorithms for solving exactly the minimum HD under rigid transformations have high computation time. The best known algorithm for rigid transformations was given by Chew et al. [4], requiring $O(m^3 n^2 \log^2 mn)$ time (for the directional HD and under L_2). To overcome the high complexity of exact algorithms, approximation schemes were considered by several authors.

Goodrich, Mitchell and Orletsky[6] approximate the minimum HD under rigid motion by approximating the transformation. They show that there is a discrete rigid transformation T such that $h(T(P), Q) \leq 4\delta^*$ where δ^* is the minimum directional HD between P and Q. Their algorithm runs in time $O(n^2 m \log n)$. Some papers consider just the *decision* version of the directional HD problem: Given $\delta > 0$, and point sets P and Q, where $|P| < |Q|$, and let G be the group of rigid transformations, find a transformation $T \in G$ for which $h(T(P), Q) \leq \delta$. Approximate algorithms for this problem under rigid transformations can be found in [2, 5]. A more detailed survey on approximate algorithms for this problem will be given in extended version of this paper.

In this paper we consider the decision version of the directional HD problem in the plane under similarity transformations and L_∞. We contribute in this paper the first attempt to work on the directional HD under similarity transformation (translation, rotation and scale) providing exact and guaranteed approximate methods. Throughout the paper we assume that the scale parameter is constrained to be above a minimum value (to avoid the trivial points in the pattern collapse to a single point). We show that if we require that the points in Q obey some separation restrictions (minimum pairwise L_∞ distance) it is possible to solve this problem efficiently. When the separation is 8δ we present an algorithm that solves the problem exactly, in $O(n^2 m \log n)$ time. For separation $c\delta$ for any $0 < c < 1$ we present a randomized algorithm for an approximate solution.

2 Preliminaries

Definition 1 *A point set Q in the plane is called δ-separated if the minimum distance under L_∞ between any pair of points in Q is greater than δ.*

Definition 2 *We define the δ-neighborhood of a point q in the plane to be all the points in R^2 that are in distance*

*This work was partly supported by the MAGNET program of the Israel Ministry of Industry and Trade (IMG4 consortium)

†Department of Computer Science, Ben Gurion University, Be'er Sheva, Israel and Orbotech LTD aiger@cs.bgu.ac.il,dror-ai@orbotech.com

‡Department of Computer Science, Ben Gurion University and Cornell University, klara@cs.bgu.ac.il

$\leq \delta$ from q. For L_∞ in the plane this is a square of side size 2δ.

Definition 3 *Given a set $\mathcal{S}$ of objects in R^d and a point $q \in R^d$, we define the depth of q in $\mathcal{S}$, $depth(q, \mathcal{S})$, to be the number of objects of $\mathcal{S}$ containing q. The depth of $\mathcal{S}$, $depth(\mathcal{S})$, is defined as $max_q depth(q, \mathcal{S})$ over all $q \in R^d$.*

We assume that the set $\mathcal{S}$ is well behaved in the sense that it induces on R^d an arrangement of complexity $O(|\mathcal{S}|^{O(d)})$. Let $\Pi(\mathcal{S})$ be the locus of points realizing $depth(\mathcal{S})$.

Lemma 1 *[1] Let $\mathcal{S}$ be a set of n objects, with $\Delta = depth(\mathcal{S})$. Let ε, $0 < \varepsilon < 1$, be fixed. Let $c > 0$ be a sufficiently large constant. Let k be an integer, with $k \geq \Delta/4$. Let $R \subset \mathcal{S}$ be a random subset formed by picking each object with probability $\rho = \rho(\varepsilon, k) = min(ck^{-1}\varepsilon^{-2}\log n, 1)$, independently. Then, for an appropriate choice of c, we have:*

(i) If $\Delta \geq 2k$, then for every $v \in \Pi(R)$ with high probability $depth(v, \mathcal{S}) \geq (1-\varepsilon)\Delta$ and therefore $depth(R) \geq 2(1-\varepsilon) \cdot k\rho$.

(ii) If $\Delta \leq k$, then with high probability $depth(R) \leq (1+\varepsilon) \cdot k\rho$.

3 The matching problem

Our goal is to find a similarity transformation T that brings every point in P to the δ-neighborhood of Q. We use the linear parametrization of similarity transformation where $(x', y') = T(x, y)$:

$$x' = ax + by + c \tag{1}$$

$$y' = -bx + ay + d \tag{2}$$

where $a = s\cos(\theta)$, $b = s\sin(\theta)$, s is the scale and θ is the rotation, and (c, d) is the translation vector. We describe the similarity transformation as the vector (a, b, c, d) in R^4. Consider the set of similarity transformations that bring a point p in P exactly to a point $q \in Q$. The transformations that correspond to Equation 1 describe a hyperplane in R^4 parallel to the d-axis. The transformations that correspond to Equation 2 describe a hyperplane in R^4 parallel to the c-axis. The set of transformations that bring p to q is thus a 2-flat which is the intersection of the two hyperplanes. Based on the parametrization as above we define a δ-stick in the transformation space:

Definition 4 *Let R_δ be a square of side size 2δ in the (c, d) plane centered at $(c, d) = (0, 0)$. A δ-stick is the polytope in R^4 determined by the Minkowski sum of R_δ with a 2-flat.*

We will also need this variant:

Definition 5 *Let r be a rectangle in the (c, d) plane centered at $(c, d) = (0, 0)$. r-stick is the polytope in R^4 determined by the Minkowski sum of r with a 2-flat.*

Lemma 2 *The set of all similarity transformations that bring p to the δ-neighborhood of q is a δ-stick in transformation space.*

Proof. Let $T = (a, b, c, d)$ be a transformation on the 2-flat that bring p exactly to q, then for every $c', d' \in R_\delta$ the point $T(p) + (c', d') = (a, b, c + c', d + d')(p)$ is in the δ-neighborhood of q. □

The same holds for r-sticks.

Lemma 3 *Let p_1 and p_2 be the points that determine the diameter of P and let q_1 and q_2 be two points in Q. Let Σ be the set of all similarity transformations that map p_1 to the δ-neighborhood of q_1, and p_2 to the δ-neighborhood of q_2. Denote by $R(\Sigma, p)$ the region in the (spatial) plane, formed by transforming a point $p \in P$ by all the transformations in Σ, then $R(\Sigma, p)$ is convex and is bounded by a square of size 6δ (see Figure 1).*

Proof. We prove first that $R(\Sigma, p)$ is convex. Σ is convex since it is an intersection of two δ-sticks each of which is a a Minkowski sum of a planar square and a 2-flat in transformation space. Each line in R^4 connecting two points in Σ corresponds to a line in $R(\Sigma, p)$ thus the convexity of $R(\Sigma, p)$ follows from the convexity of Σ.

Let T be the similarity transformation that maps p_1 and p_2 to q_1 and q_2 respectively. Let T^* be any similarity transformation that maps p_1 and p_2 to the δ-neighborhoods of q_1 and q_2 respectively. Let $q_1^* = T^*(p_1)$ and $q_2^* = T^*(p_2)$. Consider a point $p \in P$. We need to estimate $|T^*(p) - T(p)|_\infty$. We represent T^* as $R \circ S \circ I \circ T$ where I is the translation that moves q_1 to q_1^*, R and S are rotation and scale centered at q_1^* that move $I(q_2)$ to q_2^*. We first apply I on $T(p)$:

$$|(I \circ T)(p) - T(p)|_\infty = |q_1^* - q_1|_\infty \leq \delta$$

next we apply $R \circ S$ on $(I \circ T)(p)$

$$|(R\circ S\circ I\circ T)(p) - (I\circ T)(p)|_\infty \leq |(R\circ S\circ I\circ T)(p_2) - (I\circ T)(p_2)|_\infty \leq 2\delta$$

The last inequality follows from the fact that p_1 and p_2 are the points on the diameter.

thus we have

$$\begin{aligned} &|T^*(p) - T(p)|_\infty \\ &\leq |(I \circ T)(p) - T(p)|_\infty + \\ &|(R \circ S \circ I \circ T)(p) - (I \circ T)(p)|_\infty \\ &\leq \delta + 2\delta = 3\delta \end{aligned}$$

see Figure 2. Thus $R(\Sigma, p)$ is bounded by a square of size 6δ around $T(p)$ as claimed. □

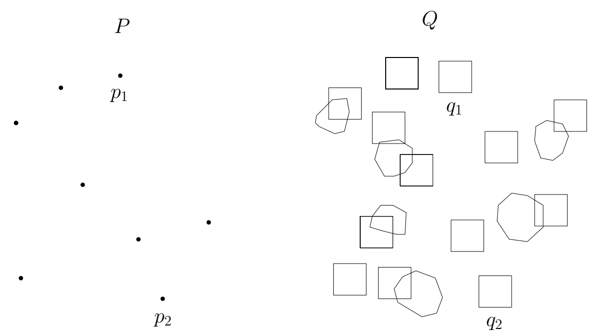

Figure 1: The point set P is on the left. The squares on the right are the δ-neighborhoods of the points in Q and the $|P|$ convex regions on the right are the mappings of all the points in P by the transformations in Σ as defined in Lemma 3. The regions of the mappings of p_1 and p_2 are exactly the δ-neighborhoods of q_1 and q_2.

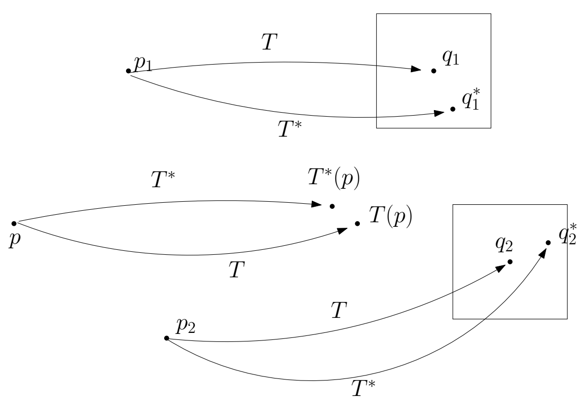

Figure 2: Bounding $|T^*(p) - T(p)|_\infty$

Corollary 4 *Let Q be a set of $c\delta$-separable points and let R be a region in the plane as specified in Lemma 3. Then the number of δ-neighborhoods of Q points whose intersection with R is not-empty is bounded by a constant.*

For $P = \{p_1, .., p_m\}$ and $Q = \{q_1, .., q_n\}$ we have mn δ-sticks which we denote by

$$\mathcal{S} = \begin{pmatrix} s_{11} & .. & s_{1n} \\ . & . & . \\ . & . & . \\ s_{m1} & .. & s_{mn} \end{pmatrix}$$

where s_{ij} is the δ-stick that brings p_i to the δ-neighborhood of q_j. Let p_1 and p_2 be the points that determine the diameter of P and let $u_{ij} = s_{1i} \cap s_{2j}$ for some $1 \leq i \neq j \leq n$. We will prove that the transformation we seek must be contained in some u_{ij}. For finding it we will compute the maximum depth in each u_{ij}.

The next theorem describes our algorithm and proves its runtime.

Theorem 5 *Given two point sets P and Q in the plane, $\delta > 0$, and a bound on the minimum scale. If Q is 8δ-separated (under L_∞) one can determine a similarity transformation T, if such exists, such that $h(T(P), Q) < \delta$. If such a transformation does not exist, then the algorithm returns 'false'. The running time of the algorithm is $O(n^2 m \log n)$.*

Proof. Find the points p_1 and p_2 that determine the diameter of P. Let q_i and q_j be any two points in Q, $1 \leq i \neq j \leq n$. Recall that u_{ij} is a convex polytope in R^4 corresponding to all similarity transformations that map p_1 and p_2 to the δ-neighborhoods of q_i and q_j respectively. By lemma 3, for any p in P, the region $R(u_{ij}, p)$ in the plane is bounded by a square of size 6δ, and since we require in this Theorem a separation of 8δ, there is at most one point from Q whose δ-neighborhood intersects $R(u_{ij}, p)$. Therefore u_{ij} is intersected by at most m δ-sticks in R^4 (one for mapping each point of P to the only point of Q in its region). For each point $p \in P$ ($p \neq p_1$ and p_2) we find the only point $q \in Q$ (if exists) that is closest to $T(p)$ where T is the transformation that brings p_1 and p_2 exactly to q_1 and q_2. Applying linear programming on all these m δ-sticks we answer 'false' if the linear programming returns infeasible, otherwise we know that there exists a transformation that maps all points of P to the δ-neighborhoods of points of Q.

As for the running time, finding the diameter of P can be done in time $O(m \log m)$. We have to solve $O(n^2)$ sub problems, one for each choice of q_i and q_j. In each of them we query the nearest neighbor from Q in the plane and solve a linear program. We first construct a Voronoi diagram of the points in Q so nearest neighbor query can be performed efficiently. This takes $O(n \log n)$ time. Then each of the m nearest neighbor queries is done in logarithmic time via point location. The linear programming is solved in linear time[9] for all the $O(m)$ δ-sticks defined by the reported points. Thus the overall time is $O(n^2(m \log n + m) = O(n^2 m \log n)$. $\square$

For the randomized approximate algorithm we first prove the following Lemma:

Lemma 6 *Let M be a subset of δ-sticks from $\mathcal{S}$ of size $O(n)$. Then one can report all intersections between M and u_{ij} for all $1 \leq i \neq j \leq n$ in time $O(n^2 \log n + K)$ where K is the number of reported intersections.*

Proof. For any δ-stick r in M, we would like to find all u_{ij} that intersect r. The δ-stick r corresponds to the transformations that bring a point p in P to the δ-neighborhood of a point q in Q. If r intersects u_{ij} for some i, j, then there exists a transformation that brings p,p_1 and p_2 to the δ-neighborhood of q,q_i and q_j respectively. In order to find all the triples of δ-sticks (r, s_{1i}, s_{2j}) for which $r \cap u_{ij}$ is not empty we do

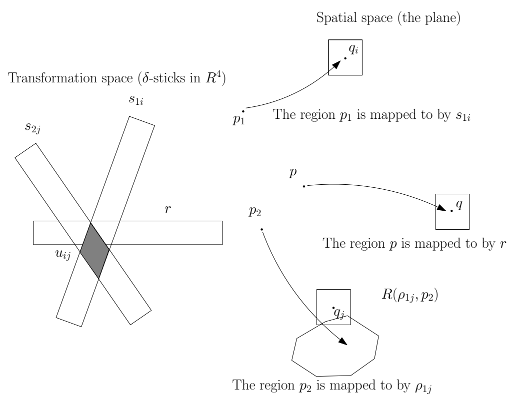

Figure 3: The transformation space is on the left. The spatial space is on the right. The non-empty intersection of the δ-neighborhood of q_j with the region $R(\rho_{1j}, p_2)$ means that r intersects the gray area $u_{i,j}$.

the following: For a given δ-stick r we construct the sets $\rho_{1i} = r \cap s_{1i}$ of transformations that bring p to the δ-neighborhood of q and p_1 to the δ-neighborhood of point $q_i \in Q$, for $i = 1, .., n$. We have n options to pick a point $q_i \in Q$ so we have n such intersection polytopes for a given r. For each ρ_{1i} we build the region $R(\rho_{1i}, p_2)$ the region to which p_2 is mapped to in the plane. Any point q_j in Q whose δ-neighborhood intersects this region implies that u_{ij} is intersected by r in transformation space (see Figure 3). We find these points in Q by prepossessing Q for simplex range searching and reporting all points of Q inside the Minkowski sum of $R(\rho_{1i}, p_2)$ with a square of size δ. Taking every r in M and every q_i in Q we thus have all the intersections in time $O(n^2 \log n + K)$ where $K = k_1 + .. + k_{n^2}$ is the number of all reported points. $\square$

We now state and prove the second theorem of this paper for randomized approximation when Q is just $c\delta$-separated for any $0 < c < 1$:

Theorem 7 *Given two point sets P and Q in R^2 of m and n points respectively, where Q is $c\delta$-separated, $0 < c < 1$, and $0 < \varepsilon < 1$, $\delta > 0$. If there exists a similarity transformation T that brings all points in P to the δ-neighborhoods of points in Q, one can determine a similarity transformation T' such that there exists a subset B of P, for which $|B| \geq (1-\varepsilon)|P|$ and $h(T'(B), Q) < \delta$. If no such T exists the algorithm either finds T' as above or returns 'false'. The expected running time is $O(n^2 c^{-4} \varepsilon^{-8} \log^4 mn)$ and the result is correct with high probability.*

Proof. Recall the matrix $\mathcal{S}$ of δ-sticks above. Here we have to slightly change the setup. Since the points in Q are just $c\delta$-separated, their δ-neighborhoods can be overlapping and so the δ-sticks above. However, we can decompose the set of n squares to a set X of $O(n)$ rectangles disjoint in their interiors in $O(n)$ time as in [3]. Thus the matrix $\mathcal{S}$ now contains $O(mn)$ $r-$sticks (as in Definition 5) built from rectangles in X instead of squares. Their longest side is smaller than δ and they do not intersect each other. Let us call the new point set constructed from all the centers of the rectangles in Q, Q'. Our original problem is now transformed to the problem of bringing the set P to the union of all rectangles in X and each point in the new set Q' is a center of some rectangle in X. Note that the number of columns of the matrix $\mathcal{S}$ is now changed but it is still $O(n)$.

All the r-sticks in one row of $\mathcal{S}$ (which corresponds to a single point in p) do not intersect each other from the construction above. Let p_1 and p_2 be the points that determine the diameter of P and let $U = \{u_{ij}\}$ be the set of polytopes that are the intersection of all pairs of r-sticks, one corresponding to p_1 and the other to p_2. Each u_{ij} is the set of all transformations that bring p_1 and p_2 to the rectangles that correspond to q_i and q_j for some $1 \leq i, j \leq |Q'|$. If there exists a transformation T such that $h(T(P), Q) < \delta$ it must be inside one of the u_{ij}'s since in particular it brings p_1 and p_2 to the neighborhoods of q_i and q_j, respectively. It follows that we only have to search for T within one of the u_{ij}'s. The existence of such a T is equivalent to the existence of a point in R^4 of depth m in the arrangement $A(\mathcal{S})$ of $\mathcal{S}$.

By Lemma 1 we can approximate $depth(\mathcal{S})$ using random sampling. As follows from the Lemma, we sample randomly R from $\mathcal{S}$ by selecting randomly each r-stick in $\mathcal{S}$ for $k = m/2$ (see the lemma), then for a point r which realizes the maximum depth in $A(R)$, $depth(r, \mathcal{S}) \geq (1-\varepsilon) depth(\mathcal{S})$. It follows that it suffices to find $depth(R)$ and by the above argument it suffices to find $depth(R \cap U)$. By corollary 4 each u_{ij} intersects at most $O(m)$ r-sticks in $\mathcal{S}$ (its size is bounded by δ), and since we select each r-stick in $\mathcal{S}$ with probability $c_1 (m/2)^{-1} \varepsilon^{-2} \log mn$ for some constant c_1, the size of the set of r-sticks in R intersecting u_{ij} for some i, j is $O(\varepsilon^{-2} \log mn)$ with high probability.

The expected size of R is $O(n\varepsilon^{-2} \log mn)$, thus, by Lemma 6, we can find these intersections in time $O(|R|^2 \log |R| + K)$ where $K = O(n^2 \varepsilon^{-2} \log mn)$. For the computation of the maximum depth in each u_{ij} we simply construct the arrangement of all the r-sticks that intersect it in time $O(\varepsilon^{-8} \log^4 mn)$ and take the maximum depth resulting from computing for all u_{ij}. The point in R^4 that realizes $depth(R)$ has the approximated depth in $\mathcal{S}$ as claimed. This point in transformation space is the desired transformation T'. As for the running time, we are now working with the set Q' whose size depends on the parameter c. We have $|Q'| = O(c^{-2}|Q|)$ so we have $O(n^2 c^{-4})$ sub problems, each one corresponding to the computation of the maximum depth in the arrangement of $O(\varepsilon^{-2} \log mn)$ polytopes in R^4. The depth is computed by constructing the whole arrangement in R^4 and thus takes $O(\varepsilon^{-8} \log^4 mn)$ time. In total, all the subproblems take expected time $O(n^2 c^{-4} \varepsilon^{-8} \log^4 mn)$. Note that if there is no transformation that brings all points in P to the δ-neighborhoods of points Q, finding the approximate transformation T' is not guaranteed.

□

4 Conclusions

We showed that for the pattern matching problem for point sets P and Q in the plane, if we apply reasonable separability restrictions on Q, we can get efficient algorithms. For 8δ-separated we showed an exact algorithm with running time $O(n^2 m \log n)$. For $c\delta$-separated $(0 < c < 1)$ we can approximate the solution in quadratic time in n.

5 Acknowledgments

The authors thank Sariel Har Peled for fruitful discussions and ideas regarding depth in arrangements. We also thank Michael Burdinov from Orbotech for useful comments on the correctness of Lemma 3

References

[1] B. Aronov, S. Har-Peled, "On approximating the depth and related problems", *In Proceedings of the Sixteenth Annual ACM-SIAM Symposium on Discrete Algorithms* (Vancouver, British Columbia, January 23 - 25, 2005).

[2] D. E. Cardoze, L. J. Schulman, "Pattern Matching for Spatial Point Sets", *FOCS* 1998: 156-165

[3] L. P. Chew, K. Kedem, "Getting around a lower bound for the minimum Hausdorff distance", *Comput. Geom.* 10(3): 197-202 (1998).

[4] P. Chew, M. Goodrich, D. Huttenlocher, K. Kedem, J. Kleinberg, D. Kravets, "Geometric Pattern Matching under Euclidean motion", *Computational Geometry: Theory and applications,* 7:113-124,1997.

[5] M. Gavrilov, P. Indyk, R. Motwani, S. Venkatasubramanian, "Combinatorial and Experimental Methods for Approximate Point Pattern Matching", *Algorithmica* 38(1): 59-90 (2003)

[6] M. T. Goodrich, J. S. B. Mitchell, M. W. Orletsky, "Approximate Geometric Pattern Matching Using Rigid Motions", *IEEE trans. Pattern Anal. Match. Intell.* 21(4):371-379(1999).

[7] P. J. Heffernan, S. Schirra, "Approximate Decision Algorithms for Point Set Congruence", *Comput. Geom.* 4: 137-156 (1994).

[8] J. Matousek, "Efficient Partition Trees". *Discrete and Computational Geometry 8*: 315-334 (1992)

[9] N. Megiddo, "Linear Programming in Linear Time When the Dimension Is Fixed". *J. ACM* 31(1): 114-127 (1984).