Lower bounds for the complexity of the Hausdorff distance

William J. Rucklidge
Department of Computer Science, Cornell University
Ithaca NY 14853, USA
wjr@cs.cornell.edu

Abstract

We describe new lower bounds for the complexity of the directed Hausdorff distance under translation and rigid motion. We exhibit lower bound constructions of $\Omega(n^3)$ for point sets under translation, for the L_1, L_2 and L_∞ norms, $\Omega(n^4)$ for line segments under translation, for any L_p norm, $\Omega(n^5)$ for point sets under rigid motion and $\Omega(n^6)$ for line segments under rigid motion, both for the L_2 norm. The results for point sets can also be extended to the undirected Hausdorff distance.

1 Introduction

The Hausdorff distance between two sets A and B is defined as

$$H(A,B) = \max(h(A,B), h(B,A))$$

where

$$h(A,B) = \sup_{a\in A} \inf_{b\in B} \|a-b\|$$

and $\|\cdot\|$ is some norm (here restricted to be some L_p norm). In this paper, A and B will be planar sets consisting of either a finite number of points or a finite number of points and non-intersecting line segments.

$h(A,B)$ is called the *directed* Hausdorff distance from the set A to the set B. $H(A,B)$ is the undirected Hausdorff distance between the sets A and B. $h(A,B)$ will be small exactly when every point in A is close to some point in B; $h(B,A)$ will be small when every point in B is close to some point in A, and $H(A,B)$ will be small when both of these are true. In particular, $h(B,A) \le \epsilon$ exactly when for any $b \in B$ there is some $a \in A$ such that $\|a-b\| \le \epsilon$. Let $A^\epsilon = A \oplus D(\epsilon)$ where $D(\epsilon)$ is the "disk" of radius ϵ (the set of all points x such that $\|x\| \le \epsilon$) and $\oplus$ is the Minkowski sum. A key observation is that $h(B,A) \le \epsilon$ iff $B \subset A^\epsilon$.

In many problems, we want to determine the transformation of one set which brings it into closest correspondence with the other set. Let G be some group of transformations. Then for any $g \in G$ define

$$D_G(g) = H(A, g(B)).$$

In other words, we transform the set B by some transformation g and compute the Hausdorff distance between this transformed set and A. This defines a function of g, and we wish to determine the minimum value of this function, as the transformation which gives rise to this minimum value will be the one bringing B into closest correspondence with A. Many approaches to determining this minimising transformation are based on searching the graph of this function (for example, by enumerating the local minima, as in [5]). It is therefore of interest to know what the geometric complexity of this graph may be. Upper bounds have been determined for some transformation groups, but few lower bounds were known. We will also be considering the graph of the function

$$d_G(g) = h(g(B), A)$$

which is the graph of the directed distance from the transformed set $g(B)$ to A.

We will exhibit lower bounds for the complexity of such graphs as follows. For each problem, we will fix some values for ϵ and n, construct sets A and B having kn elements each, for some constant k depending on the problem, and show that the set $\{g | d_G(g) \le \epsilon\}$ will have $\Omega((kn)^l) = \Omega(n^l)$ complexity, for some constant l depending on the problem. We do this by showing that this set has $\Omega(n^l)$ distinct connected components. Since each one must contain some local minimum of $d_G(g)$, this shows that there are $\Omega(n^l)$ local minima in the graph of $d_G(g)$. In some cases, we also show that the graph of $D_G(g)$ may have this complexity. Previous constructions, such as those in [4], have been for the directed Hausdorff distance alone.

We also show that the constructions for the undirected Hausdorff distance and the constructions for the directed distance on which they are based may have high complexity in a small space: for a fixed ϵ, we can make $D_G(g)$ have $\Omega(n^l)$ complexity in an arbitrarily small region of transformation space (i.e. this does not depend on just shrinking ϵ). This is motivated by the observations in [2] and [6] that for some groups G, if $D_G(g) \le \epsilon$, then g must lie in a small region in transformation space, and thus, if the undirected

Hausdorff distance could have only small complexity in a small area, we might be able to obtain efficient algorithms, as was done in [6]. The lower bounds here show that, in some cases, this is not possible. Table 1 shows the problems for which we present lower bounds, and the running times of the algorithms which search their graphs. It can be seen that in most cases, the running times are nearly tight. The exception is the bound for point sets under translation with the L_1 and L_∞ norms, in [4], where an algorithm was given which uses the structure of the problem under these norms to avoid explicitly searching the entire graph.

In this paper, we will be dealing with two transformation groups: the group T of translations and the group R of rigid motions (translations and rotations). Let t be a translation and θ an angle. Define

$$\begin{aligned} D_T(t) &= H(A_T, B_T \oplus t) \\ d_T(t) &= h(B_T \oplus t, A_T) \\ D_R(t,\theta) &= H(A_R, r_\theta(B_R) \oplus t) \\ d_R(t,\theta) &= h(r_\theta(B_R) \oplus t, A_R) \end{aligned}$$

where $r_\theta(B_R)$ denotes the set obtained by rotating B_R by θ counterclockwise about the origin.

2 Point sets under translation

This section describes two constructions of point sets A_T, B_T, each containing $O(n)$ points, for which the function $D_T(t)$ has $\Omega(n^3)$ local minima within an arbitrarily small area. The first construction is for the L_1 or L_∞ norm; the second is for the L_2 norm.

2.1 The L_1 and L_∞ example

We use the L_∞ norm throughout; rotating the point sets by 45° gives the construction for L_1.

Let A_T consist of two diagonal rows, each of n points spaced σ apart. The rows are $2\epsilon + \delta$ apart, where $\delta < \sigma/n$. A_T and A_T^ϵ are shown in Figure 1(a). The area uncovered by A_T^ϵ contains a "staircase" of $\Omega(n)$ steps, in the gap between the two sides. The width of this gap is δ. Note that by reducing σ, the two rows can be compressed inwards, thereby making the stairsteps (and thus the total length of the staircase) as small as desired.

Let B_T consist of two diagonal rows of points as shown in Figure 1(b). The points in each row are slightly more than δ apart, and are placed so that one row lies around the horizontal part of a stairstep, and the other lies around the adjacent vertical part.

Consider translating B_T slightly upwards or downwards. The points around the vertical stairstep will remain inside A_T^ϵ or outside it, as they were before, but the points around the horizontal stairstep will move into and out of A_T^ϵ as B_T moves. Similarly, as B_T moves left or right, the points around the vertical stairstep will move in and out of A_T^ϵ, but the points around the horizontal stairstep will not. We can thus see $\Omega(n^2)$ different configurations of B_T with respect to this one stairstep, since we can independently choose where the gaps lie in the two rows of B_T. B_T can also be translated so that it straddles any of the other stairsteps, each of which gives rise to $\Omega(n^2)$ configurations, for a total of $\Omega(n^3)$ configurations.

Each one of these configurations can be labelled with three numbers from 1 to n: the number of points in the bottom row of B_T that are to the right of the gap, the number of points in the top row of B_T that are to the right of the gap, and the number of the stairstep which is straddled by B_T. We will only consider configurations where there is at least one point of each row of B_T on either side of the gap. $\Omega(n^3)$ such labels are possible. Suppose t_1 and t_2 are translations representing configurations with distinct labels. Then $d_T(t_1) \le \epsilon$ and $d_T(t_2) \le \epsilon$, but any path from t_1 to t_2 must pass through a translation t where $d_T(t) > \epsilon$: either some point in one of the rows of B_T must cross the gap, in which case $d_T(t) > \epsilon$ when t is a translation placing that point inside the gap, or B_T must be moved so as to straddle another stairstep, in which case again at least one point of B_T must move through the gap. All these labels therefore label distinct regions.

Another way to visualise this is similar to that used in [4]: define $S(A_T, \epsilon, b)$ for some $b \in B_T$ to be $A_T^\epsilon \oplus -b$. Then $t \in S(A_T, \epsilon, b)$ exactly when $b + t \in A_T^\epsilon$. This set is therefore the set of all translations which map b into A_T^ϵ. Now define $S(A_T, \epsilon, B_T) = \bigcap_{b \in B_T} S(A_T, \epsilon, b)$. Then $t \in S(A_T, \epsilon, B_T)$ iff $B_T \oplus t \subset A_T^\epsilon$, or $D_T(t) \le \epsilon$; $S(A_T, \epsilon, B_T)$ is therefore the set of all translations t which make $h(A_T, B_T \oplus t) \le \epsilon$.

We can construct $S(A_T, \epsilon, B_T)$ by making a copy of A_T^ϵ for every point in B_T, translating these copies and forming their intersection. Alternately, we can consider making a copy of the *complement* of A_T^ϵ for every point in B_T, translating these copies, and forming their *union*. This union will have a "hole" for every connected component of $S(A_T, \epsilon, B_T)$. Figure 1(c) shows such a union. In this figure, each jagged line represents the "gap" staircase of some translation of A_T^ϵ. There are two sets of translations of this, corresponding to the two rows of B_T. There are other parts of the complement of A_T^ϵ which have not been drawn here, but they do not affect the $\Omega(n^3)$ holes shown here.

We now note that δ can be made as small as desired, thereby narrowing the staircase gap, and the

Problem			Lower Bound	Solution
Transformation	Set type	Norm	Complexity	Complexity
Translation	Point Sets	L_1, L_∞	$\Omega(n^3)$	$O(n^2 \log^2 n)$ [4]
		L_2	$\Omega(n^3)$	$O(n^3 \log n)$ [5]
	Points and Segments	L_1, L_∞	$\Omega(n^4)$	$O(n^4\alpha(n))$ [5]
		L_2	$\Omega(n^4)$	$O(n^4 \log^3 n)$ [1]
Rigid Motion	Point Sets	L_2	$\Omega(n^5)$	$O(n^5 \log^2 n)$ [3]
	Point and Segments	L_2	$\Omega(n^6)$	$O(n^6 \log^2 n)$ [3]

Table 1: Results for the complexity of the Hausdorff distance between two sets of size n. The lower bound results for point sets are for the undirected Hausdorff distance, and the results for sets of points and segments are for the directed distance. The results for the undirected distance also show that this complexity can occur in arbitrarily small space.

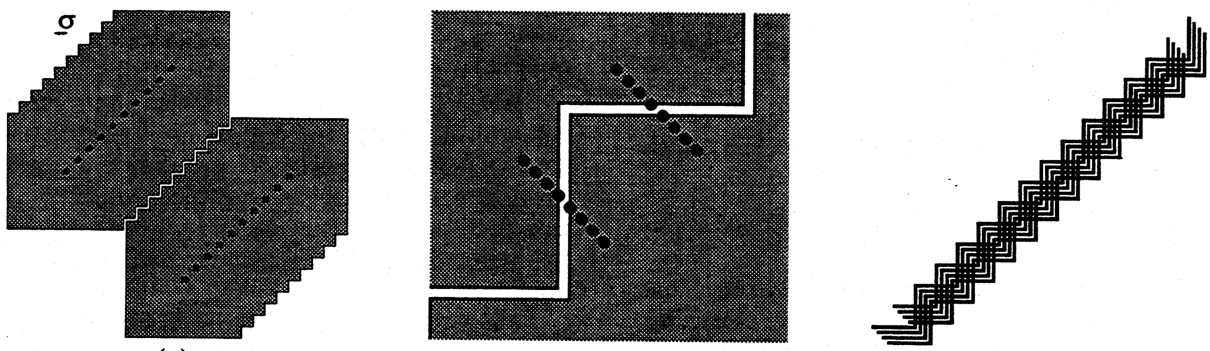

Figure 1: The L_∞ example for point sets under translation

staircase itself can be compressed as much as is desired by reducing σ. This means that we can compress the $\Omega(n^3)$-complexity region down into an arbitrarily small area.

We would also like to show that the area where $D_T(t) \leq \epsilon$ can also have large complexity in a small space. This can be done by augmenting B_T so that $h(A_T, B_T \oplus t) \leq \epsilon$ in the entire region of interest. We set $\sigma < \epsilon/n$ so that the rows of A_T have length less than ϵ, and add two points to B_T, one in the middle of each row of A_T. Then if the main body of B_T is translated anywhere on the staircase, these two extra points remain within (or close to) the rows of A_T. Since the rows have length less than ϵ, there is always some point of B_T within ϵ of any point of A_T, for any translation in the complex region. Thus, $H(A_T, B_T \oplus t) > \epsilon$ exactly where $h(B_T \oplus t, A_T) > \epsilon$ (at least in this region of interest), and so has complexity $\Omega(n^3)$; the region containing this complexity can be made arbitrarily small.

2.2 The L_2 example

In this subsection, we show how the previous example can be modified so that it works with the L_2 norm. Here, A_T consists of two vertical rows of n points. The points in each row are spaced σ apart and the rows are "staggered" by $\sigma/2$ (see Figure 2(a)). The rows are slightly less than 2ϵ apart, so that the circles of A_T^ϵ are δ apart at their closest approach. The gap between the left and right sides is not of constant width and has a complicated shape.

The set B_T again consists of two rows of n points. These rows are horizontal, and spaced $\sigma/2$ apart. The points in each row are 2δ apart. They are shown superimposed on A_T^ϵ in Figure 2(b). The idea is that, no matter what values σ, ϵ and n have, if δ is small enough, then it is possible to choose n_1 and n_2 independently, and position B_T such that $B_T \subset A_T^\epsilon$, there are n_1 points of the top row on the left side of the gap, and n_2 points of the bottom row on the left side of the gap. This would give $\Omega(n^2)$ possible configurations of B_T around a single "wobble" in the gap; as there are $\Omega(n)$ such "wobbles", there would be $\Omega(n^3)$ different configurations of B_T with $B_T \subset A_T^\epsilon$. A labelling argument, similar to that in the previous subsection, shows that these configurations are all distinct.

This is difficult to visualise, so again we will look at $S(A_T, \epsilon, B_T)$. We will construct this, as before, by taking the union of $O(n)$ copies of the "gap", translated by various amounts, and showing that the com-

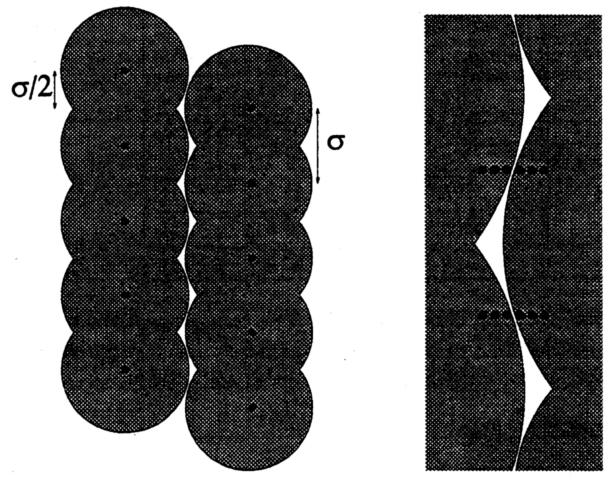

(a) (b)

Figure 2: The sets A_T, A_T^ϵ and B_T for the L_2 example

plement of this union has $\Omega(n^3)$ disjoint connected components.

Since the actual gap has such a complicated shape, we will deal only with a small part of it. In particular, we will consider only the regions where the gap's width is between δ and 2δ (recall that δ is the width of the narrowest part of the gap). There are $O(n)$ regions where this is true, each one centred around a place where the gap is at its narrowest. We bound each such region by a rectangle. These rectangles will be 2δ wide by λ long, where λ is determined by ϵ and δ, and is equal to $\sqrt{4\epsilon\delta - \delta^2}$. Now, note that $\lambda/\delta = \sqrt{4\epsilon/\delta - 1}$. Thus, for any fixed ϵ, we can make λ/δ as large as we like by making δ small enough. This means that the interesting sections of the gap can be made arbitrarily "skinny": as δ decreases, the rectangles get narrower and shorter, but their length to width ratio (i.e. $\lambda/(2\delta)$) increases.

The gap is narrowest exactly where a line from one point in the left row of A_T to one of its neighbours in the right row crosses it. The interesting rectangles are oriented perpendicular to such lines. There are two sets of such rectangles, one "leaning to the left" and the other "leaning to the right". The angle between these two sets decreases as σ decreases, but is not significantly affected by δ. Thus, for a fixed σ (and thus a basically constant angle), it is possible for one left-leaning rectangle to intersect n right-leaning rectangles which are spaced 2δ apart; this can be achieved by making δ small enough that λ/δ is large.

We will be making n copies of the gap stacked 2δ apart (corresponding to one of the rows of B_T), and having these intersect with another n copies, shifted down by $\sigma/2$ (corresponding to the other row), giving

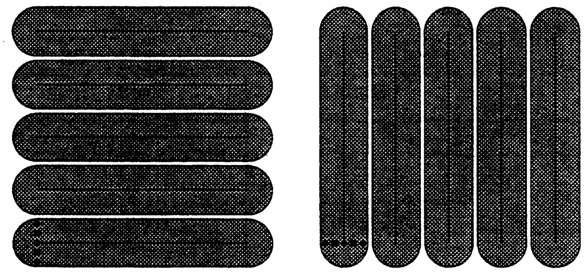

Figure 3: The sets A_T, A_T^ϵ and B_T for the line segment example

$O(n^3)$ intersections: each left-leaning rectangle from one of the copies intersects n right-leaning rectangles from other copies, and vice versa.

Since λ is a function of δ and ϵ, given values for n, ϵ and σ, we can find a value for δ such that there are $\Omega(n^3)$ distinct regions in translation space where $d_T(t) \le \epsilon$. These all lie in a space which is $O(n\sigma)$ high by $O(n\delta)$ wide, and so by a suitable choice of σ, this region of high complexity can be made arbitrarily small.

As in Subsection 2.1, if $n\sigma < \epsilon$, we may augment B_T by two points, one each in the middle of the two rows of A_T, such that $h(A_T, B_T \oplus t) \le \epsilon$ for all translations in the complex region; this construction therefore similarly shows that the undirected Hausdorff distance can have large complexity in a small area.

3 Sets of points and line segments under translation

This section describes a construction of two sets A_T and B_T, each consisting of $O(n)$ points and line segments, for which the function $D_T(t) = H(A_T, B_T \oplus t)$ has $\Omega(n^4)$ complexity.

Let $\delta < \epsilon/n$. Now let A_T consist of a group of n horizontal segments, each of length $(n-1)(2\epsilon+\delta)$, spaced $2\epsilon+\delta$ apart, plus a similar group of n vertical segments. A_T^ϵ then consists of n horizontal bars and n vertical bars, with gaps of width δ between adjacent bars. Note that this is true for any L_p norm: the caps on the ends of the bars have different shapes for different norms, but the main lengths of the gaps between the bars have the same shapes. Now, let B_T consist of a vertical row of n points, spaced 2δ apart, located at the bottom-left corner of the group of horizontal lines in A_T, plus a similar horizontal row located at the bottom-left corner of the group of vertical lines. Figure 3 shows B_T overlaid on A_T and A_T^ϵ.

There are $\Omega(n^4)$ different configurations of B_T with respect to A_T: the vertical row of B_T can be placed in any one of $\Omega(n^2)$ different configurations with respect to the horizontal segments of A_T, as it can be straddling any of the $n-1$ gaps, and from 1 to $n-1$ points can lie below the gap; similarly, the horizontal row can be placed in any one of $\Omega(n^2)$ different configurations with respect to the vertical segments of A_T. Note that sliding B_T horizontally does not affect the configuration of the vertical row, and sliding it vertically does not affect the configuration of the horizontal row (as long as these rows remain within limits); the configurations of the two rows may thus be chosen independently, for a total of $\Omega(n^4)$ different configurations. These are all clearly distinct, since any two differ in the number of points of B_T contained in one of the connected components of A_T^ϵ,so any path from one configuration to another must contain a translation of B_T where some point is crossing some gap.

4 Point sets under rigid motion

We will use the L_2 norm wherever we deal with rotation, since it is the only rotationally symmetric L_p norm.

The following construction shows that there can be $\Omega(n^5)$ distinct connected components where the undirected Hausdorff distance $D_R(t,\theta)$ is less than ϵ, in an arbitrarily small region in (t,θ) space. It is based on an augmentation of A_T and B_T from Subsection 2.2. First, note that it is possible to rotate B_T by some very small angle $\theta_{\min}$ about its centroid while still maintaining the $\Omega(n^3)$ complexity of $h(B_T \oplus t, A_T)$. This is because there must be, in the $\Omega(n^3)$ arrangement of connected components, some minimum distance between features, and so any rotation which does not move any feature of the arrangement more than half this distance cannot change the overall topology of the arrangement.

The augmentation to A_T consists of n points along a vertical line, spaced less than $\epsilon/(2n)$ apart. If we put a disk of radius ϵ about each, we will have a shape with two "scalloped" edges, as shown in Figure 4(a). We will refer to this row of points as A_A.

Now, if A_A is located sufficiently far away from the centre of rotation (and perpendicular to the line joining it to the centre of rotation), then it is possible to pass a circular arc (centred at the centre of rotation) through the inner scalloped edge so that it passes through each of the n lobes and the gaps between them. It will not pass through these lobes evenly: it will cut deeper into some of them than others. However, by moving the row farther away, we may control the magnitude of this effect, since the arc approaches a straight line. By slightly adjusting the radius of the arc, we may also control the ratio between the arc length contained inside the lobes and the arc length contained in the spaces between the lobes. We will

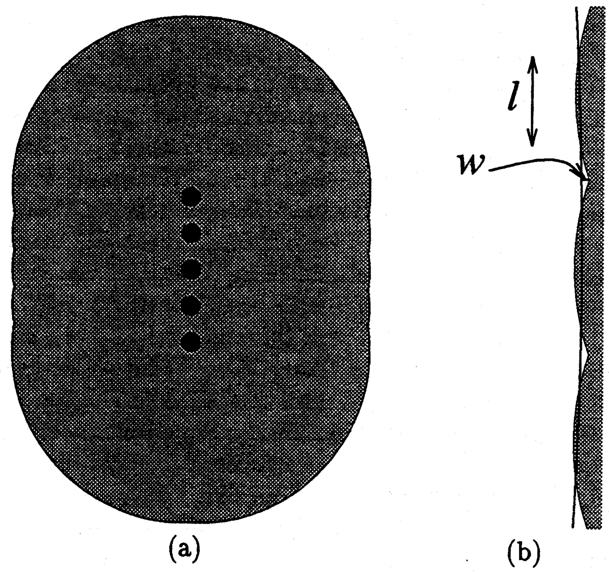

Figure 4: The sets A_A and A_A^ϵ

place A_A far enough away and position the circular arc such that the ratio between the arc length contained in any lobe and the arc length contained in the space next to that lobe is greater than $4n:1$. Let the shortest of the arc lengths contained in the lobes be l, and the smallest depth of any of the gaps be w (see Figure 4(b)).

We now construct B_A, which consists of n points positioned $l/(2n)$ apart along this circular arc, located in the lowest lobe of A_A^ϵ. We also add a point to B_A located at the lower end of A_A. Now, as B_A rotates about the centre of rotation, this row of points will move along this circular arc. Since the spaces between lobes along this path are all less than $l/(4n)$ across, and the entire arc of points will fit into a single lobe, only one point will be passing through a gap at a time, and there will be $\Omega(n^2)$ different configurations of B_A with respect to this part of A_A^ϵ. Note that all of these configurations have the property that all of the points of A_A are within ϵ of some point of B_A, specifically the extra point added at the end of the row.

We now construct A_T and B_T as in Subsection 2.2 so that the $\Omega(n^3)$ complexity region of $D_T(t)$ is smaller than $l/(8n)$ high by $w/2$ wide, and so that the centroid of B_T is at the centre of rotation. Let $A_R = A_T \cup A_A$ and $B_R = B_T \cup B_A$. Pick θ such that the points of B_A are straddling one of the spaces between lobes, and such that this straddling is "even": the two points closest to the gap are equal distances away from the edges of the gap. Then, if $|\theta| < \theta_{\min}$, there are $\Omega(n^3)$ different connected components in t space where $D_R(t,\theta) \leq \epsilon$. A labelling argument similar to that in Subsection 2.1 now shows that there are $\Omega(n^5)$

different connected components in (t,θ) space where $D_R(t,\theta) \leq \epsilon$, $\Omega(n^3)$ corresponding to each such value of θ. A key point is that it is not possible for one of the points of B_A to "sneak around" the space between the lobes, since it would have to translate at least w away from the original circular arc, which would force at least one point of B_T to cross some gap. Further, by reducing σ and δ, and by moving A_A and B_A farther out, we can make the region in (t,θ) space in which these connected components lie arbitrarily small.

There is a problem with this construction: A_A must subtend an angle of less than $\theta_{\min}$, which depends on σ and δ, which depend on l and w, which depend on the circular arc along which B_A is placed, which will have a larger radius for a smaller $\theta_{\min}$. However, as A_A and B_A are moved farther out, l and w approach limit values, as the circular arc becomes closer to a straight line. Thus, if we initially place the augmentation where l and w are within some small factor of their limit values, then determine σ and δ which work for any values of l and w between their original values and their limit values, we can move the augmentation out farther if necessary without affecting the validity of the construction.

5 Sets of points and line segments under rigid motion

This example is a modification of the example from Section 3, using the techniques from Section 4. Again, we observe that the set B_T from Section 3 may be rotated by some small angle $\theta_{\min}$ about its centroid without changing the topology of the arrangement.

Define the centre of rotation to be the centroid of B_T. We augment A_T by a group of segments A_A identical to the left-hand group of A_T, as shown in Figure 3. A_A is placed so that it subtends a total angle of less than $\theta_{\min}$ to the centroid of B_T, and lies directly to the right of it. Let $A_R = A_T \cup A_A$ We also add n points to B_T in a vertical row, in the same relative position to A_A as the first vertical row of B_T was to the left-hand group of A_T. Call this new row B_A and let $B_R = B_T \cup B_A$.

Now, any translation t for which $B_T \oplus t \in A_T^\epsilon$ will also have $B_R \oplus t \in A_R^\epsilon$. Fix such a t and consider values of θ where $|\theta| < \theta_{\min}$. As θ changes through this range, the points in B_A will sweep across the gaps in A_A^ϵ. Their spacing is such that only one point will be crossing a gap at a time. Thus, as we vary θ, the points of B_A achieve $\Omega(n^2)$ different configurations with respect to the gaps of A_A^ϵ. For this choice of t, there are thus $\Omega(n^2)$ values of θ for which $d_R(t,\theta) \leq \epsilon$, since any rotation in this range keeps $r_\theta(B_T) \oplus t \subset A_T^\epsilon$. We can choose t to represent one of the $\Omega(n^4)$ distinct configurations of B_T with respect to A_T, and so this gives $\Omega(n^6)$ different configurations of B_R with respect to the gaps of A_R^ϵ for which $d_R(t,\theta) \leq \epsilon$. These configurations are not connected in (t,θ) space, since any path from one to another must cause at least one point to cross some gap.

6 Summary

We have presented constructions which give lower bounds on the complexity of the directed and undirected Hausdorff distance in several different contexts. Several of the constructions have shown that the directed Hausdorff distance can have large complexity in small space, and this observation has led to lower bounds on the undirected Hausdorff distance.

The problems for which lower bounds on the complexity of the undirected Hausdorff distance were not shown were those involving sets of points and line segments. A remaining open problem is whether or not such bounds are possible: is it possible for the undirected Hausdorff distance under translation between such sets to have large complexity in a small space? If not, can this distance have any complexity greater than $\Omega(n^3)$? Also, is it possible to develop algorithms such as those in [4] which find the minimum Hausdorff distance under the action of some transformation group without explicitly searching the graph of the entire distance function?

7 Acknowledgments

The author would like to thank Paul Chew and Klara Kedem for the $\Omega(n^4)$ example shown in Section 3.

References

[1] P.K. Agarwal, M. Sharir, and S. Toledo. Applications of parametric searching in geometric optimization. In *Proc. Third ACM-SIAM Symposium on Discrete Algorithms*, pages 72–82, 1992.

[2] H. Alt, B. Behrends, and J. Blomer. Approximate matching of polygonal shapes. In *Proc. Seventh ACM Symposium on Computational Geometry*, 1991.

[3] L.P. Chew, M.T. Goodrich, D.P. Huttenlocher, K. Kedem, J.M. Kleinberg, and D. Kravets. Geometric pattern matching under Euclidean motion. In *Proc. Fifth Canadian Conference on Computational Geometry*, August 1993.

[4] L.P. Chew and K. Kedem. Improvements on approximate pattern matching problems. In O. Nurmi and E. Ukkonen, editors, *Proc. Third Scandinavian Workshop on Algorithm Theory*. Lecture Notes in Computer Science 621, Springer-Verlag, 1992.

[5] D.P. Huttenlocher, K. Kedem, and M. Sharir. The upper envelope of Voronoi surfaces and its applications. In *Proceedings of Seventh ACM Symposium on Computational Geometry*, pages 194–203, 1991. To appear in *Discrete Computational Geometry*.

[6] G. Rote. Computing the minimum Hausdorff distance between two point sets on a line under translation. *Information Processing Letters*, 38(3):123–127, May 1991.