Lower Bounds for the Complexity of the Hausdorff Distance

William J. Rucklidge

TR 94-1441
August 1994

Department of Computer Science
Cornell University
Ithaca, NY 14853-7501

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

The Hausdorff distance is a similarity measure defined between sets in the plane. Algorithms to find the minimum distance as one set is transformed have been described, but few lower bounds are known. We describe new lower bounds for the complexity of the directed Hausdorff distance. We exhibit lower bound constructions for both sets of points and sets of points and line segments, under translation, rigid motion, translation and scaling, and affine transformation. The results for point sets can also be extended to the undirected Hausdorff distance. As these lower bounds are for the complexity of the graph of the Hausdorff distance as a function of transformation, they do not necessarily bound functions which search this graph, but do give an indication of how complex the search may be.

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 subsets of the plane 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. The constructions presented below are all parameterised by two values, ϵ and n (and possibly other parameters). For each problem, we will fix some values for ϵ and n (and possibly other parameters), and construct sets A and B having kn elements each, for some constant k depending on the problem. We will then show that the set $\{g \mid d_G(g) \leq \epsilon\}$ has $\Omega((kn)^l) = \Omega(n^l)$ 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 [7] that for some groups G, if $D_G(g) \leq \epsilon$, then g must lie in a small region in transformation space. If the undirected Hausdorff distance could have only small complexity in a small area, even though it might have high complexity overall, we might be able to obtain efficient algorithms, as was done in [7]: the global minimum of $D_G(g)$ would be restricted to lie in a small area, and searching only this area of its graph would take less time than searching the entire graph. The lower bounds here show that, in some cases, this is not possible.

Table 1 shows the problems for which we present lower bounds; Table 2 shows the running times of the algorithms which solve those problems. It can be seen that in most cases, the running times are nearly tight with the lower bounds. 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. It may be possible to develop algorithms using similar techniques for some of the other problems, and so it should be emphasised that the lower bounds presented here are

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

Problem			Lower Bound Complexity
Transformation	Set type	Norm	
Translation	Point Sets	L_1, L_2, L_∞	$\Omega(n^3)$
	Points and Segments	Any L_p	$\Omega(n^4)$
Rigid Motion	Point Sets	L_2	$\Omega(n^5)$
	Points and Segments	L_2	$\Omega(n^6)$
Translation and x, y scale	Point Sets	L_1, L_2, L_∞	$\Omega(n^7)$
	Points and Segments	Any L_p	$\Omega(n^8)$
Affine Transformations	Point Sets	L_1, L_2, L_∞	$\Omega(n^9)$
	Points and Segments	Any L_p	$\Omega(n^{12})$

Table 2: Upper bounds for finding the exact minimum Hausdorff distance under transformation between two sets of size n.

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

for the complexity of the graph of the Hausdorff distance, and do not necessarily give lower bounds for algorithms to determine the optimal transformation.

In this chapter, we will be dealing with four transformation groups: the group T of translations, the group R of rigid motions (translations and rotations), the group S of translations and (x, y) scaling, and the group A of nondegenerate affine transformations. Let t be a translation and θ an angle. Define

$$\begin{aligned}
D_T(t) &= H(A, B \oplus t) & (1)\\
d_T(t) &= h(B \oplus t, A) & (2)\\
D_R(t, \theta) &= H(A, r_\theta(B) \oplus t) & (3)\\
d_R(t, \theta) &= h(r_\theta(B) \oplus t, A) & (4)\\
D_S(t, s_x, s_y) &= H(A, S_{s_x s_y}(B) \oplus t) & (5)\\
d_S(t, s_x, s_y) &= h(S_{s_x s_y}(B) \oplus t, A) & (6)
\end{aligned}$$

$$D_A(t, M) = H(A, M(B) \oplus t) \tag{7}$$
$$d_A(t, M) = h(M(B) \oplus t, A) \tag{8}$$

where

- $r_\theta(B)$ denotes the set obtained by rotating B by θ counterclockwise about the origin,
- $S_{s_x s_y}(B)$ denotes the set $\{(s_x x, s_y y) \mid (x, y) \in B\}$ (i.e. the set obtained by scaling B by a factor of s_x in the x direction and s_y in the y direction),
- M represents a non-singular 2×2 matrix, and
- $M(B)$ denotes the set $\{Mb \mid b \in B\}$ (i.e. the result of applying M to every point in B).

2 Point sets under translation

This section describes two constructions of point sets A and B, 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 consist of two diagonal rows, each of n points spaced σ apart (i.e. σ apart in both x and y). The rows are $2\epsilon + \delta$ apart, where $\delta < \sigma/n$. A and A^ϵ are shown in Figure 1. The area left uncovered by A^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. This is shown in Figure 2.

Let B consist of two diagonal rows of points, each of n points, as shown in Figure 3. 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. Since the stairstep lengths are all σ, and $\delta < \sigma/n$, we can make each row of B shorter (in height and width) than a stairstep segment.

Consider translating B slightly upwards or downwards. The points around the vertical stairstep will remain inside A^ϵ or outside it, as they were before, but the points around the horizontal stairstep will move into and out of A^ϵ as B moves. Similarly, as B moves left or right, the points around the vertical stairstep will move in and out of A^ϵ, but the points around the horizontal stairstep will not. We can thus see $\Omega(n^2)$ different configurations of B with respect to this one stairstep, since we can independently choose where the gaps lie in the two rows of B. B 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 $n - 1$: the number of points in the bottom row of B that are to the left of (or above) the gap, the

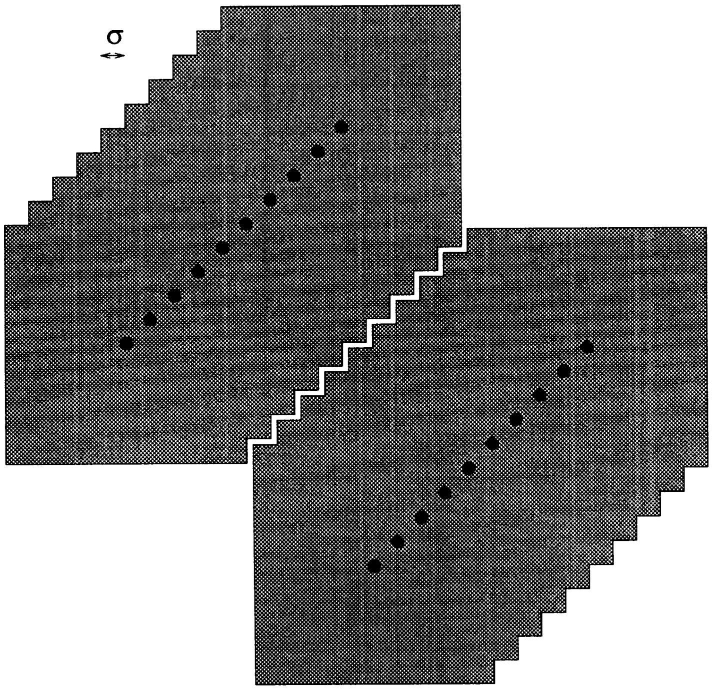

Figure 1: The sets A and A^ϵ for the L_∞ lower bound for point sets under translation.

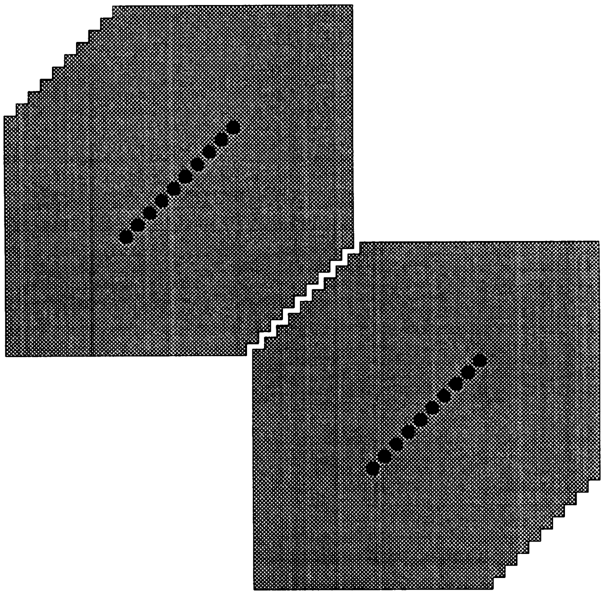

Figure 2: An illustration of the effect of reducing σ.

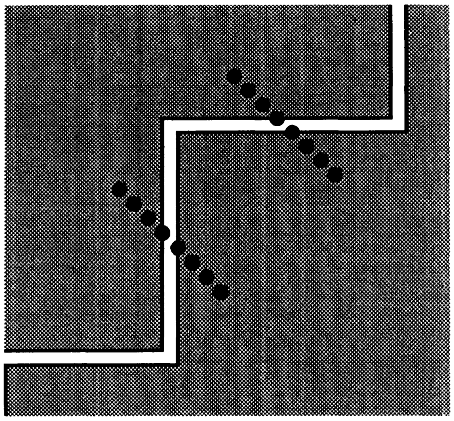

Figure 3: The sets A^ϵ and B for the L_∞ lower bound for point sets under translation.

number of points in the top row of B that are to the left of the gap, and the number of the stairstep which is straddled by B. We will only consider configurations where there is at least one point of each row of B 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 $d_T(t_1) \leq \epsilon$ and $d_T(t_2) \leq \epsilon$, 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 must cross the gap, in which case $d_T(t) > \epsilon$ when t is a translation placing that point inside the gap, or B must be moved so as to straddle another stairstep, in which case again at least one point of B 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 $S(A, \epsilon, b)$ for some $b \in B$ to be $A^\epsilon \oplus -b$. Then $t \in S(A, \epsilon, b)$ exactly when $b + t \in A^\epsilon$. This set is therefore the set of all translations which map b into A^ϵ. Now define $S(A, \epsilon, B) = \cap_{b \in B} S(A, \epsilon, b)$. Then $t \in S(A, \epsilon, B)$ iff $B \oplus t \subset A^\epsilon$, or $D_T(t) \leq \epsilon$; $S(A, \epsilon, B)$ is therefore the set of all translations t which make $h(A, B \oplus t) \leq \epsilon$.

We can construct $S(A, \epsilon, B)$ by making a copy of A^ϵ for every point in B, translating these copies and forming their intersection. Alternately, we can consider making a copy of the *complement* of A^ϵ for every point in B, translating these copies, and forming their *union.* This union will have a "hole" for every connected component of $S(A, \epsilon, B)$. Figure 4 shows such a union. Each jagged line represents the "gap" staircase of some translation of A^ϵ (i.e. portions of one copy of the complement of A^ϵ). There are two sets of translations of this staircase, corresponding to the two rows of B. There are other parts of the complement of A^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 reducing the lengths of the rows of B, and the staircase itself can be compressed as much as is desired by reducing σ (as long as δ stays smaller than σ/n). This means that, for a fixed n and ϵ, we can compress the $\Omega(n^3)$-complexity region down into an arbitrarily small area, bounded by a square $n\sigma$ on each side, since that is the length of the staircase.

We would also like to show that the area where the undirected Hausdorff distance $D_T(t)$ is no greater than ϵ can also have large complexity in a small space. This can be done by augmenting B so that $h(A, B \oplus t) \leq \epsilon$ in the entire region of interest. We set $\sigma < \epsilon/n$ so

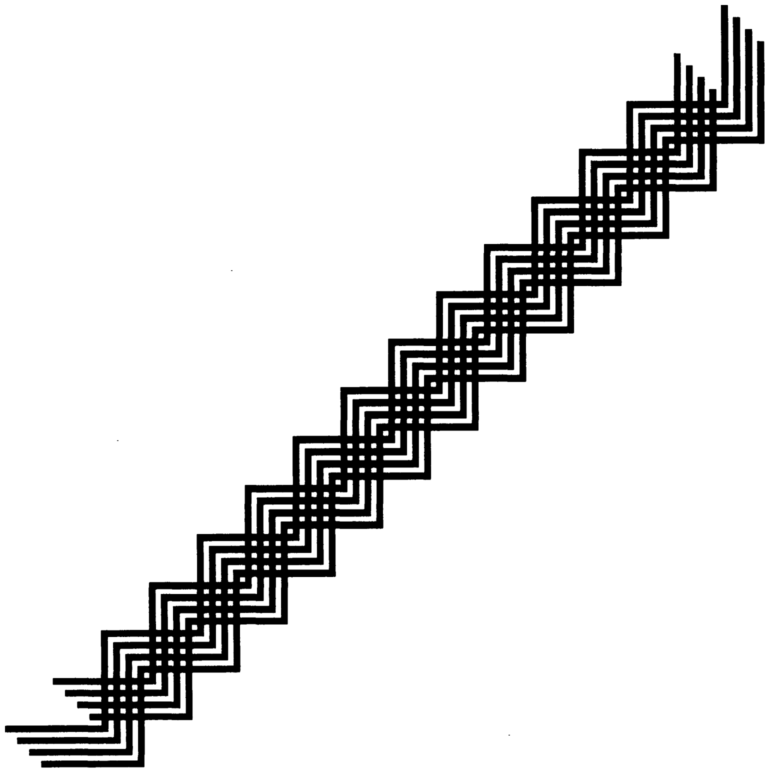

Figure 4: An illustration of $S(A, \epsilon, B)$ for the L_∞ lower bound for point sets under translation.

that the rows of A have length less than ϵ, and add two points to B, one in the middle of each row of A. Then if the main body of B is translated anywhere on the staircase, these two extra points remain within (or close to) the rows of A. Since the rows have length less than ϵ, there is always some point of B within ϵ of any point of A, for any translation in the complex region. Thus, $H(A, B \oplus t) > \epsilon$ exactly where $h(B \oplus t, A) > \epsilon$ (at least in this region of interest), since $h(A, B \oplus t)$ will always be at most ϵ. The undirected Hausdorff distance $D_T(t)$ therefore 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 again 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 5). The rows are slightly less than 2ϵ apart, so that the circles of A^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 again consists of two rows of n points. These rows are horizontal, and spaced $\sigma/2$ apart. The points in each row are slightly more than 2δ apart. They are shown superimposed on A^ϵ in Figure 6. 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 such that $B \subset A^\epsilon$, there are n_1 points of the top row on the left side of the gap, and n_2 points of

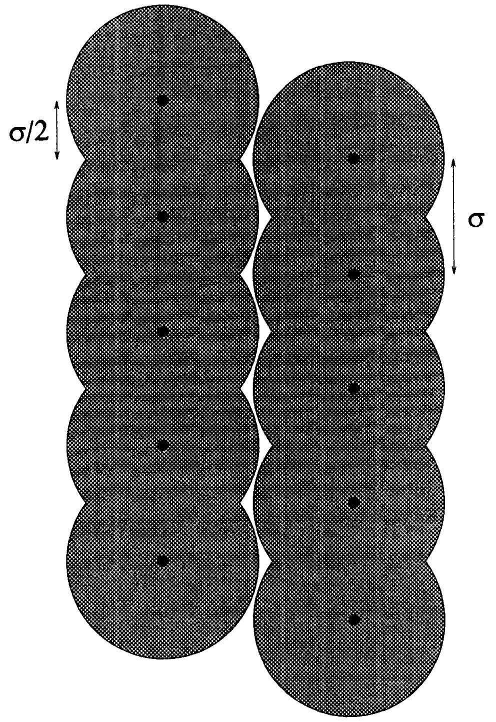

Figure 5: The sets A and A^ϵ for the L_2 lower bound for point sets under translation.

the bottom row on the left side of the gap. This would give $\Omega(n^2)$ possible configurations of B around a single "wobble" in the gap; as there are $\Omega(n)$ such "wobbles", there would be $\Omega(n^3)$ different configurations of B with $B \subset A^\epsilon$. 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 $S(A, \epsilon, B)$. 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 complement 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 $\Omega(n)$ 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 $\sqrt{4\epsilon\delta - \delta^2}$. See Figure 7 for an illustration of this region. 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 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. Suppose that we take n right-leaning rectangles, and position

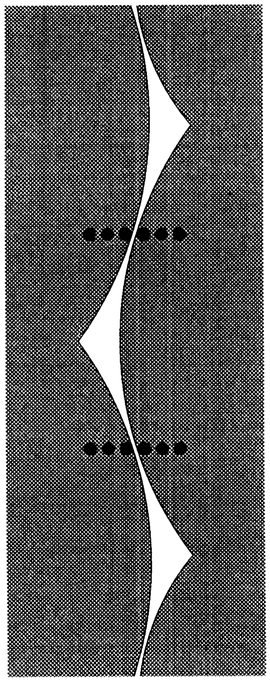

Figure 6: The sets A^ϵ and B for the L_2 lower bound for point sets under translation.

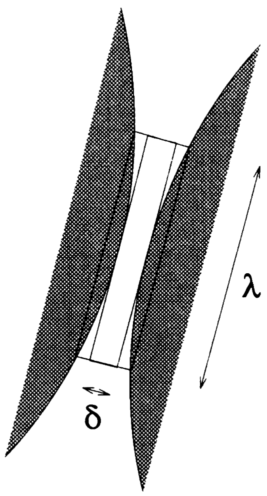

Figure 7: A closeup of the interesting region of the gap.

Figure 8: An illustration of portions of $S(A, \epsilon, B)$ for the L_2 lower bound for point sets under translation. Note the $\Omega(n^3)$ holes in the arrangement.

them slightly more than 2δ apart (the same spacing as the points of B), so that the right edge of one rectangle just touches the left edge of the next. We want to make δ small enough, and thus make the rectangles skinny enough, so that one left-leaning rectangle can intersect all n of these right-leaning rectangles. This can be achieved by making δ small enough that λ/δ is large.

In constructing $S(A, \epsilon, B)$, we will be making n copies of the gap stacked slightly more than 2δ apart (corresponding to one of the rows of B), and having these intersect with another n copies, shifted down by $\sigma/2$ (corresponding to the other row), giving $O(n^3)$ intersections: each left-leaning rectangle from one of the copies intersects n right-leaning rectangles from other copies, and vice versa. Figure 8 shows such an arrangement. In terms of B and A, this means that, as in Subsection 2.1, we can independently choose n_1 and n_2, and position B in such a way that n_1 points from its top row lie in the left half of A^ϵ, and n_2 points from its bottom row lie in the left half. There are $\Omega(n)$ such placements for any n_1 and n_2 ($1 \leq n_1, n_2 < n$), one around each "wobble" of the gap, for $\Omega(n^3)$ different configurations. Going from some configuration to another with a different n_1 or n_2 would involve some point crossing the gap. Also, since the points of B are spaced about 2δ apart and the gap becomes wider than this between two adjacent "wobbles", and all of these configurations have at least one point on either side of the gap, it is impossible to translate B from one configuration to another with the same n_1 and n_2 without some point moving outside A^ϵ. Thus, all these $\Omega(n^3)$ configurations all belong to different connected components of $S(A, \epsilon, B)$.

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 $d_T(t) \leq \epsilon$. These all are contained in an area which is $O(n\sigma)$ high by $O(n\delta)$ wide, and so by a suitable choice of

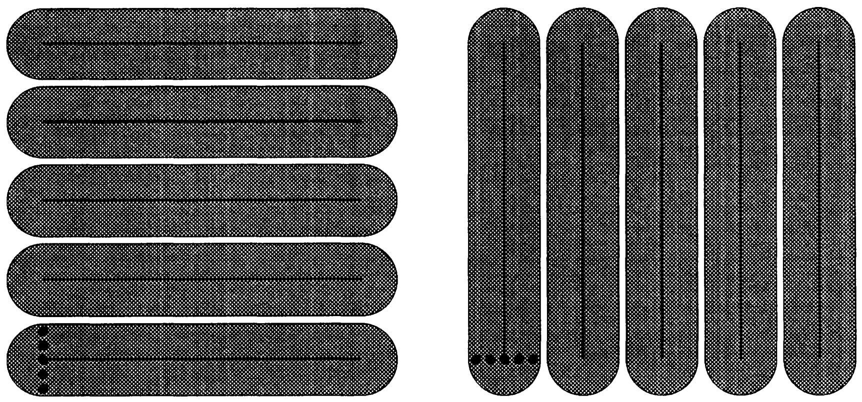

Figure 9: The sets A, A^ϵ and B for points and segments under translation.

σ, this region of high complexity can be made arbitrarily small.

As in Subsection 2.1, if $n\sigma < \epsilon$, we may augment B by two points, one each in the middle of the two rows of A, such that $h(A, B \oplus t) \leq \epsilon$ 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 and B, each consisting of $O(n)$ points and nonintersecting line segments, for which the graph of the function $d_T(t) = h(B \oplus t, A)$ has $\Omega(n^4)$ complexity.

Fix ϵ and n and let $\delta < \epsilon/n$. Now let A 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^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 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, plus a similar horizontal row located at the bottom-left corner of the group of vertical lines. Figure 9 shows B overlaid on A and A^ϵ.

There are $\Omega(n^4)$ different configurations of B with respect to A: the vertical row of B can be placed in any one of $\Omega(n^2)$ different configurations with respect to the horizontal segments of A, 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. Note that sliding B 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 contained in one of the connected components of A^ϵ, so any path from one configuration to another must

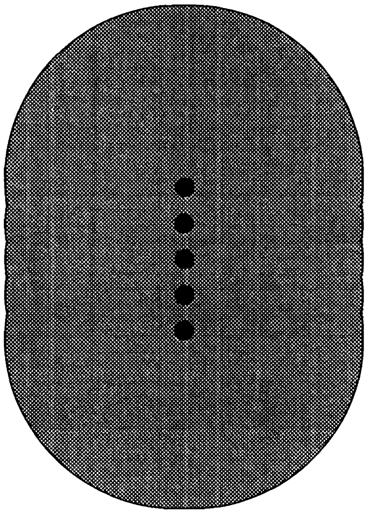

Figure 10: The sets A_2 and A_2^ϵ for point sets under rigid motion.

contain a translation of B 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 $D_R(t,\theta)$ between two sets of $\Omega(n)$ points is less than ϵ, in an arbitrarily small region in (t,θ) space. It is based on an augmentation of A and B from Subsection 2.2. For clarity, we will (both here and below) refer to the sets A and B constructed for the translational lower bounds as A_T and B_T.

First, note that it is possible to rotate B_T from that construction 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: none of the connected components merge, nor do any vanish.

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 10. We will refer to this row of points as A_2.

Now, if A_2 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 and thus increasing the radius of the arc, we may control the magnitude of this effect, since the arc approaches a straight line as A_2 moves away. By slightly adjusting the radius of the arc,

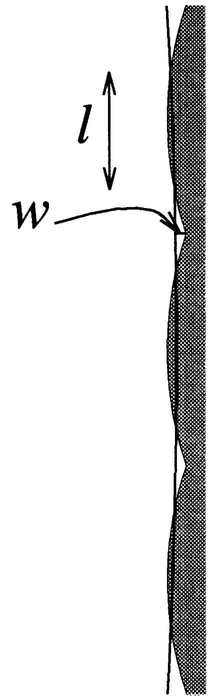

Figure 11: The interaction between A_2^ϵ and B_2.

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 place A_2 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 $8n : 1$. 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 11). As the arc becomes closer to a straight line (as A_2 is moved away from the centre of rotation), the lengths of the arc segments contained in the lobes become more similar (and, at the limit, are all equal). A_2 should be positioned far enough away that they are all within a factor of two of each other. The widest gap is then smaller than $l/(4n)$.

We now construct B_2, which consists of n points positioned $l/(2n)$ apart along this circular arc, initially located in the lowest lobe of A_2^ϵ. We also add an extra point to B_2; this point is initially located at the lower end of A_2. Now, as B_2 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_2 with respect to this part of A_2^ϵ: there are $n-1$ gaps to be straddled, and for each gap, between 1 and $n-1$ points of B_2 can be above the gap. Note that all of these configurations have the property that all of the points of A_2 are within ϵ of some point of B_2, specifically the extra point.

Pick θ such that the points of B_2 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. They will be at least $l/(8n)$ away from these edges. Now, consider translating the points of B_2 vertically up or down by up to $l/(16n)$. If the arc along

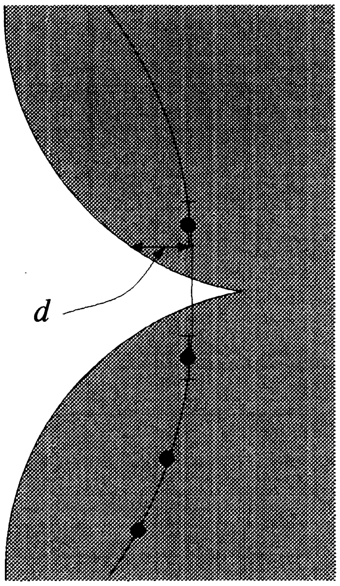

Figure 12: The minimum depth of the points of B_2.

which the points lie is close enough to a straight vertical line, then they will stay inside A_2^ϵ. Let d be the minimum horizontal "depth" achieved for any point of B_2 at any point in this translational range. Figure 12 shows this situation; the vertical bars are $l/(8n)$ high ($l/16n$ above and below the centre). Now, let d_{min} be the smallest such d value achieved for any of the $\Omega(n^2)$ possible such straddling configurations.

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 $l/(16n)$ high by $\min(d_{\text{min}}/2, w/2)$ wide, and so that the centroid of B is at the centre of rotation. Let $A = A_T \cup A_2$ and $B = B_T \cup B_2$. Let θ produce one of the straddling configurations described above. Then, if $|\theta| < \theta_{\text{min}}$, there are $\Omega(n^3)$ different connected components in t space where $D_R(t, \theta) \leq \epsilon$. This is because a slight translation of B_2 with respect to A_2 will not move any of the points of B_2 out of A_2^ϵ, and B_T and A_T have been constructed so that the range of translations required is very small. Note that $D_R(t, \theta)$ will be determined by $d_R(t, \theta)$ for all transformations in the range under consideration, as the directed distance from A to the rotated and translated B will always be less than ϵ.

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 of the $\Omega(n^2)$ such values of θ. A key point in the argument is that it is not possible for one of the points of B_2 to "sneak around" the space between the lobes (through the main body of A_2^ϵ), 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. Thus, any two configurations which differ in how the points of B_2 are straddling the gaps of A_2^ϵ must belong to different connected components. Further, by reducing σ and δ, and by moving A_2 and B_2 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 as it has been presented: A_2 must subtend an angle of less than θ_{min}, which depends on σ and δ, which depend on l and w, which depend

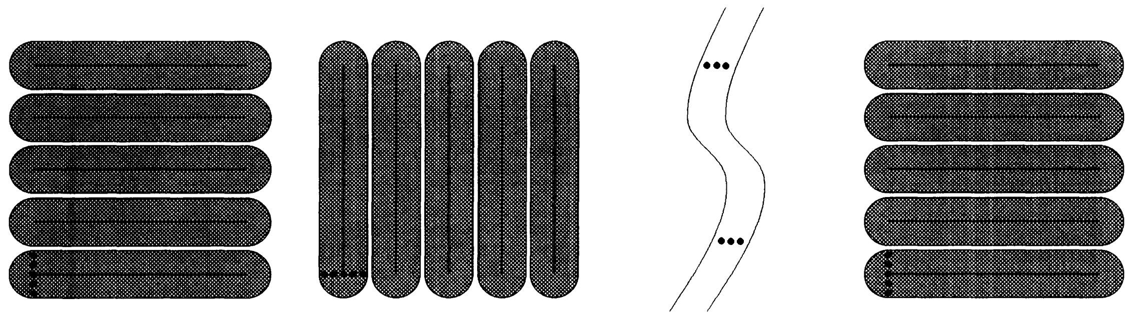

Figure 13: The sets A, A^ϵ and B for points and segments under rigid motion.

on the circular arc along which B_2 is placed, which must have a larger radius for a smaller $\theta_{\min}$ and thus depends on $\theta_{\min}$; the parameters are thus interdependent. However, as A_2 and B_2 are moved farther out, l and w approach limit values, as the circular arc becomes closer to a straight line. Thus, we can initially place A_2 and B_2 where l and w are within some small factor (say, within 1%) 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, and thus determine $\theta_{\min}$. This will give us a minimum value for the radius of the circular arc, and we know that we can move A_2 and B_2 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 will refer to the sets A and B constructed in Subsection 3 as A_T and B_T. As above, we observe that the set B_T 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_2 identical to the left-hand group of A_T, as shown in Figure 9. A_2 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 $A = A_T \cup A_2$ We also add n points to B_T in a vertical row, in the same relative position to A_2 as the first vertical row of B_T was to the left-hand group of A. Call this new row B_2 and let $B = B_T \cup B_2$. A and B are shown in Figure 13.

Now, any translation t for which $B_T \oplus t \in A_T^\epsilon$ will also have $B \oplus t \in A^\epsilon$. Fix such a t and consider values of θ where $|\theta| < \theta_{\min}$. As θ changes through this range, the points in B_2 will sweep across the gaps in A_2^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_2 achieve $\Omega(n^2)$ different configurations with respect to the gaps of A_2^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 $r_\theta(B) \oplus t \subset A^\epsilon$. We can choose t to represent one of the $\Omega(n^4)$ distinct configurations of B with respect to A, and so this gives $\Omega(n^6)$ different configurations of B with respect to the gaps of A^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 Point sets under translation and scale

In this section, we consider the complexity of the graph of the Hausdorff distance between two sets of points as one set is translated and scaled with respect to the other. This scaling is with respect to fixed x and y axes. This transformation group is of interest because it has been used in image-recognition tasks, as described in [6]. It corresponds to the transformations that the image of an object may undergo under the weak perspective projection model as a camera is moved forward or panned about a vertical axis.

We will present three different lower-bound constructions: one each for the L_1, L_2 and L_∞ norms. They are all similar in concept, and are based on the constructions in Section 2, using techniques from Section 4. All of these constructions show lower bounds for both the directed and undirected Hausdorff distance. The central idea here is that changing s_x slightly has very little effect on points which are near the y axis, while it has a largely translational effect on groups of points which are located a large distance along the x axis ("large" here means that the distance of the group from the origin is large relative to the size of the group).

6.1 The L_2 example

This example is very similar to the example for point sets under rigid motion. Again, a key observation is that the construction of A_T and B_T in Subsection 2.2 can be perturbed slightly without affecting the topology of the graph. In this case, this perturbation will take the form of a slight scaling of B_T in x and y. Suppose that the origin is placed at the lower left corner of B_T. Then let $s_{\min}$ be the valid range of such scaling: if $1 - s_{\min} \leq s_x, s_y \leq 1 + s_{\min}$, then there are $\Omega(n^3)$ translational configurations of B_T with respect to A_T for which replacing B_T by $S_{s_x s_y}(B_T)$ does not change the configuration (i.e. the same points are on the same sides of the gap of A_T^ϵ). $s_{\min}$ clearly depends on the n, ϵ and σ used to construct A_T and B_T.

First, we construct A_2 as in Section 4, and place it so that the y axis cuts through the left-hand scalloped edge in a manner similar to that described in Section 4: each lobe of the scalloped edge contains a length l of the y axis, the smallest depth of any of the gaps is w, and the ratio between the axis length contained in any lobe (i.e. l) and the length contained in the space next to that lobe is greater than $8n : 1$. Note that this construction does not depend on the location of this copy of A_2 along the y axis.

We next construct A_T and B_T as in 2.2 such that the region of $\Omega(n^3)$ complexity occupies a region less than $\min(w/2, l/(8n))$ square, and position them so that the lower left corner of this region in translation space is located at $t = (0,0)$, and the lower left corner of B_T is at the origin. This determines a value for $s_{\min}$. We then position the copy of A_2 a distance of $\epsilon/s_{\min}$ above the origin, with the y axis cutting through it as described above. We also make a copy of A_2, rotated by $90°$, to the right of the copy of A_T, with the x axis cutting through the lower scalloped edge in the same manner. A consists of A_T together with these two copies of A_2. B then consists of B_T plus two rows of n points plus an extra point per row, as described in Subsection 4; one row and its extra point are positioned along the y axis inside the uppermost lobe of the first copy of A_2; the other row and its extra point are in the corresponding position on the x axis, in the second copy of A_2. Call these two rows B_2.

Now, it is possible to independently choose seven numbers $n_1, \ldots, n_7$ from 1 to $n-1$ and construct a translation t and scale s_x, s_y of B with respect to A where $d_S(t, s_x, s_y) \leq \epsilon$ in the following manner:

1. n_1, n_2 and n_3 determine the translation. They are used to position n_1 of the points of the lower row of B_T to the left of the gap in A_T^ϵ, n_2 of the points of the upper row to the left of the gap, and with the rows straddling the n_3th "wobble" of the gap.

2. n_4 and n_5 determine s_y. They are used to position n_4 of the points of the row of B_2 lying on the y axis below one of the gaps in the scalloped edge of the upper copy of A_2. n_5 selects the gap. Note that changing s_y acts as (essentially) a translation of this row, as its distance from the origin (and therefore its y coordinate) greatly exceeds its length. It is possible to do this no matter what translation was chosen above, since the range of translations is small. Also, any s_y chosen in this manner will not exceed the range determined by $s_{\min}$.

3. n_6 and n_7 similarly determine s_x, by positioning the points of the row of B_2 lying on the x axis with respect to its copy of A_2.

Now, if two such configurations are generated with different n_i values, then it will not be possible to move from one to the other without some point crossing a gap, and so they must belong to different connected components in transformation space where $d_S(t, s_x, s_y) \leq \epsilon$. Also, due to the "extra points" added to the various parts of B, $h(A, S_{s_x s_y}(B) \oplus t)$ is no greater than ϵ for all such configurations; D_S therefore has $\Omega(n^7)$ distinct local minima. These can occur in an arbitrarily small region of transformation space for a fixed ϵ.

6.2 The L_1 example

This example is quite similar to the example in the previous subsection. The construction uses a copy of A_T and B_T from Subsection 2.1 (the translational example for L_∞), rotated 45°. Instead of A_2 being a row of points generating a scalloped edge, it is instead a row of points generating a sawtooth edge; however, it is used in the same manner.

6.3 The L_∞ example

This example must be constructed differently from the previous two, since a vertical row of points, when dilated by ϵ, generates a straight vertical edge, with no irregularities that can be exploited. In this case, however, we simply build A and B by taking three identical copies of A_T and B_T from Subsection 2.1, placing one copy at the origin (copy 1), one far out on the x axis (copy 2), and one far out on the y axis (copy 3). Now, we choose a translational configuration of one of the copies of A_T with respect to its copy of B_T (the copies are identical so which copy is used as a base is irrelevant), and initially set s_x and s_y to 1. Now, adjusting s_x slightly moves only copy 2 of B_T significantly, and moves it in a largely translational manner in the direction of the x axis. As it moves, its points will cross the gap in copy 2 of A_T^ϵ. It can move both forward and backward, as s_x can be varied both slightly above and slightly below 1. There are therefore $\Omega(n^2)$ configurations of copy 2 of B_T with respect to copy 2 of A_T^ϵ, once the overall translational configuration has been chosen.

We similarly choose one of $\Omega(n^2)$ configurations for copy 3 of A_T and B_T, by adjusting s_y. The distance of copies 2 and 3 of A_T and B_T from the origin can be made as large as desired, to reduce the non-translational effects of changing s_x and s_y until they are small enough that they have no effect on the topology of the graph of D_S. This construction method also works for the L_1 and L_2 norms.

7 Sets of points and line segments under translation and scale

This example is quite similar to the previous example; we make a copy of A_T and B_T from Section 3, shown in Figure 9, placed near the origin, together with a copy of the right-hand group of each placed some distance away along the x axis, and a copy of the left-hand group of each placed some distance away along the y axis. Similar observations give $\Omega(n^8)$ distinct local minima for d_S. Note that, as before, the exact norm being used is not relevant, since it does not affect the gaps between the segments.

8 Point sets under affine transformation

Here we are dealing with transformations which map B to $M(B) \oplus t$, where M is a 2×2 matrix defined by

$$M = \begin{bmatrix} m_{00} & m_{01} \\ m_{10} & m_{11} \end{bmatrix}$$

and $t = (t_x, t_y)$ is a translation. In other words, each point $(b_x, b_y) \in B$ is mapped to $(m_{00}b_x + m_{01}b_y + t_x, m_{10}b_x + m_{11}b_y + t_y)$.

The key observation here is that if B consists of three groups of points, one near the origin, one located a large distance along the x axis, and one located a large distance along the y axis, then

- Changing m_{00} slightly causes the second group in the transformed B to translate in x, but has little other effect.
- Changing m_{10} slightly causes the second group to translate in y, but has little other effect.
- Changing m_{01} or m_{11} slightly similarly causes the third group to translate in x or y, and has little other effect.
- If M is sufficiently close to the identity matrix, then $M(B)$ will essentially be the same as B, with the relative positions of the three groups shifted around somewhat.

Further, the magnitude of these translational motions with respect to their other effects can be increased by moving the corresponding group farther away from the origin along the appropriate axis. Thus, we have a six-parameter system (two translational and four linear parameters), which can be decomposed into three two-parameter translational systems, plus a small amount of "slop", which can be made as small as required. In order to build a $\Omega(n^9)$ example, we simply take three copies of the appropriate A_T and B_T, and arrange them as described above. Again, we can move copies 2 and 3 of A_T and B_T out from the origin until all non-translational effects are not significant, since as they move farther out, the amounts by which the m_{ij} values need to be adjusted are reduced. Since all the copies of A_T and B_T have $\Omega(n^3)$ complexity under translation, for both the directed and undirected Hausdorff distance, and are essentially independent, this construction gives $\Omega(n^9)$ complexity for both d_A and D_A.

9 Sets of points and line segments under affine transformation

This example is constructed identically to the previous example. We simply take three copies of A_T and B_T from Section 3 and position them as before. This gives $\Omega(n^{12})$ local minima of d_A, from three essentially independent $\Omega(n^4)$ translational examples.

10 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 that of determining such bounds. For example, 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 in less time than that given by the complexity of the graph of the Hausdorff distance function?

11 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*, pages 151–156, Waterloo, Ontario, 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] D.P. Huttenlocher and W.J. Rucklidge. A multi-resolution technique for comparing images using the Hausdorff distance. In *Proc. Computer Vision and Pattern Recognition*, pages 705–706, New York, NY, 1993.

[7] 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.