

Lower Bounds for the Complexity of the Graph of the Hausdorff Distance as a Function of Transformation*

W. J. Rucklidge

Xerox Palo Alto Research Center, 3333 Coyote Hill Road.
Palo Alto, CA 94304, USA
rucklidge@parc.xerox.com

Abstract. The Hausdorff distance is a measure defined between two sets in some metric space. This paper investigates how the Hausdorff distance changes as one set is transformed by some transformation group. Algorithms to find the minimum distance as one set is transformed have been described, but few lower bounds are known. We consider the complexity of the graph of the Hausdorff distance as a function of transformation, and exhibit some constructions that give lower bounds for this complexity. We exhibit lower-bound constructions for both sets of points in the plane, and sets of points and line segments; we consider the graph of the directed Hausdorff distance under translation, rigid motion, translation and scaling, and affine transformation. Many of the results can also be extended to the undirected Hausdorff distance. These lower bounds are for the complexity of the graph of the Hausdorff distance, and thus do not necessarily bound algorithms that search this graph; however, they 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). This paper examines the

* This work was supported in part by National Science Foundation PYI Grant IRI-9057928 and matching funds from General Electric, Kodak, and Xerox, and in part by Air Force Contract AFOSR-91-0328.

case where A and B are compact subsets of the plane consisting of either a finite number of points, or a finite number of points and nonintersecting 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 . $h(A, B)$ is small exactly when every point in A is close to some point in B ; $h(B, A)$ is small when every point in B is close to some point in A , and $H(A, B)$ is small when both of these are true. In particular, $h(B, A) \leq \varepsilon$ exactly when for any $b \in B$ there is some $a \in A$ such that $\|a - b\| \leq \varepsilon$. Let $A^\varepsilon = A \oplus D(\varepsilon)$ where \oplus is the Minkowski sum, and $D(\varepsilon)$ is the closed disk of radius ε (the set of all points x such that $\|x\| \leq \varepsilon$). A key observation [1], [4] is that, for compact sets such as we are considering, $h(B, A) \leq \varepsilon$ iff $B \subseteq A^\varepsilon$.

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

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

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 is the one bringing B into closest correspondence with A . Many approaches to determining this minimizing 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 [5], [4]. We also consider the graph of the function

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

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

We exhibit lower bounds for the complexity of such graphs as follows. Each construction is parametrized by two values, ε and n (and possibly other parameters). For each one, we fix some values for ε , n , and any other parameters, and construct sets A and B having kn elements each, for some constant k depending on the problem. We then show that the set $\{g \mid f_G(g) \leq \varepsilon\}$ 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 $f_G(g)$, this shows that there are $\Omega(n^l)$ local minima in the graph of $f_G(g)$. In some cases, we also show that the graph of $F_G(g)$ may have this complexity. Previous constructions, such as those in [5] and [4], have been for the directed Hausdorff distance alone.

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 $F_G(g)$ 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 $F_G(g) \leq \varepsilon$, 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, we might be able to obtain efficient algorithms: if the global minimum of $F_G(g)$ is restricted to lie in a small area, then searching only this

Table 1. Lower bounds for the complexity of the Hausdorff distance between two sets of size n . The only previously known lower bounds were those by Chew and Kedem [4] for the directed distance under translation (for both sets of points, and sets of points and line segments).

Transformation group	Point sets results for $f_G(t)$ and $F_G(t)$ except ‡ for $f_G(t)$ only			Points and segments results for $f_G(t)$ only	
	L_1	L_2	L_∞	Any L_p	L_2
Translation	$\Omega(n^3)$	$\Omega(n^3)$	$\Omega(n^3)$	$\Omega(n^4)$	
Rigid motion		$\Omega(n^5)$			$\Omega(n^6)$
Translation and x, y scale	$\Omega(n^7)$	$\Omega(n^7)$	$\Omega(n^7)^\ddagger$	$\Omega(n^8)$	
Affine transformation	$\Omega(n^9)$	$\Omega(n^9)$	$\Omega(n^9)$	$\Omega(n^{12})$	

area of its graph can take less time than searching the entire graph. The constructions here show that, in some cases, this is not possible.

In this paper we deal with four transformation groups: the group Tr of translations, the group Rig of rigid motions (translations and rotations), the group Sc of translations and (x, y) scaling, and the group Aff of nondegenerate affine transformations. 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 emphasized that the lower bounds presented here are for the complexity of the graph of the Hausdorff distance, and do not necessarily give lower bounds for algorithms that determine the optimal transformation.

2. Point Sets Under Translation

Let t be a translation. Define

$$F_{Tr}(t) = H(A, B \oplus t), \quad (1)$$

$$f_{Tr}(t) = h(B \oplus t, A) \quad (2)$$

as the (undirected and directed, respectively) Hausdorff distance between A and B as a function of translation. This section describes two constructions of point sets A and

Table 2. Time bounds for algorithms that find the exact minimum Hausdorff distance under transformation between two sets of size n .

Problem	Point sets		Points and segments	
	L_1, L_∞	L_2	L_1, L_∞	L_2
Translation	$O(n^2 \log^2 n)$ [4]	$O(n^3 \log n)$ [5]	$O(n^4 \alpha(n))$ [5]	$O(n^4 \log^3 n)$ [1]
Rigid motion		$O(n^5 \log^2 n)$ [3]		$O(n^6 \log^2 n)$ [3]

B , each containing $O(n)$ points, for which $F_{Tr}(t)$ and $f_{Tr}(t)$ have $\Omega(n^3)$ 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 $2\varepsilon + \delta$ apart, where $\delta < \sigma/n$. A and A^ε are shown in Fig. 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 inward, thereby making the stairsteps (and thus the total length of the staircase) as small as desired.

Let B consist of two diagonal rows of points, each of n points, as shown in Fig. 2. 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




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




Fig. 2. The sets A^ϵ and B for the L_∞ lower bound for point sets under translation.

part. Since the staircase lengths are all σ , and $\delta < \sigma/n$, each row of B is shorter (in height and width) than a staircase segment.

Consider translating B slightly upward or downward. The points around the vertical staircase remain inside A^ϵ or outside it, as they were before, but the points around the horizontal staircase move into and out of A^ϵ as B moves. Similarly, as B moves left or right, the points around the vertical staircase move in and out of A^ϵ , but the points around the horizontal staircase do not. We can thus see $\Omega(n^2)$ different configurations of B with respect to this one staircase, 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 staircases, each of which gives rise to $\Omega(n^2)$ configurations, for a total of $\Omega(n^3)$ configurations. We only consider configurations where there is at least one point of each row of B on either side of the gap. Each one of these configurations can be labeled with three numbers from 1 to $n - 1$: the number of points in the bottom row of B that are inside the upper-left component of A^ϵ (i.e., to the left of the gap), the number of points in the top row of B that are to the left of the gap, and the number of the staircase which is straddled by B . $\Omega(n^3)$ such labels are possible. Suppose t_1 and t_2 are translations representing configurations with distinct labels. Then $f_{Tr}(t_1) \leq \epsilon$ and $f_{Tr}(t_2) \leq \epsilon$, since each translation places B entirely inside A^ϵ , but any path from t_1 to t_2 must pass through a translation t where $f_{Tr}(t) > \epsilon$: either some point in one of the rows of B must cross the gap, in which case $f_{Tr}(t) > \epsilon$ when t is a translation placing that point inside the gap, or B must be moved so as to straddle another staircase, 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 visualize 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 $S(A, \epsilon, B) = \bigcap_{b \in B} S(A, \epsilon, b)$. Then $t \in S(A, \epsilon, B)$ iff $B \oplus t \subseteq A^\epsilon$, or $f_{Tr}(t) \leq \epsilon$; $S(A, \epsilon, B)$ is therefore the set of all translations t which make $h(B \oplus t, A) \leq \epsilon$.

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




Fig. 3. Part of $S(A, \varepsilon, B)$ for the L_∞ lower bound for point sets under translation.

union. This union has a hole for every connected component of $S(A, \varepsilon, B)$. Figure 3 shows part of such a union. Each jagged line represents one segment of 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 . These two sets intersect in $\Omega(n)$ crosshatches, each having $\Omega(n^2)$ holes; Fig. 3 shows two of the crosshatches.

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.

The area where the undirected Hausdorff distance $F_{Tr}(t)$ is no greater than ε can also have large complexity in a small space. Set $\sigma < \varepsilon/n$ so 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 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, $H(A, B \oplus t) > \varepsilon$ exactly where $h(B \oplus t, A) > \varepsilon$ (at least in this region of interest), since $h(A, B \oplus t)$ is always at most ε . The undirected Hausdorff distance $F_{Tr}(t)$ therefore has complexity $\Omega(n^3)$.

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. The set A consists of two vertical rows of n points, spaced σ apart; the two rows are staggered by $\sigma/2$ (see Fig. 4). The distance between the rows is set such




Fig. 4. The sets A and A^ϵ for the L_2 lower bound for point sets under translation.

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.

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 Fig. 5. 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 $B \subseteq A^\epsilon$, 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 gives $\Omega(n^2)$ possible configurations of B around a single “wobble” in the gap; as there are $\Omega(n)$ such wobbles, there are $\Omega(n^3)$ different configurations of B with $B \subseteq A^\epsilon$. A labeling argument, similar to that in the previous subsection, shows that these configurations are all distinct.

This is difficult to visualize, so again we look at $S(A, \epsilon, B)$. We construct this, as before, by taking the union of $O(n)$ copies of the gap, translated by various amounts, and showing that this union has $\Omega(n^3)$ disjoint holes.

Since the actual gap has such a complicated shape, we deal only with a small part of it. In particular, we 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 centered around a place where the gap is at its narrowest. We bound each such region by a rectangle. These rectangles are 2δ wide by λ long, where




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

λ is determined by ε and δ , and is equal to $\sqrt{4\varepsilon\delta - \delta^2}$. See Fig. 6 for an illustration of this region. Now, note that $\lambda/\delta = \sqrt{4\varepsilon/\delta - 1}$. Thus, for any fixed ε , we can make λ/δ as large as we like by making δ small enough: as δ decreases, the rectangles get both narrower and shorter, but their length-to-width ratio increases.

The gap is narrowest exactly where a line from one point in the left row of A to one of its neighbors 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 them slightly more than 2δ apart (the same spacing as the points of B), so that the right edge of one rectangle almost touches the left edge of the next. Then, for a small enough value of δ , λ/δ is large enough that one left-leaning rectangle can intersect all n of these right-leaning rectangles.

In constructing $S(A, \varepsilon, B)$, we make n copies of the gap stacked slightly more than 2δ apart (corresponding to one of the rows of B), and have these intersect with another n copies, shifted down by $\sigma/2$ (corresponding to the other row), giving $\Omega(n^3)$ intersections: each left-leaning rectangle from one of the copies intersects n right-leaning rectangles from other copies, and vice versa. Figure 7 shows part of such an arrangement. There are $\Omega(n^2)$ holes in the crosshatch, and $\Omega(n)$ such crosshatches in the complete arrangement. As in Section 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 involves some point crossing the gap.




Fig. 6. A closeup of the interesting region of the gap.



Fig. 7. Part of $S(A, \varepsilon, B)$ for the L_2 lower bound for point sets under translation.

Also, since the points of B are spaced about 2δ apart and the gap becomes wider than this between two adjacent wobbles, 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 $S(A, \epsilon, B)$. These components are all contained in an area 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 Section 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 $2n$ points and nonintersecting line segments, for which the graph of the directed Hausdorff distance as a function of translation $f_{Tr}(t) = h(B \oplus t, A)$ has $\Omega(n^4)$ complexity.

Fix ϵ and n and pick δ such that $\delta < \epsilon/n$. 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. Under any L_p norm, A^ϵ then consists of n horizontal bars and n vertical bars, with gaps of width δ between adjacent bars; the shape of the caps on the ends of the bars depends on the exact norm. 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 8 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 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




Fig. 8. The sets A , A^ϵ , and B for points and segments under translation.

configurations with respect to the vertical segments of A . 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 .

4. Point Sets Under Rigid Motion

Let t be a translation and let θ be an angle. Define

$$F_{Rig}(t, \theta) = H(A, r_\theta(B) \oplus t), \quad (3)$$

$$f_{Rig}(t, \theta) = h(r_\theta(B) \oplus t, A), \quad (4)$$

where $r_\theta(B)$ denotes the set obtained by rotating B by θ counterclockwise about the origin. We 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 in (t, θ) space where the undirected Hausdorff distance $F_{Rig}(t, \theta)$ between two sets of $\Omega(n)$ points is less than ε . It is based on an augmentation of A and B from Section 2.2. For clarity, we refer to the sets A and B constructed for the translational lower bounds as A_{Tr} and B_{Tr} .

First, note that it is possible to rotate B_{Tr} from that construction by some very small angle θ_{\min} about its centroid while still maintaining the $\Omega(n^3)$ complexity of $h(B_{Tr} \oplus t, A_{Tr})$. This is because there must be, in the $\Omega(n^3)$ arrangement of connected components of $S(A, \varepsilon, B)$, some minimum distance between features, and so any rotation that 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_{Tr} consists of n points along a vertical line, spaced less than $\varepsilon/(2n)$ apart. Placing a disk of radius ε about each gives a shape with two vertical scalloped edges: the left and right sides are close to vertical, but have n slight bulges or lobes. Call this row of points A_2 .

Now, if A_2 is located sufficiently far from the origin, and perpendicular to the line joining it to the origin, then it is possible to pass a circular arc (centered at the origin) through the inner scalloped edge so that it passes through each of the n lobes and the gaps between them. The arc does not pass through these lobes evenly, but cuts 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, 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 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 (see Fig. 9). As the arc becomes closer to



Fig. 9. The interaction between A_2^ϵ and B_2 .

a straight line (as A_2 is moved away from the origin), 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 origin, this row of points moves 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 fits into a single lobe, only one point passes through a gap at a time, and there are $\Omega(n^2)$ 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 , the amount by which B_2 is rotated about the origin, 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 are 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 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" inside A_2^ϵ achieved for any point of B_2 at any point in this translational range. Figure 10 shows this situation; the vertical bars are $l/(8n)$ high ($l/16n$ above and below the center). Let d_{\min} be the smallest such d value achieved for any of the $\Omega(n^2)$ possible such straddling configurations. Also, let w be the minimum distance between the circular arc and the bottom of the space between two lobes.

We now construct A_{T_r} and B_{T_r} as in Section 2.2 by choosing σ and δ so that the $\Omega(n^3)$ complexity region of $F_{T_r}(t)$ is at most $l/(16n)$ high by $\min(d_{\min}/2, w/2)$ wide, and so that the centroid of B_{T_r} is at the origin. Let $A = A_{T_r} \cup A_2$ and $B = B_{T_r} \cup B_2$.




Fig. 10. The minimum depth of the points of B_2 .

Let θ produce one of the straddling configurations described above. Then, if $|\theta| < \theta_{\min}$, there are $\Omega(n^3)$ different connected components in t space where $F_{Rig}(t, \theta) \leq \varepsilon$. This is because a slight translation of B_2 with respect to A_2 does not move any of the points of B_2 out of A_2^ε , and B_{T_r} and A_{T_r} have been constructed so that the range of translations required is very small. Note that $F_{Rig}(t, \theta)$ is determined by $f_{Rig}(t, \theta)$ for all transformations in the range under consideration, as the directed distance from A to the rotated and translated B is always less than ε .

A labeling argument similar to that in Section 2.1 now shows that there are $\Omega(n^5)$ different connected components in (t, θ) space where $F_{Rig}(t, \theta) \leq \varepsilon$, $\Omega(n^3)$ 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_r} 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.

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 , d_{\min} , and w , which depend on the circular arc along which B_2 is placed, which must have a larger radius for a smaller θ_{\min} and thus depends on θ_{\min} ; the parameters are thus interdependent. However, as A_2 and B_2 are moved farther out, l , d_{\min} , 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 θ_{\min} . This gives 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.




Fig. 11. The sets A , A^ϵ , and B for points and segments under rigid motion.

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 refer to the sets A and B constructed in Section 3 as A_{T_r} and B_{T_r} . As before, we observe that the set B_{T_r} may be rotated by some small angle θ_{\min} about its centroid without changing the topology of the arrangement.

Place B_{T_r} so that its centroid is at the origin (the center of rotation). Augment A_{T_r} by a group of segments A_2 identical to the left-hand group of A_{T_r} . A_2 is placed so that it subtends a total angle of less than θ_{\min} to the origin, and lies directly to the right of it. Let $A = A_{T_r} \cup A_2$. We also add n points to B_{T_r} in a vertical row, in the same relative position to A_2 as the first vertical row of B_{T_r} was to the left-hand group of A . Call this new row B_2 and let $B = B_{T_r} \cup B_2$. A , A^ϵ , and B are shown in Fig. 11.

Now, any translation t for which $B_{T_r} \oplus t \in A_2^\epsilon$, also has $B \oplus t \in A^\epsilon$. Fix such a t and consider values of θ where $|\theta| < \theta_{\min}$. As θ changes through this range, the points in B_2 sweep across the gaps in A_2^ϵ . Their spacing is such that only one point crosses 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 $f_{Rig}(t, \theta) \leq \epsilon$, since any rotation in this range keeps $r_\theta(B) \oplus t$ inside A^ϵ . 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 $f_{Rig}(t, \theta) \leq \epsilon$. 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 as a camera is moved forward or panned about a vertical axis, under the weak perspective projection model.

We 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. The L_1 and L_2 constructions show lower bounds for both the directed and undirected Hausdorff distance; the L_∞ construction, however, shows a lower bound for the directed Hausdorff distance only.

Define the undirected and directed Hausdorff distance as functions of a translation t and two scale parameters s_x and s_y as

$$F_{Sc}(t, s_x, s_y) = H(A, S_{s_x s_y}(B) \oplus t), \tag{5}$$

$$f_{Sc}(t, s_x, s_y) = h(S_{s_x s_y}(B) \oplus t, A), \tag{6}$$

where $S_{s_x s_y}(B)$ 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).

The central idea behind these constructions 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_r} and B_{T_r} in Section 2.2 can be perturbed slightly without affecting the topology of the graph. In this case, this perturbation takes the form of a slight scaling of B_{T_r} in x and y . Suppose that the origin is placed at the lower-left corner of B_{T_r} . 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_r} with respect to A_{T_r} for which replacing B_{T_r} by $S_{s_x s_y}(B_{T_r})$ does not change the configuration (i.e., the same points are on the same sides of the gap of $A_{T_r}^\epsilon$). s_{\min} clearly depends on the n , ϵ , and σ used to construct A_{T_r} and B_{T_r} .

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_r} and B_{T_r} as in Section 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_r} is at the origin. This determines a value for s_{\min} . We then position the copy of A_2 a distance of $2\epsilon/s_{\min}$ 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_r} , with the x axis cutting through the lower scalloped edge in the same manner. A consists of A_{T_r} together with these two copies of A_2 . B then consists of B_{T_r} plus two rows of n points plus an extra point per row, as described in Section 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 choose independently seven numbers n_1, \dots, 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 $f_{Sc}(t, s_x, s_y) \leq \varepsilon$ 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_1} to the left of the gap in $A_{T_1}^\varepsilon$, n_2 of the points of the upper row to the left of the gap, and with the rows straddling the n_3 th 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 does 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 .

If two such configurations are generated with different n_i values, then it is not 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 $f_{Sc}(t, s_x, s_y) \leq \varepsilon$. 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; F_{Sc} 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_1} and B_{T_1} from Section 2.1 (the translational example for L_∞), rotated 45° . Instead of A_{T_1} being augmented by two rows of points, each generating a scalloped edge, it is instead augmented by two rows of points, each generating a sawtooth edge; however, these augmentations are 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. Let A_2 be a vertical row of n points spaced $2\varepsilon + \delta$ apart, lying on the y axis, far away from the origin. A_2^ε is then a row of n squares with a δ -wide gap between adjacent squares. Let B_2 be a vertical row of n points spaced just over δ apart, lying close to A_2 . As s_y varies, B_2 moves mostly translationally, and its points therefore move through the $n - 1$ gaps in A_2^ε , one point at a time. $\Omega(n^2)$ configurations of the points of B_2 relative to the gaps of A_2^ε are therefore possible.

Let A be a copy of A_{T_1} from Section 2.1, plus A_2 as described above, plus a similar horizontal row of n points lying along the x axis. Similarly, let B be a copy of B_{T_1} plus B_2 plus a row of points on the x axis. We first choose a translational configuration

of B_{T_r} with respect to A_{T_r} , and initially set s_x and s_y to 1. We then choose one of the $\Omega(n^2)$ configurations of B_2 with respect to A_2 , and adjust s_y so that this configuration is achieved. We similarly choose one of the $\Omega(n^2)$ configurations of the parts of A and B lying along the x axis, and set s_x accordingly. Thus, there are $\Omega(n^7)$ different configurations of B with respect to A ; $f_{Sc}(t)$ therefore has $\Omega(n^7)$ distinct local minima.

Unlike the other constructions for point sets, it is not possible to augment B by a few points so that $h(A, t(B))$ is always below ε for the region of transformation space of interest. This construction is therefore valid for the directed distance only.

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_r} and B_{T_r} from Section 3, shown in Fig. 8; these copies are placed near the origin, together with a copy of the right-hand group of each (the group exploiting vertical gaps) 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. This construction gives $\Omega(n^8)$ distinct local minima for f_{Sc} ($\Omega(n^4)$ from A_{T_r} and B_{T_r} , and $\Omega(n^2)$ from each of the additional groups). 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 nonsingular 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 $(m_{00}b_x + m_{01}b_y + t_x, m_{10}b_x + m_{11}b_y + t_y)$. As before, we define

$$F_{Aff}(t, M) = H(A, M(B) \oplus t), \quad (7)$$

$$f_{Aff}(t, M) = h(M(B) \oplus t, A). \quad (8)$$

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 $M(B)$ is essentially the same as B , with the relative positions of the three groups shifted around somewhat.

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 an $\Omega(n^9)$ example, we simply take three copies of the appropriate A_{Tr} and B_{Tr} , and arrange them as described above. Again, we can move copies 2 and 3 of A_{Tr} and B_{Tr} out from the origin until all nontranslational 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_{Tr} and B_{Tr} 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 f_{Aff} and F_{Aff} .

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_{Tr} and B_{Tr} from Section 3 and position one at the origin, one out along the x axis, and one out along the y axis. This gives $\Omega(n^{12})$ local minima of f_{Aff} , from three essentially independent $\Omega(n^4)$ translational examples.

10. Conclusion

We have presented constructions which give lower bounds on the complexity of the directed (and, in many cases, the undirected) Hausdorff distance in several different contexts, summarized in Table 1. In a number of of the cases, we have shown that the directed and undirected Hausdorff distances can have large complexity in a small space. The large exponents in these bounds imply that, to solve these problems, either one will have to approximate the Hausdorff distance, or come up with algorithms that can somehow get around them.

We have demonstrated the complexity of the Hausdorff distance as a function of transformation by constructing cases where it has a large number of local minima. A related question is the number of *global* minima that it can have. In some cases, such as the constructions for point sets under translation with the L_1 or L_∞ norms, it is possible to make the function have equal values at all of the local minima that we have constructed, so that they are all (equal) global minima; in other cases, the number of global minima is an open question.

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, plus the problem of point sets under translation and scale with the L_∞ norm. A remaining open problem is that of determining bounds for these cases. For example, can the *undirected* Hausdorff distance under translation between sets of points and segments 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?

Acknowledgments

The $\Omega(n^4)$ example shown in Section 3 is due to Paul Chew and Klara Kedem.

References

1. P. K. Agarwal, M. Sharir, and S. Toledo. Applications of parametric searching in geometric optimization. *Proc. Third ACM-SIAM Symposium on Discrete Algorithms*, pages 72–82, 1992.
2. H. Alt, B. Behrends, and J. Blömer. Approximate matching of polygonal shapes. *Discrete and Computational Geometry*, 9:267–291, 1993.
3. L. P. Chew, M. T. Goodrich, D. P. Huttenlocher, K. Kedem, J. M. Kleinberg, and D. Kravets. Geometric pattern matching under Euclidean motion. *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*, pages 318–325. Lecture Notes in Computer Science, Vol. 621, Springer-Verlag, Berlin, 1992.
5. D. P. Huttenlocher, K. Kedem, and M. Sharir. The upper envelope of Voronoi surfaces and its applications. *Discrete and Computational Geometry*, 9(3):267–291, 1993.
6. D. P. Huttenlocher and W. J. Rucklidge. A multi-resolution technique for comparing images using the Hausdorff distance. *Proc. Computer Vision and Pattern Recognition Conference*, pages 705–706, New York, 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.

Received August 12, 1994, and in revised form November 5, 1995.