Between Shapes, Using the Hausdorff Distance

Marc van Kreveld
Department of Information and Computing Sciences, Utrecht University, The Netherlands
m.j.vankreveld@uu.nl

Tillmann Miltzow
Department of Information and Computing Sciences, Utrecht University, The Netherlands
t.miltzow@uu.nl

Tim Ophelders
Department of Computational Mathematics, Science and Engineering, Michigan State University, East Lansing, MI, USA
ophelder@egr.msu.edu

Willem Sonke
Department of Mathematics and Computer Science, TU Eindhoven, The Netherlands
w.m.sonke@tue.nl

Jordi L. Vermeulen
Department of Information and Computing Sciences, Utrecht University, The Netherlands
j.l.vermeulen@uu.nl

Abstract

Given two shapes A and B in the plane with Hausdorff distance 1, is there a shape S with Hausdorff distance $1/2$ to and from A and B? The answer is always yes, and depending on convexity of A and/or B, S may be convex, connected, or disconnected. We show a generalization of this result on Hausdorff distances and middle shapes, and show some related properties. We also show that a generalization of such middle shapes implies a morph with a bounded rate of change. Finally, we explore a generalization of the concept of a Hausdorff middle to more than two sets and show how to approximate or compute it.

2012 ACM Subject Classification Theory of computation → Computational geometry

Keywords and phrases computational geometry, Hausdorff distance, shape interpolation

Digital Object Identifier 10.4230/LIPIcs.ISAAC.2020.13

Related Version A full version of the paper is available at `https://arxiv.org/abs/2009.14719`.

Funding Research on the topic of this paper was initiated at the 4th Workshop on Applied Geometric Algorithms (AGA 2018) in Langbroek, The Netherlands, supported by the Netherlands Organisation for Scientific Research (NWO) under project no. 639.023.208.
Marc van Kreveld: supported by the NWO TOP grant no. 612.001.651.
Tillmann Miltzow: supported by the NWO Veni grant EAGER.
Jordi L. Vermeulen: supported by the NWO TOP grant no. 612.001.651.

1 Introduction

For two sets A and B in $\mathbb{R}^2$, we define the *directed Hausdorff distance* as

$$d_{\vec{H}}(A,B) = \sup_{a\in A}\inf_{b\in B} d(a,b),$$

where d denotes the Euclidean distance. The *undirected Hausdorff distance* is defined as

$$d_H(A,B) = \max(d_{\vec{H}}(A,B), d_{\vec{H}}(B,A)).$$

Figure 1 Hausdorff morphs between three shapes.

If A and B are closed sets then $d_H(A, B) = r$ is equivalent to saying that r is the smallest value such that $A \subseteq B \oplus D_r$ and $B \subseteq A \oplus D_r$, where $\oplus$ denotes the Minkowski sum, and D_r is a disk of radius r centered at the origin. Recall that the Minkowski sum of sets A and B is the set $\{a + b \mid a \in A,\ b \in B\}$. As we will use the containment property throughout the paper, we will only consider closed sets.

The Hausdorff distance has been widely used in computer vision [14] and computer graphics [6, 11] for tasks such as template matching, and error computation between a model and its simplification. At the same time, the Hausdorff distance is a classic mathematical concept. Our research motivation is to study this profound concept from a new perspective. Algorithms to compute the Hausdorff distance between two given sets are available for many types of sets, such as points, line segments, polylines, polygons, and simplices in k-dimensional Euclidean space [3, 4, 7]. However, the question whether a polynomial time algorithm exists to compute the Hausdorff distance between general semialgebraic sets remains open [13].

In this paper, we consider the natural problem of finding a set that lies "between" two or more input sets, in a Hausdorff sense. In Section 2 we investigate the Hausdorff middle of sets A and B; this is a set that has minimum undirected Hausdorff distance to A and B. Differently put, it minimizes the maximum of four directed Hausdorff distances. We show that when the Hausdorff distance between A and B is assumed to be 1, there is always a Hausdorff middle that has Hausdorff distance $1/2$ to A and B, and this is the best possible. We relate the convexity of A and/or B to the convexity and connectedness of the Hausdorff middle, and study its combinatorial complexity.

We actually treat the middle more generally, by defining a class of sets that smoothly interpolate between A and B, giving a morph between them. Figure 1 shows two examples of such morphs. We prove that this morph has a bounded rate of change. Our approach does not require any correspondence between features of the input to be calculated. However, our approach is unusual in the sense that the intermediate shapes when morphing between e.g. two polygons are not polygons themselves. Most morphing algorithms typically interpolate only the boundary, and keep all intermediate shapes polygonal [5, 17, 8, 9, 10, 18].

In Section 3 we extend to Hausdorff middles of more than two sets and generalize several results. We assume that the maximum Hausdorff distance over all pairs of input sets is 1 and examine the smallest Hausdorff distance for a middle set. That is, given sets $\mathcal{M} = \{A_1, \ldots, A_k\}$, we are interested in the value $\alpha(\mathcal{M}) = \min_S \max_{i=1,\ldots,k} d_H(A_i, S)$. This value $\alpha(\mathcal{M})$ is no longer $1/2$, but depends on the input. For convex sets, we show that a value ≈ 0.608 can always be realized and is sometimes necessary, whereas for non-convex sets it can be as bad as 1. For a given set of polygons with total complexity n, we show that $\alpha(\mathcal{M})$ and the Hausdorff middle can be computed in $O(n^6)$ time, and, for any constant $\varepsilon > 0$, $(1+\varepsilon)$-approximated in $O(n^2 \log^2 n \log 1/\varepsilon)$. We note that other interpolation methods between two shapes do not have a natural generalization to a middle of three or more sets.

Our proofs use three types of arguments. First, many of our arguments rely on simple manipulations of the formal definition of the Hausdorff distance. Those arguments immediately generalize to other normed vector spaces, for instance, $(\mathbb{R}^d, \|\cdot\|_\infty)$ or the continuous functions

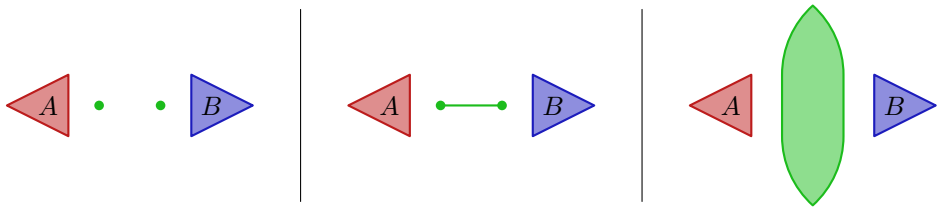

Figure 2 Three possible Hausdorff middles of A and B: two points, a line segment, and $S_{1/2}$.

endowed with the uniform norm topology $C(\mathbb{R})$. We do not state those generalizations explicitly, as this is not our focus. The second type of argument is of a topological nature. Using continuity and connectivity, we infer related properties to the output, by constructing topological structures or conclude that they cannot exist. The third type of argument uses 2-dimensional Euclidean geometry directly. We construct features, like vertices, edges and circular arcs and argue about their existence, and give distance bounds. These arguments are often intricate and do not generalize. They are of particular value, as the two-dimensional Euclidean plane is often the most interesting case in computational geometry.

2 The Hausdorff middle of two sets

Consider two closed compact sets A and B in $\mathbb{R}^2$; we are interested in computing a *Hausdorff middle*: a set C that minimizes the maximum of the undirected Hausdorff distances to A and B. That is,

$$C = \operatorname*{argmin}_{C'} \max(d_H(A, C'), d_H(B, C')).$$

Note that there may be many such sets that minimize the Hausdorff distance; see Figure 2 for a few examples. It might seem intuitive to restrict C to be the minimal set that achieves this distance, but this is ill-defined: the minimal set is not necessarily unique, and the common intersection of all minimal sets is not a solution itself (see Figure 3). However, the maximal set is well-defined and unique. Let $d_H(A, B) = 1$. Then

$$S = (A \oplus D_{1/2}) \cap (B \oplus D_{1/2})$$

is the maximal set with Hausdorff distance $1/2$ to A and B (we prove this below in Lemma 2). We want to show that $d_H(A, S) \leq 1/2$ and $d_H(B, S) \leq 1/2$. In fact, we can prove a more general statement.

We define $S_\alpha := (A \oplus D_\alpha) \cap (B \oplus D_{1-\alpha})$ for $\alpha \in [0, 1]$. We also use $\text{seg}(a, b)$ to denote the line segment connecting points a and b.

▶ **Theorem 1.** *Let A and B be two compact sets in the plane with $d_H(A, B) = 1$. Then $d_H(A, S_\alpha) = \alpha$ and $d_H(B, S_\alpha) = 1 - \alpha$.*

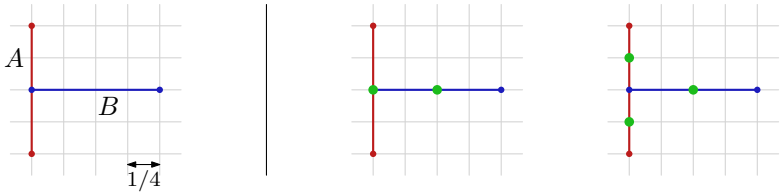

Figure 3 Two different minimal sets achieving minimal Hausdorff distance to A and B.

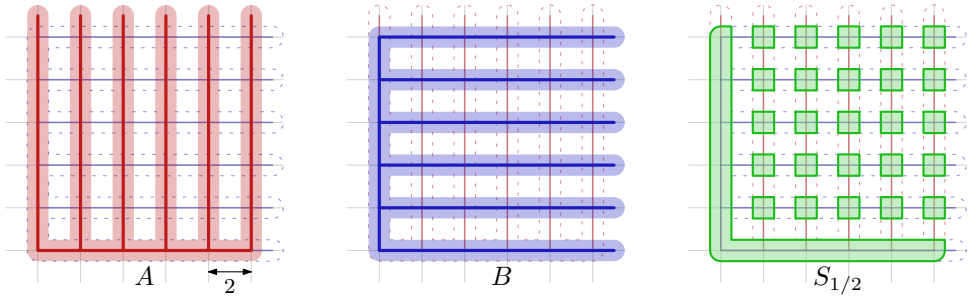

Figure 4 Sets A and B for which $S_{1/2}$ is disconnected. The shaded areas around A and B represent $A \oplus D_{1/2}$ and $B \oplus D_{1/2}$, respectively.

Proof. We first show that $d_H(A, S_\alpha) \leq \alpha$. The proof for $d_H(B, S_\alpha) \leq 1 - \alpha$ is analogous and therefore omitted. We will infer $d_H(A, S_\alpha) \leq \alpha$ from $d_{\vec{H}}(A, S_\alpha) \leq \alpha$ and $d_{\vec{H}}(S_\alpha, A) \leq \alpha$; thereafter we will show equality.

Consider any point $a \in A$; by our assumption that $d_H(A, B) = 1$, there is a point $b \in B$ with $d(a, b) \leq 1$. Now consider a point $s \in \text{seg}(a, b)$ with $d(a, s) \leq \alpha$ and $d(b, s) \leq 1 - \alpha$; clearly this point must be in S_α, as it is contained in both $A \oplus D_\alpha$ and $B \oplus D_{1-\alpha}$, and it has $d(a, s) \leq \alpha$. As this works for every $a \in A$, it holds that $d_{\vec{H}}(A, S_\alpha) \leq \alpha$. The fact that $d_{\vec{H}}(S_\alpha, A) \leq \alpha$ follows straightforwardly from S_α being a subset of $A \oplus D_\alpha$. Thus, $d_H(A, S_\alpha) \leq \alpha$.

To show equality, assume that the Hausdorff distance is realized by a point $\hat{a} \in A$ with closest point $\hat{b} \in B$, at distance 1. Consider the point $\hat{s} \in \text{seg}(\hat{a}, \hat{b})$ with $d(\hat{a}, \hat{s}) = \alpha$ and $d(\hat{b}, \hat{s}) = 1 - \alpha$. As observed, $\hat{s} \in S_\alpha$. Since $\hat{s}$ is the closest point of S_α to $\hat{a}$, and $\hat{b}$ is the closest point of B to $\hat{s}$, equality follows. ◀

▶ **Lemma 2.** *S_α is the maximal set that satisfies $d_H(A, S_\alpha) = \alpha$ and $d_H(B, S_\alpha) = 1 - \alpha$.*

Proof. Consider any set T for which we have $d_{\vec{H}}(T, A) \leq \alpha$ and $d_{\vec{H}}(T, B) \leq 1 - \alpha$. As $A \oplus D_\alpha$ contains all points with distance at most α to A, we have that $T \subseteq A \oplus D_\alpha$; similarly, we have that $T \subseteq B \oplus D_{1-\alpha}$. By the definition of S_α, this implies that $T \subseteq S_\alpha$. As this holds for any T, we conclude that S_α is maximal. ◀

2.1 Properties of S_α

In this section, we study the convexity and connectedness of S_α. Recall that a set $A \subseteq \mathbb{R}^2$ is convex if for any two points $a, b \in A$, the segment $\text{seg}(a, b)$ between them is completely contained in A. Also, recall that a set $A \subset \mathbb{R}^2$ is connected if for any two points $a, b \in A$, there exists a continuous curve $c : [0, 1] \to A$ such that $c(0) = a$ and $c(1) = b$. This type of connectedness is known as path-connectedness, but we use the term connected for simplicity. We observe the following properties:

1. If A and B are convex, S_α is convex;
2. If A is convex and B is connected, S_α is connected;
3. For some connected sets A and B, S_α is disconnected.

Property 1 is straightforward: the Minkowski sum of A and B with a disk is convex, and the intersection of convex objects is itself also convex. The example in Figure 4 demonstrates Property 3; in fact, *any* Hausdorff middle will be disconnected for those input sets.

The next lemma establishes Property 2:

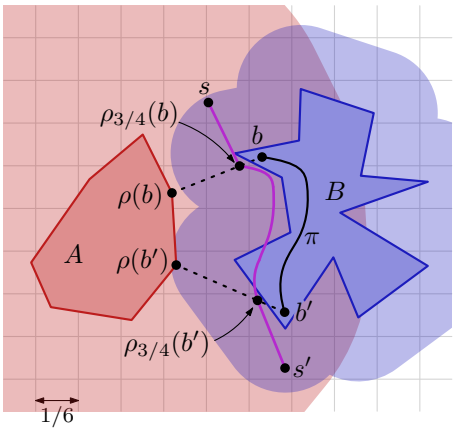

Figure 5 Illustration of the proof showing that S_α is connected if A is convex (sketched for $\alpha = 3/4$). The shaded areas around A and B represent $A \oplus D_{3/4}$ and $B \oplus D_{1/4}$, respectively, so that the doubly-shaded area is $S_{3/4}$.

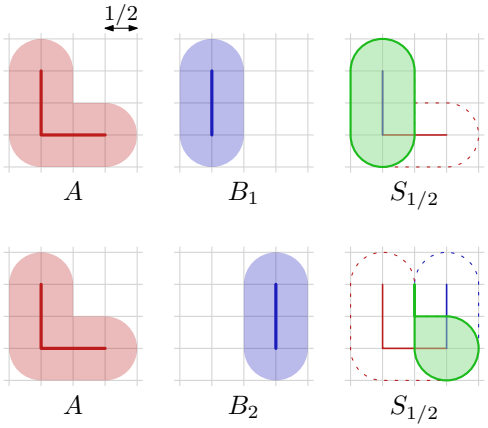

Figure 6 Although B_2 is a translate of B_1, the middle set between A and B_2 is not a translate of the middle set between A and B_1.

▶ **Lemma 3.** *Let A and B be two connected regions of the plane with Hausdorff distance 1, and A convex. Then $S_\alpha = (A \oplus D_\alpha) \cap (B \oplus D_{1-\alpha})$ is connected for $\alpha \in [0, 1]$.*

Proof. See Figure 5 for an illustration. Because A is convex, there is a continuous map $\rho\colon B \to A$ that maps each point of B to a closest point (within distance 1) in A. For $b \in B$, let $\rho_\alpha(b) = \alpha\rho(b) + (1-\alpha)b$. We have that $\rho_\alpha\colon B \to S_\alpha$ is also continuous.

Now take any two points s and s' in S_α; respectively, they have points b and $b' \in B$ within distance $1-\alpha$. The segments between s and $\rho_\alpha(b)$ and between s' and $\rho_\alpha(b')$ lie completely in S_α. Take a continuous curve π from b to b' inside B. The image of π under ρ_α connects $\rho_\alpha(b)$ to $\rho_\alpha(b')$ within S_α, so s and s' are connected inside S_α. ◀

We note that S_α may contain holes. Furthermore, S_α is not shape invariant when B is translated with respect to A. For example, let A be the union of the left and bottom sides of a unit square and let B_1 and B_2 be the left and right sides of that same unit square. Then $(A \oplus D_{1/2}) \cap (B_1 \oplus D_{1/2})$ is not a translate of $(A \oplus D_{1/2}) \cap (B_2 \oplus D_{1/2})$. See Figure 6.

2.2 Complexity of S_α

In this section, we describe the complexity of S_α in terms of the number of vertices, line segments, and circular arcs on its boundary, for several types of polygonal input sets. Recall that ∂A denotes the boundary of set A.

▶ **Lemma 4.** *Let A be a convex polygon with n vertices and B a simple polygon with m vertices. Then S_α consists of $O(n+m)$ vertices, line segments and circular arcs, and this bound is tight in the worst case.*

Proof. For brevity we let $A^\oplus = A \oplus D_\alpha$ and $B^\oplus = B \oplus D_{1-\alpha}$.

There is a trivial worst-case lower bound of $\Omega(n+m)$ by taking $\alpha = 0$ or $\alpha = 1$. Note that if the boundaries of $A^\oplus$ and $B^\oplus$ would consist of only line segments, the upper bound is easy to show: $A^\oplus$ is convex, and its boundary can therefore intersect each segment of $\partial B^\oplus$ at most twice, making ∂S_α consist of (parts of) segments from $\partial A^\oplus$ and $\partial B^\oplus$ and at most $O(m)$ intersection points. The problem is that $\partial A^\oplus$ and $\partial B^\oplus$ also contain circular arcs, in which case an arc of $\partial B^\oplus$ may intersect $\partial A^\oplus$ many times.

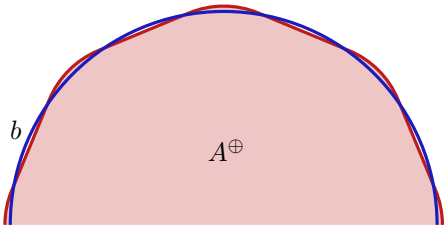

Figure 7 An arc b of $\partial B^{\oplus}$ (blue) may intersect $\partial A^{\oplus}$ (red) many times when $\alpha < 1 - \alpha$.

To show an upper bound of $O(n + m)$, we distinguish two cases. In the first case, we assume $\alpha \geq 1 - \alpha$. Note that in this case, the circular arcs that are part of the boundary of $A^{\oplus}$ have a radius larger or equal to those of $B^{\oplus}$. Additionally, $\partial A^{\oplus}$ is smooth and an alternating sequence of circular arcs and segments, as A is convex. In this case, we do in fact have that any line segment or circular arc b of $\partial B^{\oplus}$ can intersect $\partial A^{\oplus}$ at most twice. Consider two intersection points of b with $\partial A^{\oplus}$: as the curvature of $\partial A^{\oplus}$ is at most that of b, there can never be another intersection point between these two.

For the second case, we assume $\alpha < 1 - \alpha$. Again, take an arbitrary arc b of $\partial B^{\oplus}$ that intersects some arc a of $\partial A^{\oplus}$. We distinguish two cases: the center point of the disk whose boundary contains a is inside $B^{\oplus}$, or it is outside. If it is outside, b can only intersect $\partial A^{\oplus}$ in two points. If it is inside, $\partial A^{\oplus}$ may intersect b many times; see Figure 7. We charge these intersections to the arcs of $\partial A^{\oplus}$. We argue that each arc a of $\partial A^{\oplus}$ is charged at most four times: Consider any α-disk D_α and any $(1 - \alpha)$-disk $D_{1-\alpha}$ containing the center of D_α, the latter will cover at least $1/3$ of the perimeter of the former. Hence, the boundary of the union of any number of such $(1 - \alpha)$-disks intersects D_α at most four times. The circular arcs of $\partial A^{\oplus}$ cannot be charged more often because they are less than a full circle. ◀

▶ **Lemma 5.** *Let A and B be two simply connected polygons of n and m vertices, respectively. Then S_α consists of $O(nm)$ vertices, line segments and circular arcs, and this bound is tight in the worst case.*

Proof. The worst-case lower bound of $\Omega(nm)$ follows by taking A and B to be two rotated "combs"; see Figure 4. For $\alpha = 1/2$, S_α consists of $\Omega(nm)$ distinct components. The upper bound follows directly from the fact that $A \oplus D_\alpha$ and $B \oplus D_\alpha$ have complexities $O(n)$ and $O(m)$, respectively. ◀

In fact, *any* Hausdorff middle has complexity $\Theta(nm)$ for the example in Figure 4. S_α is maximal, and other middles must have at least some part of every component of S_α.

2.3 S_α as a morph

By increasing α from 0 to 1, S_α morphs from $A = S_0$ into $B = S_1$. (Examples of such morphs are presented in Figures 1 and 8.) The following lemma shows that this morph has a bounded rate of change.

▶ **Lemma 6.** *Let S_α and S_β be two intermediate shapes of A and B with $\alpha \leq \beta$. Then $d_H(S_\alpha, S_\beta) = \beta - \alpha$.*

Proof. We have $d_H(S_\alpha, S_\beta) \geq \beta - \alpha$ because, by the triangle inequality, $d_H(A, B) = 1 \leq d_H(A, S_\alpha) + d_H(S_\alpha, S_\beta) + d_H(S_\beta, B) \leq \alpha + d_H(S_\alpha, S_\beta) + 1 - \beta$.

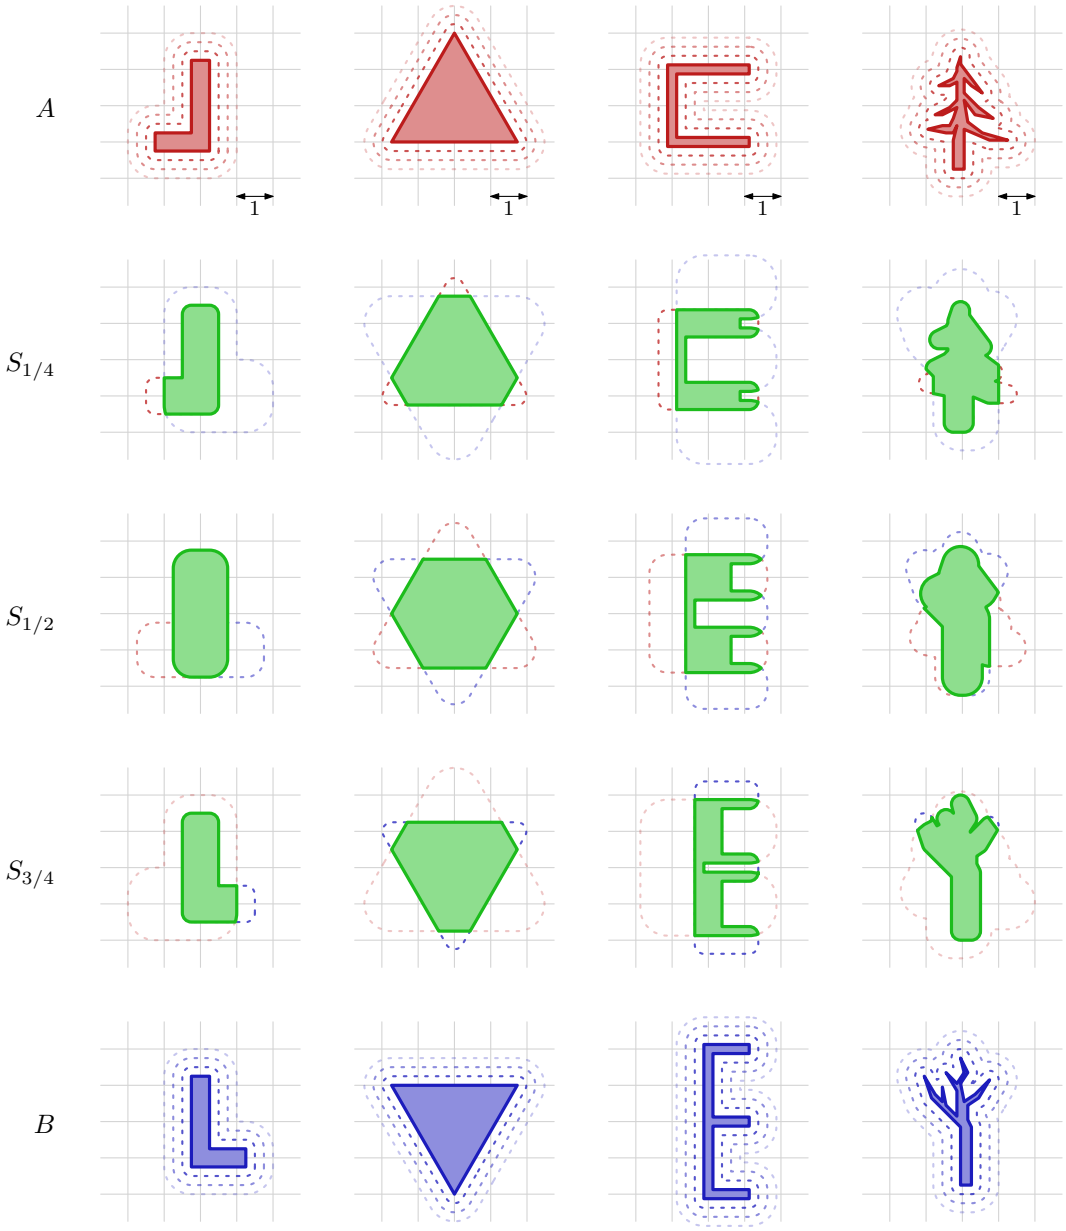

Figure 8 Example morphs.

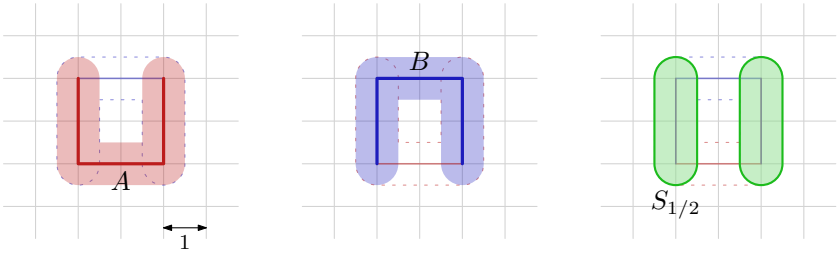

Figure 9 The left and middle figure show the offsets of A, respectively B with distance $1/2$. The right figure shows the resulting $S_{1/2}$ in green. Any connected Hausdorff middle must cross the vertical middle line or stay on one side of it. In both cases, the Hausdorff distance doubles.

It remains to show that $d_H(S_\alpha, S_\beta) \leq \beta - \alpha$. We show that $S_\beta \subseteq S_\alpha \oplus D_{\beta-\alpha}$; the proof that $S_\alpha \subseteq S_\beta + D_{\beta-\alpha}$ is analogous. Let p be some point in S_β. Then, by definition of S_β, there exist some points $a \in A$ and $b \in B$ such that $d(a,p) \leq \beta$ and $d(b,p) \leq 1-\beta$. Let $\bar{p}$ be the point obtained by moving p in the direction of a by $\beta - \alpha$. By the triangle inequality, we then have that $d(a,\bar{p}) \leq \beta - (\beta - \alpha) = \alpha$ and $d(b,\bar{p}) \leq (1-\beta) + (\beta - \alpha) = 1 - \alpha$. This implies that $\bar{p} \in S_\alpha$. As p was an arbitrary point in S_β, and $d(p,\bar{p}) \leq \beta - \alpha$, we have that $S_\beta \subseteq S_\alpha \oplus D_{\beta-\alpha}$. So $d_H(S_\alpha, S_\beta) \leq \beta - \alpha$. ◀

The lemma implies that, even though the number of connected components of S_α can change when α changes, new components arise by splitting and never "out of nothing", and the number of components can only decrease through merging and not by disappearance.

The morph from A to B has a consistent submorph property, formalized below.

▶ **Observation 7.** *If a morph from $A = S_0$ to $B = S_1$ contains a shape C, then the morph from A to C concatenated with the morph from C to B is the same as the morph from A to B: they contain the same collection of shapes in between and in the same order.*

As a corollary of this observation, $\{\alpha \in [0,1] \mid S_\alpha \text{ is convex}\}$ is a connected interval.

2.4 The cost of connectedness

For some applications, it might be necessary to insist that S_α is always connected. However, in the worst case, the cost of connecting all components of S_α can be that its Hausdorff distance to A and B becomes 1. See Figure 9 for an example where this is the case. In fact, any connected Hausdorff middle has distance 1 for this example.

3 The Hausdorff middle of more than two sets

A natural question is whether the results from the previous section extend to more than two input shapes. There are several ways to formalise the notion of a Hausdorff middle between multiple shapes. Analogous to the case of two sets, we are interested in a middle shape that minimizes the maximum Hausdorff distance to each input set. Let $\mathcal{M} = \{A_1, \ldots, A_m\}$ be a collection of m input shapes with largest pairwise Hausdorff distance 1. We define T_α as $\bigcap_i (A_i \oplus D_\alpha)$; the (maximum) middle set is then given by the smallest value α for which $T_\alpha \oplus D_\alpha$ contains all input sets. We denote this smallest α by $\alpha(\mathcal{M}) := \min\{\, \alpha \mid \max_i d_H(A_i, T_\alpha) \leq \alpha \,\}$. If α is clear from the context, we use the notation $A^\oplus$ to mean $A \oplus D_\alpha$.

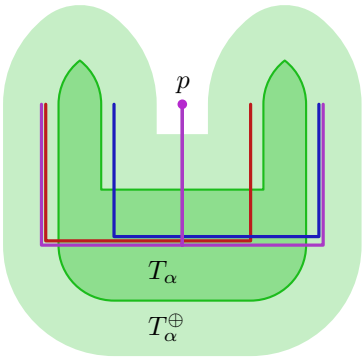

Figure 10 The pairwise Hausdorff distance in this construction is 1, and for any $\alpha < 1$, $T_\alpha^\oplus$ does not contain point p.

In this section, we first study the largest possible $\alpha(\mathcal{M})$ for general and convex input. We then study some general properties of T_α with respect to connectivity and convexity. After this, we consider whether there is some subset of $\mathcal{M}$ that requires the same value of α, and obtain a Helly-type property for convex input. Lastly, we will give various algorithms to compute or approximate $\alpha(\mathcal{M})$ efficiently.

3.1 The largest $\alpha(\mathcal{M})$

In this section, we are interested in the largest possible value of $\alpha(\mathcal{M})$. We first discuss the general case and then study the case where all sets $A \in \mathcal{M}$ are convex. In both cases, we provide an exact answer. This section relies on some tedious calculations, which turn out to be easier if we do not normalize pairwise distances of our objects to 1.

As it turns out, for some inputs it may be the case that $\alpha(\mathcal{M}) = 1$; see Figure 10. Here, there can be no shape with Hausdorff distance < 1 to all the input shapes, meaning any of the three input shapes can be chosen as "the middle". Hence, for two sets, we always have $\alpha(\mathcal{M}) = 1/2$, but for more sets, it depends on the input, and $\alpha(\mathcal{M})$ will be in $[1/2, 1]$. The example in Figure 10 requires non-convex sets, raising the question of what the range of $\alpha(\mathcal{M})$ can be when all A_i are convex.

If we have three convex sets that are points, and they form the corners of an equilateral unit-side triangle, then we can easily see that $\alpha(\mathcal{M}) = 1/\sqrt{3} \approx 0.577$ and the middle shape is exactly the point in the middle of the triangle.

An example with three line segments shown in Figure 11 surprisingly achieves (for $\lambda \approx 0.253135$, $\theta \approx 123.37°$) a larger value $\alpha^* \approx 0.6068 = r$, which we call the *magic value*. Lemma 9 shows that no three convex sets achieve $\alpha(\mathcal{M}) > \alpha^*$. Thus the magic value is the best possible upper bound for three convex sets.

We define the magic value as $\alpha^* := 1/z \approx 0.6068$, where $z := \min\{\lambda + 1 - \cos(2\theta) \mid \lambda \geq 0, \theta \in (90°, 180°), \text{and } \lambda + 1 - \cos(2\theta) = \|(-\lambda\cot(2\theta) - \sin(2\theta) + \sin(\theta), \lambda - \cos(2\theta) + \cos(\theta))\|\} \approx 1.647986325231$ (at $\lambda \approx 0.253135$, $\theta \approx 123.37°$, verified using Wolfram Cloud).

▶ **Lemma 8.** *Let $\mathcal{M} = \{A_1, \ldots, A_m\}$ be a collection of convex regions in the plane, and $\alpha := \alpha(\mathcal{M})$. There is some $A_i \in \mathcal{M}$ with $d_{\vec{H}}(A_i, T_\alpha) = \alpha$.*

Proof. By construction, we have $d_{\vec{H}}(T_\beta, A_i) \leq \beta$ for all i and all β. (Recall that this is equivalent to $T_\beta \subseteq A_i \oplus D_\beta$.) Moreover, if T_β is nonempty, then for any i, the map $\gamma \mapsto d_{\vec{H}}(T_\gamma, A_i)$ is continuous on the domain $[\beta, \infty)$, as T_γ changes continuously. We show

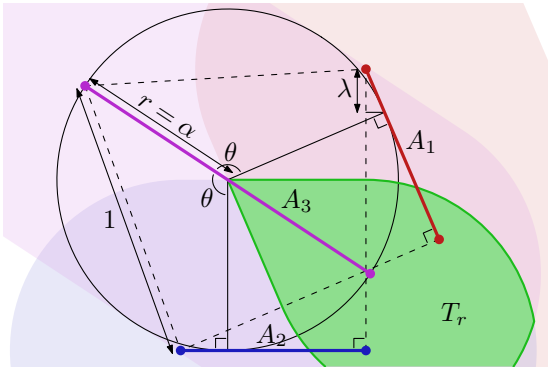

Figure 11 Three segments A_1, A_2, and A_3. Of these, A_3 is the diameter of a circle with radius r; the other two (A_1 and A_2) are tangent to the circle and are copies of one another reflected through A_3, such that all pairwise Hausdorff distances are at most 1 (length of dashed segments). The top left vertex of A_3 is furthest (at distance r) from the middle set T_r (green), so $\alpha(\{A_1, A_2, A_3\})$ is the radius r of the circle.

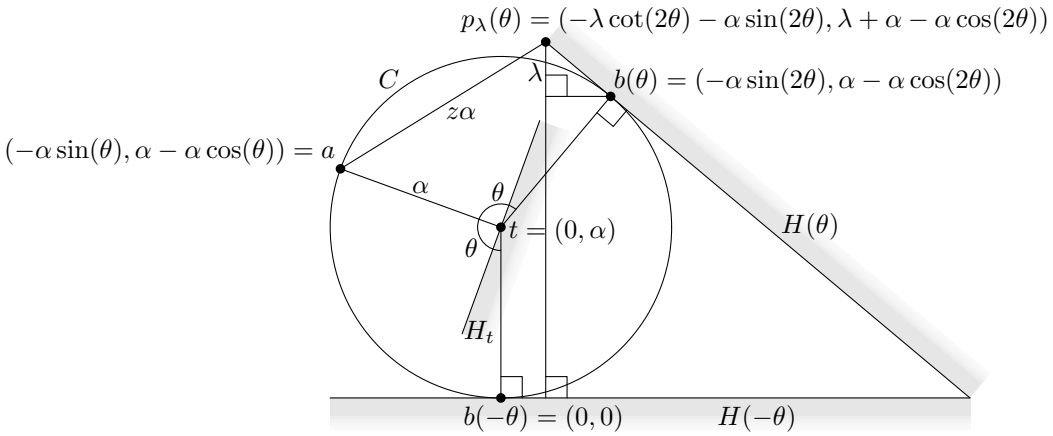

Figure 12 Derivation of the expression for z.

that for some i, we have $d_{\vec{H}}(A_i, T_\alpha) = \alpha$. If instead $d_{\vec{H}}(A_i, T_\alpha) < \alpha$ for all i, then unless T_β is empty for all $\beta < \alpha$, we can decrease α, contradicting minimality of α. If instead α is the minimal value for which T_α is nonempty, then T_α has no interior (when viewed as a subset of the plane). Because T_α is the intersection of convex sets, it is convex. If it has no interior, it is either a segment or a point, and by convexity it must lie on the boundary of $A_i^{\oplus}$ for some i, contradicting that $d_{\vec{H}}(A_i, T_\alpha) < \alpha$. ◀

▶ **Lemma 9.** *Let $\mathcal{M} = \{A_1, A_2, A_3\}$ be convex regions in the plane. Let $\alpha := \alpha(\mathcal{M})$ and $d = \max_{i,j} d_H(A_i, A_j)$, then $d \geq \alpha/\alpha^*$ (equivalently $d \geq z\alpha$).*

Proof. We refer to the full paper for full details, and only provide a proof sketch here.

We argue that in order to obtain $d \leq z\alpha$, the regions A_1, A_2, and A_3 may without loss of generality be assumed to have a particular shape. We have $d_{\vec{H}}(A_i, T_\alpha) \leq \alpha$ for all i, and by Lemma 8, also that $d_{\vec{H}}(A_i, T_\alpha) = \alpha$ for some i. So without loss assume that $d_{\vec{H}}(A_3, T_\alpha) = \alpha$, due to a point $a \in A_3$ with closest point $t \in T_\alpha$. We argue that t lies on the boundary of T_α, and since T_α does not intersect the open disk of radius α centered at a, T_α is either (a) tangent to this disk at point t, or (b) has a corner of $A_1^{\oplus} \cap A_2^{\oplus}$ there. We can exclude

case (a) because by convexity, the directed Hausdorff distance from a to the region A_1 or A_2 that gives rise to the tangent would be at least 2α, and hence $d \geq 2\alpha > z\alpha$. Because we are looking for shapes that satisfy $d \leq z\alpha$, case (b) remains, so t lies on the boundaries of $A_1^{\oplus}$ and $A_2^{\oplus}$ simultaneously. Let C be the circle of radius α centered at t. Neither A_1 nor A_2 lie inside C, but both A_1 and A_2 touch C, say at points b_1 and b_2, respectively. By convexity, the tangents to C at b_1 and b_2 define half-planes in which A_1 and A_2 must lie. Note that a also lies on C, and without loss of generality, b_1 lies (between 90° and 180°) clockwise and b_2 lies (between 90° and 180°) counter-clockwise of a on C. Since the half-planes containing A_1 and A_2 must lie sufficiently close to a for the Hausdorff distance from A_3 to be small, the actual angles are actually quite a bit less than 180° from a. On the other hand, the angle between the two half-planes cannot be too large because the point of A_1 closest to a also needs to be sufficiently close to A_2. This trade-off leads to the claimed value of z, derived from Figure 12 as shown in the full version of the paper. ◀

3.2 Properties of T_α

In this subsection, we use $\alpha := \alpha(\mathcal{M})$ for simplicity. Similar to Section 2.1, we examine the properties of T_α for different types of input. We arrive at straightforward generalizations of the results obtained for two sets:

1. If all A_i are convex, then T_α is convex.
2. If one of the A_i is connected and the rest are convex, then T_α is connected.
3. For some input where each A_i is connected, and at least two are not convex, T_α is disconnected.

Property 1 follows from the same argument as before: T_α is the intersection of convex sets, and therefore itself convex. Property 3 can be shown by extending the construction from Figure 4 with some other sets: if the intersection of two of the sets is not connected, adding more sets will not make T_α connected as long as the pairwise Hausdorff distance does not increase. We establish Property 2 with the following lemma:

▶ **Lemma 10.** *Let $\mathcal{M} = \{A_1, \ldots, A_m\}$ be a set of connected regions of the plane, with A_i convex for $i < m$. Then T_α is connected.*

Proof. Consider the set $T'_\alpha = \bigcap_{i=1}^{m-1} A_i^{\oplus}$. This set is convex, as it is the intersection of convex sets. Also note that by definition of T_α, A_m has directed Hausdorff distance at most α to T'_α. Let $A = T'_\alpha$ and $B = A_m$, normalised such that $d_{\vec{H}}(B, A) = 1$. We now apply Lemma 3 to A and B, using zero as the value for α. We obtain the result that $T_\alpha = T'_\alpha \oplus D_0 \cap A_m \oplus D_\alpha$ is connected. Note that the Hausdorff distance from A to B may be bigger than one, but this does not matter for the proof of Lemma 3. ◀

3.3 Helly-type properties

An interesting question is whether there are any sets in the input that could be removed while maintaining the same optimal value of α. To make this precise, we need some definitions. We say a collection $\mathcal{M}$ of m sets is *d-sufficient*, if there is a collection $\mathcal{M}_d \subset \mathcal{M}$ of d sets such that $\alpha(\mathcal{M}) = \alpha(\mathcal{M}_d)$.

▶ **Lemma 11.** *For every m, there is a collection $\mathcal{M}$ of m connected sets in the plane that is not $m-1$ sufficient.*

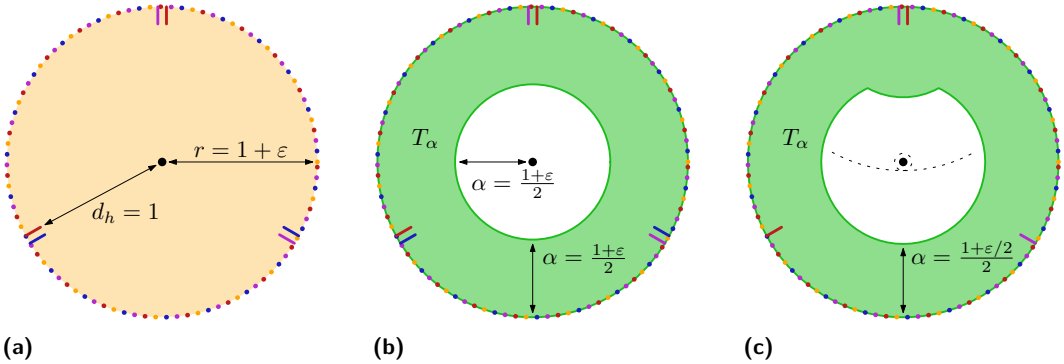

Figure 13 When the input sets are not convex, all sets may be necessary to realise the value of α. The dotted circle is part of all sets (as indicated by the colored dots); each pair of sets has an overlapping protrusion, spaced equally around the circle. (a) shows the radius of the circle and the Hausdorff distance. (b) shows that when all sets are present, the required value of α is $(1+\varepsilon)/2$. (c) shows that with the blue set removed, the required value of α is reduced to $(1+\varepsilon/2)/2$: the dilation of the protruded part of T_α fully contains the part of the disk that would otherwise not be covered, indicated by the dashed and dotted arcs.

Figure 13 depicts a collection of four sets which are not 3-sufficient. The example has one set that is a disk of radius $1+\varepsilon$ (shown in orange on the left), and $m-1$ sets that are circles on the boundary of this disk with $m-1$ protrusions of some small length ε. These protrusions are evenly spaced along the boundary of the disk, and in each location there is a distinct set out of the $m-1$ sets missing (each subset of size $m-2$ is represented by some protrusion). This way, for the case where all sets are present (Figure 13b), the protrusions do not have any influence on T_α, meaning that $\alpha \geq (1+\varepsilon)/2$ is required to let $T_\alpha^\oplus$ contain the entire disk. However, if we remove one set (other than the yellow disk), there will be one protrusion where all sets are now present, meaning it will change the shape of T_α (Figure 13c). Because of this, the center of the disk will already be covered with a smaller value of α, namely $(1+\varepsilon/2)/2$. Note that if we remove the yellow disk, it is sufficient to use a value of $\alpha = \varepsilon/2$. Further note that with a minor adaptation, all sets become polygonal and simply-connected.

We have shown that in general, we cannot remove any sets from the input while maintaining the same value of α. However, when all input sets are convex, we can show that there is always a subset of size at most three that has the same optimal value of α.

▶ **Lemma 12.** *Let $\mathcal{M} = \{A_1, \ldots, A_m\}$ be a collection of convex sets. Then there exists a subcollection $\mathcal{M}' \subseteq \mathcal{M}$ of size at most three such that $\alpha(\mathcal{M}) = \alpha(\mathcal{M}')$.*

Proof. Consider growing some value β from $1/2$ to 1. At some point, $T_\beta^\oplus$ contains all sets in $\mathcal{M}$. There are two ways in which this can happen: (1) T_β is non-empty for the first time, and immediately the condition holds, or (2) T_β grows, and its dilation now covers the last point of all sets in $\mathcal{M}$. As T_β is convex no new components can appear except for the first, and thus we have only those two cases.

In Case 1, T_β is either a segment or a point. If it is a segment, it is generated by two parallel edges of some $A_i, A_j \in \mathcal{M}$ such that we have $\alpha(\{A_i, A_j\}) = \alpha(\mathcal{M})$. If it is a point, it is the common intersection of the dilation of some number of sets from $\mathcal{M}$; we argue that you can always pick three sets for which β is optimal. Let a be the single point in T_β; consider the vectors $\mathcal{V}$ perpendicular to the boundaries of the dilated input sets intersecting in this

point. The vectors $\mathcal{V}$ must positively span the plane[1]: otherwise, all vectors would lie in the same half-plane, and a would not be the first point to appear in T_β. As we are in the plane, there must be subset $\mathcal{U} \subset \mathcal{V}$ of three vectors that positively span the plane by themselves. The three corresponding sets $A_i, A_j, A_k \in \mathcal{M}$ satisfy $\alpha(\{A_i, A_j, A_k\}) = \alpha(\mathcal{M})$.

In Case 2, as our input sets are convex, T_β itself is also convex. Let $a \in A_i$ be the last point of $\mathcal{M}$ to be covered by $T_\beta^\oplus$. As $T_\beta^\oplus$ is convex, a must be on its boundary, and therefore either on an edge or a circular arc of $T_\beta^\oplus$. Each edge can be traced back directly to an edge of some A_j, in which case A_i and A_j have Hausdorff distance 2β, and $\alpha(\{A_i, A_j\}) = \alpha(\mathcal{M})$ for any choice of k. Each circular arc is generated by a vertex of T_β, which in turn is generated by the intersection of the boundaries of some $A_j^\oplus$ and $A_k^\oplus$, in which case we also have that $\alpha(\{A_i, A_j, A_k\}) = \alpha(\mathcal{M})$. ◀

Combining the previous lemma with Lemma 9, we obtain the following result.

▶ **Theorem 13.** *Let $\mathcal{M} = \{A_1, \ldots, A_m\}$ be a collection of convex regions in the plane, and let $T_\alpha = \bigcap_i A_i^\oplus$. Then $\alpha(\mathcal{M})$ is at most the magic value $\alpha^* \approx 0.6068$.*

3.4 Algorithms

For any given collection of shapes $\mathcal{M} = \{A_1, \ldots, A_m\}$, we want to compute $\alpha(\mathcal{M})$. We present two algorithms, a simple approximation algorithm and a more complex exact algorithm. They both use the same decision algorithm as a subroutine. To be precise: given a collection of sets $\mathcal{M}$ and some α, the decision algorithm decides if $\alpha \leq \alpha(\mathcal{M})$. We first present an algorithm for the decision problem. Then we sketch how they are used in the approximation algorithm and the exact algorithm; details are deferred to the appendix. We denote all vertices and edges of the A_i as *features* of $\mathcal{M}$.

Assuming the input has total complexity n, we can test a given value of α as follows. Compute the intersection T_α of the dilations $A_1^\oplus, \ldots, A_m^\oplus$ in $O(n^2 \log n)$ time, using the construction of an arrangement of straight and circular arcs [19]. The set T_α will always have at most quadratic complexity, but it can be disconnected. Next we compute $T_\alpha^\oplus$. We take every connected component T of T_α separately, compute $T^\oplus$, and then compute their union. Since the connected components of T_α are disjoint and can be partitioned into $O(n^2)$ convex pieces, the Minkowski sums of these pieces with D_α form a set of pseudo-disks with summed complexity $O(n^2)$, see [20]. It is known that such a union has $O(n^2)$ complexity and can be computed in $O(n^2 \log^2 n)$ time [1, 20]. Thus, we can compute T_α in $O(n^2 \log^2 n)$ time.

Note that $T_\alpha \subseteq A_i^\oplus$, by definition. It remains to test $A_i \subseteq T_\alpha^\oplus$, for each A_i. We test all those containments by a standard plane sweep [12] in $O(n^2 \log n)$ time. As soon as we find any proper intersection between an arc of $\partial(T_\alpha^\oplus)$ and some edge of some ∂A_i, we can stop the sweep and conclude that α needs to be larger. If there were no proper intersections of this type, there were only $O(n^2)$ events (and not $O(n^3)$), including the ones between edges of different ∂A_i. When there are no proper intersections, each shape A_i lies fully inside or outside $T_\alpha^\oplus$. We can test this in $O(n^2 \log n)$ time (replace each A_i by a single point and then test by a plane sweep or planar point location [12]), and conclude that α must be larger or smaller than the one tested. Thus this decision algorithm takes $O(n^2 \log^2 n)$ total time.

The decision algorithm leads to a simple approximation algorithm to find a value of α that is at most a factor $1 + \varepsilon$ from the optimum. We can perform $\lceil \log 1/\varepsilon \rceil$ steps of binary search in the range $[1/2, 1]$, testing if $T_\alpha^\oplus$ contains all A_i using the above decision algorithm. This takes $O(n^2 \log^2 n \log 1/\varepsilon)$ time in total.

[1] We say $v_i \in \mathbb{R}^2$ *span the plane positively*, if for every point $p \in \mathbb{R}^2$ there are some numbers $a_i \in \mathbb{R}^+$ such that $\sum a_i v_i = p$.

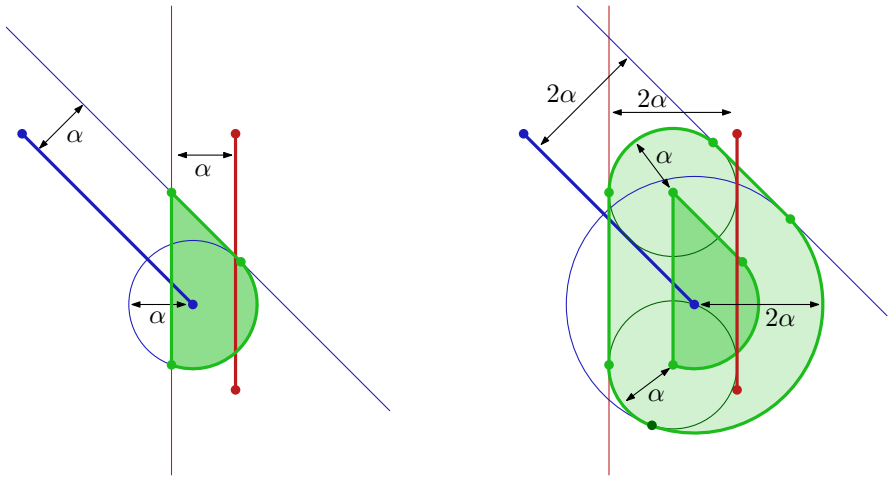

Figure 14 Left, two sets shown by red and blue line segments, and the construction of T_α from lines parallel to edges of $\mathcal{M}$ and circles centered at vertices of $\mathcal{M}$. Right, construction of $T_\alpha^\oplus$ from lines at distance 2α from edges of $\mathcal{M}$, circles of radius 2α centered at vertices of $\mathcal{M}$, and circles of radius α centered at certain vertices of T_α.

We can compute an exact value of $\alpha(\mathcal{M})$ in polynomial time. To this end, we imagine a continuous process where we grow α from $1/2$, and keep track of $T_\alpha^\oplus$. The first time (smallest α) $T_\alpha^\oplus$ covers all A_i, we have found the Hausdorff distance $\alpha(\mathcal{M})$ corresponding to the Hausdorff middle, and we can construct T_α explicitly as the Hausdorff middle. Such an approach is sometimes called *wavefront propagation* or *continuous Dijkstra*; it has been used before to compute Voronoi diagrams [12, 16], straight skeletons [2] and shortest paths on terrains [22]. This approach is combinatorial if there are finitely many events and we can determine each on time, before it occurs. Instead of explicitly maintaining $T_\alpha^\oplus$ when α grows, we will determine a polynomial-size set of critical α values that contains the sought one, and find it by binary search, using the decision algorithm described above.

The value $\alpha(\mathcal{M})$ that we aim to compute occurs when $T_\alpha^\oplus$ has grown just enough to cover all A_i. This can happen in three ways, roughly corresponding to a vertex of A_i becoming covered, an edge of A_i becoming covered at some point "in the middle", or a hole of $T_\alpha^\oplus$ collapsing and disappearing interior to A_i. We call the vertices, edges, and arcs of $\mathcal{M}$ and $T_\alpha^\oplus$ the *features* (of their boundaries). The three ways of covering all A_i, expressed in the features of $\mathcal{M}$ and $T_\alpha^\oplus$, are now: (1) a feature of $T_\alpha^\oplus$ coincides with a vertex of some A_i, (2) a vertex of $T_\alpha^\oplus$ lies on a feature of some A_i, or (3) features of $T_\alpha^\oplus$ collapse and cause a hole of $T_\alpha^\oplus$ to disappear. In the last case, when that hole was inside some A_i, this can be the event where A_i is covered fully for the first time. In all cases, one, two, or three features of $T_\alpha^\oplus$ and zero or one feature of some A_i are involved, and at most three features in total. When three edge or circular arc features pass through a single point for some value of α, we say that these features are *concurrent*. Similarly, when an edge or circular arc passes through a vertex for some α, we say they are concurrent.

It can be that more than three features of $T_\alpha^\oplus$ pass through the point where e.g. a hole in $T_\alpha^\oplus$ disappears, but then we can still determine this critical value by examining just three features of $T_\alpha^\oplus$, and computing the α value when the curves of these three features are concurrent.

Let us analyze which features make up the boundary of $T_\alpha^\oplus$, see Figure 14. There are four types: (1) straight edges, which are at distance 2α from an edge of $\mathcal{M}$, and parallel to it, (2) circular arcs of radius 2α, which are parts of circles centered at vertices of $\mathcal{M}$, (3)

circular arcs of radius α, centered at a vertex of T_α, and (4) vertices where features of types (1)–(3) meet. Every one of the features of the boundary of $T_\alpha^\oplus$ is determined by one or two features of $\mathcal{M}$. In particular, each arc of type (3) is centered on an intersection point which is a vertex of T_α, of which there can be $\Theta(n^2)$ in the worst case (Figure 4). Depending on the type of intersection point, its location may change linearly in α, or according to a low-degree algebraic curve (when the intersection has equal distance α to an edge and a vertex of $\mathcal{M}$).

Since any critical value can be determined as a concurrency of two (vertex and edge or arc) or three features (three edges or arcs) from $\mathcal{M}$ and $T_\alpha^\oplus$, and features of $T_\alpha^\oplus$ in turn are determined by up to two features of $\mathcal{M}$, every critical value depends on at most six features of the input, $\mathcal{M}$. If we choose any tuple with up to six features of $\mathcal{M}$, and compute the α values that may be critical, we obtain a set of $O(n^6)$ values that contain all critical α values, among which $\alpha(\mathcal{M})$. We can compute this set in $O(n^6)$ time, as it requires $O(1)$ time for each tuple of up to six features of $\mathcal{M}$.

▶ **Theorem 14.** *Let $\mathcal{M}$ be a collection of m polygonal shapes in the plane with total complexity n, such that the Hausdorff distance between any pair is at most 1, and let $\varepsilon > 0$ be a constant. The Hausdorff middle can be computed exactly in $O(n^6)$ time, and approximated within ε in $O(n^2 \log^2 n \log 1/\varepsilon)$ time.*

Parametric search could result in a faster exact algorithm, but for this one would need to express whether input features are close to a given S_α in terms of low degree polynomials. This is nontrivial given that S_α as function of α varies in a complex manner.

4 Discussion and future research

We have defined and studied the Hausdorff middle of two planar sets, leading to a new morph between these sets. We also considered the Hausdorff middle for more than two sets. While we assumed that the input sets are simply-connected, our definition of middle and the morph immediately generalize to more general sets, like sets with multiple components and holes. In this sense our definition of middle is very general. Other interpolation methods between shapes do not generalize to more than two input sets and cannot easily handle sets with multiple components.

There are many interesting open questions. For example, when both input sets are one-dimensional curves, is there a natural way to define a Hausdorff middle curve that is also 1-dimensional?

Besides the maximal middle set, there are other options for a Hausdorff middle. For example, we can choose S_α clipped to the convex hull of $A \cup B$, which is also a valid Hausdorff middle. In Figure 9, the green shape would be reduced to the part inside the square, which may be more natural. This Hausdorff middle can also be used in a morph.

Another interesting question could be if, for two shapes A and B, we can find a translation or rigid motion of A such that some measure on the Hausdorff middle (e.g. area, perimeter, diameter) is minimised.

For two or more shapes in the plane, we could also define a middle based on area-of-symmetric-difference. Here we may want to average the areas for the middle shape, and possibly choose the middle that minimizes perimeter. This problem is related to minimum-length area bisection [21].

Similarly, for a set of curves, we could define a middle curve based on the Fréchet distance. This appears related to the Fréchet distance of a set of curves rather than just a pair [15].

References

1 P.K. Agarwal, J. Pach, and M. Sharir. State of the union (of geometric objects). Technical report, American Mathematical Society, 2008.

2 O. Aichholzer, F. Aurenhammer, D. Alberts, and B. Gärtner. A novel type of skeleton for polygons. *Journal of Universal Computer Science*, 1(12):752–761, 1995.

3 H. Alt, B. Behrends, and J. Blömer. Approximate matching of polygonal shapes. *Annals of Mathematics and Artificial Intelligence*, 13(3):251–265, September 1995.

4 H. Alt, P. Braß, M. Godau, C. Knauer, and C. Wenk. Computing the Hausdorff distance of geometric patterns and shapes. In Boris Aronov, Saugata Basu, János Pach, and Micha Sharir, editors, *Discrete and Computational Geometry - The Goodman-Pollack Festschrift*, pages 65–76. Springer, 2003.

5 H. Alt and L.J. Guibas. Discrete geometric shapes: Matching, interpolation, and approximation. In *Handbook of Computational Geometry*, pages 121–153. Elsevier, 2000.

6 Nicolas Aspert, Diego Santa-Cruz, and Touradj Ebrahimi. Mesh: Measuring errors between surfaces using the Hausdorff distance. In *Proc. IEEE International Conference on Multimedia and Expo*, volume 1, pages 705–708, 2002.

7 M.J. Atallah. A linear time algorithm for the Hausdorff distance between convex polygons. *Information Processing Letters*, 17(4):207–209, 1983.

8 G. Barequet, M.T. Goodrich, A. Levi-Steiner, and D. Steiner. Contour interpolation by straight skeletons. *Graphical Models*, 66(4):245–260, 2004.

9 G. Barequet and A. Vaxman. Nonlinear interpolation between slices. *International Journal of Shape Modeling*, 14(01):39–60, 2008.

10 J.-D. Boissonnat. Shape reconstruction from planar cross sections. *Computer Vision, Graphics, and Image Processing*, 44(1):1–29, 1988.

11 Paolo Cignoni, Claudio Rocchini, and Roberto Scopigno. Metro: measuring error on simplified surfaces. *Computer Graphics Forum*, 17(2):167–174, 1998.

12 M. de Berg, O. Cheong, M. van Kreveld, and M. Overmars. *Computational Geometry: Algorithms and Applications*. Springer, 3rd edition, 2008.

13 M.G. Dobbins, L. Kleist, T. Miltzow, and P. Rzążewski. $\forall\exists\mathbb{R}$-completeness and area-universality. In *International Workshop on Graph-Theoretic Concepts in Computer Science*, pages 164–175. Springer, 2018.

14 M.-P. Dubuisson and Anil K. Jain. A modified hausdorff distance for object matching. In *Proceedings of 12th IEEE International Conference on Pattern Recognition*, volume 1, pages 566–568, 1994.

15 A. Dumitrescu and G. Rote. On the Fréchet distance of a set of curves. In *CCCG*, pages 162–165, 2004.

16 S. Fortune. A sweepline algorithm for Voronoi diagrams. *Algorithmica*, 2(1-4):153, 1987.

17 Barequet G. and M. Sharir. Piecewise-linear interpolation between polygonal slices. *Computer Vision and Image Understanding*, 63(2):251–272, 1996.

18 C. Gotsman and V. Surazhsky. Guaranteed intersection-free polygon morphing. *Computers & Graphics*, 25(1):67–75, 2001.

19 D. Halperin and M. Sharir. Arrangements. In C.D. Tóth, J. O'Rourke, and J.E. Goodman, editors, *Handbook of Discrete and Computational Geometry*, chapter 28, pages 723–762. Chapman and Hall/CRC, 3rd edition, 2018.

20 K. Kedem, R. Livne, J. Pach, and M. Sharir. On the union of Jordan regions and collision-free translational motion amidst polygonal obstacles. *Discrete & Computational Geometry*, 1(1):59–71, 1986.

21 E. Koutsoupias, C. H. Papadimitriou, and M. Sideri. On the optimal bisection of a polygon. *ORSA Journal on Computing*, 4(4):435–438, 1992.

22 J.S.B. Mitchell, D.M. Mount, and C.H. Papadimitriou. The discrete geodesic problem. *SIAM Journal on Computing*, 16(4):647–668, 1987.