

Abstract Morphing Using the Hausdorff Distance and Voronoi Diagrams

Lex de Kogel

Utrecht University, The Netherlands

Marc van Kreveld

Utrecht University, The Netherlands

Jordi L. Vermeulen

Utrecht University, The Netherlands

Abstract

This paper introduces two new abstract morphs for two 2-dimensional shapes. The intermediate shapes gradually reduce the Hausdorff distance to the goal shape and increase the Hausdorff distance to the initial shape. The morphs are conceptually simple and apply to shapes with multiple components and/or holes. We prove some basic properties relating to continuity, containment, and area. Then we give an experimental analysis that includes the two new morphs and a recently introduced abstract morph that is also based on the Hausdorff distance [23]. We show results on the area and perimeter development throughout the morph, and also the number of components and holes. A visual comparison shows that one of the new morphs appears most attractive.

2012 ACM Subject Classification Theory of computation → Computational geometry

Keywords and phrases Morphing, Hausdorff distance, Voronoi diagrams

Digital Object Identifier 10.4230/LIPIcs.ESA.2022.74

Related Version *Full Version*: <https://arxiv.org/abs/2206.15339>

Supplementary Material *Interactive Resource*: <https://hausdorff-morphing.github.io/>

Funding Research was funded by NWO TOP grant no. 612.001.651.

1 Introduction

Morphing, also referred to as shape interpolation, is the changing of a given shape into a target shape over time. Applications include animation and medical imaging. Animation is often motivated by the film industry, where morphing can be used to create cartoons or visual effects. In medical imaging, the objective is a 3D reconstruction from cross-sections, such as those from MRI or CT scans. Reconstruction between two 2D slices is essentially 2D interpolation between shapes, which is a form of morphing. We regard morphing itself as the change of one shape into another shape by a parameter, or, more precisely, a function from the interval $[0, 1]$ to shapes in a space, such that the image at 0 is the one input shape and the image at 1 is the other input shape. It is often convenient to see the morphing parameter as time. In the rest of this paper, we will refer to the shape of the morph at any particular time value as an *intermediate shape*. See Figure 1 for an example of two halfway shapes between polygons resembling a butterfly and a spider.

The quality of a morph depends on the application. For medical imaging, the implied 3D reconstruction must be anatomically plausible. For morphing between two drawings of a cartoon character, the shapes in between must keep the dimensions of the limbs, for example. Furthermore, semantically meaningful features (nose, chin) should morph from their position in the one shape to their position in the other shape.



© Lex de Kogel, Marc van Kreveld, and Jordi L. Vermeulen;
licensed under Creative Commons License CC-BY 4.0

30th Annual European Symposium on Algorithms (ESA 2022).

Editors: Shiri Chechik, Gonzalo Navarro, Eva Rotenberg, and Grzegorz Herman; Article No. 74; pp. 74:1–74:16



Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



■ **Figure 1** The intermediate shapes of two different morphing methods at time value $1/2$ when morphing between the input shapes on the left. The middle shows the dilation morph (introduced in [23]), the right shows the Voronoi morph (introduced in this paper).

In this paper we concentrate on abstract morphing of shapes. A morphing task is abstract if there is no (semantic) reason to transform certain parts of a starting shape into certain parts of a goal shape. In a recent paper, van Kreveld et al. [23] presented a new type of abstract morphing based on the Hausdorff distance. It takes any two compact planar shapes A and B as input, and produces a morphed shape that interpolates smoothly between them. For a time value $\alpha \in [0, 1]$, this morph is equal to A at $\alpha = 0$ and to B at $\alpha = 1$. For any value of α it has Hausdorff distance α to A and Hausdorff distance $1 - \alpha$ to B , if the initial Hausdorff distance is 1 (the input can be scaled to make this true without changing intermediate shapes). Morphs with this property are called *Hausdorff morphs* [23]. The Hausdorff morph introduced by van Kreveld et al. is based on Minkowski sums with a disk, and hence we refer to this specific one as the *dilation morph*.

While the dilation morph has nice theoretical properties, in practice it will often grow intermediate shapes from A until $\alpha = 1/2$, at which point the greatly dilated shape will shrink back towards B . For α close to $1/2$, the morphed shape typically resembles neither of the input shapes unless they already looked alike. We can see this in Figure 1.

In this paper we present a new Hausdorff morph called *Voronoi morph* that gives a more visually convincing morph, while maintaining many of the properties of the previous work. Our morph uses Voronoi diagrams to partition each input shape into regions with the same closest point on the other shape, and then scales and moves each such region to that closest point based on the value of α . We show that the Voronoi morph is also a Hausdorff morph. It interpolates smoothly between A and B , but does not have the same problem of significantly increasing the area during the morph. We also present a variant called *mixed morph* that reduces the problem of unnecessarily increasing the perimeter of the interpolated shape. It uses dilation and erosion to overcome some shortcomings of the Voronoi morph.

Related work. The Hausdorff distance is a widely used distance metric that can be used for any two subsets of a space. It is a bottleneck measure: only a maximum distance determines the Hausdorff distance. Efficient algorithms to compute the Hausdorff distance between two simple polygons or their higher-dimensional equivalents exist [3, 4, 7]. The Hausdorff distance is used in computer vision [19] and computer graphics [6, 17] for template matching, and the computation of error between a model and its simplification.

Several algorithmic approaches to morphing have been described. Many of these are motivated by shape interpolation between slices (e.g., [2, 9, 10, 12], an overview can be found in [11]). Other papers discuss morphing explicitly and not as an interpolation problem. Many of these results use compatible triangulations [20, 24], in particular those that avoid self-intersections. It is beyond the scope of this paper to give a complete overview of morphing methods. For (not so recent) surveys of shape matching, interpolation, and correspondence, see [5, 22]. Our paper builds upon the morphing approach given in [23], which introduced Hausdorff morphs as a new technique for abstract morphing, and the dilation morph as a specific example of a Hausdorff morph.

Another shape similarity measure than the Hausdorff distance, the *Fréchet* distance, can also be used to define a morph. In particular, *locally correct Fréchet matchings* [14] immediately imply a smooth transition of one shape outline into another, because they match all pairs of points on the two curves. Similar approaches were given in [25, 26]. During the transition, however, the outline may be self-intersecting. This problem was addressed in [15, 16]. A more important shortcoming of morphing using the Fréchet distance is that it is unclear how to morph between shapes with different numbers of components and holes.

Much of the commercial software for morphing applies to images, with or without additional human control. Other software is meant as toolkits for designers to design their own morphs, most notably Adobe After Effects.

Our results. We introduce two new abstract morphs based on the Hausdorff distance. They are – just like the dilation morph – conceptually simple and easy to implement if one has code for Minkowski sum and difference with a disk, Voronoi diagrams, and polygon intersection and union. We examine basic properties of the two new morphs and compare how they relate to the dilation morph. In particular, we show that the Voronoi morph is a Hausdorff morph and that it is 1-Lipschitz continuous. We also show that for any morphing parameter (time), the Voronoi-morph intermediate shape is a subset of the mixed-morph intermediate shape, which in turn is a subset of the dilation-morph intermediate shape.

We then proceed with an extensive experimental analysis where we compare four basic quantities: area, perimeter, number of components, and number of holes. We show how these quantities develop throughout the three morphs. We also present visual results. As data we use simple drawings of animals, country outlines, and text (letters and whole words).

2 Preliminaries

Given two sets A and B , we can define the directed Hausdorff distance from A to B 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* between A and B is then defined as the maximum of both directed distances:

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

When A and B are closed sets, we can alternatively define the Hausdorff distance using Minkowski sums. Recall that the Minkowski sum $A \oplus B$ is defined as $\{a + b \mid a \in A, b \in B\}$; the directed Hausdorff distance between A and B is then the smallest value r for which $A \subseteq B \oplus D_r$, where D_r is a disk of radius r .

Van Kreveld et al. [23] then define a function that interpolates between two shapes in a Hausdorff sense: For any time parameter $\alpha \in [0, 1]$, they define the dilation morph

$$S_\alpha(A, B) := (A \oplus D_\alpha) \cap (B \oplus D_{1-\alpha}),$$

and prove that this shape has Hausdorff distance α to A and $1 - \alpha$ to B , and that it is the maximal shape with this property. Additionally, they show that this morph is 1-Lipschitz continuous: for two time parameters α and β , $d_H(S_\alpha(A, B), S_\beta(A, B)) \leq |\beta - \alpha|$. Note that we will omit the arguments A, B when they are clear from context.

Structurally, it turns out that the intermediate shapes may have quadratic complexity, even when the input is two simple polygons with n vertices each. For instance, if the input consists of a horizontal comb and an overlapping vertical comb, each with $n/4$ prongs, S_α will consist of $\Omega(n^2)$ components for any $\alpha \in (0, 1)$. In fact, this is not limited to the dilation morph: *any* intermediate shape with Hausdorff distance α to A and $1 - \alpha$ to B will have $\Omega(n^2)$ components [23], so every Hausdorff morph has this feature.

Note that both S_α and the morphing methods described below change when we translate one of the input shapes. That is, if we write $t + B$ to be the translation of B over a vector t , it is not true that $S_\alpha(A, t + B) = \alpha * t + S_\alpha(A, B)$, as one might expect. This is because the positions of the input shapes are important both for the Hausdorff distance and for the shape of S_α . However, we can simply calculate $S_\alpha(A, B)$ with A and B aligned in some way, and then generate $S_\alpha(A, t + B)$ by explicitly translating $S_\alpha(A, B)$ by $\alpha * t$. Sensible alignment methods include aligning the centroids, maximising the overlap of the shapes, and minimising the Hausdorff distance.

3 Voronoi morph

As demonstrated in Figure 1, one of the problems with the dilation morph is that the intermediate shapes tend to lose any resemblance to the input during the morphing process. The main reason for this is that the dilated shapes we are intersecting contain many points that do not influence the Hausdorff distance in any way, because they are not on the shortest path from a point on one shape to the closest point on the other. In other words, much of S_α can be removed without changing the Hausdorff distance to and from the input. That said, there is no obvious “correct” way to determine which parts should be removed to obtain the greatest resemblance to the input.

We propose a morph in which we only take the points of S_α that are on the shortest path between points in one input shape and the closest point on the other. Specifically, we only take the points where the ratio of distances to the one shape and the closest point on the other is $\alpha : 1 - \alpha$. More formally, we define our new morph T_α as follows:


$$T_\alpha(A, B) := \{a + \alpha(c(a, B) - a) \mid a \in A\} \cup \{b + (1 - \alpha)(c(b, A) - b) \mid b \in B\},$$

where $c(a, B)$ denotes the point on B closest to a . In other words, we move each point in A closer to the closest point in B by a fraction α of that distance, and each point in B closer to the closest point in A by a fraction $1 - \alpha$, and take the union of those two shapes. If a point is equidistant to multiple points in the other shape, we include all options. We can prove that this morph has the desired Hausdorff distances to the input.

► **Theorem 1.** *Let A and B be two compact sets in the plane with $d_H(A, B) = 1$. Then for any $0 \leq \alpha \leq 1$, we have $d_H(A, T_\alpha) = \alpha$ and $d_H(B, T_\alpha) = 1 - \alpha$.*

Proof. We first show that $d_H(A, T_\alpha) \leq \alpha$, and then show strict equality. The case for $d_H(B, T_\alpha)$ is analogous and therefore omitted.

By construction, any point $a \in A$ has a point at distance $\leq \alpha$ in T_α , showing that $d_H(A, T_\alpha) \leq \alpha$. Similarly, by construction, for each point $b \in B$ there is a point $t_b \in T_\alpha$ such that $t_b = (1 - \alpha)(c(b, A) - b)$. As $d(b, c(b, A)) \leq 1$, it must be the case that t_b has distance



■ **Figure 2** On the left, A is partitioned by the Voronoi diagram $V(B)$ of B . On the right, each partitioned part of A , shown in green, is scaled towards the closest point on B by a factor α .

at most α to $c(b, A)$. It follows that all points in T_α have distance at most α to a point in A , thereby showing that $d_{\bar{H}}(T_\alpha, A) \leq \alpha$. As we have both $d_{\bar{H}}(A, T_\alpha) \leq \alpha$ and $d_{\bar{H}}(T_\alpha, A) \leq \alpha$, it follows that $d_H(A, T_\alpha) \leq \alpha$.

To show strict equality, assume the Hausdorff distance is realised by some point $\hat{a} \in A$ with closest point $\hat{b} \in B$, i.e., $d(\hat{a}, \hat{b}) = 1$. By construction, there is a point $\hat{t} \in T_\alpha$ at distance α from \hat{a} and at distance $1 - \alpha$ from \hat{b} . As \hat{t} is the closest point to \hat{a} in T_α , we have $d_H(A, T_\alpha) = \alpha$, and as \hat{b} is the closest point to \hat{t} in B , we have $d_H(B, T_\alpha) = 1 - \alpha$. If the Hausdorff distance is realised by a point on B , we use a symmetric argument. ◀

We can additionally show that the Voronoi morph, like the dilation morph, is 1-Lipschitz continuous in a Hausdorff sense:

► **Lemma 2.** *Let $\alpha, \beta \in [0, 1]$. Then $d_H(T_\alpha, T_\beta) \leq |\beta - \alpha|$.*

Proof. Let t_α be any point on T_α . Assume without loss of generality that there is some $a \in A$ such that $t_\alpha = a + \alpha(c(a, B) - a)$ (the case for t_α being included due to a point in B is analogous). Now consider the point $t_\beta = a + \beta(c(a, B) - a)$: t_α and t_β are on the same straight line segment between a and $c(a, B)$, and have distance $|\beta - \alpha| \cdot |c(a, B) - a|$ to each other. As $d_H(A, B) = 1$, we know that $|c(a, B) - a| \leq 1$, and therefore that $|t_\beta - t_\alpha| \leq |\beta - \alpha|$. This holds for any $t_\alpha \in T_\alpha$, and the argument is symmetric for T_β . ◀

Note that this type of continuity implies that components of T_α can only form or disappear by merging with or splitting from another component.

In addition to the Hausdorff distance-related properties, it is also interesting to study the general geometric and topological properties of T_α . We first show that the number of components $\#C(T_\alpha)$ of T_α does not change during the morph, except possibly at $\alpha = 0$ and $\alpha = 1$. We prove this for the case of polygonal input; the proof can likely be generalised, but the formalisation is somewhat tedious and not particularly interesting.

Let $V(A)$ be the Voronoi diagram of the vertices, open edges and the interior components of A . We now define $\text{Par}(A, B)$ to be the input shape A partitioned into regions by $V(B)$. Note that $\text{Par}(A, B)$ is a set of regions of A that each have the closest point of B on the same vertex, edge or face of B . For some region $P \in \text{Par}(A, B)$, let P_α be the region obtained by scaling P towards the site of the Voronoi cell of B it is in by a factor α . If this site is a vertex, we simply scale P uniformly towards it; if it is an edge, we scale it perpendicular to the supporting line of that edge; and if it is a face, it does not scale or move at all; see Figure 2 for an illustration. Now let $\text{Par}_\alpha(A, B) := \{P_\alpha \mid P \in \text{Par}(A, B)\}$. Note that T_α is the union of all elements of $\text{Par}_\alpha(A, B)$ and $\text{Par}_{1-\alpha}(B, A)$.

► **Lemma 3.** *Let $0 < \alpha < \beta < 1$. Then $\#C(T_\alpha) = \#C(T_\beta)$.*

Proof. Assume that $\#C(T_\alpha) \neq \#C(T_\beta)$. We can assume without loss of generality that $\#C(T_\alpha) > \#C(T_\beta)$, as in the other case we can take $T_\alpha(B, A)$ instead of $T_\alpha(A, B)$ and get the same morph, but parametrised in reverse. We can also assume that for fixed α, β is the smallest value such that $\#C(T_\alpha) > \#C(T_\beta)$. In this case, there are two regions P and Q of $\text{Par}(A, B)$ or $\text{Par}(B, A)$ that are disjoint and in different components of T_α , but intersect and are in the same component of T_β . This is because, as a consequence of Lemma 2, components cannot appear or disappear. In the following we assume $P, Q \in \text{Par}(A, B)$; the arguments for when one or both are in $\text{Par}(B, A)$ are identical.

As $P_\beta \cap Q_\beta \neq \emptyset$, there must be some point p in both P_β and Q_β . As both P_β and Q_β are formed by regions moving towards the closest point on the other shape, this point is then on the intersection of two shortest paths between A and B . Let a_1, b_1, a_2 and b_2 be the endpoints of these paths intersecting in p . One of the two segments $\overline{pb_1}, \overline{pb_2}$ will be the shortest; assume without loss of generality that it is $\overline{pb_1}$. In this case the path a_2pb_1 is shorter than a_2pb_2 , and by the triangle inequality b_1 must be closer to a_2 than b_2 .

This contradicts the assumption that b_2 was the closest point to a_2 . We conclude that such shortest paths can never intersect, and therefore $P_\alpha \cap Q_\alpha = \emptyset$ for any $\alpha \in (0, 1)$. As such, components can never merge or split for $\alpha \in (0, 1)$, and as they also cannot appear or disappear by Lemma 2, the statement in the lemma follows. ◀

Note that the number of components can change at $\alpha = 0$ or $\alpha = 1$, as in these limit cases elements of $\text{Par}_\alpha(A, B)$ and $\text{Par}_\alpha(B, A)$ turn into points or line segments. Using the strategy from this proof, it also follows that $\text{Par}_\alpha(A, B)$ and $\text{Par}_{1-\alpha}(B, A)$ are interior-disjoint. An interesting corollary of this observation is that the area $|T_\alpha|$ of T_α is bounded from below by $(1 - \alpha)^2 |A| + \alpha^2 |B|$, which is attained when both shapes are disjoint and all parts are moving to a finite number of points (vertices) on the other shape.

3.1 A variant morph

One problem with the Voronoi morph is that it can introduce many slits into the boundary, thereby greatly increasing the perimeter of the shape. This is because parts of the input that have different closest points on the other shape will tend to move away from each other. We present a variant of the Voronoi morph that tries to reduce these problems. As it uses both the Voronoi morph and the dilation morph, we call this variant the *mixed morph*. The mixed morph $M_{\alpha, \varphi}$ is defined as follows:

$$M_{\alpha, \varphi}(A, B) := ((T_\alpha(A, B) \oplus D_\varphi) \ominus D_\varphi) \cap S_\alpha,$$

where \ominus is the Minkowski difference, defined as $A \ominus B := (A^c \oplus B)^c$, where A^c is the complement of A . Taking a Minkowski sum with a disk is also known as *dilation*, and the Minkowski difference with a disk is known as *erosion*. Performing first a dilation and then an erosion with disks of the same radius is known as *closing*, and can be used to close small gaps and holes in a shape without modifying the rest too much. The closing operator is widely used and studied in the field of image analysis [21].

The resulting morph may no longer be a Hausdorff morph: we may have increased the Hausdorff distance by closing certain gaps or holes. We therefore intersect the closed version of T_α with the dilation morph S_α , so that gaps that are necessary to obtain the appropriate Hausdorff distance are maintained. This results in the mixed morph $M_{\alpha, \varphi}$.

The mixed morph has a new parameter, φ , being the radius of the disk used in the closing. Note that $M_{\alpha,0} = T_\alpha$. We can show that $M_{\alpha,\varphi}$ contains all shapes obtained with the same α but smaller value of φ :

► **Lemma 4.** *Let $\varphi, \psi \in \mathbb{R}^+$ and $\varphi \leq \psi$. Then $M_{\alpha,\varphi} \subseteq M_{\alpha,\psi}$.*

Proof. Let us assume that $M_{\alpha,\varphi} \supset M_{\alpha,\psi}$ instead. Then there is some point p such that $p \in M_{\alpha,\varphi}$, but $p \notin M_{\alpha,\psi}$. There are two reasons why p might not be in $M_{\alpha,\psi}$: either $p \notin T_\alpha \oplus D_\psi$, or $p \in T_\alpha \oplus D_\psi$ but $p \notin (T_\alpha \oplus D_\psi) \ominus D_\psi$.

It can clearly not be the case that $p \in M_{\alpha,\varphi}$ but $p \notin T_\alpha \oplus D_\psi$: $M_{\alpha,\varphi}$ is a subset of $T_\alpha \oplus D_\varphi$, and as $\varphi \leq \psi$, we have that $T_\alpha \oplus D_\varphi \subseteq T_\alpha \oplus D_\psi$.

It must then be the case that $p \in T_\alpha \oplus D_\psi$ but $p \notin (T_\alpha \oplus D_\psi) \ominus D_\psi$. In this case, the distance between p and the boundary ∂T_α^\oplus of $T_\alpha \oplus D_\psi$ must be less than ψ . Let $q \in \partial T_\alpha^\oplus$ be the point on the boundary closest to p . As $p \in T_\alpha \oplus D_\varphi$ and $T_\alpha \oplus D_\varphi \subseteq T_\alpha \oplus D_\psi$, the segment \overline{pq} must intersect the boundary of $T_\alpha \oplus D_\varphi$ in some point q' . We must have that $d(p, q') \geq \varphi$, or p would not be in $M_{\alpha,\varphi}$, and we must have $d(q, q') \geq \psi - \varphi$, as $T_\alpha \oplus D_\psi = (T_\alpha \oplus D_\varphi) \oplus D_{\psi-\varphi}$. But then, by the triangle inequality, $d(p, q) \leq d(p, q') + d(q, q') \geq \psi$, which is a contradiction. Hence, $p \in M_{\alpha,\psi}$. As this holds for all $p \in M_{\alpha,\varphi}$, the statement in the lemma follows. ◀

Note that this means we now have the following hierarchical containment of morphs: $T_\alpha \subseteq M_{\alpha,\varphi} \subseteq M_{\alpha,\psi} \subseteq S_\alpha$, for $\varphi \leq \psi$. As T_α is a Hausdorff morph, and S_α is the maximal Hausdorff morph, this shows that $M_{\alpha,\varphi}$ is a Hausdorff morph as well. However, $M_{\alpha,\varphi}$ is not 1-Lipschitz continuous: components may suddenly merge when their distance falls below 2φ .

3.2 Algorithm

To give an algorithm for computing T_α , we assume A and B are (sets of) polygons, possibly with holes. As T_α is based on moving all points on the one shape to the closest point on the other shape, we can compute the Voronoi diagram of each input shape, and then use these to partition the other shapes. This gives us a partitioning of A into pieces that overlap B , or have the same closest point or edge on B , and vice versa. Pieces of A completely inside B are unchanged, pieces with a vertex as closest element are scaled uniformly towards that vertex by a factor α , and pieces with an edge as closest element are scaled perpendicular to the supporting line of that edge by a factor α . For pieces of B we do the same, except that we scale them with a factor $1 - \alpha$. Figure 2 shows an example of how a shape A is partitioned by the Voronoi diagram $V(B)$ of B , and each piece is scaled towards the closest point on B .

Given this algorithm, we can also straightforwardly compute $M_{\alpha,\varphi}$ by computing T_α and S_α , dilating and eroding T_α by a distance φ , and then intersecting the result with S_α .

Our computations rely solely on Voronoi diagrams, Minkowski sums and differences with disks, intersections and unions of polygons, all of which can be found in standard books or surveys [1, 8, 18] and an intermediate shape can be calculated in $O(n^2 \log n)$ time.

4 Experiments

We compare the dilation, Voronoi and mixed morphs experimentally on three data sets. The first data set is a collection of outlines of animals taken from [13]. The second is a selection of the outlines of European countries obtained from the Thematic Mapping World Borders data set;¹ we use the outlines of Austria, Belgium, Croatia, Czechia, France, Germany, Greece, Ireland, Italy, the Netherlands, Poland, Spain and Sweden. For these two sets we compute

¹ http://www.thematicmapping.org/downloads/world_borders.php

the morphs for all pairs of animals and all pairs of countries in the sets. None of the three morphs is translation-invariant or scale-invariant, so it matters where we place the shapes with respect to each other and what sizes they initially have. We choose to scale the shapes to have the same area and translate them to have a common centroid.

The third data set is a small collection of words and letters manually traced as polygons. We use three pairs of words (wish/luck, kick/stuff, try/it), and the letters f, i and u in a serif and a sans serif font. Observe that our morphs could in theory be used to define an infinite family of fonts by interpolating between the glyphs of each element. For these experiments we do not scale the shapes but use the font size, and we align them manually.

For each experiment, we measure the area, perimeter, number of components and number of holes of the morph for α values starting at zero and increasing in steps of $1/8$. The parameter φ of the mixed morph was universally set to 0.02 based on preliminary experimentation.

It is not necessarily insightful to compare areas and especially perimeters between experiments. To make the results more comparable, we make the assumption that an ideal morph linearly interpolates the area and perimeter between those of the input shapes. For each experiment, we can then give the ratio between the measured area and perimeter and these “ideal” values. For the number of components and holes this is less meaningful, as these are discrete values, so we simply record the numbers directly.

Each morphing method was implemented in C++, using Boost² to calculate intersections and unions of polygons, Voronoi diagrams, and Minkowski sums. Although efficiency is not the focus of this paper, running all our experiments only took a few minutes in total.

5 Results

A summary of our measurements of area and perimeter can be seen in Tables 1 and 2. A summary of the number of components and holes for only the animals data set can be seen in Table 3; we exclude the other data sets because the inputs have different numbers of components. Topological measurements for the animals data set can be viewed in Table 4; the measurements for the other experiments can be viewed in the full version. We note that the Voronoi and mixed morphs sometimes have spurious holes caused by numerical precision issues (e.g., the Voronoi morph should not have an intermediate shape with five holes in our experiment with the letter i). Animations of the different morphs for each experiment can be viewed online.³

In Figure 3 we can see that the average area of the dilation morph quickly grows as α increases, until reaching its peak at $\alpha = 1/2$, to about three times the desired size. For the perimeter we see the opposite trend, with the dilation morph typically having a lower perimeter than desired. This is a consequence of the dilation erasing details in the boundary of the input shapes. We can see in Figure 4 that this happens more quickly in the experiments with country shapes. This is expected, as most of the country shapes have more sharp coastline features and islands that quickly disappear, whereas the animal shapes are generally smoother and only have one component.

Our Voronoi morph on average has an area that is much closer to the desired value, and with much lower variance than the dilation morph. However, we see that on average the perimeter is much higher than the desired value. This is because points on opposite sides of a Voronoi edge move in different directions, causing new boundaries to appear in the interior

² <https://www.boost.org>

³ <https://hausdorff-morphing.github.io>

■ **Table 1** The distributions of areas for each morphing method over all experiments for all nine tested values of α , separated by experiment category.

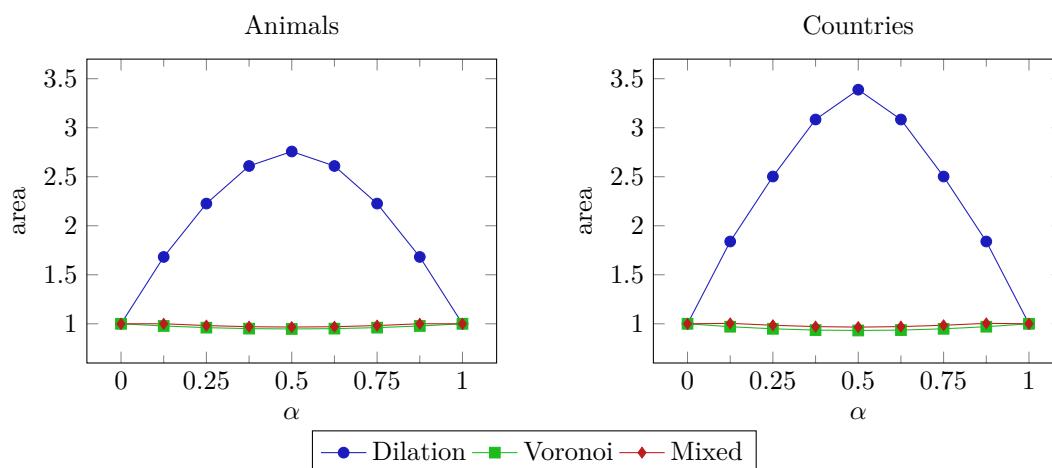
Category	Dilation		Voronoi		Mixed	
	Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.
Animals	1.977	0.763	0.969	0.024	0.986	0.019
Countries	2.249	1.498	0.960	0.039	0.987	0.039
Text	2.118	1.046	0.980	0.035	0.989	0.028

■ **Table 2** The distributions of perimeters for each morphing method over all experiments for all nine tested values of α , separated by experiment category.

Category	Dilation		Voronoi		Mixed	
	Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.
Animals	0.857	0.137	1.725	0.432	1.183	0.155
Countries	0.876	0.237	1.610	0.471	1.129	0.184
Text	0.955	0.142	1.401	0.418	1.155	0.192

■ **Table 3** The distributions of the number of components and holes for each morphing method for all tested values of α except 0 and 1. This only includes the animals data set, as these shapes have only one component and no holes.

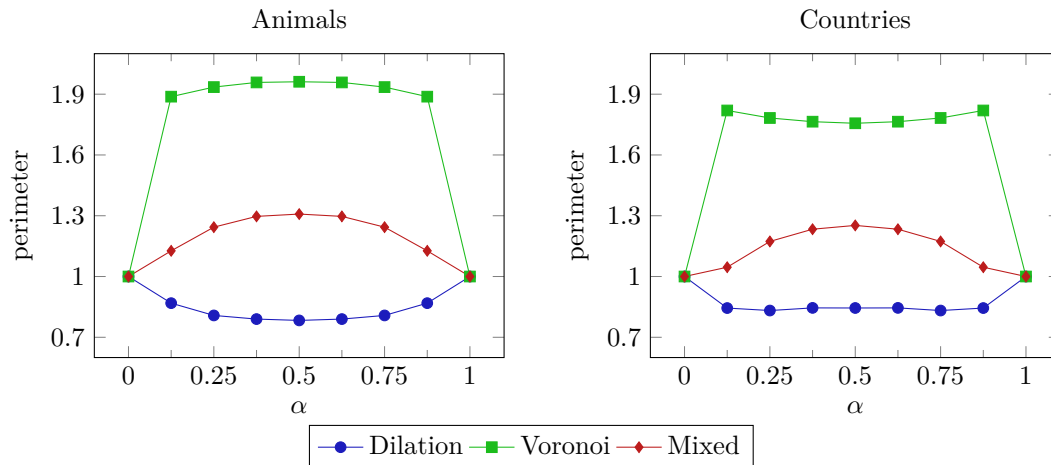
Category	Dilation		Voronoi		Mixed	
	Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.
Components	1.004	0.063	18.556	8.089	5.317	3.213
Holes	0.218	0.602	2.544	2.699	0.218	0.532



■ **Figure 3** The average area over all experiments as a function of α , for both the animals and countries data sets.

■ **Table 4** The minimum and maximum number of components and the maximum number of holes for the experiments with animal shapes, separated by morph type.

Experiment	Dilation			Voronoi			Mixed		
	min	max	holes	min	max	holes	min	max	holes
bird → butterfly	1	1	0	1	11	6	1	4	1
bird → cat	1	1	1	1	12	6	1	5	0
bird → dog	1	1	1	1	23	4	1	7	1
bird → horse	1	2	1	1	27	6	1	8	2
bird → ostrich	1	1	0	1	26	5	1	12	0
bird → shark	1	1	0	1	14	4	1	6	0
bird → spider	1	1	0	1	30	8	1	9	1
bird → turtle	1	1	0	1	21	3	1	9	1
butterfly → cat	1	1	1	1	6	2	1	3	1
butterfly → dog	1	1	1	1	9	4	1	4	1
butterfly → horse	1	1	1	1	28	3	1	9	2
butterfly → ostrich	1	1	1	1	11	7	1	4	2
butterfly → shark	1	1	0	1	8	4	1	3	1
butterfly → spider	1	1	1	1	30	9	1	10	2
butterfly → turtle	1	1	0	1	13	6	1	3	2
cat → dog	1	1	1	1	17	1	1	3	1
cat → horse	1	1	1	1	17	2	1	5	1
cat → ostrich	1	1	0	1	13	4	1	4	0
cat → shark	1	1	0	1	9	0	1	5	0
cat → spider	1	1	2	1	29	3	1	7	3
cat → turtle	1	1	0	1	14	1	1	7	0
dog → horse	1	1	2	1	31	4	1	9	1
dog → ostrich	1	1	1	1	22	2	1	7	1
dog → shark	1	1	1	1	17	3	1	6	1
dog → spider	1	1	3	1	32	2	1	14	2
dog → turtle	1	1	1	1	16	1	1	6	0
horse → ostrich	1	1	2	1	27	7	1	15	1
horse → shark	1	1	1	1	22	5	1	7	1
horse → spider	1	1	4	1	38	3	1	9	1
horse → turtle	1	1	1	1	22	4	1	12	0
ostrich → shark	1	1	0	1	15	8	1	5	0
ostrich → spider	1	1	4	1	38	9	1	17	0
ostrich → turtle	1	1	0	1	21	12	1	4	0
shark → spider	1	1	0	1	23	2	1	5	2
shark → turtle	1	1	0	1	11	3	1	3	0
spider → turtle	1	1	4	1	25	14	1	8	3



■ **Figure 4** The average perimeter over all experiments as a function of α , for both the animals and countries data sets.

of a shape as soon as $\alpha > 0$. We can see this happen in the middle column of Figure 5, and this is reflected in Figure 4, where we see the perimeter sharply increase and then stay mostly the same, before sharply dropping back down.


Our mixed morph achieves its purpose of reducing the perimeter of the Voronoi morph: the measured perimeters are close to the desired values, while the measured areas stay comparable to those of the Voronoi morph. In Figure 4, we see that the perimeter typically still increases during the morphing process, but does not jump up sharply as soon as $\alpha > 0$. This is because the small value of φ lets us close only the narrow gaps that appear around the edges of the Voronoi diagram, but not the gaps that develop as pieces of the shapes move apart significantly. We can see this when comparing the middle and right columns of Figure 5: fewer gaps are closed at $\alpha = 1/2$ than at the other time values.

In addition to area and perimeter, we also tracked the number of components and holes for each morph type. We observe that for the dilation morph, there is an intermediate shape with only one component in all but one of our experiments, showing that this morph tends to turn everything into a blob during the morphing process. On the other hand, the Voronoi morph tends to have an intermediate shape with a number of components much larger than either of the input shapes. The mixed morph exhibits neither of these behaviours. This is illustrated in Figure 5.

Inspecting the morphs visually (see both the full version online for more figures, as well as <https://hausdorff-morphing.github.io> for all experiments), our mixed morph looks quite reasonable, especially when the area of symmetric difference between the input shapes is small. In many cases, the intermediate shape at $\alpha = 1/2$ is a recognisable mix of the two input shapes. This is not the case for the dilation morph, where the Hausdorff distance needs to be very small compared to the size of the input shapes for it to look good. For instance, when one shape has some small islands far away, the dilation morph will grow to have a very large area, whereas with the Voronoi and mixed morphs, the islands just slowly move towards the closest point on the other shape; see the full version online. However, both the Voronoi and mixed morph can still look bad when the area of symmetric difference is large. It may therefore be best to align the input shapes such that the area of symmetric difference is minimised, rather than simply aligning the centroids.



■ **Figure 5** Intermediate shapes for $\alpha \in \{0, 1/4, 1/2, 3/4, 1\}$ when morphing between the outlines of Germany and Italy. The columns show the dilation morph, Voronoi morph and mixed morph from left to right.



■ **Figure 6** Intermediate shapes for $\alpha \in \{0, 1/4, 1/2, 3/4, 1\}$ when morphing between the outlines of the words try and it. The columns show the dilation morph, Voronoi morph and mixed morph from left to right. Note that some artefacts in the Voronoi and mixed morphs, such as in the curved part of the letter r, are caused by having polygonal input instead of smooth curves.

The morphs generally look less convincing on our experiments with text, as the shapes can be very different. For single letters the morphs can look convincing, but when morphing between words, especially of different numbers of letters, the intermediate shape at $\alpha = 1/2$ does not necessarily resemble both input shapes (see Figure 6). However, the intermediate shapes at $\alpha = 1/4$ and $\alpha = 3/4$ still do clearly resemble input shapes A and B , respectively, for the Voronoi and mixed morphs, but less so for the dilation morph. A better approach to morphing text may be to morph on a per-letter basis, rather than treating the whole text as a single shape. Some strategy would then have to be devised that determines which letter will morph to which, and how to deal with different Hausdorff distances between the letter pairs.

Both the Voronoi morph and the mixed morph often have small parts separating, moving, and then merging somewhere else (for example, the beak in the bird-to-ostrich morphs on <https://hausdorff-morphing.github.io>). Such artifacts may be circumvented by choosing a slightly warped Voronoi diagram, but this upsets the simplicity of the current methods. We can sometimes notice in the animations that the mixed morph is indeed not Lipschitz continuous, but since φ is rather small, this does not show clearly.

6 Conclusion

We introduced a new abstract morphing method based on Voronoi diagrams. This new method satisfies the same bounds on the Hausdorff distance as the previously introduced dilation morph, and is also 1-Lipschitz continuous. We have shown experimentally that the intermediate shapes of the Voronoi morph have areas that more closely match those of the input shapes than the dilation morph, but tends to have a perimeter that is larger than desired. To remedy this, we introduced a variant morph, the mixed morph, that we experimentally show to reduce this problem of increasing the perimeter. This mixed morph still satisfies the bounds on the Hausdorff distance, but is no longer 1-Lipschitz continuous. Our experimental analysis is the first we are aware of that analyses the development of area, perimeter, number of components and number of holes throughout the morphs.

An interesting open question is whether we can prevent the increase in perimeter caused by the Voronoi morph without losing 1-Lipschitz continuity. This would require somehow anticipating the moment when two pieces of boundary will meet, and smoothly bridging the gap between them over time, instead of just instantly filling it. To optimise the mixed morph, we can study the effects of choosing different φ , or even changing φ throughout the morph. Another direction is to develop other morphs that guarantee a smooth change of some distance measure other than the Hausdorff distance; we noted that it is unclear how to employ the Fréchet distance for morphing in the presence of multiple components.

A more practically oriented direction for further research would be to develop a less naive method of filling gaps than the mixed morph. It does not necessarily make sense to use the same radius for the closing operator everywhere, which sometimes closes gaps that will be opened again. However, any adaptation of this type will disrupt the conceptual simplicity of the Voronoi and mixed morphs.

References

- 1 P. K. Agarwal, J. Pach, and M. Sharir. State of the union (of geometric objects). In J. Goodman, J. Pach, and R. Pollack, editors, *Surveys on Discrete and Computational Geometry*, volume 453 of *Contemporary Mathematics*, pages 9–48. American Mathematical Society, 2008.

- 2 A. B. Albu, T. Beugeling, and D. Laurendeau. A morphology-based approach for interslice interpolation of anatomical slices from volumetric images. *IEEE Transactions on Biomedical Engineering*, 55(8):2022–2038, 2008.
- 3 H. Alt, B. Behrends, and J. Blömer. Approximate matching of polygonal shapes. *Annals of Mathematics and Artificial Intelligence*, 13(3):251–265, 1995.
- 4 H. Alt, P. Braß, M. Godau, C. Knauer, and C. Wenk. Computing the Hausdorff distance of geometric patterns and shapes. In B. Aronov, S. Basu, J. Pach, and M. 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 N. Aspert, D. Santa-Cruz, and T. Ebrahimi. Mesh: Measuring errors between surfaces using the Hausdorff distance. In *Proceedings of the 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 Franz Aurenhammer, Rolf Klein, and Der-Tsai Lee. *Voronoi diagrams and Delaunay triangulations*. World Scientific Publishing Company, 2013.
- 9 G. Barequet, M. T. Goodrich, A. Levi-Steiner, and D. Steiner. Contour interpolation by straight skeletons. *Graphical Models*, 66(4):245–260, 2004.
- 10 G. Barequet and M. Sharir. Piecewise-linear interpolation between polygonal slices. *Computer Vision and Image Understanding*, 63(2):251–272, 1996.
- 11 G. Barequet and A. Vaxman. Reconstruction of multi-label domains from partial planar cross-sections. *Computer Graphics Forum*, 28(5):1327–1337, 2009.
- 12 J.-D. Boissonnat. Shape reconstruction from planar cross sections. *Computer Vision, Graphics, and Image Processing*, 44(1):1–29, 1988.
- 13 Q. W. Bouts, I. Kostitsyna, M. van Kreveld, W. Meulemans, W. Sonke, and K. Verbeek. Mapping polygons to the grid with small Hausdorff and Fréchet distance. In *Proceedings of the 24th Annual European Symposium on Algorithms*, pages 22:1–22:16, 2016.
- 14 K. Buchin, M. Buchin, W. Meulemans, and B. Speckmann. Locally correct Fréchet matchings. *Computational Geometry*, 76:1–18, 2019.
- 15 K. Buchin, E. W. Chambers, T. Ophelders, and B. Speckmann. Fréchet isotopies to monotone curves. In *Proceedings of the 33rd European Workshop on Computational Geometry*, pages 41–44, 2017.
- 16 E. W. Chambers, D. Letscher, T. Ju, and L. Liu. Isotopic Fréchet distance. In *Proceedings of the 23rd Canadian Conference on Computational Geometry*, 2011.
- 17 P. Cignoni, C. Rocchini, and R. Scopigno. Metro: measuring error on simplified surfaces. *Computer Graphics Forum*, 17(2):167–174, 1998.
- 18 M. de Berg, O. Cheong, M. van Kreveld, and M. Overmars. *Computational Geometry: Algorithms and Applications*. Springer, 3rd edition, 2008.
- 19 M.-P. Dubuisson and A. 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.
- 20 C. Gotsman and V. Surazhsky. Guaranteed intersection-free polygon morphing. *Computers & Graphics*, 25(1):67–75, 2001.
- 21 R. M. Haralick, S. R. Sternberg, and X. Zhuang. Image analysis using mathematical morphology. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-9(4):532–550, 1987.
- 22 O. van Kaick, H. Zhang, G. Hamarneh, and D. Cohen-Or. A survey on shape correspondence. *Computer Graphics Forum*, 30(6):1681–1707, 2011.
- 23 M. van Kreveld, T. Miltzow, T. Ophelders, W. Sonke, and J. L. Vermeulen. Between shapes, using the Hausdorff distance. *Computational Geometry*, 100:101817, 2022.
- 24 Z. Liu, L. Zhou, H. Leung, and H. P.-H. Shum. High-quality compatible triangulations and their application in interactive animation. *Computers & Graphics*, 76:60–72, 2018.

74:16 Abstract Morphing Using the Hausdorff Distance and Voronoi Diagrams

- 25 A. Maheshwari, J.-R. Sack, and C. Scheffer. Approximating the integral Fréchet distance. *Computational Geometry*, 70:13–30, 2018.
- 26 G. Rote. Lexicographic Fréchet matchings. In *Proceedings of the 30th European Workshop on Computational Geometry*, 2014.